



HAL
open science

Online Stochastic Matching: A Polytope Perspective

Céline Comte, Fabien Mathieu, Sushil Mahavir Varma, Ana Bušić

► **To cite this version:**

Céline Comte, Fabien Mathieu, Sushil Mahavir Varma, Ana Bušić. Online Stochastic Matching: A Polytope Perspective. 2024. hal-03502084v5

HAL Id: hal-03502084

<https://hal.science/hal-03502084v5>

Preprint submitted on 13 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License


Online Stochastic Matching: A Polytope Perspective

Céline Comte¹ ✉ 🏠 

Eindhoven University of Technology, Eindhoven, The Netherlands
CNRS and LAAS, Toulouse, France

Fabien Mathieu ✉ 🏠 

LINCS, Paris, France
Swapcard, Paris, France

Sushil Mahavir Varma ✉ 🏠 

Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, US

Ana Bušić ✉ 🏠 

Inria, Paris, France
DI ENS, CNRS, PSL University, Paris, France

Abstract

Stochastic dynamic matching problems have recently gained attention in the stochastic-modeling community due to their diverse applications, such as supply-chain management and kidney exchange programs. In this paper, we study a matching problem where items of different classes arrive according to independent Poisson processes. Unmatched items are stored in a queue, and compatibility between items is represented by a simple graph, where items can be matched if their classes are connected. We analyze matching policies in terms of stability, delay, and long-term matching rate optimization. Our approach relies on the conservation equation, which ensures a balance between arrivals and departures in any stable system. Our main contributions are as follows. We establish a link between the existence of stable policies, the dimensionality of the solution set of the conservation equation, and the compatibility graph's structure. We describe the convex polytope formed by non-negative solutions to the conservation equation, and we design policies that can achieve or closely approximate the vertices of this polytope. Lastly, we discuss potential extensions of our results beyond the main assumptions of this paper.

2012 ACM Subject Classification Mathematics of computing → Queueing theory; Mathematics of computing → Markov processes; Mathematics of computing → Matchings and factors

Keywords and phrases stochastic dynamic matching, graph theory, linear algebra, stability, matching rates, conservation equation

Supplementary Material Stochastic Matching Package

Software (Python, PyPi): Documentation [34]

Acknowledgements The work presented in this paper has been partly carried out at LINCS (<https://www.lincs.fr>).


1 Introduction.

Stochastic dynamic matching problems, in which items arrive at random instants to be matched with other items, have recently attracted much attention in the stochastic-modeling community. These challenging control problems are highly relevant in many applications, including supply-chain management, pairwise kidney exchange programs, and online marketplaces. In pairwise kidney exchange programs for example, each item represents a

¹ Corresponding author.

donor-receiver pair, and two pairs can be matched if the donor of each pair is compatible with the receiver of the other pair. In online marketplaces, items are typically divided into two categories, called demand and supply, and the goal is to maximize a certain long-term performance criteria by appropriately matching demand items with supply items.

In this paper, we consider the following dynamic matching problem². Items of different classes arrive according to independent Poisson processes. Compatibility constraints between items are described by a simple graph on their classes, such that two items can be matched if their classes are neighbors in the graph. Unmatched items are stored in the queue of their class, and the matching policy decides which matches are performed and when. All in all, a stochastic matching model is described by a triplet (G, λ, Φ) , where $G = (V, E)$ is the compatibility graph, $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ is the vector of per-class arrival rates, and Φ is the matching policy. In Figure 1 for instance, there are four item classes numbered from 1 to 4; classes 2 and 3 are compatible with all classes, while classes 1 and 4 are compatible only with classes 2 and 3.



■ **Figure 1** Illustration of a matching model (G, λ, Φ) on the diamond graph.

A matching model (G, λ, Φ) is said to be stable if the associated Markov chain is positive recurrent; roughly speaking, this means that the queue length of each class does not explode as time elapses. Assuming that this matching model is stable, the matching rate μ_k along an edge $k \in E$ with endpoints $i, j \in V$ is the rate at which class- i items and class- j items are matched with one another. Our end goal is to characterize the matching rates that are achievable under stable matching policies and to use this characterization to design matching policies that achieve certain matching rate vectors. We focus specifically on two performance indicators, namely the matching rates along edges and the delay (proportional to the mean queue length of each class by Little's law). Our results show that these metrics are sometimes incompatible, in the sense that optimizing a function of the matching rate vector can lead to instability.

1.1 Motivation.

Besides providing us with criteria to compare the long-term impact of different matching policies, stability, delay, and matching rates are relevant performance criteria in many instances of matching models.

² Detailed definitions of the concepts discussed here will be given in Section 2.

Linear optimization.

First assume a reward r_k is earned each time a match is performed between items of the classes i and j that are endpoints of edge $k \in E$, and that we look for a policy Φ that maximizes the long-run reward rate $\sum_{k \in E} r_k \mu_k$ while making the system stable. This optimization problem is relevant to applications in organ exchange programs, where rewards capture the desirability of a match, accounting for metrics such as the quality of life after transplant and the survival rates of the recipients and donors [7].

Perhaps the most natural policy to consider is a priority policy whereby each incoming item is immediately matched with an unmatched item of the compatible class leading to the highest reward, if any. Unfortunately, it was observed in [33, Section 5] that such a priority policy may lead to instability. Furthermore, using a matching policy with a specific priority order for edges does not guarantee that the ordering of the resulting matching rates is consistent with the edge priorities. Depending on the properties of the compatibility graph G and the arrival rate vector λ , we will introduce either a stable matching policy that yields the optimal reward (if such a stable policy exists) or a family of stable policies that get arbitrarily close to the optimal.

Chained matching.

Matching rates become also crucial when the output of a stable “first-level” matching model (G, λ, Φ) is the (non-Poisson) arrival process of a “second-level” matching model (or, more generally, of another stochastic system). The second-level model consists of a compatibility graph $G' = (E, E')$ whose nodes are the edges of G , and the matching rate vector μ of the first-level model is the arrival rate vector of this second-level model. The ability to control this vector μ becomes instrumental in stabilizing or optimizing performance in the second-level model. Such two-level matching systems are found in various applications such as quantum switches [40, 43].

1.2 Contributions.

We propose a unified approach to studying stability, delay, and matching rates in stochastic matching problems. The starting point is the conservation equation, which any stable policy must satisfy, expressed as $A\mu = \lambda$, where A is the incidence matrix of the compatibility graph G , λ gives the arrival rates, and μ gives the matching rates achieved by the policy. As we will see, what can be achieved in terms of stability, delay, and matching rates is closely related to properties of the linear mapping $\mu \mapsto A\mu$, particularly injectivity and surjectivity. We establish simple equivalences between these linear properties and the structure of the graph (Definitions 2–5 and Proposition 6). With a slight abuse of notation, we say that a graph has a given property if the associated linear map does, e.g., G is said to be bijective if $\mu \mapsto A\mu$ is bijective. Building on these definitions, our main contributions are as follows.

First, we establish a direct relationship between the solution space of the conservation equation, the structure of the compatibility graph, and the existence of a stable policy. Specifically, we show that a compatibility graph G is stabilizable (i.e., there exists a vector λ and a policy Φ such that the matching model (G, λ, Φ) is stable) if and only if the graph G is surjective (Proposition 7). We also prove that a matching problem (G, λ) is stabilizable (i.e., there exists a policy Φ such that the matching model (G, λ, Φ) is stable) if and only if G is surjective and the conservation equation has a solution with positive coordinates (Proposition 8). This result provides a new direct method to verify stabilizability, in polynomial time with respect to the number of classes and edges.

Second, we describe the space Π of solutions to the conservation equation. When G is bijective, Π reduces to a unique solution, for which we provide a closed-form expression (Proposition 9). When G is surjective-only, we give a parametric expression for Π (Proposition 10 and Definition 11) and characterize the polytope $\Pi_{\geq 0}$ of non-negative solutions. We show that the vertices of $\Pi_{\geq 0}$, i.e., its extreme points, correspond to injective solutions in the sense that the subgraph restricted to their support is injective (Proposition 13). Moreover, we show that the vertices are bijective with probability 1 when the vector λ is sampled randomly from a reasonable distribution (Proposition 14).

Third, we apply these results to the linear optimization of matching rates achievable by stable policies. This optimization problem is equivalent to reaching a vertex of $\Pi_{\geq 0}$ (Propositions 15 and 16). A solution is achievable if and only if the vertex is bijective, in which case we propose an optimal policy (Proposition 17). Otherwise, we demonstrate a fundamental trade-off between regret with respect to the optimal solution and the experienced delay (Proposition 18), and we introduce a sequence of policies that we prove to have arbitrarily small regret, at the cost of increasing mean queue sizes (Proposition 19). These results are supported by extensive simulations, where we benchmark our policies against the state-of-the-art.

Finally, we outline possible extensions of our work. We show via a convexity result that optimizing a non-linear function of μ is also possible, in the sense that a target μ^* in the interior of $\Pi_{\geq 0}$ can always be achieved by a stable policy (Propositions 20 and 21). We also explore the optimization of μ using greedy policies (Propositions 23 and 24 and Conjecture 25), and we investigate what happens when G is a hypergraph.

1.3 State of the art.

We now review the relevant work related to (static or dynamic) matching problems.

Non-bipartite or general stochastic matching.

Our work is part of a broader research effort on the stochastic matching model that was briefly discussed earlier and will be described in details in Section 2, see [5, 6, 12, 15, 25, 33, 35]. Among these works, the following are particularly relevant because directly related to our results on stability. The work in [33] is the earliest on this matching model. It derives necessary and sufficient stability conditions that are instrumental in several of our results, in particular Propositions 7 and 8. This work also proves that the Match-the-Longest (ML) policy is maximally stable (in the sense that it always leads to stability whenever the matching problem (G, λ) is stabilizable), a result that is also applied in Proposition 8. [15, 35] focus on the First-Come-First-Match (FCFM) policy. In particular, [35] proves that the FCFM policy is maximally stable, and [15] provides a new sufficient stability condition we prove to be also necessary in Proposition 8.

The recent work in [5] is perhaps the closest to ours, so we provide a detailed discussion to highlight the relation with our paper. The first statement of [5, Theorem 1] is synonym to the equivalence of statements ii and iii in Proposition 8. Our proof is significantly shorter because it relies more heavily on existing results. Our observation at the beginning of Section 4 that the conservation equation has a unique solution if and only if the graph is bijective (and not surjective-only) summarizes [5, Theorem 3]. Some formulas derived in [5, Section 9] are special cases of the formulas derived in Proposition 9. The non-bipartite matching model in [5] is slightly more general because it considers graphs with self-loops, that is, an edge can have identical endpoints, and this paper also considers the bipartite model that will be

discussed in the next paragraph. However, [5] does not adopt the polytope approach that allows us to derive the necessary *and sufficient* conditions of Sections 4–6. Furthermore, although we decided to focus on nonbipartite matching models to alleviate the discussion, the general approach we develop in Section 3 equally applies to bipartite graphs and can be used to derive results for bipartite matching model such as those studied in [5].

Other variants of the model were studied recently, and an interesting future work would consist of generalizing our results to these variants. In particular, [25] considers item abandonment, [6] considers graphs with self-loops, and [23, 36, 28, 22, 41, 38] allow matches not limited to two items by replacing the graph with a hypergraph, which is also called multi-way matching in the literature.

A series of papers on multi-way [28, 22, 41] and two-way matching [29] are closely related to our analysis of matching rates optimization (see Section 6). Under a certain general position condition, these papers propose reward rate maximizing matching policies achieving a regret of $O(1)$ at all times simultaneously. In our setting, the general position condition corresponds to approaching a bijective vertex of $\Pi_{\geq 0}$, and so, Proposition 17 is an equivalent steady-state result. Additionally, we consider the case of injective-only vertices in Section 6.2 and propose a matching policy that ensures stability and arbitrarily small regret in the steady-state at the cost of increasing mean queue length size. Such a policy that ensures stability while optimizing the reward rate seems to be new in the literature, as the results of [28, 22, 41, 29] cannot be directly generalized to the steady-state for the injective-only case. In addition, we explicitly characterize when a vertex is either bijective or injective-only, which also appears to be a new contribution.

Bipartite stochastic matching.

To the best of our knowledge, the first example of a stochastic matching model with an infinite time horizon in the literature, which predated the model that we consider, is the bipartite matching model introduced in [13] and studied in [1, 2, 3, 9, 10, 11, 16]. In this model, the compatibility graph is bipartite, with two parts that correspond to supply and demand items, respectively. This bipartite model differs from ours by its arrival process: time is slotted and, during each time slot, one demand item *and* one supply item arrive. [2, 9] have made contributions about stability and [2] focused on matching rates, and they obtained results similar to those derived in the literature on our model. The bipartite nature of the graph simplifies some calculations, for instance by allowing the application of flow-maximization algorithms to calculate optimal matching rates.

Static and fractional matching.

The static matching problem, in which the nodes of the graph represent items (rather than classes), has been extensively studied in mathematics, computer science, and economy, see [32]. Although the questions raised in static and dynamic matching are often different, the conservation equation that we obtain is reminiscent of several results in static matching. For example, finding a maximum-cardinality matching in the graph G (that is, a maximum-cardinality set of edges without common endpoints) is equivalent to finding integers $\mu_k \in \{0, 1\}$ for each edge $k \in E$ that maximize $\sum_{k \in E} \mu_k$ while satisfying the conservation equation with $\lambda_i = 1$ for each $i \in V$. The relaxation of this integer linear program leads to the so-called fractional matching problem, which has been studied in the literature, see [32, Section 7.2]. Therefore, the fractional matching polytope defined in [32, Section 7.5] is a special case of the convex polytope that we consider in Section 5.2, and our characterization of this convex

polytope is a natural generalization of existing characterizations of the fractional polytope ³.

1.4 Outline.

The remainder of the paper is organized as follows. **Section 2** gives a formal definition of the model. **Section 3** introduces the conservation equation and defines the notions of surjectivity, injectivity, and bijectivity for a graph. We use these definitions to formulate new necessary and sufficient stability conditions. In particular, we show that stability requires the compatibility graph to be surjective (that is, either bijective or surjective-only). In **Section 4**, we focus on bijective graphs and give a closed-form expression of the unique solution to the conservation equation. **Section 5** characterizes the solution set of the conservation equation for surjective-only graphs. **Section 6** focuses on the linear optimization of the matching rates (that is, optimizing a linear combination of the matching rates) under a stable policy. The performance of the proposed policies is validated in **Section 7** through simulations. Lastly, in **Section 8**, we consider several extensions of our framework: non-linear optimization, restriction to greedy policies, and hypergraph matchings.

2 Stochastic dynamic matching.

We consider a stochastic dynamic matching system in which items arrive at random times to be matched with other items. Each incoming item may be matched with any unmatched item of a compatible class; in this case, both items disappear immediately. Unmatched items are gathered in a waiting queue. In this paper, such a stochastic dynamic matching system will be described by a triplet (G, λ, Φ) , where G is the *compatibility graph*, λ the vector of *arrival rates*, and Φ the *matching policy*. We now review each component in details. To facilitate reference and understanding, notation is summarized in Table 1.

2.1 Compatibility graph.

Compatibility constraints between items are described by a graph $G = (V, E)$, called the *compatibility graph* of the model, which is simple (undirected and without self-loop). The number of nodes is represented by n , while m denotes the number of edges.

We use $V = \{v_1, v_2, \dots, v_n\}$ to denote the set of nodes, where each node corresponds to a class in the matching model. In cases where there is no confusion, we may refer to a class v_i simply by its index i . Following the intuition conveyed in Figure 1, we will use the terms “class i ” and “queue i ” interchangeably, and refer for instance to the number of unmatched class- i items as the size of queue i ; this is a convenience of terminology, and this does not preclude the matching policy from using information not captured by this state representation, such as the arrival order of items of different classes (see Section 2.3 for more details).

The set of edges is denoted by $E = \{e_1, e_2, \dots, e_m\}$. These edges represent compatibility constraints between item classes, in the sense that a class- i item and a class- j item can be matched with one another if and only if their classes are adjacent, that is, if there is an edge with endpoints i and j . When there is no ambiguity, we may refer to an edge $e_k \in E$ with endpoints $i, j \in V$ by its index k or its set $\{i, j\}$ of endpoints. In Figure 1 for instance, there are four item classes numbered from 1 to 4. Classes 2 and 3 are compatible with all classes,

³ The fractional matching polytope is actually defined using non-strict inequalities rather than equalities. However, one can verify that these two convex polytopes have the same non-zero vertices.

but classes 1 and 4 are only compatible with classes 2 and 3. The absence of self-loop means that an item of a given class cannot be matched with other items of the same class.

Lastly, we let \mathbb{I} denote the family of independent sets of the compatibility graph G , where an independent set of G is a non-empty set of nodes that are pairwise non-adjacent. The family of independent sets in the compatibility graph of Figure 1 is $\mathbb{I} = \{\{1\}, \{2\}, \{3\}, \{4\}, \{1, 4\}\}$.

2.2 Arrival process.

Class- i items arrive according to an independent Poisson process with rate $\lambda_i > 0$, for each $i \in V$. The vector of arrival rates is denoted by $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n) \in \mathbb{R}_{>0}^n$. Scaling all coordinates of λ by the same positive constant is equivalent to changing the time unit, so we can renormalize λ without changing the dynamics. For example, we will sometimes use the *unit* normalization, in which $\sum_{i \in V} \lambda_i = 1$. We also let $I = (I_t, t \in \mathbb{N})$ denote the sequence of independent and identically distributed (i.i.d.) item classes, so that I_t is the class of the $(t + 1)$ -th item, equal to i with probability $\lambda_i / (\sum_{j \in V} \lambda_j)$, for each $t \in \mathbb{N}$. The couple (G, λ) is called a *matching problem* or simply a *problem*. Occasionally, when we need to specify the sequence of incoming items and not merely its distribution, we will also refer to the couple (G, I) as a (matching) problem.

2.3 Policy and matching dynamics.

Most of the paper focuses on deterministic size-based policies, that is, policies whereby matching decisions are deterministic functions of the queue-size vector. However, as we will briefly explain at the end of this section (and discuss in more details in the supplementary material), our results also apply to a more general definition of a policy. Throughout the paper, we assume that the system is initially empty, meaning that it starts with no unmatched item.

2.3.1 Deterministic size-based policies.

A (deterministic) size-based *matching policy* is defined formally as a function $\Phi : \mathcal{Q} \times V \rightarrow V \cup \{\perp\}$, where \mathcal{Q} is an infinite subset of \mathbb{N}^n that contains the reachable states of the system. In-depth discussion on \mathcal{Q} will come in Sections 2.3.2 and 2.3.3. For each $q \in \mathcal{Q}$ and $i \in V$, an incoming class- i item that finds the system in state q is matched with an item of class $\Phi(q, i)$ if $\Phi(q, i) \in V$ and is added to class- i queue if $\Phi(q, i) = \perp$. The matching policy is assumed to be *adapted* to the compatibility graph G in the sense that

$$\Phi(q, i) \in \{j \in V_i : q_j \geq 1\} \cup \{\perp\}, \quad q \in \mathcal{Q}, \quad i \in V, \quad (1)$$

where V_i is the neighbor set of node i in G . The system dynamics are described by a Markov chain $Q = (Q_t, t \in \mathbb{N})$, called the *queue-size process*. For each $t \in \mathbb{N}$, $Q_t = (Q_{t,1}, Q_{t,2}, \dots, Q_{t,n})$ is an n -dimensional vector giving the number of unmatched items of each class right after the arrival of the t -th item, with the assumption that the system is initially empty, that is, $Q_0 = 0$. The system dynamics satisfy the recursion

$$Q_{t+1} = \begin{cases} Q_t + \mathbb{1}_{I_t} & \text{if } J_t = \perp, \\ Q_t - \mathbb{1}_{J_t} & \text{if } J_t \neq \perp, \end{cases} \quad (2)$$

where $J_t = \Phi(Q_t, I_t)$ for each $t \in \mathbb{N}$, and $\mathbb{1}_i$ is the n -dimensional vector with one in coordinate i and zero elsewhere, for each $i \in V$. We assume that the policy Φ is such that

the Markov chain Q has state space \mathcal{Q} and is irreducible. By unfolding the recursion (2), we obtain that, for each $t \in \mathbb{N}$,

$$Q_{t,i} = L_{t,i} - \sum_{k \in E_i} M_{t,k}, \quad t \in \mathbb{N}, \quad i \in V, \quad (3)$$

where $E_i \subseteq E$ is the set of edges that are incident to node i in the graph G , for each $i \in V$, $L_{t,i}$ is the number of class- i items among the first t arrivals, for each $t \in \mathbb{N}$ and $i \in V$, and $M_{t,k}$ is the number of times that classes i and j are matched over the first t arrivals, for each $t \in \mathbb{N}$ and $\{i, j\} = e_k \in E$:

$$L_{t,i} = \sum_{s=0}^{t-1} \mathbb{1}_{\{I_s=i\}}, \quad t \in \mathbb{N}, \quad i \in V, \quad (4)$$

$$M_{t,k} = \sum_{s=0}^{t-1} \mathbb{1}_{\{\{I_s, J_s\}=e_k\}}, \quad t \in \mathbb{N}, \quad k \in E, \quad (5)$$

with the convention that the sums are zero if $t = 0$. The triplet (G, λ, Φ) is called a *matching model*, or simply a *model*. Occasionally, when specifying the sequence of incoming item classes is useful, we will also refer to the triplet (G, I, Φ) as a (matching) model.

2.3.2 Greedy policies.

A policy Φ is called *greedy* if an incoming item is matched whenever possible, that is, if there is an unmatched item that is compatible. More formally, the policy Φ is greedy if

$$\Phi(q, i) \neq \perp \text{ for each } (q, i) \in \mathcal{Q} \times V \text{ such that } \{j \in V_i : q_j \geq 1\} \neq \emptyset. \quad (6)$$

Equivalently, a policy Φ is greedy if the set of unmatched item classes under this policy is an independent set of the compatibility graph, meaning that the state space \mathcal{Q} of the queue-size process is equal to

$$\mathcal{Q}_G = \{q \in \mathbb{N}^m : q_i q_j = 0 \text{ for each } i, j \in V \text{ such that } \{i, j\} \in E\}. \quad (7)$$

Here are two examples of greedy matching policies that will appear later in the paper:

- **Match-the-Longest (ML)**: For each $(q, i) \in \mathcal{Q} \times V$ such that $\sum_{j \in V_i} q_j \geq 1$, we choose $\Phi(q, i) \in \arg \max_{j \in V_i} (q_j)$ (ties are broken arbitrarily). This policy was considered by [33, 6, 25, 5].
- **Highest-Reward-First (HRF)**: This policy selects matches according to rewards defined on edges. If r denotes a vector of \mathbb{R}^m (indexed by the edges) with distinct coordinates, we let $\Phi(q, i) = \arg \max_{j \in V_i: q_j \geq 1} r_{i,j}$ for each $(q, i) \in \mathcal{Q} \times V$ such that $\sum_{j \in V_i} q_j \geq 1$. This definition remains valid if the coordinates of r have ties that do not impact the decision, i.e., if all pairs of edges that are incident to the same node and are not part of a triangle have distinct rewards⁴. Note that the order induced by r fully defines an HRF policy.

⁴ If three classes form a triangle in the compatibility graph, at most one of them can be non-empty under a greedy policy, so at most one edge of the triangle can be in the input of $\arg \max$.

2.3.3 Non-greedy policies.

The state space \mathcal{Q} under non-greedy policies is a strict superset of the set \mathcal{Q}_G defined in (7). In Section 6, non-greedy policies are obtained by applying the following modifications to greedy policies:

- **Filtering:** Given a subset $E^* \subsetneq E$ of edges, replace V_i (the neighbors of node i) with $V_i^* = \{j \in V_i : \{i, j\} \in E^*\}$ in the definition (1) of Φ . Intuitively, we eliminate the edges of $E \setminus E^*$ and follow a (greedy or non-greedy) policy on the subgraph $G^* = (V, E^*)$.
- **Semi-filtering:** We consider a variant of filtering policies in which the restriction of selecting matches from a subset E^* is not always enforced. Examples of semi-filtering policies will be given in Sections 6 and 7.

2.3.4 Other policies.

Although our results are stated for deterministic size-based policies because they are notationwise convenient, our results apply to a broader family of policies that are either random or require a more complex state descriptor, or both.

The extension to random policies is standard. A random (size-based) policy Φ can be defined as a function $\Phi : \mathcal{Q} \times V \times (V \cup \{\perp\}) \rightarrow [0, 1]$ such that, for each $t \in \mathbb{N}$, J_t is sampled according to the distribution $\Phi(q, i, \cdot)$ given that $Q_t = q$ and $I_t = i$. Saying that the policy is adapted to the compatibility graph G is then equivalent to saying that, for each $q \in \mathcal{Q}$ and $i \in V$, the support of $\Phi(q, i, \cdot)$ is included into $\{j \in V_i : q_j \geq 1\} \cup \{\perp\}$. The policy is greedy if $\Phi(q, i, \perp) = 0$ for each $(q, i) \in \mathcal{Q} \times V$ such that $\{j \in V_i : q_j \geq 1\} \neq \emptyset$ or, equivalently, if $\mathcal{Q} = \mathcal{Q}_G$.

An even broader of policies that fit into our framework is described in Appendix A. In a nutshell, the extended definition of a policy starts with a pair $(\mathcal{S}, |\cdot|)$, where \mathcal{S} is a countably infinite set and $|\cdot| : \mathcal{S} \rightarrow \mathbb{N}^n$ is a function that maps any state $s \in \mathcal{S}$ to the corresponding queue-size vector. The policy is then a function $\Phi : \mathcal{S} \times V \times (V \cup \{\perp\}) \times \mathcal{S} \rightarrow [0, 1]$ such that $\Phi(s, i, j, s')$ is the conditional probability that, given an incoming class- i item finds the system in state s , the matching decision is j and the new state is s' . The policy is assumed to be such that the Markov chain $S = (S_t, t \in \mathbb{N})$ defined on \mathcal{S} by the evolution of the system state is irreducible. The stochastic process $Q = (Q_t = |S_t|, t \in \mathbb{N})$ is called the *queue-size process*. It no longer satisfies the Markov property in general, but it does satisfy the evolution equations (2) and (3), with $L_i = (L_{t,i}, t \in \mathbb{N})$ and $M_k = (M_{t,k}, t \in \mathbb{N})$ defined by (4) and (5) for each $i \in V$ and $k \in E$. The state space of the queue-size process is given by $\mathcal{Q} = \{|s|, s \in \mathcal{S}\}$. The policy is greedy if $\mathcal{Q} = \mathcal{Q}_G$ and non-greedy if $\mathcal{Q} \supsetneq \mathcal{Q}_G$, where \mathcal{Q}_G is still given by (7). First-Come-First-Match (FCFM) (e.g., see [35, 15]) is a classical example of a deterministic policy that requires an expanded state descriptor (remembering the arrival order of items of different classes).

This extended policy definition can initially be ignored, with S_t (resp. \mathcal{S}) understood as Q_t (resp. \mathcal{Q}). All results for extended policies remain valid if restated for deterministic size-based policies, except for the convexity result in Proposition 20, which specifically relies on extended policies.

2.4 Performance criteria.

We now define the main performance indicators. Later, we will show that these indicators can sometimes be antagonistic, meaning that certain matching problems may require to make trade-offs between them.

Stability.

We now define the notions of *stability* and *stabilizability*, which will be explored in Section 3.

► **Definition 1** (Stability and stabilizability).

- (i) A model (G, λ, Φ) is called *stable* if the Markov chain $(S_t, t \in \mathbb{N})$ is positive recurrent. In this case, we say that the policy Φ stabilizes the problem (G, λ) .
- (ii) A problem (G, λ) is called *stabilizable* if there exists a policy that stabilizes it.
- (iii) A compatibility graph G is called *stabilizable* if there exists a vector $\lambda \in \mathbb{R}_{>0}^n$ such that the problem (G, λ) is stabilizable.

The ML greedy policy, discussed in Section 2.3.2, has the property of stabilizing all stabilizable matching problems ([33]). The First-Come-First-Match (FCFM) greedy policy has the same property (see [35] for details). When the matching problem (G, λ) is clear from the context, we will refer to a policy (adapted to G) as *stable* if the corresponding model (G, λ, Φ) is stable.

Delay.

In a stable matching model (G, λ, Φ) , the delay is defined as the long-term average of the waiting times of the items (where the waiting time is zero if the item is matched immediately upon its arrival). According to Little's law [30, 31], the delay is equal to $\frac{\sum_{i \in V} \mathbb{E}[Q_i^\infty]}{\sum_{i \in V} \lambda_i}$, where $Q^\infty = (Q_1^\infty, Q_2^\infty, \dots, Q_n^\infty)$ is a random vector following the stationary distribution of the Markov chain $Q = (Q_t, t \in \mathbb{N})$ under policy Φ .

Matching rates.

Consider a stable matching model (G, λ, Φ) . We define the *matching rate* μ_k along an edge $e_k \in E$, with endpoints i and j , as the long-run average number of matches between a class- i item and a class- j item per unit of time, given by:

$$\left(\sum_{i \in V} \lambda_i \right) \times \frac{1}{t} M_{t,k} \xrightarrow[t \rightarrow +\infty]{\text{almost surely}} \mu_k, \quad k \in E. \quad (8)$$

This quantity is uniquely defined according to the ergodic theorem, see [37, Theorem 1.10.2]. Note that $M_{t,k}/t$ represents the average number of matches *per arrival* (out of the first t arrivals), and the factor $(\sum_{i \in V} \lambda_i)$ converts this into the average number of matches *per unit of time*. The vector of matching rates associated with the model (G, λ, Φ) is denoted as $\mu = (\mu_1, \dots, \mu_m) \in \mathbb{R}_{\geq 0}^m$. If ambiguity arises, we will use the notation $\mu(\Phi)$ or $\mu(G, \lambda, \Phi)$. Sections 4 and 5 will investigate the constraints that a vector of matching rates must satisfy.

► **Remark 1** (Zero-rate edges). *The following remark will play a key role in Section 6.1. The matching rate along an edge $e_k \in E$ with endpoints i and j is given by*

$$\mu_k = \lambda_i \sum_{s, s' \in \mathcal{S}} \pi(s) \Phi(s, i, j, s') + \lambda_j \sum_{s, s' \in \mathcal{S}} \pi(s) \Phi(s, j, i, s'),$$

where π is the equilibrium distribution of the Markov chain $(S_t, t \in \mathbb{N})$. Since $\lambda_i > 0$, $\lambda_j > 0$, and the distribution π is positive on its support \mathcal{S} , it follows that the matching rate μ_k is zero if and only if $\Phi(s, i, j, s') = \Phi(s, j, i, s') = 0$ for each $s, s' \in \mathcal{S}$, that is, items of classes i and j are never matched with one another under policy Φ . In this case, the policy Φ is also adapted to the subgraph $G^* = (V, E^*)$ with $E^* = E \setminus \{e_k\}$.

Regret.

To turn a matching rate vector into a scalar performance metric, we assume that the *quality* of a matching rate vector can be evaluated by some *reward function* $f : \mu \in \mathbb{R}^m \rightarrow f(\mu) \in \mathbb{R}$. It is then convenient to define the *regret* of a policy Φ as $R(\Phi) = f_{\text{sup}} - f(\mu(\Phi))$, where $f_{\text{sup}} = \sup_{\Psi: (G, \lambda, \Psi) \text{ stable}} f(\mu(\Psi))$ is the supremum of the rewards achievable by any stable policy. Unless otherwise stated, we consider *linear* reward functions of the form $f : \mu \in \mathbb{R}^m \rightarrow r^\top \mu$, where $r = (r_1, \dots, r_m) \in \mathbb{R}^m$ is a vector of rewards associated with the edges of G .

Delay and regret will be the main focus of Sections 6–8.

General notation	
$\mathbb{N}, \mathbb{N}_{>0}, \mathbb{R}, \mathbb{R}_{\geq 0}, \mathbb{R}_{>0}$	Sets of non-negative integers, positive integers, real numbers, non-negative real numbers, positive real numbers.
$\geq, \leq, >, <$	Coordinate-wise comparison in \mathbb{R}^n .
$ \mathcal{A} $	Cardinality of the set \mathcal{A} .
Graph notation	
$G = (V, E)$	Simple graph G with $ V = n$ vertices and $ E = m$ edges.
v_i	Vertex indexed by i (denoted i if there is no ambiguity).
e_k, k , or $\{i, j\}$	Edge indexed by k , with endpoint vertices i and j .
Matching notation	
$\lambda = (\lambda_i)_{1 \leq i \leq n}$	Vector of arrival rates of the item classes.
Φ	A matching policy.
$\mu = (\mu_k)_{1 \leq k \leq m} = (\mu_{i,j})_{\{i,j\} \in E}$	Vector of matching rates along the edges.
Linear-algebra notation	
$A = (a_{i,k})_{i \in V, k \in E}$	Incidence matrix of the graph G .
$A^\top = (a_{k,i})_{k \in E, i \in V}$	Transpose of the matrix A .
$\ker(A) = \{y \in \mathbb{R}^m : Ay = 0\}$	Right kernel of the matrix A . Its dimension is called the nullity of A .
$\ker(A^\top) = \{x \in \mathbb{R}^n : A^\top x = 0\}$	Left kernel of the matrix A . Its dimension is the nullity of A^\top .
$d = m - n$	Dimension of the right kernel of the matrix A if G is surjective.

■ **Table 1** Table of key notation.

3 Graph theory and linear algebra.

In this section, we present a central result of this paper, namely, the connection between stability and the structure of the compatibility graph. Section 3.1 introduces the conservation equation, a system of linear equations satisfied by all matching rate vectors. Section 3.2 introduces the related concepts of surjectivity, injectivity, and bijectivity that will play a key role throughout the paper. In Section 3.3, we combine these concepts to formulate new necessary and sufficient conditions under which a compatibility graph G or a matching problem (G, λ) is *stabilizable* in the sense of Definition 1. Section 3.4 illustrates these results with early examples.

3.1 Conservation equation.

Computing the matching rate vector achieved by a given stable policy or characterizing the set of matching rate vectors that can be achieved by stable policies is a difficult problem *a priori*. To circumvent this difficulty, we first establish a necessary condition known as the *conservation equation*. This equation is satisfied by the matching rate vectors achieved by all stable policies. It asserts that, in a stable system, the arrival of items and their departure due to matches balance each other in the long run. More formally, given a stable matching model (G, λ, Φ) , the conservation equation (CE) can be derived by dividing (3) by t and taking the limit as t tends to infinity, and it can be written in two equivalent forms, either as a system of linear equations (CE-1) or in matrix form (CE-2):

$$\sum_{k \in E_i} \mu_k = \lambda_i, \quad i \in V, \quad (\text{CE-1})$$

$$A\mu = \lambda, \quad (\text{CE-2})$$

where the $n \times m$ matrix $A = (a_{i,k})_{i \in \{1,2,\dots,n\}, k \in \{1,2,\dots,m\}}$ is the *incidence* matrix of the graph G , defined by $a_{i,k} = 1$ if edge e_k is incident to node v_i and $a_{i,k} = 0$ otherwise. Using a conservation equation is common in queueing theory, but our contributions in this paper primarily stem from the novel mixed approach that combines graph theory and linear algebra, which we will elaborate on in the following sections.

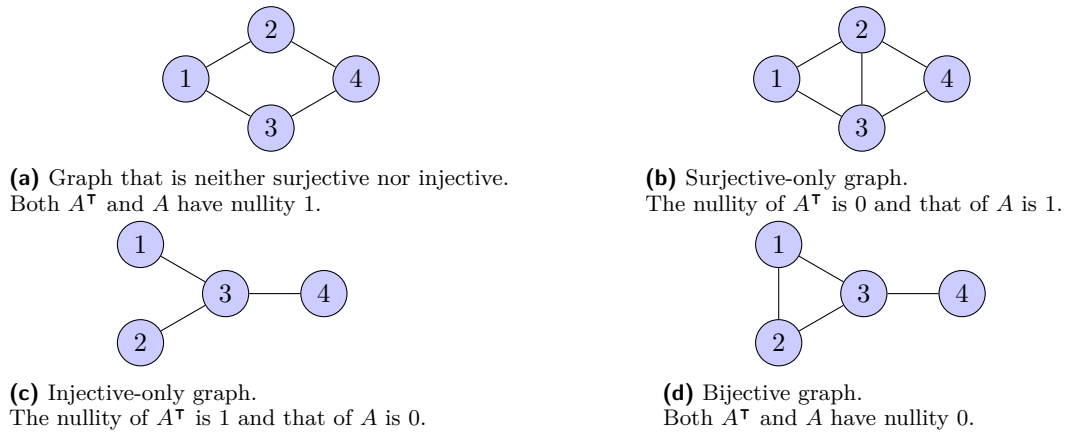
The conservation equation (CE) holds significant importance throughout this paper. While our primary focus is to understand the matching rate vector, we often find it beneficial to temporarily depart from interpreting λ and μ as vectors of arrival and matching rates in a matching model. Instead, we view the conservation equation as a linear equation with λ as a free variable and μ as an unknown. This perspective allows us to explore various aspects of the equation and its implications. In this context, we will sometimes allow the coordinates of λ and μ to be negative, even if the vectors of arrival and matching rates in a matching model have non-negative coordinates.

3.2 Surjectivity, injectivity, and bijectivity.

Definitions 2–5 below introduce the notions of *surjectivity*, *injectivity*, and *bijectivity* of a graph. In a nutshell, a compatibility graph G is said to be surjective (resp. injective, bijective) if the linear application $\mu \mapsto A\mu$ defined by its incidence matrix A is surjective (resp. injective, bijective). Interestingly, simple equivalent conditions exist in terms of the graph structure. As we will see later, these notions are fundamental to study the stability of matching models and the associated matching rate vector. In particular, we will see that (i) a compatibility graph G is stabilizable if and only if G is surjective, and (ii) the matching rates in a stabilizable matching problem (G, λ) are independent of the policy Φ if and only if G is bijective. Examples are shown in Figure 2. The equivalence of the conditions given in Definitions 2–4 and the implications in Proposition 6 are proved in Appendix B.1.

► **Definition 2 (Surjective graph).** *Consider a simple graph $G = (V, E)$ with n nodes and m edges. Let A denote the $n \times m$ incidence matrix of G . The graph G is called surjective if one of the following equivalent conditions is satisfied:*

- (i) *The function $\mu \in \mathbb{R}^m \mapsto A\mu \in \mathbb{R}^n$ is surjective.*
- (ii) *For each $\lambda \in \mathbb{R}^n$, the equation $A\mu = \lambda$ of unknown $\mu \in \mathbb{R}^m$ has at least one solution.*
- (iii) *The left kernel of the matrix A is trivial.*
- (iv) *Each connected component of the graph G is non-bipartite.*



■ **Figure 2** Examples of graphs. The nullity of a matrix is the dimension of its kernel.

► **Definition 3** (Injective graph). Consider a simple graph $G = (V, E)$ with n nodes and m edges. Let A denote the $n \times m$ incidence matrix of G . The graph G is called injective if one of the following equivalent conditions is satisfied:

- (i) The function $\mu \in \mathbb{R}^m \mapsto A\mu \in \mathbb{R}^n$ is injective.
- (ii) For each $\lambda \in \mathbb{R}^n$, the equation $A\mu = \lambda$ of unknown $\mu \in \mathbb{R}^m$ has at most one solution.
- (iii) The right kernel of the matrix A is trivial.
- (iv) Each connected component of the graph G contains at most one odd cycle and no even cycle; in other words, each connected component of G is either a tree or a unicyclic graph with an odd cycle.

► **Definition 4** (Bijective graph). Consider a simple graph $G = (V, E)$ with n nodes and m edges. Let A denote the $n \times m$ incidence matrix of G . The graph G is called bijective if the following equivalent conditions are satisfied:

- (i) The function $\mu \in \mathbb{R}^m \mapsto A\mu \in \mathbb{R}^n$ is bijective.
- (ii) For each $\lambda \in \mathbb{R}^n$, the equation $A\mu = \lambda$ of unknown $\mu \in \mathbb{R}^m$ has exactly one solution.
- (iii) The matrix A is invertible.
- (iv) Each connected component of the graph G contains one cycle and this cycle is odd.

► **Definition 5** (Surjective-only graph, injective-only graph). A simple graph G is called surjective-only (resp. injective-only) if G is surjective but not injective (resp. injective but not surjective).

The following proposition gives necessary conditions for surjectivity and injectivity in terms of the number of nodes and edges in the graph.

► **Proposition 6.** Consider an undirected graph $G = (V, E)$ with n nodes and m edges.

- (i) If G is surjective, then $n \leq m$.
- (ii) If G is injective, then $n \geq m$.
- (iii) If G is bijective, then $n = m$.
- (iv) If G is surjective, then G is also injective if and only if $n = m$.
- (v) If G is injective, then G is also surjective if and only if $n = m$.

3.3 Necessary and sufficient conditions for stabilizability

The definitions and results of Sections 3.1 and 3.2 allow us to formulate new necessary and sufficient conditions for stabilizability.

3.3.1 Stabilizable compatibility graph.

The following proposition gives necessary and sufficient conditions for a graph G to be stabilizable, in terms of either its structure or its incidence matrix.

► **Proposition 7.** *Let G be a compatibility graph. The following conditions are equivalent:*

- (i) *The graph G is stabilizable.*
- (ii) *The graph G is surjective.*

Proof. Equivalence between Proposition 7i and Definition 2iv has been proved by [33, Theorem 1]. ◀

Although the equivalence between Proposition 7i and Definition 2iv is a known result, no prior literature on stochastic matching models (to the best of our knowledge) has explored the connection between Proposition 7i and the alternative definitions of surjectivity introduced in Definition 2. As we will see later, this new characterization of the stabilizability of a graph G will be useful to analyze the relationship between delay and matching rates.

In the remainder, we will use the words “stabilizable” and “surjective” interchangeably. Furthermore, unless stated otherwise, we will assume that the graph G is surjective.

3.3.2 Stabilizable matching problem.

We now turn to the stabilizability of a matching problem (G, λ) . As recalled in Section 2.4, two examples of greedy policies that stabilize the model whenever this matching problem is stabilizable are ML and FCFM, as shown by [33] and [35], respectively. Proposition 8 below provides necessary and sufficient conditions for the matching problem (G, λ) to be stabilizable; condition ii was already derived by [33], but condition iii is new⁵.

► **Proposition 8.** *Consider a matching problem (G, λ) with a surjective graph G . The following conditions are equivalent:*

- (i) *The matching problem (G, λ) is stabilizable.*
- (ii) *For each independent set $\mathcal{I} \in \mathbb{I}$, we have $\sum_{i \in \mathcal{I}} \lambda_i < \sum_{i \in V(\mathcal{I})} \lambda_i$.*
- (iii) *The conservation equation (CE) admits a solution $\mu \in \mathbb{R}_{>0}^m$ (i.e., a solution with strictly positive components).*

Proof. This result is proved in Appendix B.2 using results from [33, 15]. ◀

One might expect that the time complexity to verify condition ii in Proposition 8 is exponential in the general case, considering that the number of independent sets grows exponentially fast with the number n of classes. However, [33, Proposition 1] proved that there exists a polynomial algorithm for verifying this condition. It is worth noting that this verification process is indirect in that it involves constructing a bipartite double cover of G . In contrast, condition iii offers a more direct, polynomial-time method to verify the stabilizability of a matching model (G, μ) . To present this approach more explicitly, we differentiate between two cases based on whether the graph G , assumed to be surjective, is surjective-only or bijective.

► **Remark 2.** *As observed by [15, Lemma 12], if the graph G is surjective, condition iii in Proposition 8 gives a simple way of generating vectors $\lambda \in \mathbb{R}_{\geq 0}^n$ such that the problem (G, λ) is stabilizable: it suffices to take $\lambda = A\mu$ for some $\mu \in \mathbb{R}_{>0}^m$. For instance, if $\mu = (\beta, \dots, \beta)$ for some $\beta > 0$, then the coordinates of λ are proportional to the degree of each node.*

⁵ [35] concurrently proved a similar result, see Section 1.3.

Verify stabilizability when G is bijective.

If the graph G is bijective, then the matrix A is invertible, and (CE) has a unique solution, namely $A^{-1}\lambda$. Proposition 8 implies that the matching problem (G, λ) is stabilizable if and only if all coordinates of $A^{-1}\lambda$ are positive. The special case of bijective graphs will be investigated in detail in Section 4, where we will provide a more direct expression for $A^{-1}\lambda$.

Verify stabilizability when G is surjective-only.

If the compatibility graph G is surjective-only, (CE) has multiple solutions. To determine if one of these solutions is positive, it suffices to solve a linear optimization problem that searches for a solution to (CE) whose smallest coordinate is as large as possible:

$$\begin{aligned} & \underset{z=(z_1, z_2, \dots, z_{m+1}) \in \mathbb{R}^{m+1}}{\text{Maximize}} && z_{m+1}, \\ & \text{Subject to} && A(z_1, z_2, \dots, z_m)^\top = \lambda, \\ & && z_i \geq z_{m+1}, \quad i \in \{1, 2, \dots, m\}. \end{aligned} \tag{10}$$

Here, the first m coordinates of the vector z are the coordinates of a vector $\mu \in \mathbb{R}^m$ that satisfies (CE), and the last coordinate of z is a lower bound of the coordinates of this vector μ . Indeed, the equality constraint means that μ satisfies (CE), and the inequality constraint means that the last coordinate of z is less than or equal to its other coordinates. The value to maximize is the last coordinate of the vector z .

If z is a solution to (10), we call the corresponding vector $\mu \in \mathbb{R}^m$ a *maximin* solution to (CE). The optimization problem (10) is a textbook linear optimization problem. It can be solved with a time complexity that is polynomial in the number n of nodes and the number m of edges using many methods, for instance the interior-point-method, see [26].

The linear optimization problem (10) has a solution with positive coordinates if and only if (CE) has a solution with positive coordinates. According to Proposition 8, this is equivalent to saying that the matching problem (G, λ) is stabilizable. Therefore, to verify if a matching problem (G, λ) is stabilizable, it suffices to find a solution to the linear optimization problem (10) and to check if its last coordinate is positive.


► **Remark 3.** *Observe that the optimization problem (10) always has solutions with finite coordinates. Indeed, the set of vectors that satisfy the constraints of (10) contains at least one valid solution with real-valued coordinates (this is again a consequence of the surjectivity of G). We just need to consider an arbitrary solution μ of (CE) (see Section 5.1.2 for a concrete example using the Moore-Penrose inverse) and to let $z_\mu = (\mu_1, \mu_2, \dots, \mu_m, \min(\mu))$, where $\min(\mu)$ is the smallest coordinate of the vector μ . Any solution better than z_μ has all its coordinates lower-bounded by $\min(\mu)$ and upper-bounded by $\max(\lambda) - \min(0, (n-2)\min(\mu))$. The latter bound is obtained by observing that, if edge k is incident to node i and if $(\mu'_1, \mu'_2, \dots, \mu'_m, x')$ is a solution to (10) such that $x' \geq \min(\mu)$, then $\mu'_k = \lambda_i - \sum_{\ell \in E_i \setminus k} \mu'_\ell$ by (CE). We then use the inequalities $\lambda_i \leq \max(\lambda)$ and $\sum_{\ell \in E_i \setminus k} \mu'_\ell \geq \min(0, (n-2)x') \geq \min(0, (n-2)\min(\mu))$ (this latter inequality is obtained by distinguishing two cases, depending on whether $\min(\mu) \geq 0$ or $\min(\mu) < 0$). Therefore, the solutions better than z_μ belong to a compact set of \mathbb{R}^{m+1} , which ensures the existence of an optimal solution with finite coordinates.*

3.4 Early examples.


We now provide illustrative examples that illustrate the stabilizability results of Propositions 7 and 8 as well as our definitions of surjectivity, injectivity, and bijectivity. These examples will also introduce useful notions that will be further explored in Sections 4–6.

3.4.1 Bijective graphs.


We first consider connected compatibility graphs G that are both surjective and injective: the graph contains exactly one cycle, and this cycle is odd. According to Definition 4, (CE) has a unique solution for each vector $\lambda \in \mathbb{R}^n$ of arrival rates. Proposition 8 implies that the matching problem (G, λ) is stabilizable if and only if the coordinates of this solution are positive, in which case this solution gives the matching rates achieved by *all* stable matching policies. By Remark 2, one can always exhibit a vector $\lambda \in \mathbb{R}_{>0}^n$ of arrival rates such that the matching problem (G, λ) is stable.




(a) Matching rates in the triangle graph \mathcal{C}_3 .




(b) Matching rates in the paw graph. $\bar{\lambda}_3 = \lambda_3 - \lambda_4$ denotes the residual rate that class 3 can provide to classes 1 and 2.



(c) Matching rates in the square graph \mathcal{C}_4 with the normalization $\lambda_1 + \lambda_4 = \lambda_2 + \lambda_3 = \frac{1}{2}$. This graph is not stabilizable.



(d) Matching rates in the diamond graph with the normalization $\lambda_1 + \lambda_4 = \frac{1}{2}$. $2\beta = \lambda_2 + \lambda_3 - \lambda_1 - \lambda_4 = \lambda_2 + \lambda_3 - \frac{1}{2}$ is the difference between the arrival rates of the inner part $\{2, 3\}$ and the outer part $\{1, 4\}$. $\bar{\lambda}_2 = \lambda_2 - \beta$ and $\bar{\lambda}_3 = \lambda_3 - \beta$ represent the residual rates that classes 2 and 3 can provide to classes 1 and 4, and they are such that $\lambda_1 + \lambda_4 = \bar{\lambda}_2 + \bar{\lambda}_3 = \frac{1}{2}$, as in Example 3.



(e) Matching rates in the kayak paddle $KP_{3,3,1}$.

■ **Figure 3** Examples of Section 3.4.

► **Example 1** (Triangle). If the graph G is a triangle graph \mathcal{C}_3 , and the vector $\lambda = (\lambda_1, \lambda_2, \lambda_3)$ is given, the solution to (CE) is unique and shown in Figure 3a. According to Proposition 8iii, (G, λ) is stabilizable if and only if all coordinates of this solution are positive. This condition is indeed equivalent to Proposition 8ii, which states that $\lambda_1 < \lambda_2 + \lambda_3$, $\lambda_2 < \lambda_1 + \lambda_3$, and $\lambda_3 < \lambda_1 + \lambda_2$. Alternatively, this condition can be expressed as λ_1 , λ_2 , and λ_3 being the lengths of the sides of a non-degenerate triangle (i.e., they satisfy the triangular inequality).

Under these conditions, the model (G, λ, Φ) is stable if Φ is the (unique) greedy policy adapted to G (indeed, Proposition 34 shows that a complete graph G admits a unique greedy policy Φ).

► **Example 2** (Paw graph). If G is a paw graph, the solution to (CE) is again unique and shown in Figure 3b. Here, $\bar{\lambda}_3 = \lambda_3 - \lambda_4$ is the remaining rate of class 3 after accounting for the needs of class 4. After this subtraction, the matching rates along edges $\{1, 2\}$, $\{1, 3\}$, and $\{2, 3\}$ are defined as in the triangle graph of Figure 3a.

It is important to note that, while the existence of a matching rate vector with positive coordinates guarantees that a greedy policy like ML is stable, there may still exist greedy policies that are unstable. Consider, for instance, the paw graph of Figure 3b and a Highest-Reward-First policy where the edges $\{1, 3\}$ and $\{2, 3\}$ offer higher rewards than $\{3, 4\}$. Even if the problem is stabilizable, this policy might select the edge $\{3, 4\}$ at a slower rate than class-4 items arrive, causing instability (see [33, Section 5]).

3.4.2 Bipartite graph (neither injective nor surjective).

Besides explaining intuitively why bipartite graphs cannot be stabilized, the following example will prepare the ground for Example 4.

► **Example 3** (Square graph). Figure 3c shows a square graph C_4 . This graph is not surjective because it is bipartite with parts $\{1, 4\}$ (called the outer part) and $\{2, 3\}$ (inner part). Therefore, according to Proposition 7, this graph is not stabilizable. Yet, given a vector $\lambda = (\lambda_1, \lambda_2, \lambda_3, \lambda_4)$ of arrival rates, (CE) may still have a solution with positive coordinates. This does not contradict Proposition 8, as Item iii is equivalent to Items i and ii only if the graph G is surjective. Assuming unit normalization, the conservation equation (CE-2) has a solution if and only if

$$\lambda_1 + \lambda_4 = \lambda_2 + \lambda_3 = \frac{1}{2}. \quad (11)$$

If (11) is satisfied, the solutions to (CE) can be described with a parameter α as shown in Figure 3c: starting from the particular solution $\mu = (2\lambda_1\lambda_2, 2\lambda_1\lambda_3, 2\lambda_2\lambda_4, 2\lambda_3\lambda_4)$, all solutions can be generated by alternately adding and subtracting α along the (even) cycle 1-2-4-3. The positive solutions correspond to values of α such that $-2 \min(\lambda_1\lambda_2, \lambda_3\lambda_4) < \alpha < 2 \min(\lambda_1\lambda_3, \lambda_2\lambda_4)$.

To understand why the matching problem (C_4, λ) is not stabilizable even when (CE) has a solution with positive coordinates, let us focus on the system dynamics. We can use (3) to show that the difference in total queue size between the outer part $\{1, 4\}$ and the inner part $\{2, 3\}$ satisfies $Q_{t,1} + Q_{t,4} - (Q_{t,2} + Q_{t,3}) = L_{t,1} + L_{t,4} - (L_{t,2} + L_{t,3})$ for each $t \in \mathbb{N}$. Therefore, the Markov chain $(Q_{t,1} + Q_{t,4} - (Q_{t,2} + Q_{t,3}), t \in \mathbb{N})$ is a random walk on the integer number line $\{\dots, -2, -1, 0, 1, 2, \dots\}$, with transition probability proportional to $\lambda_1 + \lambda_4$ in the increasing direction and to $\lambda_2 + \lambda_3$ in the decreasing direction. If (11) is not satisfied, this random walk is transient, and the difference between the queue sizes of the two parts grows linearly with time. On the other hand, if (11) is satisfied, then the random walk does not have this bias, but the model is still unstable because the corresponding Markov chain is null recurrent⁶.

⁶ Some existing studies of matching in bipartite graphs solve this issue by coupling arrivals in the two components, see [2, 9, 13], or by assuming that items have a finite patience time, as in [25].

3.4.3 Surjective-only graphs.

We finally consider a compatibility graph G that is surjective but not injective. In other words, the graph G is stabilizable and (CE) has an infinite number of solutions. The achievability of these solutions by a stable matching policy will be discussed in Section 6.

► **Example 4** (Diamond (double-fan) graph). *Figure 3d shows the diamond graph D , that is, a square graph with an additional edge between nodes 2 and 3. Compared to Example 3, this additional edge makes the graph non-bipartite, and therefore surjective, so that the graph is stabilizable according to Proposition 7. For ease of computation, we assume that the vector $\lambda = (\lambda_1, \lambda_2, \lambda_3, \lambda_4)$ of arrival rates is normalized so that $\lambda_1 + \lambda_4 = \frac{1}{2}$. Under this assumption, and with $\beta = \frac{1}{2}(\lambda_2 + \lambda_3 - \lambda_1 - \lambda_4) = \frac{1}{2}(\lambda_2 + \lambda_3) - \frac{1}{4}$, $\bar{\lambda}_2 = \lambda_2 - \beta$, and $\bar{\lambda}_3 = \lambda_3 - \beta$, the general solution to (CE) can be described with a parameter α as shown in Figure 3d. After subtracting β from λ_2 and λ_3 , the solutions to (CE) for all edges but $\{2, 3\}$ are exactly the same as in the square graph of Example 3.*

According to Proposition 8ii, the matching problem (D, λ) is stabilizable if and only if

$$\lambda_2 < \lambda_1 + \lambda_3 + \lambda_4, \quad \lambda_3 < \lambda_1 + \lambda_2 + \lambda_4, \quad \lambda_1 + \lambda_4 < \lambda_2 + \lambda_3, \quad (12)$$

that is, $\bar{\lambda}_3 > 0$, $\bar{\lambda}_2 > 0$, and $\beta > 0$. If these inequalities are satisfied, the positive solutions correspond to values of α such that $-2 \min(\lambda_1 \bar{\lambda}_2, \bar{\lambda}_3 \lambda_4) < \alpha < 2 \min(\lambda_1 \bar{\lambda}_3, \bar{\lambda}_2 \lambda_4)$. Intuitively, compared to the square graph, stabilizable matching problems (D, λ) have a positive difference of 2β between the arrival rates of the inner part $\{2, 3\}$ and the outer part $\{1, 4\}$. This difference is absorbed by the central edge $\{2, 3\}$, which has matching rate β . Like Example 1 and unlike Example 2, the matching model (D, λ, Φ) is stable for every greedy policy Φ provided that (12) is satisfied (as shown in Corollary 27 of Appendix B.3).

► **Example 5** (Kayak paddle graph). *Figure 3e shows a kayak paddle $KP_{3,3,1}$, consisting of two triangles linked by an edge. According to Proposition 8, the matching problem (G, λ) is stabilizable if and only if there exists $\alpha > 0$ such that $(\lambda_1, \lambda_2, \lambda_3 - \alpha)$ and $(\lambda_4 - \alpha, \lambda_5, \lambda_6)$ are the arrival rate vectors of two stabilizable triangle graphs \mathcal{C}_3 . The solutions to (CE) can be described by varying α as shown in Figure 3e. Assuming the matching problem (G, λ) is stabilizable, the solutions to (CE) with positive coordinates correspond to the values of α such that*

$$0 < \alpha < \min(\lambda_3 - |\lambda_2 - \lambda_1|, \lambda_4 - |\lambda_5 - \lambda_6|).$$

Intuitively, solutions with positive coordinates have a positive matching rate α along edge $\{3, 4\}$. After subtracting this rate from λ_3 and λ_4 , the subgraphs restricted to nodes 1, 2, and 3 and to nodes 4, 5, and 6, respectively, behave like the triangle of Figure 3a. Like Example 2 and unlike Examples 1 and 4, the fact that (G, λ) is stabilizable does not guarantee the stability of all greedy policies.

4 Matching rates in bijective graphs.

In the remainder, we consider exclusively compatibility graphs G that are stabilizable. According to Definition 2 and Proposition 7, this implies that the graph G is surjective, i.e., that each connected component of G is non-bipartite. By Proposition 6iv and Proposition 8, there are only two possible cases:

1. If $n = m$, the graph G is also bijective: each connected component of G is a unicyclic graph, and its (only) cycle is odd. The conservation equation (CE) has a unique solution given by $\mu = A^{-1}\lambda$. The matching problem (G, λ) is stabilizable if and only if all

components of this solution are positive, in which case the solution provides the matching rates achievable by any stable policy.

2. If $n < m$, the graph G is surjective-only: each connected component of G is non-bipartite, and at least one of these connected components contains an even cycle or a pair of odd cycles. The conservation equation (CE) has an infinite number of solutions. The matching problem (G, λ) is stabilizable if and only if one of them has all-positive components.

Case 1 is studied in this section, while case 2 will be studied in Sections 5 and 6.

In Proposition 9 below, we give a simpler expression for the unique solution $\mu = A^{-1}\lambda$ of (CE) in terms of the arrival rate vector λ , under the assumption that the graph G is bijective. We assume without loss of generality that G is connected, as otherwise we can consider each connected component independently. Compared to the expression $\mu = A^{-1}\lambda$, the advantage of Proposition 9 is twofold: it does not require calculating a matrix inversion, and it highlights the monotonicity of the matching rates with respect to the arrival rates. This result will be illustrated in Examples 6 and 7.

► **Proposition 9.** *Consider a problem (G, λ) with a compatibility graph $G = (V, E)$ that is connected and bijective, and consider an edge $k \in E$.*

- (i) *If edge k does not belong to the (unique odd) cycle of the graph G , then edge k separates the graph G into two parts, namely a tree and a unicyclic graph. If $V_k \subset V$ denotes the set of nodes that belong to the tree (including one endpoint of edge k), then the matching rate along edge k is given by*

$$\mu_k = \sum_{i \in V_k} (-1)^{d_{i,k}} \lambda_i, \quad (13)$$

where $d_{i,k}$ is the distance between node i and edge k , defined as the minimum distance between node i and an endpoint of edge k .

- (ii) *If edge k belongs to the (unique odd) cycle of the graph G , then the matching rate along edge k is*

$$\mu_k = \frac{1}{2} \left(\sum_{i \in V} (-1)^{d_{i,k}} \lambda_i \right). \quad (14)$$

Proof. We first prove (13) for every edge k that does not belong to the cycle. As observed in the proposition, each edge k that does not belong to the cycle separates the graph into two parts, one of which is a tree with node set V_k ; the rooted tree associated with k is obtained by designating the corresponding endpoint of edge k as the root. We now prove (13) by induction on the height of this rooted tree. Equation (13) is true if the height of this tree is zero. Indeed, in this case, the endpoint of edge k that belongs to the tree, say node i , has no other incident edge, so that applying (CE-1) to node i yields $\mu_k = \lambda_i$, which is consistent with (13). Now assume that the assumption is satisfied for each edge whose associated rooted tree has height at most $h - 1$ for some $h \geq 0$, and consider an edge k whose associated rooted tree has height h . By applying (CE-1) to the root i of this associated rooted tree, we obtain

$$\mu_k = \lambda_i - \sum_{\ell \in E_i \setminus \{k\}} \mu_\ell. \quad (15)$$


The induction hypothesis guarantees that (13) is satisfied for every $\ell \in E_i \setminus \{k\}$ (as the height of the associated rooted tree is at most $h - 1$). After injecting this observation to (15), the result for edge k follows by observing that $d_{j,k} = d_{j,\ell} + 1$ for each $\ell \in E_i \setminus \{k\}$ and $j \in V_\ell$, and that $V_k = \{i\} \cup (\bigcup_{\ell \in E_i \setminus \{k\}} V_\ell)$ (all sets being disjoint).

We now prove (14) for each edge k that belongs to the cycle. Since the graph G is unicyclic, deleting edge k from G yields a (connected) tree, and therefore a bipartite graph. We let V_+ denote set of nodes in the part that contains both endpoints of edge k (that both endpoints belong to the same part follows from the fact that the cycle is odd) and V_- the set of nodes in the other part. We obtain

$$\sum_{i \in V_+} \lambda_i - \sum_{i \in V_-} \lambda_i = \sum_{i \in V_+} \sum_{\ell \in E_i} \mu_\ell - \sum_{i \in V_-} \sum_{\ell \in E_i} \mu_\ell = 2\mu_k.$$

The first equality follows from (CE-1). The second equality holds because each edge $\ell \in E \setminus \{k\}$ has one endpoint in V_+ and another in V_- , so that μ_ℓ appears once in the first nested sum and once in the second; on the contrary, since both endpoints of edge k belong to V_+ , μ_k appears twice in the first nested sum and zero times in the second. Equation (14) follows by observing that $d_{i,k}$ is even if and only if $i \in V_+$. ◀

We remark that the influence of the arrival rate of a node on the matching rate along an edge only depends on the parity of the distance between the edge and the node. The actual distance does not. In particular, even in a very large (bijective) graph, a node far away from an edge has the same (although possibly reversed) impact on this edge's matching rate as an endpoint of this edge.



(a) Matching rates in the pentagon graph C_5 . Only rate $\mu_{1,2}$ is shown for ease of display. The other rates are deduced by permutation.

(b) Matching rates in a “lying puppet” graph with $n = 9$ nodes and $m = 9$ edges. The differences $\bar{\lambda}_7 = \lambda_7 - \lambda_8 - \lambda_9$, $\bar{\lambda}_4 = \lambda_4 - \lambda_5 - \lambda_6 - \lambda_7$, and $\bar{\lambda}_3 = \lambda_3 - \lambda_4$ are the residual rates that classes 7, 4, and 3 provide to their neighbors of lower index.

■ **Figure 4** Matching rates in bijective graphs.

► **Example 6** (Cycle graph with 5 nodes). A cycle graph is the simplest bijective graph that we can consider, as it contains an odd cycle and no other edges. In the cycle graph of Figure 4a, a direct application of Proposition 9ii yields $\mu_{1,2} = \frac{1}{2}(\lambda_1 + \lambda_2 - \lambda_3 + \lambda_4 - \lambda_5)$. Matching rates along other edges follow by symmetry. From the point of view of edge $\{1, 2\}$, we can partition nodes into two sets, namely $\{1, 2, 4\}$ and $\{3, 5\}$. The former (resp. latter) set contains nodes at an even (resp. odd) distance of edge $\{1, 2\}$, and increasing the arrival rate of these nodes increases (resp. decreases) the matching rate along edge $\{1, 2\}$.

► **Example 7** (Lying puppet). We now consider the graph of Figure 4b. Edges $\{1, 2\}$, $\{1, 3\}$, and $\{2, 3\}$ belong to the cycle, and the other edges do not. According to Proposition 9, we have

$$\mu_{1,2} = \frac{\lambda_1 + \lambda_2 - \bar{\lambda}_3}{2}, \quad \mu_{1,3} = \frac{\lambda_1 - \lambda_2 + \bar{\lambda}_3}{2}, \quad \mu_{2,3} = \frac{-\lambda_1 + \lambda_2 + \bar{\lambda}_3}{2},$$

where $\bar{\lambda}_3 = \lambda_3 - \mu_{3,4}$, and

$$\mu_{4,5} = \lambda_5, \quad \mu_{4,6} = \lambda_6, \quad \mu_{7,8} = \lambda_8,$$

$$\mu_{7,9} = \lambda_9, \quad \mu_{4,7} = \lambda_7 - \mu_{7,8} - \mu_{7,9}, \quad \mu_{3,4} = \lambda_4 - \mu_{4,5} - \mu_{4,6} - \mu_{4,7}.$$

This second set of equations can be obtained either by a direct application of (13) or by applying (CE-1) recursively from the leaves. Indeed, applying (CE-1) to nodes 5, 6, 8, and 9 gives directly the values of $\mu_{4,5}$, $\mu_{4,6}$, $\mu_{7,8}$, and $\mu_{7,9}$, then applying (CE-1) to node 7 gives the value of $\mu_{4,7}$, and finally applying (CE-1) to node 4 gives the value of $\mu_{3,4}$. The values of $\mu_{1,2}$, $\mu_{1,3}$, and $\mu_{2,3}$ are similar to Example 2, where the arrival rate λ_3 is again replaced with the effective arrival rate $\bar{\lambda}_3$ from the point of view of classes 1 and 2.

5 Polytope of solutions in surjective-only graphs.

Let G be a surjective compatibility graph. In Section 4, we considered the case where G was bijective, so that the vector of matching rates for any stable policy was the unique solution to (CE). Let us now assume that G is surjective-only, meaning that each connected component of G is non-bipartite and that at least one component contains either an even cycle or a pair of odd cycles (cf. Definitions 2, 3, and 5). In this case, the solution set of (CE) is infinite. The matching rate vector under any stable policy belongs to this set and has non-negative coordinates.

To better control the matching rates, we first need to understand the structure of the solution set of (CE). To this end, Section 5.1 generalizes the earlier examples from Section 3.4 by systematically characterizing the affine space of all real-valued solutions to (CE). Following this, Section 5.2 describes the convex polytope formed by the solutions with non-negative coordinates, in particular its vertices, that is its extreme points. Whether these solutions can be achieved through a stable matching policy is explored in Section 6.

5.1 Affine space of real-valued solutions.

We first consider the set of solutions to (CE) with real-valued (positive, zero, or negative) coordinates:

$$\Pi = \{\mu \in \mathbb{R}^m : A\mu = \lambda\}. \quad (16)$$

We now delve into the properties of Π . In Section 5.1.1, we recall that Π is an affine space of dimension $d = m - n$ that can be described as a translation of the kernel of the incidence matrix A by a particular solution to (CE). Section 5.1.2 gives examples of particular solutions that can be used. Section 5.1.3 gives an algorithm to construct a basis for the kernel of the incidence matrix using a spanning tree of the graph G .

5.1.1 Edge basis, kernel basis.

The following proposition characterizes the solution set Π of (CE) using the incidence matrix of the compatibility graph. Equation (17) shows in particular that, up to translation, this solution set depends only on the structure of the compatibility graph G , while the arrival rate vector λ impacts only the translation vector.

► **Proposition 10.** *Consider a matching problem (G, λ) with a surjective-only compatibility graph G , and let A denote the incidence matrix of G . The solution set Π of (CE) is the affine space obtained by translating the kernel $\ker(A)$ of the matrix A by a particular solution μ° of (CE), that is,*

$$\Pi = \{\mu^\circ + \mu : \mu \in \ker(A)\}. \quad (17)$$

Furthermore, the vector space $\ker(A)$ and the affine space Π have dimension $d = m - n$.

Proof. That the set Π is of the form (17) is a well-known result in linear algebra. Definition 2 about surjectivity implies that the rank of A is n , and we conclude from the rank-nullity theorem that the nullity of A is $d = m - n$. The affine space Π has the same dimension according to (17). \blacktriangleleft

Thanks to Proposition 10, given a particular solution μ° of (CE) and a basis $\mathcal{B} = (b_1, b_2, \dots, b_d)$ of $\ker(A)$, we can rewrite the affine space Π as

$$\Pi = \{\mu^\circ + \alpha_1 b_1 + \alpha_2 b_2 + \dots + \alpha_d b_d : \alpha_1, \alpha_2, \dots, \alpha_d \in \mathbb{R}\}.$$

In fact, we can define the following affine isomorphism between the coordinate space \mathbb{R}^d and the d -dimensional affine space Π :

$$\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d) \in \mathbb{R}^d \mapsto \mu = \mu^\circ + \alpha_1 b_1 + \alpha_2 b_2 + \dots + \alpha_d b_d \in \Pi. \quad (18)$$

This allows us to define two coordinate systems, *edge coordinates* and *kernel coordinates*.

► **Definition 11** (Edge basis, kernel basis). *Consider a matching problem (G, λ) with a surjective-only graph G . Given a particular solution μ° of (CE) and a basis $\mathcal{B} = (b_1, b_2, \dots, b_d)$ of $\ker(A)$, there are two natural bases to represent vectors in Π :*

- *Edge basis: A vector of Π is described by its canonical coordinates $\mu = (\mu_1, \mu_2, \dots, \mu_m) \in \mathbb{R}^m$, solution to (CE), where μ_k represents a candidate matching rate along edge k , for each $k \in E$.*
- *Kernel basis: A vector of Π is described by its coordinates $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d) \in \mathbb{R}^d$ in the basis \mathcal{B} , where $d = m - n$ is the dimension of the affine space Π .*

If B is the $m \times d$ matrix giving the coordinates of the vectors of the basis \mathcal{B} in the edge basis, the change-of-basis formulas are as follows:

- *A vector of Π with coordinates α in kernel basis has coordinates $\mu = \mu^\circ + B\alpha$ in edge basis;*
- *A vector of Π with coordinates μ in edge basis has coordinates $\alpha = B^+(\mu - \mu^\circ)$ in kernel basis, where B^+ is the pseudo-inverse ⁷ (or Moore-Penrose inverse) of B .*

Both bases have advantages. The edge basis, by definition, gives directly the candidate matching rates. The kernel basis allows us to work in lower dimension (d instead of m) and to ignore the conservation equation (CE), which is implicitly enforced. In the remainder, we will often use interchangeably the edge coordinates $\mu = (\mu_1, \mu_2, \dots, \mu_m)$ and the kernel coordinates $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d)$ to describe a given vector in Π . Which basis we are actually using will be made clear by our choice of letters (either μ or α).

For graphs that have a low kernel dimension d , it is convenient to mix both approaches by representing a generic vector of Π , i.e., a generic solution to (CE), in the form $\mu^\circ + \alpha_1 b_1 + \alpha_2 b_2 + \dots + \alpha_d b_d$. The solutions in Examples 4 and 5 displayed in Figures 3d and 3e follow this convention. This representation, along with the possibility to switch between edge basis and kernel basis, will be used extensively in Sections 5.2 and 6.

5.1.2 Particular solution.

We propose two ways of computing a particular solution to (CE).

⁷ The columns of the matrix B are linearly independent because \mathcal{B} is a basis, so that the pseudo-inverse B^+ has the simple expression $B^+ = (B^\top B)^{-1} B^\top$, where the $d \times d$ matrix $B^\top B$ is invertible because $\ker(B^\top B) = \ker(B) = \{0\}$.

Maximin solution.

A solution to the linear optimization problem (10) from Section 3.3.2 is a particular solution. Recall that such a solution allows us to determine whether the matching problem (G, λ) is stabilizable by checking if all its coordinates are positive.

Pseudoinverse.

Alternatively, a standard approach to simultaneously finding a particular solution μ° and characterizing $\ker(A)$ makes use of the pseudoinverse (or Moore-Penrose inverse) of the matrix A . Since Definition 2 on surjectivity implies that the rows of A are linearly independent, the pseudoinverse A^+ of A has the following simple form:

$$A^+ = A^\top(AA^\top)^{-1},$$

where the $n \times n$ matrix AA^\top is invertible because $\ker(AA^\top) = \ker(A^\top) = \{0\}$. We can then describe a particular solution μ^\perp and the kernel $\ker(A)$ as follows:

$$\mu^\perp = A^+\lambda, \quad \ker(A) = \{(\text{Id}_{m \times m} - A^+A)\mu : \mu \in \mathbb{R}^m\}, \quad (19)$$

where $\text{Id}_{m \times m}$ denotes the m -dimensional identity matrix. The vector μ^\perp is the *least-squares solution* to (CE), and it is orthogonal to $\ker(A)$. In general, some coordinates of this solution can be negative even if non-negative solutions exist. For example, if G is the diamond graph D of Example 4, then the matching problem (D, λ) with $\lambda = (4, 5, 2, 1)$ is stabilizable ($\mu = (\frac{7}{2}, \frac{1}{2}, 1, \frac{1}{2}, \frac{1}{2})$ is a solution to (CE) with positive coordinates), but the solution given by the pseudoinverse is $\mu^\perp = (\frac{11}{4}, \frac{5}{4}, 1, \frac{5}{4}, -\frac{1}{4})$.


Equation (19) shows that the pseudoinverse also provides an implicit characterization of $\ker(A)$, though this is not very practical, as it relies on a projection from \mathbb{R}^m to $\ker(A)$. Section 5.1.3 offers a more direct characterization by constructing a basis for $\ker(A)$ based on the structure of the compatibility graph G .

5.1.3 Basis of the kernel of the incidence matrix.

Recall that a vector $\mu \in \mathbb{R}^m$ belongs to $\ker(A)$ if and only if $A\mu = 0$, which reads $\sum_{k \in E_i} \mu_k = 0$ for each $i \in V$. In other words, a vector $\mu \in \mathbb{R}^m$ belongs to $\ker(A)$ if and only if, for each $i \in V$, the sum of the coordinates of μ associated with the edges that are incident to node i is zero. Using this observation, we first give examples of vectors that belong to $\ker(A)$, and then we give an algorithm that generates a basis $\mathcal{B} = (b_1, b_2, \dots, b_d)$ of $\ker(A)$.

First observe that an even cycle, if it exists, is the support of a vector in $\ker(A)$: it suffices to assign alternatively the values $+1$ and -1 to the edges of this cycle and the value 0 to all other edges. In the diamond graph of Example 4, if edges are numbered in lexicographical order, then $y = (1, -1, 0, -1, 1)$ is a vector of the unidimensional kernel ($d = m - n = 1$), with support the even cycle 1-2-4-3 (see Figure 5a). Intuitively, even cycles can be used to move weight between “odd” and “even” edges of the cycle without modifying the value of the product $A\mu$. Actually, in this example, Figure 3d shows that the only way to increase the matching rate along edges $\{1, 2\}$ and $\{3, 4\}$ is if we reduce the matching rate along edges $\{1, 3\}$ and $\{2, 4\}$, and conversely.

Apart from even cycles, other structures of interest are kayak paddles $KP_{\ell, r, p}$, made of two odd cycles (of lengths ℓ and r) connected by a path (of length p). Such graphs have a unidimensional kernel, and a base vector can be found by assigning properly the values $+1$



(a) Vector of the kernel space of the diamond graph.

(b) Vector of the kernel space of the kayak paddle $KP_{3,5,2}$.

■ **Figure 5** Base vectors of kernel spaces of two toy graphs.

and -1 along the cycles and the values $+2$ and -2 along the path. Figure 5b shows such an assignment for $KP_{3,5,2}$.

Surprisingly, for any surjective graph G , one can build a basis of $\ker(A)$ using only subgraphs of G that are even cycles and kayak paddles. Algorithm 1, derived from [20], describes such a construction. Appendix C.1 gives a more detailed description of this algorithm and proof that it terminates and returns the desired result.


Figure 6 shows a possible run of Algorithm 1 for the codomino graph: in Figure 6a, a spanning tree \mathcal{T} (plain back edges) and a pivotal edge a such that $\mathcal{T} \cup \{a\}$ contains an odd cycle (dotted edge) are selected. Then, for each edge s_i that does not belong to $\mathcal{T} \cup \{a\}$, we build a base vector of the kernel space with support included into $\mathcal{T} \cup \{a, s_i\}$, after observing that $\mathcal{T} \cup \{a, s_i\}$ contains either an even cycle or a kayak paddle similar to those shown in Figure 5. Other examples are shown in Appendix C.1.

Data: A connected surjective-only compatibility graph $G = (V, E)$

Result: A basis \mathcal{B} of the kernel of the incidence matrix A of G

- 1 $\mathcal{T} \leftarrow$ Edges of a spanning tree of G
- 2 $a \leftarrow$ An edge in $E \setminus \mathcal{T}$ such that $\mathcal{T} \cup \{a\}$ contains an odd cycle
- 3 $\mathcal{B} \leftarrow \emptyset$
- 4 **for** $s \in E \setminus (\mathcal{T} \cup \{a\})$ **do**
- 5 Select from $\mathcal{T} \cup \{a, s\}$ an even cycle or a kayak paddle made of two odd cycles
- 6 Weight the selected edges like in Figure 5 (unselected edges have weight 0)
- 7 Add resulting kernel vector to \mathcal{B}
- 8 **return** \mathcal{B}

■ **Algorithm 1** High-level description of the construction of a basis of the kernel of the incidence matrix A of a compatibility graph G . See Appendix C.1 for detailed description.



(a) Selection of \mathcal{T} and a .

(b) Kernel vector derived from s_1 ($KP_{3,3,1}$).

(c) Kernel vector derived from s_2 (\mathcal{C}_4).

■ **Figure 6** A possible execution of Algorithm 1 for the codomino graph.

► **Remark 4.** Equation (17) implies that, given an edge $k \in E$, all solutions to (CE) have the same value along edge k if and only if edge k does not belong to the support of any basis vector. According to Algorithm 1, this is equivalent to saying that edge k belongs neither to an even cycle nor to a kayak paddle. In the diamond graph of Example 4 for instance, the edge $\{2, 3\}$ is the only one that does not belong to the even cycle $1-2-4-3$, and it is indeed the only one with a fixed rate β . In general, if an edge $k \in E$ satisfies this unicity condition, then the matching rate along edge k in a stable matching model (G, λ, Φ) is independent of the policy Φ . Note that there is no straightforward relation between the number of edges with uniquely-defined matching rates and the dimensionality d of the affine space Π .

5.2 Polytope description

We continue to focus exclusively on matching problems (G, λ) that are stabilizable. Let $\Pi_{\geq 0}$ denote the set of solutions to (CE) that have non-negative coordinates, defined as

$$\Pi_{\geq 0} = \Pi \cap \mathbb{R}_{\geq 0}^m = \{\mu \in \mathbb{R}^m : A\mu = \lambda, \mu \geq 0\}. \quad (20)$$

The set $\Pi_{\geq 0}$ is a d -dimensional convex polytope in \mathbb{R}^m , as it is the intersection of a d -dimensional affine space with the positive orthant $\mathbb{R}_{\geq 0}^m$, both of which are convex. This set $\Pi_{\geq 0}$ is neither empty nor degenerated to a dimension lower than d because the matching problem (G, λ) is assumed to be stabilizable, which by Proposition 8 means that Π contains a vector with positive coordinates (i.e., in the interior of the positive orthant). It is bounded because each $\mu \in \Pi_{\geq 0}$ satisfies $0 \leq \mu_k \leq \min_{i \in V_k} (\lambda_i)$ for each $k \in E$.

Equation (20) describes $\Pi_{\geq 0}$ in the edge basis. As $\Pi_{\geq 0}$ is a subset of Π , we can also express its elements in the kernel basis introduced in Section 5.1.1. In the kernel basis, $\Pi_{\geq 0}$ is defined by the vectors whose coordinates belong to

$$\tilde{\Pi}_{\geq 0} = \{\alpha \in \mathbb{R}^d : \mu^\circ + \alpha_1 b_1 + \alpha_2 b_2 + \dots + \alpha_d b_d \geq 0\}. \quad (21)$$

As (20) and (21) basically represent the same polytope up to the change-of-basis formulas of Definition 11, in the remainder, we will use the same notation $\Pi_{\geq 0}$ to describe both sets; the underlying basis will be made clear by our choice of letters (as before, μ for the edge basis and α for the kernel basis).


5.2.1 Vertex characterization

The vertices of a convex polytope, also known as its *corners* or extreme points, are instrumental in optimization, which will be the object of Section 6. Definition 12 provides the formal definitions of vertices, along with faces and facets.


► **Definition 12** (Vertices, faces, and facets; adapted from [42]). *Let Υ denote a convex polytope of dimension $d \in \mathbb{N}_{>0}$. A (non-empty) face of Υ is a non-empty intersection of Υ with a hyperplane such that Υ is included into one of the two halfspaces defined by the hyperplane. A vertex of Υ is a face of dimension 0. Equivalently, a vector $\mu \in \Upsilon$ is a vertex of Υ if, and only if, it cannot be written as a convex combination of points in $\Upsilon \setminus \{\mu\}$. A facet of Υ is a face of dimension $d - 1$.*

Proposition 13 below gives a simple yet powerful characterization of the vertices of $\Pi_{\geq 0}$.

► **Proposition 13.** *Consider a vector $\mu \in \Pi_{\geq 0}$. Let $E^* = \{k \in E : \mu_k > 0\}$ denote the support of the vector μ and $G^* = (V, E^*)$ its support graph. The following statements are equivalent:*



■ **Figure 7** Matching problem (G, λ) , where G is the codomino graph and $\lambda = (2, 3, 3, 2, 3, 3) \in \mathbb{R}^6$.



■ **Figure 8** Matching problem (G, λ) , where G is the codomino graph and $\lambda = (4, 4, 3, 4, 3, 4) \in \mathbb{R}^6$.

- (i) The vector μ is a vertex of $\Pi_{\geq 0}$.
- (ii) The graph G^* is injective.

In particular, if μ is a vertex, we can distinguish two cases depending on the value of $|E^*|$:

1. If $|E^*| = n$ then G^* is bijective.
2. If $|E^*| < n$ then G^* is injective-only.

Proof (borrowed from [14]). See Appendix C.2. ◀

With a slight abuse of notation, we say that a vertex $\mu \in \Pi_{\geq 0}$ is bijective (resp. injective-only) if the support graph G^* of μ is bijective (resp. injective-only).

Figures 7 and 8 shows $\Pi_{\geq 0}$ for the codomino graph. In each figure, the left part displays a generic solution to (CE) on the graph, while the right part shows the polytope $\Pi_{\geq 0}$ and its vertices in the kernel basis. In Figure 7, $\Pi_{\geq 0}$ has five vertices. The vertex $\alpha = (0, 1)$ is bijective, as $|E^*| = |E \setminus \{\{2, 3\}, \{5, 6\}\}| = 6$. One can verify that the remaining four vertices are injective-only (as at least three edge coordinates are null). In Figure 8, $\Pi_{\geq 0}$ has three vertices, all of which are bijective.

For further details, we encourage the reader to consult Appendix C.3, which provides additional results on the relationships between vertex properties and the edge positivity inequalities from Equation (21), along with more detailed examples of polytopes and their vertices.

5.2.2 Probability of bijectivity

Section 6 will show that the bijectivity of vertices is central for optimizing the matching rates achieved by a stable policy. It is thus natural to wonder how frequent (or rare) bijective vertices are. Proposition 14 gives a part of the answer.

► **Proposition 14.** *Let G be a surjective-only graph, and assume that the arrival rate vector λ is drawn according to a positive Lebesgue density over Δ^{n-1} , the standard simplex of \mathbb{R}^n . Then,*

$$\mathbb{P}(\text{All vertices of } \Pi_{\geq 0} \text{ are bijective} | (G, \lambda) \text{ stabilizable}) = 1.$$

Proof. See Appendix C.4. ◀

Proposition 14 suggests that under random arrival rates, the polytope $\Pi_{\geq 0}$ is likely to have only bijective vertices. This result also holds if each coordinate of λ follows a continuous, independent probability distribution, provided that the distribution's projection on Δ^{n-1} has a positive measure and includes at least one vector λ for which (G, λ) is stabilizable. In practice, injective-only vertices may still occur, especially when the coordinates of λ display certain regularities, such as being proportional to node degrees (as in Figure 7). Moreover, injective-only vertices offer unique theoretical challenges worth exploring.


In a sense, Proposition 14 resembles the statement *with probability 1, a square matrix is invertible*: while generally true under suitable conditions, this does not imply that non-invertible matrices can be ignored.

6 Matching rates optimization.

In Section 5, we identified the matching rate vectors μ that *may* be achieved by a stable policy when the compatibility graph G is surjective-only. We now use these results to characterize and optimize the matching rates *actually* achieved by a stable policy. As stated in Section 2.4, given a stabilizable problem (G, λ) and a reward vector $r = (r_1, \dots, r_m) \in \mathbb{R}^m$ giving the immediate reward yielded by a match, our goal is to design a stable policy Φ adapted to G and that maximizes $r^\top \mu(G, \lambda, \Phi)$, the long-run average reward.

► **Remark 5.** *The constraint of stability is important, as it forces $\mu \in \Pi_{\geq 0}$, which can impact reward optimization. Consider for example the matching problem from Figure 9 with associated rewards $(1, 1, 1, 0, 0)$ (matches involving class 4 are not rewarded). One can verify that r is orthogonal to $\ker(A)$, so that all stable policies yield the same reward, which is 5 here. On the other hand, if we accept to never match class-4 items, the resulting policy is obviously unstable (class 4 has unbounded queue length) but can achieve a reward of 6, e.g. by applying the ML policy restricted to the triangle $\{1, 2, 3\}$. Optimizing unstable policies is outside the scope of this paper, but in Section 8.3.2, we briefly investigate how the paradox from Figure 9 can be solved using the formalism of hypergraphs.*

The remainder of this section is structured as follows: Section 6.1 recalls that optimizing linear rewards amounts to reaching a vertex of the polytope. We show that: if the vertex is bijective, the optimal solution is trivially achieved by a stable policy; if the vertex is injective only, no stable policy can achieve the vertex. Section 6.2 focuses on the non-trivial case where the vertex is injective-only. We demonstrate that such a vertex can be approached at the price of an increase in the average queue sizes, and we introduce a family of vertex-approaching policies with a guaranteed trade-off in performance.



■ **Figure 9** A matching problem where all stable policies have the same reward if $r = (1, 1, 1, 0, 0)$.

6.1 Vertex optimality

We first recall a classical result from convex optimization; namely, a convex optimization problem defined on a convex polytope always admits a vertex of the polytope as solution. More specifically, Proposition 15 states that for optimizing a linear reward function, it is sufficient to be able to reach the vertices of $\Pi_{\geq 0}$. Proposition 16 shows that, most of the time, it is also necessary.

► **Proposition 15.** *Let (G, λ) a stabilizable matching problem with a surjective-only compatibility graph G , and let $r = (r_1, \dots, r_m) \in \mathbb{R}^m$ be a vector of rewards associated with the edges of G . Consider the problem of finding a solution to (CE) with non-negative components that maximizes the reward rate $r^\top \mu$, and let $F = \{\mu \in \Pi_{\geq 0} : r^\top \mu = \max_{z \in \Pi_{\geq 0}} r^\top z\}$. Then F is a non-empty face of $\Pi_{\geq 0}$. In particular, there exists a vertex $\mu \in \Pi_{\geq 0}$ that maximizes the reward (i.e., $\mu \in F$).*

Proof. This is a standard result from convex optimization. Since $\Pi_{\geq 0}$ is closed and bounded, there is a maximum r_{\max} among the rewards associated with vectors inside $\Pi_{\geq 0}$. The set F is the intersection of the hyperplane $\{y \in \mathbb{R}^m : r^\top y = r_{\max}\}$ with $\Pi_{\geq 0}$. Therefore, it is a non-empty face of $\Pi_{\geq 0}$. That any non-empty face contains a vertex follows from the lattice structure of the faces of polytopes. ◀

► **Proposition 16.** *Let (G, λ) a stabilizable matching problem with a surjective-only compatibility graph G , and let $r \in \Delta^{m-1}$ be a vector of rewards drawn according to a positive Lebesgue density over Δ^{m-1} , the standard simplex of \mathbb{R}^m . Let $F = \{\mu \in \Pi_{\geq 0} : r^\top \mu = \max_{z \in \Pi_{\geq 0}} r^\top z\}$ be the set of the optimal matching rate vectors for r . Then, with probability 1, $F = \{\mu_r\}$, where μ_r is a vertex of $\Pi_{\geq 0}$.*

Proof. As in Proposition 15, the non-empty face F is the intersection of the hyperplane $\{y \in \mathbb{R}^m : r^\top y = \max_{z \in \Pi_{\geq 0}} r^\top z\}$ with $\Pi_{\geq 0}$. If its dimension is more than 0, it contains a 1-dimensional face of $\Pi_{\geq 0}$ (an *edge*). Let $b \in \mathbb{R}^m \setminus \{0\}$ be a direction of that edge. We have $b^\top r = 0$. The set $\{r \in \Delta^{m-1} : b^\top r = 0\}$, if it exists, has a dimension at most $m - 2$, so the probability that $b^\top r = 0$ is 0. As the number of edges of $\Pi_{\geq 0}$ is finite, we conclude that, with probability 1, F does not contain any edge of $\Pi_{\geq 0}$, which means it is reduced to a single vertex. ◀

Building on Propositions 15 and 16, a natural question is: *Can a given vertex be achieved by a stable policy?* It is answered by Proposition 17.

► **Proposition 17.** *Let (G, λ) be a stabilizable matching problem with a surjective-only compatibility graph G . Let $\Pi_{\mathcal{P}}$ be the set of matching rate vectors achieved by stable policies adapted to (G, λ) , μ a vertex of the polytope $\Pi_{\geq 0}$ associated to (G, λ) , and $E^* \subset E$ the*

support of μ . We distinguish two cases, depending on whether μ is bijective or injective-only (cf Proposition 13).

- (i) If μ is bijective (meaning that $|E^*| = n$), then $\mu \in \Pi_{\mathcal{P}}$. In fact, $\mu = \mu(\Phi_{E^*})$, where Φ_{E^*} is the ML policy adapted to G with a filter on E^* .
- (ii) If μ is injective-only (meaning that $|E^*| < n$), then $\mu \notin \Pi_{\mathcal{P}}$: no stable policy can achieve μ .

Proof. Let $G^* = (V, E^*)$ be the subgraph of G associated with μ , $p = |E^*|$ be the number of positive coordinates of μ , and let A^* denote the $n \times p$ incidence matrix of G^* . As G^* is injective (Proposition 13), the restriction $\tilde{\mu}$ of the vector μ to its positive coordinates is the only solution of the conservation equation $A^* \tilde{z} = \lambda$, of unknown $\tilde{z} \in \mathbb{R}^p$. We now consider the two cases separately:

- (i) If G^* is bijective, Proposition 8 implies that the matching problem (G^*, λ) is stabilized by the ML policy: G^* is surjective and $\tilde{\mu}$ is a solution to the conservation equation $A^* z = \lambda$ with positive coordinates. To prove $\mu \in \Pi_{\mathcal{P}}$, we consider the ML policy with a filter on E^* on the matching problem (G, λ) , denoted Φ_{E^*} . Φ_{E^*} behaves exactly like the greedy ML policy on (G^*, λ) , which is stable with matching rate $\tilde{\mu}$ as we just saw. Hence, the model (G, λ, Φ_{E^*}) is stable, and its matching rate is necessarily equal to μ by injectivity of G^* . This proves that $\mu \in \Pi_{\mathcal{P}}$.
- (ii) If G^* is injective-only, Proposition 7 implies that the matching problem (G^*, λ) is not stabilizable. We prove that $\mu \notin \Pi_{\mathcal{P}}$ by contradiction. Suppose that $\mu \in \Pi_{\mathcal{P}}$, and let Φ be a stable policy on the matching problem (G, λ) such that $\mu(\Phi) = \mu$. Since $\mu_k(\Phi) = 0$ for each $k \in E \setminus E^*$, we know from Remark 1 that Φ never performs a match supported by an edge in $E \setminus E^*$. Hence, Φ also defines a stable policy on the matching problem (G^*, λ) , which contradicts the instability of the matching problem (G^*, λ) . Therefore, $\mu \notin \Pi_{\mathcal{P}}$. \blacktriangleleft

Proposition 17 tells that optimizing linear rewards is easy when the vertex is bijective: we just need to forbid the edges outside the support of the vertex. We now focus on the injective-only case.

6.2 Approaching injective-only vertices

Consider an injective-only vertex μ of $\Pi_{\geq 0}$, and let its support graph be denoted by $G^* = (V, E^*)$, where $E^* = \{k \in E : \mu_k > 0\}$. As shown in Proposition 17, the matching problem (G^*, λ) is not stabilizable. In this section, we present a stronger result: no sequence of matching policies can converge to an injective-only vertex of $\Pi_{\geq 0}$ while maintaining finite queue lengths. On the other hand, we also provide an achievability result. We propose a simple variant of the ML policy and demonstrate that it achieves matching rates that are $O(\epsilon)$ close to μ , with an expected queue length of $O(1/\epsilon)$. Note that the distance to μ and the expected queue length are directly tied to our main scalar performance metrics, namely regret and delay⁸.

⁸ Queue length and delay are related by Little's law (see Section 2.4). Consider a stable policy Φ with matching rate $\mu(\Phi) \in \Pi_{\geq 0}$ and regret $R(\Phi)$. We have $R(\Phi) = r^\top(\mu - \mu(\Phi)) = \cos(\theta) \|r\|_2 \cdot \|\mu - \mu(\Phi)\|_2$, where θ denotes the angle between r and $\mu - \mu(\Phi)$. In particular, $R(\Phi) \leq \|r\|_2 \cdot \|\mu - \mu(\Phi)\|_2$: if you are close to μ , your regret is low. Conversely, if α denotes the minimum possible cosine between r and $\mu - \mu'$ for $\mu' \in \Pi_{\geq 0} \setminus \{\mu\}$ (i.e., $\alpha = \min_{\mu' \in \Pi_{\geq 0} \setminus \{\mu\}} \frac{r^\top(\mu - \mu')}{\|r\|_2 \cdot \|\mu - \mu'\|_2}$), we obtain $R(\Phi) \geq \alpha \|r\|_2 \cdot \|\mu - \mu(\Phi)\|_2$. Therefore, $\alpha \|r\|_2 \cdot \|\mu - \mu(\Phi)\|_2 \leq R(\Phi) \leq \|r\|_2 \cdot \|\mu - \mu(\Phi)\|_2$. Note that if μ is the unique optimal solution (as is often the case; see Proposition 16), then $\alpha > 0$: if you are far from μ , your regret is high.

6.2.1 Impossibility result

The next proposition captures the fundamental trade-off between approaching an injective-only vertex and keeping the queue length short.

► **Proposition 18.** *Consider an injective-only vertex μ of $\Pi_{\geq 0}$. For each policy Φ adapted to G and such that the matching model (G, λ, Φ) is stable, if $\|\mu(\Phi) - \mu\|_1 \leq \epsilon$, then $\mathbb{E}[\|Q\|_2^2] \geq \Omega(1/\epsilon)$.*

Proof. See Appendix D.1. ◀

The above proposition provides a bound of the order $\Omega(1/\epsilon)$ on the second moment of the queue length. We believe such a lower bound is not tight. In particular, we expect an $\Omega(1/\epsilon)$ lower bound on the first moment. Nonetheless, the result demonstrates that the trade-off between approaching an injective-only vertex and keeping the queue lengths small is fundamental. Remind, however, that such trade-off does not exist for bijective vertices, as demonstrated in Proposition 17.

► **Remark 6.** *Appendix D.1 actually proves a more general version of Proposition 18 where we consider matching policies that can form at most M pairs of items at each time step, for some $M \in \mathbb{N}_{>0}$.*

6.2.2 Achievability result

Given an injective-only vertex μ of $\Pi_{\geq 0}$, we now construct a sequence of policies, that we call ϵ -filtering policies, which converge to the vertex μ as $\epsilon \downarrow 0$. The principle of these policies is as follows: when a new item arrives, it is marked with low probability as eligible for all matches in E . For a match outside E^* to be made, the selected pair must contain at least one such marked item. Marks allow some matches to be made outside E^* , which ensures stability. Their low probability ensures that the resulting matching rate vector is close to the target vertex μ .

Formally, given $0 < \epsilon < 1$, we label each incoming arrival by ‘-’ (marked) with probability ϵ and ‘+’ (unmarked) otherwise, independently of everything else. An item with class i and label $\ell \in \{+, -\}$ is said to have type- i^ℓ . We denote the number of items of type i^ℓ at time t by Q_{t,i^ℓ} . Thus, the total number of class- i items at time t is $Q_{t,i} = Q_{t,i^+} + Q_{t,i^-}$. Now, we restrict our attention to policies that do *not* match type- i^+ items to type- j^+ items for all $\{i, j\} \in E \setminus E^*$. The policies that follow this restriction in the original graph G can be viewed as policies operating on the augmented graph $G' = (V', E')$, with $V' = \{i^\ell : i \in V, \ell \in \{+, -\}\}$ and $E' = E^+ \cup E^\pm \cup E^-$, where

$$E^+ = \{\{i^+, j^+\} : \{i, j\} \in E^*\}, \quad E^\pm = \{\{i^+, j^-\} : \{i, j\} \in E\}, \quad E^- = \{\{i^-, j^-\} : \{i, j\} \in E\}.$$

The arrival-rate vector in the augmented graph is given by $\lambda'_\epsilon \in \mathbb{R}_{\geq 0}^{2n}$ with $\lambda'_{\epsilon, i^+} = (1 - \epsilon)\lambda_i$ and $\lambda'_{\epsilon, i^-} = \epsilon\lambda_i$ for each $i \in V$. An ϵ -filtering policy adapted to (G, λ) with filter E^* refers to a policy adapted to (G', λ'_ϵ) . For example, an ϵ -filtering ML policy adapted to (G, λ) with filter E^* mimics the ML policy adapted to (G', λ'_ϵ) . To be more precise, for a given state $q \in \mathbb{R}_{\geq 0}^{2n}$, an incoming item of type i^ℓ is matched with a type in the set $\arg \max_{j^{\tilde{\ell}} \in V' : \{i^\ell, j^{\tilde{\ell}}\} \in E'} q_{j^{\tilde{\ell}}}$.

Note that, for any greedy policy adapted to (G', λ'_ϵ) , most of the arrivals are matched using edges in E^+ , or equivalently E^* , as items are labeled ‘+’ with probability $1 - \epsilon$. As G^* is the support graph of the vertex μ , we can then show that such a policy ensures that the matching rate on $\{i^+, j^+\}$ is $O(\epsilon)$ close to $\mu_{i,j}$ for all $\{i, j\} \in E^*$. We require $\epsilon > 0$, as $\epsilon = 0$ corresponds to operating a greedy policy on (G^*, λ) , and the matching problem (G^*, λ)

is not stabilizable. We show that $\epsilon > 0$ is sufficient to ensure stability as long as (G, λ) is stabilizable (even when (G^*, λ) is not). In particular, whenever $\epsilon > 0$, we use the edges in $E^\pm \cup E^-$, or equivalently, all the edges in E , to match an $\Omega(\epsilon)$ fraction of items that ensures stability and also results in an $O(1/\epsilon)$ bound on the expected number of items queued in the steady state. We make this argument formal in the following proposition for the ϵ -filtering ML policy.

► **Proposition 19.** *Consider a vertex μ of $\Pi_{\geq 0}$ with support graph $G^* = (V, E^*)$, where $E^* = \{k \in E : \mu_k > 0\}$. For each $\epsilon > 0$, let Φ_ϵ denote the ϵ -filtering ML policy adapted to (G, λ) with filter E^* . Then, $(G, \lambda, \Phi_\epsilon)$ is stable and there exist $C_1, C_2 \in \mathbb{R}_{\geq 0}$ and $\epsilon_0 \in (0, 1)$ such that, for all $\epsilon \leq \epsilon_0$, we have*

$$\|\mu(\Phi_\epsilon) - \mu\|_1 \leq C_2\epsilon, \quad \mathbb{E} \left[\sum_{i \in V} Q_i \right] \leq \frac{C_1}{\epsilon}.$$

The first part of the proposition asserts that the matching rate vector under the ϵ -filtering ML policy is $O(\epsilon)$ close to the vertex μ . The second part proves an upper bound of the order $O(1/\epsilon)$ on the expected sum of queue lengths. While Proposition 18 formalized the fundamental trade-off between approaching an injective-only vertex and maintaining small queue lengths, the above proposition demonstrates that such a trade-off is achievable. Now we present a sketch of the proof that highlights the intuition behind the result.

Sketch of the Proof (See Appendix D.2 for the complete proof). We first consider an arbitrary matching problem (G, λ) . Define the CRP gap $\delta(G, \lambda)$ as follows, where CRP stands for Complete Resource Pooling:

$$\delta(G, \lambda) = \min_{\mathcal{I} \in \mathcal{I}} \left\{ \sum_{j \in V(\mathcal{I})} \lambda_j - \sum_{i \in \mathcal{I}} \lambda_i \right\}. \quad (22)$$

Note that, by Proposition 8, the matching problem (G, λ) is stabilizable if and only if $\delta(G, \lambda) > 0$. Intuitively, $\delta(G, \lambda)$ characterizes the minimum slack between the arrival rate of any independent set \mathcal{I} and its neighbors $V(\mathcal{I})$. The arrival rate of the neighbors acts as a service rate for \mathcal{I} , and the CRP gap $\delta(G, \lambda)$ is reminiscent of the heavy-traffic parameter for a single-server queue. We make this intuition rigorous by showing that, under the ML policy, we have

$$\mathbb{E} \left[\sum_{i \in V} Q_i \right] = O \left(\frac{1}{\delta(G, \lambda)} \right). \quad (23)$$

The above result holds for any matching problem (G, λ) operating under the ML policy as long as $\delta(G, \lambda) > 0$ (equivalently, whenever the matching problem (G, λ) is stabilizable). Now, we use the above result to analyze the ϵ -filtering ML policy adapted to (G, λ) with filter E^* . Recalling that this is exactly the ML policy adapted to (G', λ'_ϵ) , we turn our attention to the augmented matching problem (G', λ'_ϵ) and prove the following:

Expected queue length is $O(1/\epsilon)$: We show that $\delta(G', \lambda'_\epsilon) = \Omega(\epsilon)$, and so, the expected queue length is $O(1/\epsilon)$ by (23). The idea is that, while the edges in E^+ may not contribute to a non-zero CRP gap because E^+ mimics E^* , matching $\Theta(\epsilon)$ arrivals using the edges in $E^\pm \cup E^-$ ensures a CRP gap of $\Omega(\epsilon)$ because $E^\pm \cup E^-$ does mimic E . Thus, we have $\delta(G', \lambda'_\epsilon) = \Omega(\epsilon)$, which by (23) implies the expected queue length in steady state is at most $O(1/\epsilon)$.

Matching rate is $O(\epsilon)$ close to μ : We show that we are $O(\epsilon)$ close to the vertex μ by bounding the frequency of “bad” matches, i.e., the matches that do not correspond to the support graph G^* . Such edges are of the form (i^+, j^-) or (i^-, j^+) . We use the conservation equations given by (CE-1) along with the fact that the arrival rate to vertices $\{i^- : i \in V\}$ is $O(\epsilon)$ to prove this statement.

Combining the two steps above completes the proof of the proposition. \blacktriangleleft

► **Remark 7.** In [36], Nazari and Stolyar introduce the *Extended Greedy Primal-Dual (EGPD) policies*, a family of reward-based stable policies designed to achieve (near)-optimal rewards. EGPD provides an alternate solution to approaching injective-only vertices. The main benefit of ϵ -filtering policies is that the trade-off between queue size and reward is proved, while for EGPD it was observed only empirically in [36]. Also, our approach allows us to identify the cases where the trade-off is not necessary (bijective vertices).

7 Practical performance

Let (G, λ) be a stabilizable problem with a surjective-only compatibility graph G , $r \in \mathbb{R}^m$ a vector of rewards, and μ a vertex of $\Pi_{\geq 0}$ that optimizes the average reward $r^\top \mu$. Let $E^* \subsetneq E$ denote the edges of G that support μ . In this section, we use simulations to quantitatively evaluate the performance of various policies in reaching μ . Section 7.1 introduces the policies under consideration, while Section 7.2 outlines the simulation methodology. Sections 7.3 and 7.4 analyze the injective-only and bijective cases, respectively. Finally, Section 7.5 presents a discussion of the results.

7.1 Considered policies

For ease of display, we use the same notation Φ to designate all policies; which policy is considered will be made clear by the letter in subscript (e.g., ϵ for the ϵ -filtering, β for EGPD). More specifically, we consider the following policies:

- **Filtering ML policy:** The ML policy restricted to E^* , denoted by Φ_{E^*} , is stable only if μ is bijective (see Proposition 17). In this case, it is optimal (i.e., incurs no regret), and its delay can be used as a benchmark for evaluating other policies.
- **ϵ -filtering ML policy:** The ϵ -filtering ML policy with filter E^* , denoted by Φ_ϵ for $\epsilon \in (0, 1)$, is stable as long as $\epsilon > 0$ and converges to μ as $O(\epsilon)$, with a bounded delay in $O(1/\epsilon)$ (Proposition 19).
- **k -filtering ML policy:** The k -filtering ML policy with filter E^* , denoted by Φ_k for $k \in \mathbb{N}$, applies the filter E^* depending on the state: it applies Φ_{E^*} if the length of the longest queue is less than k , otherwise it applies the ML policy without filter.
- **EGPD policies:** The Extended Greedy Primal-Dual (EGPD) algorithm, denoted by Φ_β for $\beta \in (0, +\infty)$, is the state-of-the-art policy for minimizing regret. Introduced in [36], the algorithm is based on two ingredients: a virtual queue of matches to be performed in the future, and a score function that is mostly driven by r , with the addition of a small term proportional to β to account for queue lengths and stabilize the system. An informal definition of EGPD is as follows. After each arrival:
 - (i) The system updates scores and selects the highest-scoring potential match, regardless of immediate feasibility. If scores are too low, no edge is selected.
 - (ii) If an edge is selected, its label is added to the virtual queue.

- (iii) If a queued label becomes feasible, meaning all required items are available, the match is executed, and the label is removed. If multiple labels are feasible, the oldest one is chosen (FCFM selection).

Compared to our approach, EGPD has two main differences: first, it is oblivious to the structure of the polytope $\Pi_{\geq 0}$; second, it is designed to work for arbitrary hypergraphs, whereas we focus on simple graphs. These two aspects will be discussed in Sections 7.5 and 8.3, respectively.

► **Remark 8.** *The motivation for introducing Φ_k is our belief that it should outperform Φ_ϵ in practical applications. Specifically, while Φ_ϵ consistently permits a small proportion of items to be matched outside E^* , Φ_k relaxes the constraint on allowed edges only when necessary—namely, in cases where delay is significantly impacted. We provide in Appendix E.1 elements to substantiate the conjecture that Φ_k offers theoretical guarantees similar to those of Φ_ϵ .*

7.2 Methodology

For measuring the performance criteria presented in Section 2.4, we used the Python package *Stochastic Matching* of [34]. Among other things, the package leverages our theoretical results to tell if a matching problem is stabilizable, compute an optimal vertex for a given reward, and measure the regret and delay of policies through simulations. A matching model is evaluated by simulating 10^{10} arrivals and measuring regret and delay. For policies that depend on a parameter, we choose the following parameter values:

- $k \in [1, 2^{11}]$ for k -filtering;
- $\epsilon \in [10^{-4}, 10^{-1}]$ for ϵ -filtering;
- $\beta \in [10^{-3}, 10]$ for EGPD.

Unless otherwise stated, our evaluation protocol is as follows: after choosing G and λ , we select a vertex μ of $\Pi_{\geq 0}$. Then we choose a reward vector r such that μ is optimal, using one of the following methods:

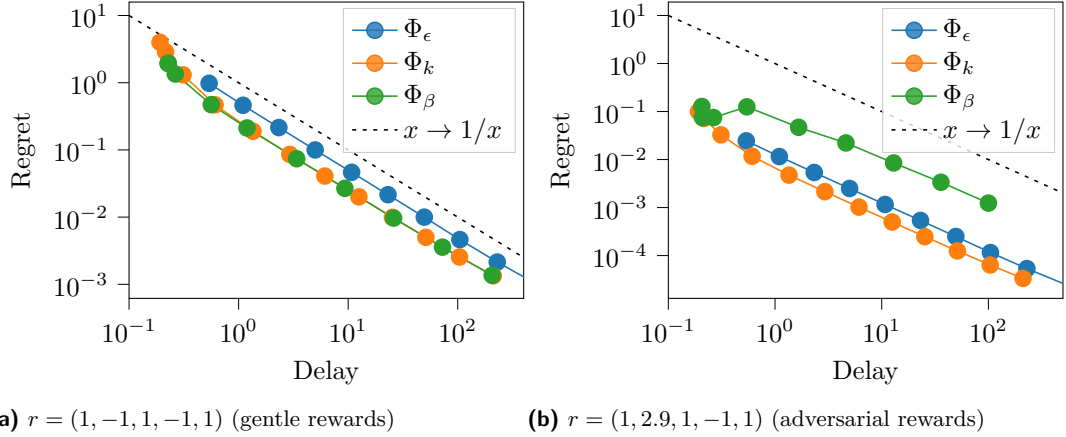
- **Gentle rewards:** An edge e_k receives a reward of 1 if $\mu_k > 0$ and -1 otherwise. This clearly highlights the support of μ .
- **Adversarial rewards:** While ensuring that μ remains the unique optimum for r , we manually adjust the weights to make it more difficult to identify. This is typically achieved by assigning a large reward to an edge outside the support of μ , or by reducing the reward on an edge within the support.

Note that changing the rewards primarily affects the behavior of the EGPD policy. Indeed, since the other policies rely solely on knowledge of the support E^* of μ to operate, the choice of reward vector r does not impact the matching-rate vector they achieve (although it still impacts the regret).

7.3 Approaching an injective-only vertex

To begin, we consider the diamond graph from Figure 1 with degree-proportional arrival rates, i.e., $\lambda = (2, 3, 3, 2)$. Our goal is to approach the injective-only vertex $\mu = (2, 0, 1, 0, 2)$, which disables edges $\{1, 3\}$ and $\{2, 4\}$. The results are presented in Figure 10.

When the rewards associated with μ are gentle (Figure 10a), all policies clearly demonstrate the $(\epsilon, 1/\epsilon)$ trade-off between delay and regret, as predicted by Proposition 19. Specifically, Φ_k and Φ_β exhibit very similar performance, while Φ_ϵ lags slightly behind.

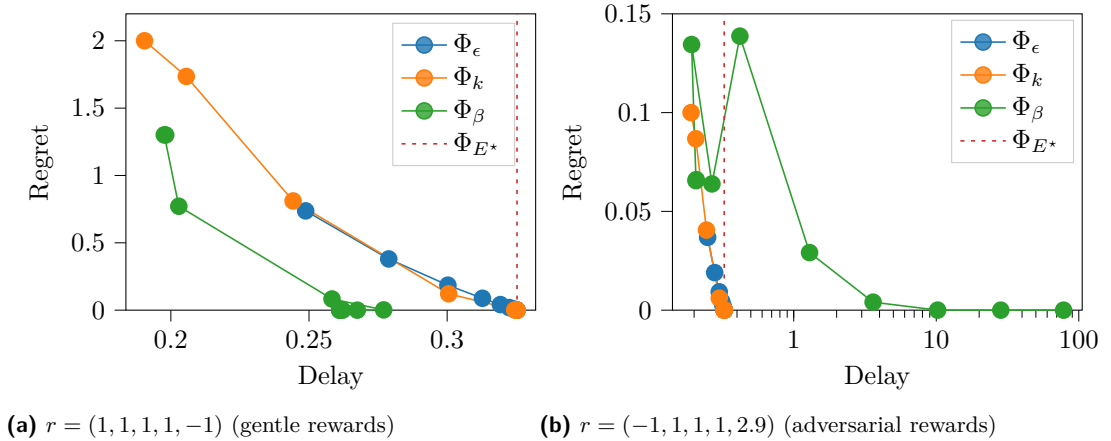


■ **Figure 10** G : diamond graph; λ : $(2, 3, 3, 2)$ (degree-proportional); μ : $(2, 0, 1, 0, 2)$ (injective-only).

When the adversarial reward vector $r = (1, 2.9, 1, -1, 1)$ is considered (Figure 10b), both Φ_ϵ and Φ_k remain largely unaffected, as expected. However, the performance of Φ_β significantly deteriorates. The performance gap between Φ_k and Φ_β exceeds a factor of 10 — meaning that Φ_β incurs ten times more regret than Φ_k for the same target delay, or, equivalently, Φ_β requires a delay ten times longer to achieve the same regret as Φ_k .

7.4 Approaching a bijective vertex

Still focusing on the diamond graph, we now take $\lambda = (4, 4, 4, 2)$ and examine the vertex $\mu = (1, 3, 1, 2, 0)$, which disables edge $\{3, 4\}$. Since μ is bijective, we know from Proposition 17 that Φ_{E^*} is optimal in terms of regret. However, it is interesting to assess how vertex-approaching policies perform in this case. The results are presented in Figure 11, where the dotted line shows the delay of Φ_{E^*} .



■ **Figure 11** G : diamond graph; λ : $(4, 4, 4, 2)$; μ : $(1, 3, 1, 2, 0)$ (bijective).

When the rewards associated with μ are gentle (Figure 11a), both Φ_ϵ and Φ_k converge to Φ_{E^*} : their regret approaches zero, and their delay matches that of Φ_{E^*} . The performance of Φ_β is superior, as it achieves zero regret with a shorter delay than Φ_{E^*} .

The situation changes drastically when the adversarial reward vector $r = (-1, 1, 1, 1, 2.9)$ is considered (Figure 11b). Although the regret approaches zero for all policies, the delay behaves differently: the delays of Φ_ϵ and Φ_k remain contained and are bounded by that of Φ_{E^*} , while the delay of Φ_β appears to grow as $1/\beta$ and becomes virtually unbounded.

7.5 Discussion

The observations made in Sections 7.3 and 7.4 are not specific to the diamond graph. Appendix E.2 displays other examples (up to 100 nodes) that are qualitatively similar. From them we conclude that:

- ϵ -filtering ML is a robust vertex-approaching policy that works well on injective-only and bijective vertices. It has the advantage of offering theoretical guarantees, but its numerical performance can often be bested by other (state-dependent) policies.
- EGPD has the advantage that it does not require the knowledge of the vertex μ that corresponds to the considered arrival-rate and reward vectors. Empirically, its performance is optimal when the rewards and the support E^* of μ are aligned enough, but it degrades strongly when it not the case.
- While the theoretical guarantees of k -filtering ML are only conjectured for now, simulations indicate that it offers the best compromise between performance and robustness. It is the policy we recommend if the arrival rates are steady enough for the vertex to be computed.

8 Extensions of our results

The central focus of this article is the linear optimization of the matching rate vector achieved by a stable policy adapted to a matching problem based on a simple graph. Yet, our approach applies to a broader scope. In this section, we discuss some of these extensions. Section 8.1 considers the case of a non-linear reward function. Section 8.2 revisits the results above when only *greedy* policies are considered. Section 8.3 briefly describes how to extend our results on hypergraphs.

8.1 Non-linear optimization

Proposition 17 introduced $\Pi_{\mathcal{P}}$, the set of matching rates achievable by a stable policy, and compared it to the vertices of $\Pi_{\geq 0}$, which contain the optimal solutions for linear reward optimization.

However, when considering a non-linear reward function $r(\mu)$, an optimal solution can lie anywhere within $\Pi_{\geq 0}$. For example, one might want a stable policy to be as close as possible to a target matching rate vector μ_0 and use the reward function $r(\mu) = -\|\mu - \mu_0\|_1$. Ideally, we would like $\Pi_{\mathcal{P}}$ to be as close to $\Pi_{\geq 0}$ as possible. The following proposition demonstrates that this is indeed the case.

► **Proposition 20.** *Let $\Pi_{\mathcal{P}}$ be the set of matching rate vectors achievable by a stable policy adapted to a given stabilizable matching problem (G, λ) . Then, $\Pi_{\mathcal{P}}$ is convex. Furthermore, any positive solution to the conservation equation (CE) can be obtained by a stable policy, that is,*

$$\Pi_{>0} \subseteq \Pi_{\mathcal{P}} \subseteq \Pi_{\geq 0}, \text{ where } \Pi_{>0} = \{\mu \in \mathbb{R}_{>0}^m : A\mu = \lambda\}.$$

In the particular case where all vertices of $\Pi_{\geq 0}$ are bijective, we have $\Pi_{\mathcal{P}} = \Pi_{\geq 0}$.

Proof. Convexity of $\Pi_{\mathcal{P}}$ is shown in Appendix F.1. The inclusion $\Pi_{\mathcal{P}} \subseteq \Pi_{\geq 0}$ follows from the discussion in Section 3.1: any stable policy yields a matching rate vector that satisfies (CE) with non-negative coordinates. We now prove that $\Pi_{>0} \subseteq \Pi_{\mathcal{P}}$. Since $\Pi_{\mathcal{P}}$ is convex, its closure is also convex. By Proposition 19 and the convexity of $\Pi_{\mathcal{P}}$, the closure of $\Pi_{\mathcal{P}}$ contains all vertices of $\Pi_{\geq 0}$, and, by convexity, encompasses $\Pi_{\geq 0}$ as well. As a result, by virtue of convexity, $\Pi_{\mathcal{P}}$ contains the interior⁹ of its own closure, which includes $\Pi_{>0}$, the interior of $\Pi_{\geq 0}$. In the particular case where all vertices are bijective, they all belong to $\Pi_{\mathcal{P}}$, leading to the equality $\Pi_{\mathcal{P}} = \Pi_{\geq 0}$. ◀

Proposition 20 is essentially an existence result, as the policy constructed to show convexity of $\Pi_{\mathcal{P}}$ in Appendix F.1 may be both hard to implement (as it requires computing the mean return time to a particular state of two Markov chains) and undesirable in practice (as the matching rate will have high variance). An interesting follow-up question is thus to explore practical methods to reach an arbitrary vector of $\Pi_{\mathcal{P}}$.

8.2 Greedy policies

Greedy policies are appealing candidates for controlling matching problems, and they have been extensively studied in the literature. Some, such as ML and FCFM, are direct to implement and stabilize all stabilizable matching problems. Moreover, greedy policies may be more socially acceptable, particularly in scenarios involving human participants. However, as we will demonstrate in this section, greedy policies are generally not well-suited for optimizing matching rate vectors. They can never achieve a vertex, and examples suggest that in most cases, they are unable to even approach one.

In what follows, consider a stabilizable problem (G, λ) with a surjective-only graph G , and let Π_G denote the set of matching rate vectors achieved by stable *greedy* policies adapted to the problem (G, λ) . As in Proposition 20, let $\Pi_{>0} = \{\mu \in \mathbb{R}_{>0}^m : A\mu = \lambda\}$ denote the set of the positive solutions to the conservation equation (CE), i.e., those with positive coordinates, which form the interior of $\Pi_{\geq 0}$.

8.2.1 Convexity and impossibility result

We first demonstrate that, unlike stable policies in general, stable *greedy* policies can never reach the boundary of the convex polytope $\Pi_{\geq 0}$, irrespective of whether the vertices of $\Pi_{\geq 0}$ are bijective or injective-only.

► **Proposition 21.** *Let Π_G denote the set of matching rate vectors achievable by a stable greedy policy for a given stabilizable matching problem (G, λ) , where G is surjective-only. Then, Π_G forms a non-empty convex subset of $\Pi_{>0}$.*

Proof. The set Π_G is non-empty because, as recalled in Section 2.4, the greedy ML and FCFM policies are stable. Its convexity is shown in Appendix F.1. We now prove that $\Pi_G \subseteq \Pi_{>0}$. Consider a stable greedy policy Φ and let μ denote the matching-rate vector in the model (G, λ, Φ) . Consider an edge $e_k = \{i, j\}$. Since the policy Φ is greedy, two items of classes i and j are always matched if the following sequence of events occurs: the system is in the empty state \emptyset , then a class- i item arrives, and then a class- j item arrives. Let

⁹ Here, the notion of *interior* refers to the canonical \mathbb{R}^d topology of the d -dimensional affine space of solutions to the conservation equation.

p_\emptyset denote the stationary probability that the model (G, λ, Φ) is in the empty state¹⁰ \emptyset . We know that $p_\emptyset > 0$ because the model is stable, and the previous remark implies that $\mu_k \geq p_\emptyset \lambda_i \lambda_j / (\sum_{\ell \in V} \lambda_\ell)^2 > 0$. Since this is true for each edge $e_k \in E$ and each $\mu \in \Pi_G$, we conclude that $\Pi_G \subseteq \Pi_{>0}$. ◀

Proposition 21 has the following consequence regarding greedy policies and linear optimization.

► **Corollary 22.** *Let $r_{\max} = \max_{\mu \in \Pi_{\geq 0}} r^\top \mu$ be the optimal reward in the linear optimization problem defined by the reward vector r in Proposition 15. One of the following must hold:*

- (i) *All stable policies (greedy or not) are optimal, i.e., $r^\top \mu = r_{\max}$ for each $\mu \in \Pi_G$.*
- (ii) *All stable greedy policies are suboptimal, i.e., $r^\top \mu < r_{\max}$ for each $\mu \in \Pi_G$.*

Proof. We know from Proposition 15 that the set of vectors $\mu \in \Pi_{\geq 0}$ maximizing the reward forms a non-empty face F of $\Pi_{\geq 0}$. If $F = \Pi_{\geq 0}$, then we are in case i, meaning all $\mu \in \Pi_G \subseteq \Pi_{\geq 0}$ are optimal. This occurs when the vector r is orthogonal to $\ker(A)$ ¹¹. Otherwise, by the lattice structure of polytope faces, F is contained in a facet of $\Pi_{\geq 0}$, implying there is at least one edge $k \in E$ where the k -th coordinate of all vectors in F is zero. As stable greedy policies produce matching rate vectors with all positive coordinates, no greedy policy can be optimal in this case. ◀

8.2.2 Achievability results

We now explore the relationship $\Pi_G \subseteq \Pi_{>0}$ through several examples. In particular, Propositions 23 and 24 illustrate situations where Π_G is a strict subset of $\Pi_{>0}$, suggesting that the greedy constraint imposes significant limitations on the set of achievable matching rates. However, this is not universally true, as Conjecture 25 provides a (carefully chosen) example where $\Pi_G = \Pi_{>0}$.

► **Proposition 23.** *Let (K_n, λ) be a stabilizable matching problem, where K_n is the complete graph with $n \geq 3$ nodes. All greedy policies adapted to K_n are stable and yield the same matching-rate vector, denoted μ_G . In particular, we have $\Pi_G = \{\mu_G\} \subsetneq \Pi_{>0}$ whenever $n \geq 4$.*

Proof. See Appendix F.2.1 for proof and discussion. ◀

► **Proposition 24.** *Let (D, λ) be a stabilizable matching problem, where D is the diamond graph from Figure 3d. All greedy policies adapted to D are stable. In kernel coordinates, there exist $\alpha_- < \alpha_+$ such that $\Pi_G = [\alpha_-, \alpha_+] \subsetneq \Pi_{>0}$.*


Proof. See Appendix F.2.2 for proof, discussion, and numerical results. ◀

► **Conjecture 25.** *Let (G, λ) be the stabilizable matching problem depicted in Figure 12 (the Fish matching problem). For this problem, we have $\Pi_G = (-1/2, 1/2) = \Pi_{>0}$.*

Proof. See Appendix F.2.3 for sketch of proof, discussion, and numerical results. ◀

¹⁰ As mentioned in Appendix A, we assume that there exists a unique state $s \in \mathcal{S}$ such that $|s| = 0$. This state is called the *empty state* and denoted by \emptyset . This assumption guarantees that the intuitive notion of system stability is captured by the positive recurrence of the Markov chain describing the evolution of the system state. If the policy is queue-based (so that $\mathcal{S} = \mathcal{Q}$), this empty state is simply the n -dimensional zero vector.

¹¹ An example of this is shown in Figure 9, or when all coordinates of r are equal, making all edges equivalent.



■ **Figure 12** Generic solution to (CE) in the Fish matching ($\lambda = (4, 4, 3, 2, 3, 2)$).

8.3 Hypergraphs

A hypergraph $G = (V, E)$ consists of a set of nodes V and a set of hyperedges E , where each hyperedge is a multi-subset (i.e. potentially allowing redundancy) of V of arbitrary size, starting from one. This contrasts with simple graphs, where each edge is a pair of distinct nodes. Stochastic matching models can be extended from simple graphs to hypergraphs in the following natural manner: when a matching decision is made, items corresponding to the selected hyperedge are removed from the system.

A full extension of our results to hypergraphs is beyond the scope of this paper, but the key ideas are:

- The incidence matrix A for hypergraphs is defined similarly to simple graphs, with $a_{i,k}$ indicating whether node i belongs to hyperedge k (or counting occurrences in case of multiplicity). We believe that the connection between properties of A and stability still holds, particularly Proposition 7. However, some graph-specific relations, such as the connection between instability and bipartite structures or the description of $\ker(A)$ in terms of cycles and kayak paddles, do not generalize easily to hypergraphs.
- We believe that Proposition 8 remains valid. Specifically, the stability condition iii can still be checked by computing $A^{-1}\lambda$ for bijective hypergraphs, or solving (10) for surjective-only hypergraphs.
- The policies discussed in Sections 6 and 7 for optimizing the matching-rate vector, namely Φ_{E^*} , Φ_ϵ , and Φ_k , rely on the ML policy and its maximal stability in simple graphs. This property does not extend to hypergraphs¹², making it necessary to adopt a different maximally stable policy. As [36] prove that Φ_β is maximally stable for each $\beta \in (0, +\infty)$, one can use its version with null rewards as a substitute for ML. We refer to this unbiased version as Φ_U .


We now illustrate this adaptation through numerical results on a few example cases. We consider the policies from Section 7.1, adapted to hypergraphs as necessary, and we follow the simulation methodology outlined in Section 7.2. These cases will highlight how the proposed adaptations perform in the context of hypergraphs.

8.3.1 Original example from [36]

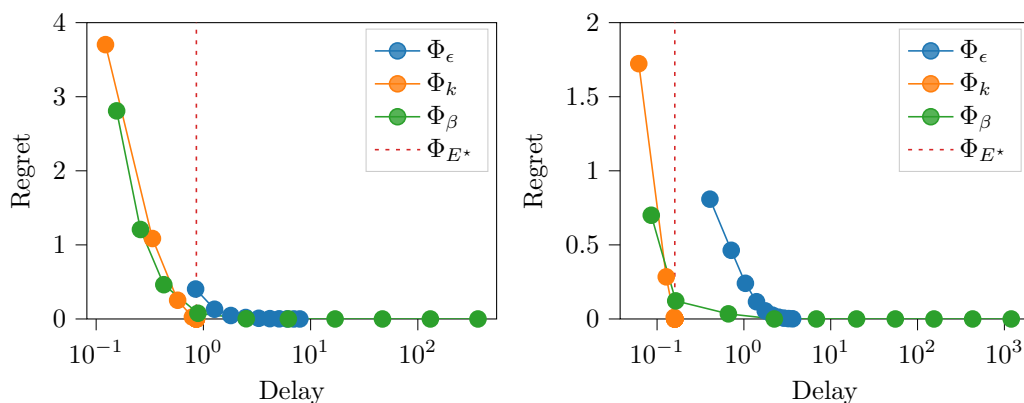
[36] considered the hypergraph shown in Figure 13 for the numerical evaluation of their policy Φ_β . This hypergraph consists of four nodes, each with a self-loop (1-edge), two regular 2-edges, and one 3-edge. It can be verified that this hypergraph is surjective-only, and the solution space to (CE) has dimension 3.

The authors consider one unique reward vector (shown in Figure 13) and two distinct arrival-rate vectors: $\lambda = (1.2, 1.5, 2, 0.8)$ and $\lambda = (1.8, 0.8, 1.4, 1)$. Both vectors are stabi-

¹²In hypergraphs, stabilizable matching problems exist where no greedy policy is stable.



■ **Figure 13** Hypergraph example studied in [36]. The rewards associated with each hyperedge are indicated.



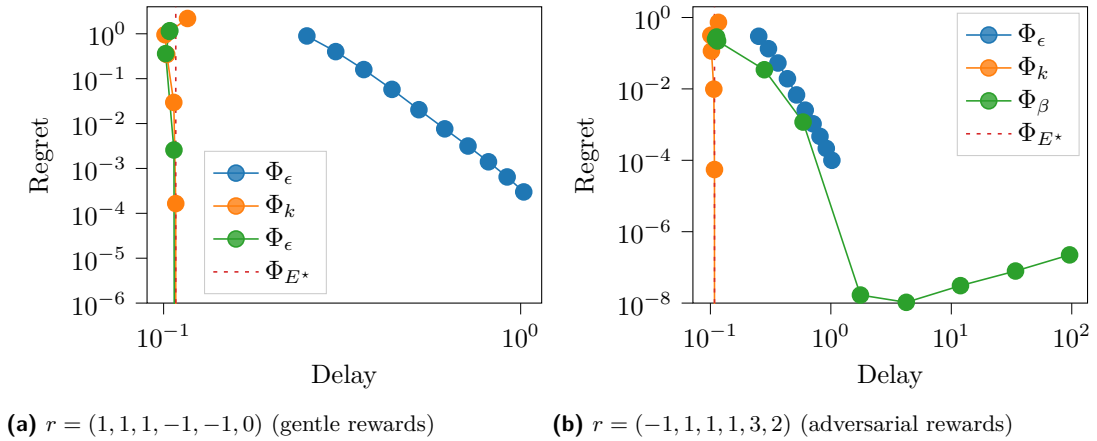
■ **Figure 14** Reaching the vertex corresponding to $r = (-1, -1, 1, 2, 5, 4, 7)$ in the hypergraph from Figure 13.

lizable. The first generates a polytope $\Pi_{\geq 0}$ with 8 vertices, while the second generates a polytope with 4 vertices. All vertices are bijective, meaning that the reward can always be optimized using a filtering policy Φ_{E^*} adapted to the optimal vertex. The performance of the considered policies under this setting is shown in Figure 14. The most noticeable findings are:

- All parameterized policies manage to achieve optimal rewards when the parameter is aggressive enough.
- As noted in [36], the delay experienced by Φ_{β} continues to grow as $1/\beta$, even after rewards have been optimized.
- The delay experienced by Φ_{ϵ} appears to remain bounded, yet above that of Φ_{E^*} . This differs from the convergence expected (as in Section 7.4) and may be due to interactions between the expanded graph, on which Φ_{ϵ} dispatches arrivals in a unbalanced way, and the virtual queue reservation mechanism used by Φ_U .
- The delay of Φ_k converges to that of Φ_{E^*} , as observed in Section 7.4.

8.3.2 Unstable policies

Figure 9 in Remark 5 introduced a toy example demonstrating that imposing stability can limit the achievable reward. With the formalism of hypergraphs, such issues can be resolved by adding self-loops to absorb unmatched items. In the toy example of Figure 9, adding a self-loop to class 4 creates a third vertex in the polytope, which zeroes edges $\{2, 4\}$ and $\{3, 4\}$. This vertex is bijective, so Φ_{E^*} is stable without regret.



■ **Figure 15** Optimizing rewards with a policy that tolerates abandonment.

The performance of the considered policies under this setting is shown in Figure 15. The most noteworthy findings are:

- All parameterized policies have regret that diminishes to zero as the parameter becomes aggressive enough.
- The delay of Φ_β converges to that of Φ_{E^*} when adapted rewards are considered (Figure 15a). However, for adversarial rewards (Figure 15b), the delay of Φ_β grows as $1/\beta$.
- The performance of Φ_ϵ is lower than that of the other policies.
- The delay of Φ_k converges to that of Φ_{E^*} .

References

- 1 Ivo Adan, Ana Bušić, Jean Mairesse, and Gideon Weiss. Reversibility and Further Properties of FCFS Infinite Bipartite Matching. *Mathematics of Operations Research*, 43(2):598–621, December 2017. Publisher: INFORMS. doi:10.1287/moor.2017.0874.
- 2 Ivo Adan and Gideon Weiss. Exact FCFS matching rates for two infinite multitype sequences. *Operations Research*, 60(2):475–489, 04 2012. Publisher: INFORMS. doi:10.1287/opre.1110.1027.
- 3 Angelos Aveklouris, Levi DeValve, and Amy R. Ward. Matching Impatient and Heterogeneous Demand and Supply. *arXiv:2102.02710 [cs, eess, math]*, December 2021. arXiv: 2102.02710. URL: <http://arxiv.org/abs/2102.02710>.
- 4 Egon Balas and William R. Pulleyblank. The perfectly matchable subgraph polytope of a bipartite graph. *Networks*, 13(4):495–516, 1983. doi:10.1002/net.3230130405.
- 5 Jocelyn Begot, Irène Marcovici, and Pascal Moyal. Stability regions of systems with compatibilities and ubiquitous measures on graphs. *Queueing Systems*, 103(3):275–312, 2023. doi:10.1007/s11134-023-09872-0.
- 6 Jocelyn Begot, Irène Marcovici, Pascal Moyal, and Youssef Rahme. A general stochastic matching model on multigraphs. *arXiv:2011.05169 [math]*, November 2020. arXiv: 2011.05169. URL: <http://arxiv.org/abs/2011.05169>.
- 7 Péter Biró, Joris van de Klundert, David Manlove, William Pettersson, Tommy Andersson, Lisa Burnapp, Pavel Chromy, Pablo Delgado, Piotr Dworczak, Bernadette Haase, Aline Hemke, Rachel Johnson, Xenia Klimentova, Dirk Kuypers, Alessandro Nanni Costa, Bart Smeulders, Frits Spieksma, María O. Valentín, and Ana Viana. Modelling and optimisation in european kidney exchange programmes. *European Journal of Operational Research*, 291(2):447–456, 2021. doi:10.1016/j.ejor.2019.09.006.

- 8 Pierre Bremaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Texts in Applied Mathematics. Springer-Verlag, New York, 1999. URL: <https://www.springer.com/gp/book/9780387985091>, doi:10.1007/978-1-4757-3124-8.
- 9 Ana Bušić, Varun Gupta, and Jean Mairesse. Stability of the Bipartite Matching Model. *Advances in Applied Probability*, 45(2):351–378, June 2013. Publisher: Cambridge University Press. doi:10.1239/aap/1370870122.
- 10 Ana Bušić and Sean Meyn. Approximate optimality with bounded regret in dynamic matching models. *arXiv:1411.1044 [cs, math]*, June 2016. arXiv: 1411.1044. URL: <http://arxiv.org/abs/1411.1044>.
- 11 Arnaud Cadas, Ana Bušić, and Josu Doncel. Optimal control of dynamic bipartite matching models. In *Proceedings of the 12th EAI International Conference on Performance Evaluation Methodologies and Tools, VALUETOOLS 2019, Universitat de les Illes Balears, Palma de Mallorca, Spain, March 12-15, 2019*, pages 39–46. ACM, 2019. doi:10.1145/3306309.3306317.
- 12 Arnaud Cadas, Josu Doncel, Jean-Michel Fourneau, and Ana Bušić. Flexibility can hurt dynamic matching system performance. *arXiv:2009.10009 [cs, math]*, September 2020. arXiv: 2009.10009. URL: <http://arxiv.org/abs/2009.10009>.
- 13 René Caldentey, Edward H. Kaplan, and Gideon Weiss. FCFS infinite bipartite matching of servers and customers. *Advances in Applied Probability*, 41(3):695–730, 09 2009. Publisher: Cambridge University Press. doi:10.1239/aap/1253281061.
- 14 Céline Comte and Jaap Storm. Global and local stability in parallel service systems with redundancy, 2022. Unpublished.
- 15 Céline Comte. Stochastic non-bipartite matching models and order-independent loss queues. *Stochastic Models*, 38(1):1–36, January 2022. doi:10.1080/15326349.2021.1962352.
- 16 Céline Comte and Jan-Pieter Dorsman. Performance Evaluation of Stochastic Bipartite Matching Models. In *Performance Engineering and Stochastic Modeling*, Lecture Notes in Computer Science, pages 425–440. Springer International Publishing, 2021. doi:10.1007/978-3-030-91825-5_26.
- 17 Dragoš Cvetkovic, Peter Rowlinson, and Slobodan Simic. *Spectral Generalizations of Line Graphs: On Graphs with Least Eigenvalue -2*. London Mathematical Society Lecture Note Series. Cambridge University Press, 2004. doi:10.1017/CB09780511751752.
- 18 J. G. Dai. On positive harris recurrence of multiclass queueing networks: A unified approach via fluid limit models. *The Annals of Applied Probability*, 5(1):49–77, 1995. Publisher: Institute of Mathematical Statistics. doi:10.1214/aop/1177004828.
- 19 Reinhard Diestel. *Graph Theory*. Springer Publishing Company, Incorporated, 5th edition, August 2017.
- 20 Michael Doob. An interrelation between line graphs, eigenvalues, and matroids. *Journal of Combinatorial Theory, Series B*, 15(1):40–50, 08 1973. doi:10.1016/0095-8956(73)90030-0.
- 21 David Gregory. Hall conditions for edge-weighted bipartite graphs. Technical report, Queen’s University, 2010.
- 22 Varun Gupta. Greedy algorithm for multiway matching with bounded regret. *Operations Research*, 72(3):1139–1155, 2024.
- 23 Itai Gurvich and Amy Ward. On the Dynamic Control of Matching Queues. *Stochastic Systems*, 4(2):479–523, October 2014. Publisher: INFORMS. doi:10.1287/13-SSY097.
- 24 B. Hajek. *Random Processes for Engineers*. Cambridge university press, 2015.
- 25 Matthieu Jonckheere, Pascal Moyal, Claudia Ramírez, and Nahuel Soprano-Loto. Generalized max-weight policies in stochastic matching. *arXiv:2011.04535 [math]*, November 2020. arXiv: 2011.04535. URL: <http://arxiv.org/abs/2011.04535>.
- 26 N. Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, STOC ’84, pages 302–311, New York, NY, USA, December 1984. Association for Computing Machinery. doi:10.1145/800057.808695.
- 27 F. P. Kelly. *Reversibility and Stochastic Networks*. Cambridge University Press, 2011.

- 28 Süleyman Kerimov, Itai Ashlagi, and Itai Gurvich. On the optimality of greedy policies in dynamic matching. *Operations Research*, 2023.
- 29 Süleyman Kerimov, Itai Ashlagi, and Itai Gurvich. Dynamic matching: Characterizing and achieving constant regret. *Management Science*, 70(5):2799–2822, 2024.
- 30 John D. C. Little. A proof for the queuing formula: $l = \lambda w$. *Operations Research*, 9(3):383–387, 1961. doi:10.1287/opre.9.3.383.
- 31 John D. C. Little. Little’s law as viewed on its 50th anniversary. *Operations Research*, 59(3):536–549, 2011. doi:10.1287/opre.1110.0940.
- 32 László Lovász and M. D. Plummer. *Matching Theory*. American Mathematical Soc., 2009.
- 33 Jean Mairesse and Pascal Moyal. Stability of the stochastic matching model. *Journal of Applied Probability*, 53(4):1064–1077, 12 2016. Publisher: Cambridge University Press. doi:10.1017/jpr.2016.65.
- 34 Fabien Mathieu. Stochastic Matching, 2022. URL: https://balouf.github.io/stochastic_matching/index.html.
- 35 Pascal Moyal, Ana Bušić, and Jean Mairesse. A product form for the general stochastic matching model. *Journal of Applied Probability*, 58(2):449–468, 06 2021. doi:10.1017/jpr.2020.100.
- 36 Mohammadreza Nazari and Alexander L. Stolyar. Reward maximization in general dynamic matching systems. *Queueing Systems*, 91(1):143–170, February 2019. doi:10.1007/s11134-018-9593-y.
- 37 J. R. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 1997. doi:10.1017/CB09780511810633.
- 38 Youssef Rahme and Pascal Moyal. A stochastic matching model on hypergraphs. *Advances in Applied Probability*, 53(4):951–980, December 2021. doi:10.1017/apr.2021.8.
- 39 R. Srikant and Lei Ying. *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*. Cambridge University Press, New York, NY, USA, 2014.
- 40 Thirupathaiiah Vasantam and Don Towsley. A throughput optimal scheduling policy for a quantum switch. In *Quantum Computing, Communication, and Simulation II*, volume 12015, pages 14–23. SPIE, 2022. doi:10.1117/12.2616950.
- 41 Yehua Wei, Jiaming Xu, and Sophie H Yu. Constant regret primal-dual policy for multi-way dynamic matching. In *Abstract Proceedings of the 2023 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pages 79–80, 2023.
- 42 Günter M. Ziegler. *Lectures on polytopes*. Springer-Verlag, New York, 1995. URL: http://www.worldcat.org/search?qt=worldcat_org_all&q=9780387943657.
- 43 Martin Zubeldia, Prakirt R. Jhunjhunwala, and Siva Theja Maguluri. Matching queues with abandonments in quantum switches: Stability and throughput analysis, 2022. URL: <http://arxiv.org/abs/2209.12324>, arXiv:2209.12324[cs], doi:10.48550/arXiv.2209.12324.

A Supplementary material of Section 2.3.4 (Other policies).

Our results are not limited to deterministic size-based policies but apply to a broader family of policies that are either random or require a more complex state descriptor, or both. We now introduce this family of policies, with the goal of being as general as possible.

A.1 Extended definition.

Under this more general definition, the match-maker makes decisions based not only on the vector of queue sizes, but also (possibly) on additional information that is captured by the system state. The state space is a couple $(\mathcal{S}, |\cdot|)$, where \mathcal{S} is a countably infinite set and $|\cdot| : \mathcal{S} \rightarrow \mathbb{N}^n$ is a function that maps any state $s \in \mathcal{S}$ to the vector giving the number of unmatched items of each class in that state, denoted by $|s| = (|s|_1, |s|_2, \dots, |s|_n)$. The existence of the function $|\cdot|$ guarantees that the system state contains enough information to retrieve the number of unmatched items of each class, which is a classical assumption in queueing theory (see for instance [27, Section 3.2]). We assume that there exists a unique state $s \in \mathcal{S}$ such that $|s| = 0$. This state is called the *empty state* and denoted by \emptyset . This assumption guarantees that the intuitive notion of system stability will be captured by the positive recurrence of the Markov chain describing the evolution of the system state¹³.

The policy is now a function $\Phi : \mathcal{S} \times V \times (V \cup \{\perp\}) \times \mathcal{S} \rightarrow [0, 1]$ such that $\Phi(s, i, j, s')$ is the conditional probability that, given an incoming class- i item finds the system in state s , the matching decision is j and the new state is s' . More formally, the dynamics are described by a Markov chain $((S_t, I_t, J_t), t \in \mathbb{N})$, where $I = (I_t, t \in \mathbb{N})$ is the sequence of incoming item classes and, for each $t \in \mathbb{N}$, $i \in V$, $j \in V \cup \{\perp\}$, and $s, s' \in \mathcal{S}$, we have

$$\mathbb{P}(J_t = j, S_{t+1} = s' \mid S_t = s, I_t = i) = \Phi(s, i, j, s').$$

The stochastic process $S = (S_t, t \in \mathbb{N})$ is also a Markov chain, with transition probabilities

$$\mathbb{P}(S_{t+1} = s' \mid S_t = s) = \frac{\sum_{i \in V} \lambda_i \sum_{j \in V \cup \{\perp\}} \Phi(s, i, j, s')}{\sum_{i \in V} \lambda_i}, \quad t \in \mathbb{N}, \quad s, s' \in \mathcal{S}.$$

We assume that the Markov chain S has state space \mathcal{S} and is irreducible, and that $S_0 = \emptyset$. The policy is assumed to be *adapted* to the compatibility graph G and *consistent* in the sense that, for each $(s, i, j, s') \in \mathcal{S} \times V \times (V \cup \{\perp\}) \times \mathcal{S}$, we have $\Phi(s, i, j, s') > 0$ only if

$$j \in \{j' \in V_i : |s|_{j'} \geq 1\} \cup \{\perp\}, \quad \text{and} \quad |s'| = \begin{cases} |s| + \mathbb{1}_i & \text{if } j = \perp, \\ |s| - \mathbb{1}_j & \text{if } j \neq \perp. \end{cases}$$

Using this extended definition, the previously defined policy models can be easily expressed. For example, the matching policy Φ is called *size-based* if $|\cdot|$ is the identity (implying $\mathcal{S} \subset \mathbb{N}^n$) and *deterministic* if, for each $s \in \mathcal{S}$ and $i \in V$, there exists $(j, s') \in (V \cup \{\perp\}) \times \mathcal{S}$ such that $\Phi(s, i, j, s') = 1$. The policy is called *greedy* if $\sum_{s' \in \mathcal{S}} \Phi(s, i, \perp, s') = 0$ for each $(s, i) \in \mathcal{S} \times V$ such that $\{j \in V_i : |s|_j \geq 1\} \neq \emptyset$, and *non-greedy* otherwise. FCFM (e.g., see [15, 35]) is a classical example of a deterministic policy that is not size-based: its state space is

¹³Without this assumption, one may construct two Markov chains associated with the same system, one positive recurrent and the other transient, for example if the state of the latter Markov chain embeds the time t . This assumption is used only in the proofs of Propositions 20 and 26. In both cases, we can verify that the same conclusion holds as long as the set of states $s \in \mathcal{S}$ such that $|s| = 0$ is finite. Assuming that this set is reduced to a singleton is merely a notational convenience.

a couple $(\mathcal{S}, |\cdot|)$ where \mathcal{S} is a subset of the set of sequences $c = (c_1, c_2, \dots, c_p)$ made of a finite but arbitrarily large number p of elements of \mathcal{I} , and $|c|_i$ is the cardinality of the set $\{q \in \{1, 2, \dots, p\} : c_q = i\}$, for each $i \in \mathcal{I}$. A policy that is neither size-based nor deterministic will appear in the proof of Proposition 20.

The stochastic process $Q = (Q_t, t \in \mathbb{N})$ defined by $Q_t = |S_t|$ for each $t \in \mathbb{N}$ is called the *queue-size process*. This process does not satisfy the Markov property in general, but it does satisfy the evolution equations (2) and (3), with $L_i = (L_{t,i}, t \in \mathbb{N})$ and $M_k = (M_{t,k}, t \in \mathbb{N})$ defined by (4) and (5) for each $i \in V$ and $k \in E$. The state space of the queue-size process is given by $\mathcal{Q} = \{|s|, s \in \mathcal{S}\}$. The policy is greedy if $\mathcal{Q} = \mathcal{Q}_G$ and non-greedy if $\mathcal{Q} \supsetneq \mathcal{Q}_G$, where \mathcal{Q}_G is still given by (7).

► **Remark 9** (Arrival rates vs. arrival sequence). *We will often identify the matching model (G, λ, Φ) with the Markov chain S . This is a slight abuse of language: the triplet (G, λ, Φ) specifies the transition diagram of this Markov chain but, even if Φ is deterministic, characterizing its sample paths requires specifying the sequence I , sampled according to $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$. This slight abuse of language will not cause confusion when discussing stability and matching rates, but the distinction will matter in Section 6.*

► **Remark 10** (Discrete time vs. continuous time). *The discrete-time Markov chain S gives the sequence of states observed by incoming items, and it was analyzed under various policies [33, 25]. Yet, in queueing theory, it is more common to consider the continuous-time Markov chain describing the system state over time. However, as observed in [15, Section 2.2.2], S is the jump chain of this continuous-time Markov chain, and both Markov chains have the same stationary measures because the departure rate from each state in the continuous-time Markov chain is constant equal to $\sum_{i \in V} \lambda_i$. Therefore, our results are equally relevant to study performance metrics like the mean queue size or the mean waiting time of items.*

A.2 Equivalent policies.

With our extended definition, a decision rule can be associated with an infinite number of policies. For instance, it is always possible to artificially expand the state definition, resulting in an unlimited range of policies. We define here an equivalence relation between policies that captures the intuitive concepts of *yielding identical distributions of matching decisions* (for random policies) and *making the same decisions* (for deterministic policies). This discussion will also prepare the ground for Propositions 34 and 36.

Consider a policy Φ_1 adapted to a compatibility graph $G = (V, E)$, and let $(\mathcal{S}_1, |\cdot|_1)$ denote its state space. Assume that the function $|\cdot|_1 : \mathcal{S}_1 \rightarrow \mathbb{N}^n$ can be written as a composition of two functions, $\langle \cdot \rangle : \mathcal{S}_1 \rightarrow \mathcal{S}_2$ and $|\cdot|_2 : \mathcal{S}_2 \rightarrow \mathbb{N}^n$, such that \mathcal{S}_2 is the image of \mathcal{S}_1 through $\langle \cdot \rangle$. Moreover, assume that there exists a policy Φ_2 with state space $(\mathcal{S}_2, |\cdot|_2)$, adapted to the graph G , such that for each $s_2, s'_2 \in \mathcal{S}_2$, $i \in V$, and $j \in V \cup \{\perp\}$, we have

$$\sum_{s'_1 \in \mathcal{S}_1 : \langle s'_1 \rangle = s'_2} \Phi_1(s_1, i, j, s'_1) = \Phi_2(s_2, i, j, s'_2) \quad \text{for each } s_1 \in \mathcal{S}_1 \text{ such that } \langle s_1 \rangle = s_2. \quad (24)$$

We say that policy Φ_1 can be *reduced* to policy Φ_2 and that $\langle \cdot \rangle$ is a *reduction* function. If $((S_{1,t}, I_t, J_{1,t}), t \in \mathbb{N})$ and $((S_{2,t}, I_t, J_{2,t}), t \in \mathbb{N})$ denote the Markov chains associated with policies Φ_1 and Φ_2 , respectively, under the same sequence $(I_t, t \in \mathbb{N})$ of incoming item classes, then for each $t \in \mathbb{N}$, (i) the conditional distribution of $(J_{1,t}, \langle S_{1,t+1} \rangle)$ given that $S_{1,t} = s_1$ and $I_t = i$ is the same for all states $s_1 \in \mathcal{S}_1$ that have the same image $s_2 = \langle s_1 \rangle$, and (ii) $(\langle S_{1,t} \rangle, I_t, J_{1,t})$ and $(S_{2,t}, I_t, J_{2,t})$ have the same distribution. Conclusion (ii) follows from an

inductive argument and implies that policies Φ_1 and Φ_2 are stable or unstable under the same conditions and, if stable, yield the same matching rate vector.

The special case where the policy Φ_2 is deterministic will be useful in Propositions 34 and 36. In this case, (24) says that, for each $s_2 \in \mathcal{S}_2$ and $i \in V$, there exist $j \in V \cup \{\perp\}$ and $s'_2 \in \mathcal{S}_2$ such that $\sum_{s'_1 \in \mathcal{S}_1: \langle s'_1 \rangle = s'_2} \Phi_1(s_1, i, j, s'_1) = \Phi_2(s_2, i, j, s'_2) = 1$ for each $s_1 \in \mathcal{S}_1$ such that $\langle s_1 \rangle = s_2$. This condition implies that the Markov chains under Φ_1 and Φ_2 are equivalent pathwise (i.e., $\langle S_{1,t} \rangle = S_{2,t}$ and $J_{1,t} = J_{2,t}$ for each $t \in \mathbb{N}$) and not just in distribution.

In general, we say that two policies Φ_1 and Φ_2 adapted to the graph G are *equivalent* if there exists a policy Φ adapted to the graph G such that both Φ_1 and Φ_2 can be reduced to Φ . This equivalence between Φ_1 and Φ_2 can be interpreted as indicating that the two policies, when they are in equivalent states, i.e. states that have the same reduction, yield identical distributions of matching decisions. It follows that, if we let $(Q_{1,t}, t \in \mathbb{N})$ and $(Q_{2,t}, t \in \mathbb{N})$ denote the queue-size processes under Φ_1 and Φ_2 with the same sequence $(I_t, t \in \mathbb{N})$ of incoming item classes, then we have $\mathbb{P}(Q_{1,t} = q) = \mathbb{P}(Q_{2,t} = q)$ for each $t \in \mathbb{N}$ and $q \in \mathcal{Q}$.

B

 Supplementary material of Section 3 (Graph theory and linear algebra).

B.1 Proofs of Section 3.2 (Surjectivity, injectivity, and bijectivity)

Proof of Definition 2. The equivalence of i, ii, and iii is a well-known result in linear algebra. We prove that conditions iii and iv are equivalent. This proof is adapted from [17, Lemma 2.2.3].

The key argument consists of observing that a vector $x = (x_1, x_2, \dots, x_n) \in \mathbb{R}^n$ belongs to the left kernel of the matrix A if and only if

$$\sum_{i=1}^n x_i a_{i,k} = 0, \quad k \in \{1, 2, \dots, m\}.$$

For each $k \in \{1, 2, \dots, m\}$, the k -th equation reads $x_j = -x_i$, where i and j are the endpoints of edge k . An induction argument shows that, for every path i_1, i_2, \dots, i_k in the graph G , we have $x_{i_p} = (-1)^{p-1} x_{i_1}$ for each $p \in \{1, 2, \dots, k\}$.

First assume that condition iv is satisfied. Let $x \in \mathbb{R}^n$ be a vector of the left kernel of the matrix A . Since each connected component of G is non-bipartite, for each $i \in V$, there exists a path of length say ℓ that connects node i to a cycle $i_1, i_2, \dots, i_p, i_{p+1} = i_1$ consisting of an odd number p of nodes. We then obtain $x_i = (-1)^{\ell+p+\ell} x_i = -x_i$, which implies that $x_i = 0$. Therefore, the left kernel of A is trivial, meaning that condition iii is satisfied.

On the contrary, if condition iv is not satisfied, then there exists a connected component of G that is bipartite with parts V_+ and V_- . We build a non-zero vector in the left kernel of A by choosing $x_i = 1$ for each $i \in V_+$, $x_i = -1$ for each $i \in V_-$, and $x_i = 0$ for each $i \in V \setminus (V_+ \cup V_-)$. This implies that condition iii is not satisfied. ◀

Proof of Definition 3. The equivalence of conditions i, ii, and iii is a well-known result in linear algebra. We now prove that conditions iii and iv are equivalent.

We first assume that the graph G is connected, and we distinguish the following two cases:

- If G is non-bipartite, according to Definition 2, the nullity of A^\top is 0. The rank-nullity theorem implies that the rank of A^\top is n , so that the rank of A is also n . A second application of the rank-nullity theorem implies that the nullity of A is $m - n$. In particular, $\ker(A) = \{0\}$ if and only if $m = n$.
- If G is bipartite, any non-zero vector of the left kernel of A must be parallel (collinear) to the non-zero vector x constructed in the proof of Definition 2. This parallelism is due to the constraints $x_i = -x_j$ for all edges $\{i, j\}$. Based on this, the nullity of A^\top is 1, and we conclude from another double application of the rank-nullity theorem that the nullity of A is $m - n + 1$. In particular, $\ker(A) = \{0\}$ if and only if $m = n - 1$.

All in all, we obtain that condition iii is true if and only if either the graph G is non-bipartite and contains as many edges as nodes, or the graph G is bipartite and contains one less edge than it contains nodes. This, in turn, is equivalent to condition iv.

If the graph G is not connected, we can rewrite A as a bloc matrix in which each bloc corresponds to a connected component, and we can then use the previous argument to prove the equivalence for each connected component. ◀

Proof of Definition 4. The function $\mu \in \mathbb{R}^m \mapsto A\mu \in \mathbb{R}^n$ is bijective if and only if it is both surjective and injective. Hence, the equivalence of conditions i to iv follows directly from Definitions 2 and 3. ◀

Proof of Proposition 6. These statements (transposed to A) are again well-known in linear algebra. ◀

B.2 Proof of Section 3.3.2 (Stabilizable matching problem)

Proof of Proposition 8. Equivalence of i and ii follows from [33, Proposition 2 and Theorem 2]. For completeness, we observe that [33, Proposition 2] is proved under the assumption that the matching policy Φ is greedy and deterministic and that the state space $(\mathcal{S}, |\cdot|)$ has a particular form, but we can verify that the argument remains valid under the assumptions of Section 2. We now prove that ii and iii are equivalent. Condition ii implies condition iii because (a) according to [33], under condition ii, (G, λ, Φ) is stable when Φ is the ML policy, and (b) the associated vector μ of matching rates satisfies condition iii by ergodicity. That condition iii implies condition ii was proved by [15, Lemma 12]. ◀

B.3 Minimal stability region for greedy matching policies (Section 3.4.3 and Appendices F.2.1 and F.2.2).

The following result gives a sufficient stability condition for greedy matching policies. The proof relies on a linear Lyapunov function. This result can be seen as the counterpart of [9, Proposition 5.1] for non-bipartite matching models.

► **Proposition 26.** *Consider a matching problem (G, λ) with a connected graph G . If*

$$\sum_{i \in V(\mathcal{I})} \lambda_i > \frac{1}{2} \sum_{i \in V} \lambda_i, \quad \mathcal{I} \in \mathbb{I}, \quad (25)$$

then the matching model (G, λ, Φ) is stable for every greedy matching policy Φ .

Proof. Consider a matching problem (G, λ) that satisfies (25) and a greedy matching policy Φ adapted to the graph G . Since the Markov chain $(S_t, t \in \mathbb{N})$ associated with the matching model (G, λ, Φ) depends on the vector λ only up to a positive multiplicative constant, we can assume without loss of generality that $\sum_{i \in V} \lambda_i = 1$. Let \mathcal{S} denote the state space of this Markov chain and $|\cdot|$ the corresponding queue-size function (as defined in Appendix A). We consider the Lyapunov function $F : \mathcal{S} \rightarrow \mathbb{R}$ defined by $F(s) = \sum_{i \in V} |s|_i$ (that is, $F(s)$ is the number of unmatched items in state s) for each $s \in \mathcal{S}$. For each $t \in \mathbb{N}$ and $s \in \mathcal{S}$, we have

$$\mathbb{E}(F(S_{t+1}) \mid S_t = s) - F(s) = \sum_{i \in V \setminus V(\mathcal{I})} \lambda_i - \sum_{i \in V(\mathcal{I})} \lambda_i = - \left(\sum_{i \in V(\mathcal{I})} \lambda_i - \sum_{i \in V \setminus V(\mathcal{I})} \lambda_i \right),$$

where $\mathcal{I} = \{i \in V : |s|_i \geq 1\}$ is the set of classes of unmatched items in state s . Importantly, if $F(s) > 0$ (that is, $s \neq \emptyset$), then \mathcal{I} is an independent set of the compatibility graph G because it is non-empty and the policy Φ is greedy. It follows that, for each $s \in \mathcal{S} \setminus \{\emptyset\}$,

$$\mathbb{E}(F(S_{t+1}) \mid S_t = s) - F(s) \leq -\varepsilon, \quad \text{with } \varepsilon = \min_{\mathcal{I} \in \mathbb{I}} \left(\sum_{i \in V(\mathcal{I})} \lambda_i - \sum_{i \in V \setminus V(\mathcal{I})} \lambda_i \right).$$

Equation (25) implies that $\varepsilon > 0$. Using the Lyapunov-Foster theorem, see [8, Theorem 1.1 in Chapter 5], we conclude that the matching model (G, λ, Φ) is stable. ◀

As one would expect, any matching problem (G, λ) that satisfies (25) is stabilizable in the sense of Definition 1. Indeed, (25) implies Proposition 8ii because $\mathcal{I} \subseteq V \setminus V(\mathcal{I})$ for each

$\mathcal{I} \in \mathbb{I}$. Corollary 27 below shows that, conversely, whether a stabilizable matching problem satisfies (25) depends on the structure of the graph G : conditions i and ii exhibit compatibility graphs G such that (25) is satisfied whenever the matching problem (G, λ) is stabilizable, while conditions iii and iv exhibit stabilizable compatibility graphs G for which (25) is never satisfied.

► **Corollary 27.** *Consider a matching problem (G, λ) .*

Under the following two conditions, the stabilizability of the matching model (G, λ) implies that (25) is satisfied, and therefore that the matching model (G, λ, Φ) is stable under every greedy policy Φ adapted to G :

(i) *G is a complete graph with $n \geq 3$ nodes.*

(ii) *G is the diamond graph of Example 4.*

Under the following conditions, (25) is never satisfied:

(iii) *The graph G has diameter greater than or equal to 3.*

(iv) *The graph G contains a leaf (that is, a node with degree 1).*

Proof. We first need to prove that, under either condition i or condition ii, the matching model (G, λ) is stabilizable if and only if (25) is satisfied. We proceed by verifying that, under either of these two conditions, Proposition 8ii and (25) are equivalent:

(i) First assume that condition i is satisfied. The independent sets of a complete graph K_n are the singletons. Using this observation, we can verify that Proposition 8ii and (25) are both equivalent to $\lambda_i < \frac{1}{2} \sum_{i \in V} \lambda_i$ for each $i \in V$.

(ii) Now assume that condition ii is satisfied, that is, G is the diamond graph. The conclusion follows by recalling that Proposition 8ii simplifies to (12), and then by observing that (12) and (25) are equivalent.

To prove that (25) cannot be satisfied under either condition iii or iv, we proceed by contradiction:

(iii) First assume that condition iii is satisfied, and let i and j denote two nodes that are at distance 3 or more. In particular, the sets V_i and V_j are disjoint. If (25) is satisfied, then applying this equation to both $\{i\}$ and $\{j\}$ and summing the inequalities yields $\sum_{i' \in V_i \cup V_j} \lambda_{i'} > \sum_{i' \in V} \lambda_{i'}$, which is a contradiction since $V_i \cup V_j \subseteq V$. Hence, (25) cannot be satisfied by both $\{i\}$ and $\{j\}$.

(iv) Now assume that condition iv is satisfied. Let i denote a leaf node of G and j the (only) neighbor of i . Then again, applying (25) to both $\{i\}$ and $\{j\}$ and summing the inequalities yields $\sum_{i' \in V_j \cup \{j\}} \lambda_{i'} > \sum_{i' \in V} \lambda_{i'}$, which is again a contradiction. ◀

C Supplementary material of Section 5 (Polytope of solutions in surjective-only graphs).

C.1 Algorithm of Section 5.1.3 (Basis of the kernel of the incidence matrix).

Data: A connected surjective-only compatibility graph $G = (V, E)$

Result: A basis \mathcal{B} of the kernel of the incidence matrix A of G

```

1  $\mathcal{T} \leftarrow$  the set of edges of a spanning tree of  $G$ 
2  $a \leftarrow$  an edge in  $E \setminus \mathcal{T}$  such that  $\mathcal{T} \cup \{a\}$  contains an odd cycle
3  $\mathcal{B} \leftarrow \emptyset$ 
4 for  $s \in E \setminus (\mathcal{T} \cup \{a\})$  do
5    $b \leftarrow (0, 0, \dots, 0) \in \mathbb{R}^m$ 
6   if  $\mathcal{T} \cup \{a, s\}$  contains an even cycle  $C_\ell$  then
7      $c_1, \dots, c_\ell \leftarrow$  consecutive edges of  $C_\ell$ 
8     for  $d \in \{1, \dots, \ell\}$  do
9        $k \leftarrow$  index of  $c_d$  in  $E$ 
10       $b_k \leftarrow (-1)^d$ 
11   else
12      $\mathcal{T} \cup \{a, s\}$  contains a kayak paddle  $KP_{\ell, r, p}$  with  $\ell$  odd,  $r$  odd, and  $p \geq 0$ 
13      $v_i \leftarrow$  node connecting the kayak's cycle  $C_\ell$  to the kayak central path  $P_p$ 
14      $v_j \leftarrow$  node connecting the kayak's cycle  $C_r$  to the kayak central path  $P_p$ 
15      $c_1, \dots, c_\ell \leftarrow$  consecutive edges of  $C_\ell$ , starting and ending at node  $v_i$ 
16     for  $d \in \{1, \dots, \ell\}$  do
17        $k \leftarrow$  index of  $c_d$  in  $E$ 
18        $b_k \leftarrow (-1)^d$ 
19      $c_1, \dots, c_p \leftarrow$  consecutive edges of  $P_p$ , starting at node  $v_i$  and ending at node  $v_j$ 
20     for  $d \in \{1, \dots, p\}$  do
21        $k \leftarrow$  index of  $c_d$  in  $E$ 
22        $b_k = 2(-1)^{d+1}$ 
23      $c_1, \dots, c_r \leftarrow$  consecutive edges of  $C_r$ , starting and ending at node  $v_j$ 
24     for  $d \in \{1, \dots, r\}$  do
25        $k \leftarrow$  index of  $c_d$  in  $E$ 
26        $b_k \leftarrow (-1)^{d+p+1}$ 
27    $\mathcal{B} \leftarrow \mathcal{B} \cup \{b\}$ 
28   return  $\mathcal{B}$ 

```

■ **Algorithm 1** Construction of a basis of the kernel of the incidence matrix A of the compatibility graph G . This algorithm was initially introduced by [20, Section 3] to build a basis of the eigenspace associated with the eigenvalue -2 of the adjacency matrix of a line graph (i.e., a graph whose nodes and edges represent, respectively, the edges and their incidence relations in another graph).

Given a surjective-only graph G with an incidence matrix denoted by A , Algorithm 1 builds a basis of the kernel of A as follows. The algorithm first identifies a spanning tree \mathcal{T} of G (Line 1) and an edge $a \in E \setminus \mathcal{T}$ such that the set $E \setminus (\mathcal{T} \cup \{a\})$ (of cardinality $d = m = n$) contains an odd cycle (Line 2). Then, for each edge $s \in E \setminus (\mathcal{T} \cup \{a\})$, the algorithm builds (Lines 5–26) a base vector b whose support (i) is either an even cycle or a kayak paddle and (ii) contains s and is included into $\mathcal{T} \cup \{a, s\}$. We assume without loss of generality that

the graph G is connected (in addition to being surjective-only). If not, we can apply the algorithm to each connected component separately, and then we embed the obtained vectors to \mathbb{R}^n via zero padding.

We now verify that Algorithm 1 terminates and yields the desired result.

► **Proposition 28.** *Algorithm 1 terminates and returns a basis of the kernel of the incidence matrix A of the compatibility graph G .*

Proof. The proof is based mainly on the notion of *cycle space* in a graph. We briefly summarize the concepts that are useful to understand the proof (see [19, Section 1.9] for details).

A *spanning subgraph* of a graph $G = (V, E)$ is a subgraph $G^\circ = (V, E^\circ)$ with $E^\circ \subseteq E$. Importantly, G and G° have the same set of nodes. A subgraph is *Eulerian* if every vertex has an even degree (possibly zero). In particular, if E° is a set of edges that form a cycle in G , then the graph (V, E°) is Eulerian. The *cycle space* of G is the vector space made of all Eulerian spanning subgraphs of G , using the symmetric difference of the edge sets for addition and the two-element field for scalar multiplication. Equivalently, the cycle space can be described as a vector space of the finite field $\mathbb{Z}/2\mathbb{Z}$: each vector $g = (g_1, g_2, \dots, g_m)$ in this vector space satisfies $\sum_{k \in E_i} g_k = 0$ (modulo 2) for each $i \in V$, and the addition and multiplication are the usual operations in $\mathbb{Z}/2\mathbb{Z}$. For example, if G is the codomino graph of Figure 17, the spanning subgraphs G_1 , G_2 , and G_3 of G with edge sets $E_1 = \{\{1, 2\}, \{1, 6\}, \{2, 3\}, \{3, 4\}, \{4, 5\}, \{5, 6\}\}$, $E_2 = \{\{2, 3\}, \{2, 6\}, \{3, 5\}, \{5, 6\}\}$, and $E_3 = \{\{1, 2\}, \{1, 6\}, \{2, 6\}\}$, respectively, belong to the cycle space of G . The edge set of the addition $G_1 + G_2$ is the set $\{\{1, 2\}, \{1, 6\}, \{2, 6\}, \{3, 4\}, \{3, 5\}, \{4, 5\}\}$ of edges that are either in E_1 or in E_2 , but not in both. Similarly, the edge set of $G_1 + G_3$ is $\{\{2, 3\}, \{2, 6\}, \{3, 4\}, \{4, 5\}, \{5, 6\}\}$, and the edge set of $G_2 + G_3$ is $\{\{1, 2\}, \{1, 6\}, \{2, 3\}, \{3, 5\}, \{5, 6\}\}$. One can verify that $\{G_1, G_2, G_3\}$ forms a basis of the cycle space of G . Importantly, in general, the dimension of the cycle space is $m - n + 1$.

Algorithm 1 terminates. We first prove the existence of edge a defined on Line 2 of Algorithm 1. By definition of a spanning tree, \mathcal{T} contains $n - 1$ edges and, for each edge $a \in E \setminus \mathcal{T}$, $\mathcal{T} \cup \{a\}$ contains a unique cycle. The $m - n + 1 = |E \setminus \mathcal{T}|$ cycles thus obtained are independent in the sense that each cycle contains at least one edge (a) that is not contained in the other cycles. Therefore, these $m - n + 1$ cycles form a basis of the cycle space of G . Since a linear combination of even cycles cannot produce a subgraph consisting of a single odd cycle¹⁴, and since G contains an odd cycle (as it is non-bipartite), then at least one of the $m - n + 1$ basis cycles is odd.

We now verify that, for each $s \in E \setminus (\mathcal{T} \cup \{a\})$, $\mathcal{T} \cup \{a, s\}$ contains either (i) an even cycle C_ℓ or (ii) a kayak paddle $KP_{\ell, r, p}$ with two odd cycles. By construction, $\mathcal{T} \cup \{a\}$ contains a unique cycle C_r , which is odd, and $\mathcal{T} \cup \{s\}$ contains a unique cycle C_ℓ . $\mathcal{T} \cup \{a, s\}$ contains both C_r and C_ℓ . We now proceed by elimination:

- If C_ℓ is even, then C_ℓ is an even cycle included into $\mathcal{T} \cup \{a, s\}$, and we are in case (i).
- If C_ℓ is odd and shares at least one edge with C_r , then the symmetric difference of C_r and C_ℓ is an even cycle, and it is again included into $\mathcal{T} \cup \{a, s\}$, so we are again in case (i).
- If C_ℓ is odd and C_r and C_ℓ have no edge in common, then we are in case (ii).

¹⁴A linear combination of even cycles may produce a subgraph consisting of an even number of disjoint odd cycles, as illustrated by $G_1 + G_2$ in the example above, but not a subgraph consisting of one odd cycle. Indeed, the symmetric difference of two edge sets contains an even number of edges if both edge sets contain an even number of edges.

Algorithm 1 returns the correct result. We finally prove that the family \mathcal{B} returned by Algorithm 1: (i) has cardinality $m - n$, (ii) is linearly independent, and (iii) is included into the kernel of A . We prove each item one after another:

- (i) The family \mathcal{B} has cardinality $m - n$: It suffices to observe that \mathcal{B} has same cardinality as $E \setminus (\mathcal{T} \cup \{a\})$, which we already mentioned has cardinality $m - n$.
- (ii) The family \mathcal{B} is linearly independent: For each $s \in E \setminus (\mathcal{T} \cup \{a\})$, the basis vector constructed from edge s is the only vector in \mathcal{B} whose support contains this edge.
- (iii) The family \mathcal{B} is included into the kernel of A : Let $b \in \mathcal{B}$. Our goal is to prove that b belongs to the kernel of A , i.e., that $\sum_{k \in E_i} b_k = 0$ for each $i \in V$. First observe that, for each $i \in V$, we have $\sum_{k \in E_i} b_k = \sum_{k \in E_i \cap S} b_k$, where S is the support of the vector b . In particular, we have immediately $\sum_{k \in E_i} b_k = 0$ for each $i \in V$ such that $E_i \cap S = \emptyset$. Now consider a node $i \in V$ such that $E_i \cap S \neq \emptyset$. We make a case disjunction depending on the support S of b :

- If S is an even cycle C_ℓ , then $E_i \cap S = \{k_1, k_2\}$, where k_1 and k_2 are two consecutive edges of the cycle C_ℓ . Line 10 in the algorithm implies that $b_{k_1} = -b_{k_2} \in \{1, -1\}$. It follows that $\sum_{k \in E_i} b_k = b_{k_1} + b_{k_2} = 0$.
- If S is a kayak paddle $KP_{\ell,r,p}$ with odd cycles C_ℓ and C_r and central path P_p , $p \in \mathbb{N}$, we again distinguish several cases:

Node i does not belong to the central path: If $E_i \cap S \subseteq C_\ell$ or $E_i \cap S \subseteq C_r$, we conclude as before.

Node i does not belong to a cycle: If $E_i \cap S \subseteq P_p$, then $E_i \cap S = \{k_1, k_2\}$ where k_1 and k_2 are two consecutive edges of the path P_p . Line 22 of the algorithm implies that $b_{k_1} = -b_{k_2} \in \{2, -2\}$. It follows that $\sum_{k \in E_i} b_k = b_{k_1} + b_{k_2} = 0$.


Node i belongs to a cycle and the central path: The only remaining case is when $E_i \cap S$ intersects several sets among C_ℓ , C_r , and P_p . If $p = 0$, that is, if the central path is the node i , then $E_i \cap S = \{k_1, k_2, k_3, k_4\}$, where k_1 and k_2 (resp. k_3 and k_4) are two consecutive edges in C_ℓ (resp. C_r). Lines 18 and 26 yield $b_{k_1} = b_{k_2} = -1$ and $b_{k_3} = b_{k_4} = 1$, which implies that $\sum_{k \in E_i} b_k = b_{k_1} + b_{k_2} + b_{k_3} + b_{k_4} = 0$. If $p \geq 1$, then $E_i \cap S = \{k_1, k_2, k_3\}$, where k_1 and k_2 are two consecutive edges of either C_ℓ or C_r , and k_3 is an edge in P_p . Lines 18, 22, and 26 of the algorithm imply that $b_{k_1} = b_{k_2} \in \{1, -1\}$ and that $b_{k_3} = -2b_{k_1}$, so that we conclude again that the desired sum is zero. ◀

Figures 16 and 17 shows possible runs of Algorithm 1 on the triamond and codomino graphs, which both have a two-dimensional kernel. Note that the basis is not unique and depends on our choice of the spanning tree \mathcal{T} and the augmenting edge a (see Lines 1 and 2 in Algorithm 1).


C.2 Proof of Section 5.2.1 (Vertex characterization).

Proof of Proposition 13 (borrowed from [14]). We prove that the negations of i and ii are equivalent. Let A^* denote the incidence matrix of G^* .

By Definition 12, if μ is not a vertex of $\Pi_{\geq 0}$, there exist $z_1, z_2 \in \Pi_{\geq 0} \setminus \{\mu\}$ and $0 < \theta < 1$ such that $\mu = \theta z_1 + (1 - \theta)z_2$. The coordinates of the vectors z_1 and z_2 are non-negative, so this equality implies that their supports are included into the support E^* of the vector μ . In particular, if $\tilde{\mu}$ and \tilde{z}_1 denote the restrictions of μ and z_1 to coordinates in E^* , respectively, then $A^* \tilde{\mu} = A\mu = \lambda = Az_1 = A^* \tilde{z}_1$ with $\tilde{\mu} \neq \tilde{z}_1$, which means that $G^* = (V, E^*)$ is not injective.



■ **Figure 16** Two possible constructions of a kernel basis for the triamond graph. Construction A yields the basis vectors $b_1 = (1, -1, -1, -1, 1, 1, 0)$ and $b_2 = (0, 0, 1, 0, -1, -1, 1)$. Construction B yields the basis vectors $b_1 = (-1, 1, 0, 1, 0, 0, -1)$ and $b_2 = (0, 0, -1, 0, 1, 1, -1)$.



■ **Figure 17** Two possible constructions of a kernel basis for the codomino graph. Construction A yields the vectors $b_1 = (1, -1, 0, -1, 1, -1, -1, 2)$ and $b_2 = (0, 0, 1, -1, 0, -1, 0, 1)$. Construction B yields the vectors $b_1 = (-1, 1, 1, 0, -1, 0, 1, -1)$ and $b_2 = (0, 0, -1, 1, 0, 1, 0, -1)$.

Conversely, if G^* is not injective, there exists a non-zero vector \tilde{z} in $\mathbb{R}^{|E^*|}$ such that $A^* \tilde{z} = 0$. If we embed \tilde{z} into $\mathbb{R}^{|E|}$ with zero-padding, we obtain a non-zero vector z such that $Az = 0$, and whose support is included into that of the vector μ . This implies that there exists $\varepsilon > 0$ such that both $\mu - \varepsilon z$ and $\mu + \varepsilon z$ belong to $\Pi_{\geq 0}$. The convex combination $\mu = \frac{1}{2}(\mu - \varepsilon z) + \frac{1}{2}(\mu + \varepsilon z)$ proves that the vector μ is not a vertex of $\Pi_{\geq 0}$.


The last part directly derives from Items ii and v in Proposition 6. ◀

C.3 Additional examples and results for Section 5.2.1 (Vertex characterization).

Following Proposition 13, the bijectivity of a vertex is determined by the number of its coordinates that are positive in edge coordinates, that is, by the cardinality of the set of edges

that form its support. Recall that the d -dimensional polytope $\Pi_{\geq 0}$ is actually characterized by the m inequalities $\mu_k \geq 0$ for each $k \in E$. In particular, this polytope has at most m facets, one for each inequality, but it typically has fewer. Indeed, some inequalities may be redundant and/or not tight, in a sense that will be defined in Definition 29 below. For example, by looking more closely at the general solution obtained for the diamond graph in Figure 3d, we conclude that:

- The inequality $\mu_{2,3} \geq 0$ is satisfied trivially by every vector $\mu \in \Pi$, as we have $\mu_{2,3} = \beta > 0$. Therefore, this inequality does not define a facet of $\Pi_{\geq 0}$.
- If $\lambda_1 \bar{\lambda}_2 < \bar{\lambda}_3 \lambda_4$, the inequality $\mu_{1,2} \geq 0$ supersedes the inequality $\mu_{3,4} \geq 0$, and conversely. If $\lambda_1 \bar{\lambda}_2 = \bar{\lambda}_3 \lambda_4$, these two inequalities are equivalent. In both cases, the inequalities $\mu_{1,2} \geq 0$ and $\mu_{3,4} \geq 0$ lead to a single facet of $\Pi_{\geq 0}$.
- If $\lambda_1 \bar{\lambda}_3 < \bar{\lambda}_2 \lambda_4$, the inequality $\mu_{1,3} \geq 0$ supersedes the inequality $\mu_{2,4} \geq 0$, and conversely. If $\lambda_1 \bar{\lambda}_3 = \bar{\lambda}_2 \lambda_4$, these two inequalities are equivalent. In both cases, the inequalities $\mu_{1,3} \geq 0$ and $\mu_{2,4} \geq 0$ lead to a single facet of $\Pi_{\geq 0}$.



■ **Figure 3d (repeated)** Matching rates in the diamond graph with the normalization $\lambda_1 + \lambda_4 = \frac{1}{2}$. $2\beta = \lambda_2 + \lambda_3 - \frac{1}{2}$ is the difference between the arrival rates of $\{2, 3\}$ and $\{1, 4\}$. $\bar{\lambda}_2 = \lambda_2 - \beta$ and $\bar{\lambda}_3 = \lambda_3 - \beta$ represent the residual rates that classes 2 and 3 can provide to classes 1 and 4.

All in all, the 1-dimensional convex polytope $\Pi_{\geq 0}$ associated with the diamond graph of Figure 3d has two facets, even if it is defined by five inequalities. Definition 29 below will help us relate these notions to the number of zero coordinates of the vertices of the convex polytope $\Pi_{\geq 0}$.

► **Definition 29** (Adapted from [4, 42]). *Let $k \in E$.*

- (i) *The inequality $\mu_k \geq 0$ is said to be tight if there exists a vector $\mu \in \Pi_{\geq 0}$ such that $\mu_k = 0$, in which case we also say that this inequality is tight for the vector μ .*
- (ii) *The inequality $\mu_k \geq 0$ is said to be redundant if removing this inequality does not change the polytope $\Pi_{\geq 0}$, in the sense that*

$$\Pi_{\geq 0} = \{\mu \in \mathbb{R}^m : A\mu = \lambda \text{ and } \mu_\ell \geq 0 \text{ for each } \ell \in E \setminus \{k\}\}.$$

Otherwise, this inequality is called irredundant.

- (iii) *The matching problem (G, λ) is called essential if all tight inequalities are irredundant.*
- (iv) *The polytope $\Pi_{\geq 0}$ is said to be simple if every vertex of $\Pi_{\geq 0}$ belongs to exactly d facets, which is the minimal number of facets a vertex belongs to.*

Importantly, the number of positive coordinates of a vertex μ (considered in Proposition 17) is the number of inequalities that are not tight for this vertex. More generally, Definition 29 has the following intuitive interpretation. An inequality is tight if the convex polytope $\Pi_{\geq 0}$ intersects the hyperplane obtained by transforming this inequality into an equality. Non-tight inequalities are “useless” (and redundant) because they are never satisfied as equalities by any vector in $\Pi_{\geq 0}$. The matching problem (G, λ) is essential if each tight inequality defines

a distinct facet of the convex polytope $\Pi_{\geq 0}$. Under this condition, the number of facets that contain a vertex is equal to the number of inequalities that are tight for this vertex. In particular, as we will see in Proposition 30, if the matching problem (G, λ) is essential and the polytope $\Pi_{\geq 0}$ is simple, then every vertex satisfies exactly d (tight) inequalities as equalities, which means that this vertex has d zero coordinates, and therefore $n = m - d$ positive coordinates, so that this vertex is bijective.

All these notions are illustrated in Examples 8–10 below, which show in particular that a matching problem (G, λ) may be essential even if the polytope $\Pi_{\geq 0}$ is not simple, and conversely. Consistently with the discussion above on Figure 3d, these examples use a kernel basis to verify effortlessly whether an inequality is tight and/or irredundant.

► **Example 8 (Essential matching problem).** *Figure 18 considers a codomino graph with the vector of arrival rates $\lambda = (4, 5, 3, 2, 3, 5)$. A particular solution to (CE) is $\mu^\circ = (2, 2, 1, 2, 1, 1, 1, 1) \in \mathbb{R}^8$, and the basis of $\ker(A)$ consists of the vectors $b_1 = (-1, 1, 1, 0, -1, 0, 1, -1)$ and $b_2 = (0, 0, -1, 1, 0, 1, 0, -1)$ obtained in construction B of Figure 17. The generic solution to (CE) is shown in Figure 18b.*

The inequalities are listed in Figure 18a. The 2-dimensional polytope $\Pi_{\geq 0}$, shown in Figure 18c in kernel basis, is characterized by five tight inequalities which are also irredundant:

$$-1 \leq \alpha_1 \leq 1, \quad \alpha_2 \geq -1, \quad \alpha_1 - \alpha_2 \geq -1, \quad \alpha_1 + \alpha_2 \leq 1.$$

The matching problem (G, λ) is essential. In kernel basis, the vertices of the convex polytope $\Pi_{\geq 0}$ are $(0, 1)$, $(-1, 0)$, $(1, 0)$, $(-1, -1)$, and $(1, -1)$, and we can verify on Figure 18c that each vertex belongs to exactly 2 facets. Therefore $\Pi_{\geq 0}$ is simple (more generally, all 2-dimensional polytopes are simple). All in all, each vertex of $\Pi_{\geq 0}$ has 2 zero coordinates and 6 positive coordinates in edge coordinates, so that this vertex is bijective. These vertices are represented in edge basis in Figures 18d–18h.

► **Example 9 (Non-essential matching problem).** *Figure 19 shows the same codomino graph as in Example 8, with the same basis of $\ker(A)$, but with the vector of arrival rates $\lambda = (2, 4, 4, 2, 2, 2)$. A particular solution to (CE) is $\mu^\circ = (1, 1, 2, 1, 1, 1, 1, 0)$, and the general solution is shown in Figure 19b.*


The inequalities are listed in Figure 19a. The 2-dimensional convex polytope $\Pi_{\geq 0}$ is shown in kernel basis in Figure 19c. All inequalities are tight, but only one is irredundant, so we conclude that the matching problem (G, λ) is not essential, even if the polytope $\Pi_{\geq 0}$ is still simple. Correspondingly, even if each vertex belongs to exactly two facets, they all have more than 2 zero coordinates, so none of them is bijective. For example, the vertex $(1, -1)$ in kernel basis has coordinates $(0, 2, 4, 0, 0, 0, 2, 0)$ in edge basis (Figure 19f). This vertex has 5 zero coordinates in edge coordinates (and only 3 positive coordinates) even if it belongs to only 2 facets.

► **Example 10 (Non-simple polytope).** *We finally exhibit an essential matching problem with a non-simple associated polytope. As 2-dimensional polytopes are simple, we need to consider a more complex example. We consider the matching problem of Figure 20a. The arrival rate is $\lambda = (3, 3, 6, 3, 4, 4, 6, 3, 4, 4) \in \mathbb{R}^{10}$. The particular solution and kernel basis are shown on the edges. The set $\Pi_{\geq 0}$, shown in Figure 20b in kernel basis, is an Egyptian pyramid characterized by the following tight inequalities:*


$$\alpha_3 \geq 0, \quad 1 + \alpha_1 - \alpha_3 \geq 0, \quad 1 - \alpha_1 - \alpha_3 \geq 0, \quad 1 + \alpha_2 - \alpha_3 \geq 0, \quad 1 - \alpha_2 - \alpha_3 \geq 0.$$

Edge basis	Kernel basis	Tight?	Irredundant?
$\mu_{1,2} \geq 0$	$\alpha_1 \leq 2$	✗	✗
$\mu_{1,6} \geq 0$	$\alpha_1 \geq -2$	✗	✗
$\mu_{2,3} \geq 0$	$\alpha_1 - \alpha_2 \geq -1$	✓	✓
$\mu_{2,6} \geq 0$	$\alpha_2 \geq -2$	✗	✗
$\mu_{3,4} \geq 0$	$\alpha_1 \leq 1$	✓	✓
$\mu_{3,5} \geq 0$	$\alpha_2 \geq -1$	✓	✓
$\mu_{4,5} \geq 0$	$\alpha_1 \geq -1$	✓	✓
$\mu_{5,6} \geq 0$	$\alpha_1 + \alpha_2 \leq 1$	✓	✓


(a) Inequalities.




(b) Generic solution to (CE).




(d) Edge coordinates of (0, 1).




(e) Edge coordinates of (-1, 0).




(f) Edge coordinates of (1, 0).



(c) Polytope $\Pi_{\geq 0}$ in kernel basis.



(g) Edge coordinates of (-1, -1).



(h) Edge coordinates of (1, -1).



■ **Figure 18** An essential matching problem (G, λ) with a simple polytope $\Pi_{\geq 0}$. The vector of arrival rates is $\lambda = (4, 5, 3, 2, 3, 5) \in \mathbb{R}^6$, a particular solution to (CE) is $\mu^\circ = (2, 2, 1, 2, 1, 1, 1, 1) \in \mathbb{R}^8$, and the chosen base vectors for $\ker(A)$ are $b_1 = (-1, 1, 1, 0, -1, 0, 1, -1)$ and $b_2 = (0, 0, -1, 1, 0, 1, 0, -1)$.

These five inequalities are irredundant (each one corresponds to exactly one of the five facets), so we conclude that the matching problem (G, λ) is essential. In kernel basis, the vertices of this convex polytope are $(-1, -1, 0)$, $(1, -1, 0)$, $(1, 1, 0)$, $(-1, 1, 0)$, and $(0, 0, 1)$. These vertices are shown in edge basis in Figures 20c–20g. The polytope $\Pi_{\geq 0}$ is not simple because the vertex $(0, 0, 1)$ (the “top” of the pyramid) belongs to 4 facets, while the polytope has dimension 3. Consistently, we can see in Figure 20g that this vertex has 4 zero coordinates and only 9 positive coordinates in edge basis; the subgraph defined by the support of this vertex is injective-only.

In light of these examples, we can give the following characterization of the bijective vertices of $\Pi_{\geq 0}$.




Edge basis	Kernel basis	Tight?	Irredundant?
$\mu_{1,2} \geq 0$	$\alpha_1 \leq 1$	✓	✗
$\mu_{1,6} \geq 0$	$\alpha_1 \geq -1$	✓	✗
$\mu_{2,3} \geq 0$	$\alpha_1 - \alpha_2 \geq -2$	✓	✗
$\mu_{2,6} \geq 0$	$\alpha_2 \geq -1$	✓	✗
$\mu_{3,4} \geq 0$	$\alpha_1 \leq 1$	✓	✗
$\mu_{3,5} \geq 0$	$\alpha_2 \geq -1$	✓	✗
$\mu_{4,5} \geq 0$	$\alpha_1 \geq -1$	✓	✗
$\mu_{5,6} \geq 0$	$\alpha_1 + \alpha_2 \leq 0$	✓	✓

(a) Inequalities



(b) Generic solution to (CE).

(c) Polytope $\Pi_{\geq 0}$ in kernel basis. Dashed lines show tight redundant inequalities.



(d) Edge coordinates of $(-1, 1)$. (e) Edge coordinates of $(-1, -1)$. (f) Edge coordinates of $(1, -1)$.

■ **Figure 19** Non-essential matching problem (G, λ) with a simple polytope $\Pi_{\geq 0}$. The vector of arrival rates is $\lambda = (2, 4, 4, 2, 2, 2) \in \mathbb{R}^6$, a particular solution to (CE) is $\mu^\circ = (1, 1, 2, 1, 1, 1, 1, 0) \in \mathbb{R}^8$, and the chosen base vectors for $\ker(A)$ are $b_1 = (-1, 1, 1, 0, -1, 0, 1, -1)$ and $b_2 = (0, 0, -1, 1, 0, 1, 0, -1)$.


► **Proposition 30.** Let μ be a vertex of $\Pi_{\geq 0}$. The following statements are equivalent:

- (i) μ is bijective.
- (ii) μ belongs to exactly d facets of $\Pi_{\geq 0}$ and none of the inequalities tight for μ is redundant. In particular, all vertices of $\Pi_{\geq 0}$ are bijective if, and only if, the matching problem (G, λ) is essential and the polytope $\Pi_{\geq 0}$ is simple.

Proof. To prove the equivalence of i and ii, we first remark that the number of zero edge coordinates of a vector $\mu \in \Pi_{\geq 0}$ is by definition the number of inequalities that are tight for μ . It is in particular at least the number of facets that intersect μ , with equality if, and only if, none of the inequalities tight for μ is redundant.

If a vector μ is a vertex of $\Pi_{\geq 0}$, then μ belongs to at least d facets of $\Pi_{\geq 0}$. Now, the vector μ is bijective if, and only if, d of its coordinates are zero, which in view of the remark above is equivalent to say that μ belongs to exactly d facets of $\Pi_{\geq 0}$ and none of the inequalities tight for μ is redundant.

As for the last statement, it follows directly from Definition 29. ◀



■ **Figure 20** Essential matching problem (G, λ) with a non-simple polytope $\Pi_{\geq 0}$. The arrival rate is $\lambda = (3, 3, 6, 3, 4, 4, 6, 3, 4, 4) \in \mathbb{R}^{10}$. A particular solution and the chosen base vectors for $\ker(A)$ are implicitly shown on the edges of Figure 20a.

► **Remark 11.** We can further explore the relationship between vertices and their support graph. Again consider a vertex μ of $\Pi_{\geq 0}$, and let G^* denote the support graph of μ and A^* its incidence matrix. According to Definitions 2, 3, and 5, the subgraph G^* is injective if and only if each connected component of G^* is either a tree or a unicyclic graph with an odd cycle. For each connected component of G^* that is a tree, and therefore a bipartite graph with parts V_+ and V_- , the existence of vertex μ implies that

$$\sum_{i \in V_+} \lambda_i = \sum_{i \in V_-} \lambda_i. \quad (26)$$

This equation follows by summing (CE-1) over the nodes in V_+ on the one hand, summing (CE-1) over the nodes in V_- on the other hand, and verifying that the left-hand sides of both equations are equal. In fact, one can verify that the vector λ belongs to the image of A^* if and only if λ satisfies (26) for each connected component of G^* that is a tree. Note that this condition is void if G^* is bijective because, in this case, none of the connected components of G^* is a tree.

Conversely, one can wonder which injective subgraph G^* of G defines a vertex of $\Pi_{\geq 0}$. Satisfying (26) for each tree connected component of G^* only guarantees the existence of a (unique) solution $z \in \mathbb{R}^P$ to the conservation equation $A^*z = \lambda$. If each coordinate of z is

positive, then we indeed obtain a vertex of $\Pi_{\geq 0}$ by embedding z in \mathbb{R}^m with zero padding; otherwise, G^* does not define a vertex of $\Pi_{\geq 0}$.

C.4 Proof of Section 5.2.2 (Probability of bijectivity).

Proof of Proposition 14. Let C denote the set of normalized arrival rate vectors λ such that (G, λ) is stabilizable:

$$C = \{\lambda \in \Delta^{n-1} : (G, \lambda) \text{ stabilizable}\}$$

To prove Proposition 14, we proceed as follows:

- We prove that C is non-empty and convex, hence measurable.
- We prove that C is an open set (for the Lebesgue measure), so its probability is positive and the conditional probability given C is well-defined.
- We prove that if a vertex is injective-only, the vector λ must respect a constraint that reduces the dimension of the possible solutions, meaning that the measure of the event is null.

We first prove the different properties of C .

C is not empty. We use Remark 2: consider $\lambda = A\mu$, with $\mu = (\frac{1}{2m}, \dots, \frac{1}{2m})$. The matching problem (G, λ) is stabilizable as all coordinates of μ are positive. Moreover, one can check that $\lambda = (\frac{d_1}{2m}, \dots, \frac{d_n}{2m})$, where d_i is the degree of node i , so $\lambda \in \Delta^{n-1}$ ($\sum_i d_i = 2m$). Hence, $\lambda \in C$.

C is convex. Let λ_a and λ_b be two vectors of C . Let $\mu_a \in \mathbb{R}_{>0}^m$ and $\mu_b \in \mathbb{R}_{>0}^m$ be positive solutions to $A\mu_a = \lambda_a$ and $A\mu_b = \lambda_b$ respectively (such solutions exist as stated in Proposition 8). For $0 < x < 1$, consider $\lambda = x\lambda_a + (1-x)\lambda_b$. $\lambda \in \Delta^{n-1}$ because Δ^{n-1} is convex. Moreover, we have $A\mu = \lambda$, with $\mu = x\mu_a + (1-x)\mu_b$. As $\mathbb{R}_{>0}^m$ is convex, $\mu \in \mathbb{R}_{>0}^m$, so (G, λ) is stabilizable (Proposition 8), hence $\lambda \in C$.

C is open in Δ^{n-1} . Let $\lambda \in C$. For some $\epsilon > 0$, let λ_ϵ be an element of Δ^{n-1} such that $\|\lambda - \lambda_\epsilon\|_1 < \epsilon$. To prove that C is open, we just need to show that $\lambda_\epsilon \in C$ if ϵ is small enough. Let $\mu \in \mathbb{R}_{>0}^m$ be a positive solution to $A\mu = \lambda$. Using Equations (18) and (19), we can write $\mu = A^+\lambda + \mu_\alpha$, with $\mu_\alpha = \mu - A^+\lambda = (\text{Id}_{m \times m} - A^+A)\mu \in \ker(A)$. Consider now $\mu_\epsilon = A^+\lambda_\epsilon + \mu_\alpha$. We have $A\mu_\epsilon = AA^+\lambda_\epsilon + A\mu_\alpha = \lambda_\epsilon$ so $(\lambda_\epsilon, \mu_\epsilon)$ checks (CE). Moreover, $\|\mu - \mu_\epsilon\|_1 = \|A^+(\lambda - \lambda_\epsilon)\|_1 < \|A^+\|_1 \epsilon$, where $\|A^+\|_1 < +\infty$ is the operator norm of A^+ associated to the 1-norm. All coordinates of μ are positive, so if ϵ is small enough, it also the case for μ_ϵ , meaning that (G, λ_ϵ) is stabilizable (Proposition 8). In other words, $\lambda_\epsilon \in C$.

This establishes that C (i.e., the stabilizability of (G, λ)) defines a proper conditional probability. We now examine the conditions under which the polytope $\Pi_{\geq 0}$ of a stabilizable problem (G, λ) possesses an injective-only vertex.

Consider $\lambda \in \Delta^{n-1}$. Let μ be a solution to $A\mu = \lambda$. Using Definition 11 and the pseudoinverse approach from Section 5.1.2, the kernel coordinates of μ are $\alpha = B^+(\text{Id}_{m \times m} - A^+A)\mu$, so we can write

$$\mu = B\alpha + A^+\lambda = D\gamma,$$

where $D = (B|A^+)$ is the $m \times m$ horizontal concatenation of B and A^+ , and $\gamma = (\alpha_1, \dots, \alpha_d, \lambda_1, \dots, \lambda_n)$ is the concatenation of α and λ .

Notice that D depends solely on G and is invertible. Its inverse can be verified as the vertical concatenation of $B^+(\text{Id}_{m \times m} - A^+A)$ and A : any matching rate vector μ uniquely defines λ and α .

We now assume that one vertex of the polytope is injective-only and look for the implications for λ .

The existence of an injective-only vertex implies that there exists a subset of edges $K \subset E$, with $|K| > d$, such that μ is null on K (remind that $d = m - n$ is the dimension of $\ker(A)$). Let $D|_K$ denote the $|K| \times m$ submatrix of D formed by the rows corresponding to K , and let $B|_K$ denote the d first columns of $D|_K$, that is the part of $D|_K$ that handles α . As $|K| > d$, we have $\ker((B|_K)^\top) \neq \{0\}$, meaning that there exists a non-trivial linear combination $v \in \mathbb{R}^{|K|} \setminus \{0\}$ of the rows of $B|_K$ such that $vB|_K = 0$.

Let $A^+|_K$ denotes the n last columns of $D|_K$, that is the part of $D|_K$ that handles λ . We know that $vD|_K \neq 0$ because D is invertible, so its row are linearly independent. As $vB|_K = 0$, this means that $vA^+|_K \neq 0$. The vector $u = vA^+|_K$ can be seen as a non-trivial linear combination of the coordinates of λ .

We can now conclude: if the polytope $\Pi_{\geq 0}$ associated to some $\lambda \in C$ admits an injective-only vertex, we have a subset of edges $K \subset E$, with $|K| > d$ such that $D|_K \gamma = 0$. In particular, we must have $vD|_K \gamma = 0$, which simplifies as $u\lambda = 0$. In other words, for μ to be null over K , the vector λ must belong to some hyperplane L_u of dimension $n - 1$ that can be built solely from G and K . As $\Delta^{n-1} \neq L_u$ (one is affine, the other is not), the intersection $\Delta^{n-1} \cap L_u$, if it exists, has a dimension at most $n - 2$. In particular, the Lebesgue measure of $\Delta^{n-1} \cap L_u$ (over Δ^{n-1}) is null, meaning that the probability to nullify all edges of K is null.

The number of subsets $K \subset E$ such that $|K| > d$ is finite, so the probability to nullify all edges of any of them is also null, so the probability that the polytope $\Pi_{\geq 0}$ associated to some $\lambda \in C$ admits an injective-only vertex is null. \blacktriangleleft

D Supplementary material of Section 6 (Matching rates optimization).

D.1 Proof of Proposition 18

We consider an arbitrary Markovian, stationary matching policy that matches at most $\mathcal{M} \in \mathbb{N}$ pairs of items in one time step. In particular, we consider a matching policy that only depends on the current state of the underlying Markov chain of the matching problem. To be more precise, consider an irreducible, aperiodic, and positive recurrent DTMC $(S_t, t \in \mathbb{N})$ defined in a similar way as in Appendix A, such that Q_t and Φ_t are functions of S_t , where $\Phi_{t,i}$ is the number of class- i items matched at time t for all $i \in V$. If $Q_t = S_t$, then the matching policy simply depends on the current queue lengths, e.g., ML policy. Now, given an injective-only vertex μ , the following proposition formalizes the trade-off between closeness to μ and the queue lengths in the steady-state (denoted by Q_∞) for any such matching policy. Proposition 18 is a special case of Extended Proposition 18 with $\mathcal{M} = 1$. The proof of Extended Proposition 18 is given in the rest of the appendix.

► **Extended Proposition 18.** *Consider an injective-only vertex μ of $\Pi_{\geq 0}$. Then, under any Markovian, stationary matching policy Φ that matches at most $\mathcal{M} \in \mathbb{N}$ pairs of items in one-time step, if $\|\mu(\Phi) - \mu\|_1 \leq \epsilon$, then $\mathbb{E}[\|Q_\infty\|_2^2] \geq \Omega(1/\epsilon)$, where Q_∞ is the queue length under the steady state of $(S_t, t \in \mathbb{N})$.*

Proof. As μ is injective-only, there exists a connected component of $G^* = (V, E^*)$ that is bipartite by Definitions 3 and 4. Denote such a connected component by $G^\circ = (V_+ \cup V_-, E^\circ)$, where V_+ and V_- are the (disjoint) parts of the connected component, and $E^\circ \subseteq E^*$. Now, consider the following test function:

$$W(q) = \left\langle \mathbf{1}^{(V_+, V_-)}, q \right\rangle^2, \quad q \in \mathcal{Q},$$

where $\mathbf{1}_i^{(V_+, V_-)} = 1$ for $i \in V_+$, -1 for $i \in V_-$, and 0 otherwise. For each $t \in \mathbb{N}$, let A_t denote the class of the incoming item at time t , so that $A_t = \mathbf{1}_{I_t}$, where I_t is the class of $(t+1)$ -th item. Thus, $\lambda_i / \sum_{j \in V} \lambda_j$ is the probability that $A_t = \mathbf{1}_i$. We also define $\delta M_{t,k} = M_{t,k} - M_{t-1,k}$ to be the number of items matched using edge $k \in E$ at time t . Thus, we have $\sum_{k \in E_i} \delta M_{t,k} = \Phi_{t,i}$ for all $i \in V$. For the rest of the proof, we consider the steady-state of the Markov chain underlying the matching problem. In particular, we consider the steady-state of the (discrete-time) Markov chain $(S_t, t \in \mathbb{N})$, of which $Q = (Q_t, t \in \mathbb{N})$ and $\Phi = (\Phi_t, t \in \mathbb{N})$ are functions.

If $\mathbb{E}[W(Q)] = \infty$, then we have $\mathbb{E}[\sum_{i \in V} Q_i^2] = \infty$ which completes the proof of the proposition. Indeed, for each $q \in \mathbb{N}^n$,

$$\begin{aligned} W(q) &= \left(\sum_{i \in V_+} q_i \right)^2 + \left(\sum_{i \in V_-} q_i \right)^2 - 2 \left(\sum_{i \in V_+} q_i \right) \left(\sum_{i \in V_-} q_i \right), \\ &\leq \left(\sum_{i \in V_+} q_i \right)^2 + \left(\sum_{i \in V_-} q_i \right)^2, \\ &\leq |V_+| \sum_{i \in V_+} q_i^2 + |V_-| \sum_{i \in V_-} q_i^2. \end{aligned}$$

Now, assume instead that $\mathbb{E}[W(Q)] < \infty$. In stationarity, we have $\mathbb{E}[W(Q_{t+1})] = \mathbb{E}[W(Q_t)]$. In particular, we take the expectation under the steady state distribution of $(S_t, t \in \mathbb{N})$.

Using this, we get

$$\begin{aligned}
0 &= \mathbb{E} \left[\left\langle \mathbf{1}^{(V_+, V_-)}, Q_{t+1} \right\rangle^2 - \left\langle \mathbf{1}^{(V_+, V_-)}, Q_t \right\rangle^2 \right], \\
&= \mathbb{E} \left[\left\langle \mathbf{1}^{(V_+, V_-)}, Q_t + A_t - \Phi_t \right\rangle^2 - \left\langle \mathbf{1}^{(V_+, V_-)}, Q_t \right\rangle^2 \right], \\
&= \mathbb{E} \left[\left\langle \mathbf{1}^{(V_+, V_-)}, A_t - \Phi_t \right\rangle^2 \right] + 2\mathbb{E} \left[\left\langle \mathbf{1}^{(V_+, V_-)}, Q_t \right\rangle \left\langle \mathbf{1}^{(V_+, V_-)}, A_t - \Phi_t \right\rangle \right], \\
&\stackrel{(a)}{=} \mathbb{E} \left[\left\langle \mathbf{1}^{(V_+, V_-)}, A_t - \Phi_t \right\rangle^2 \right] - 2\mathbb{E} \left[\left\langle \mathbf{1}^{(V_+, V_-)}, Q_t \right\rangle \left\langle \mathbf{1}^{(V_+, V_-)}, \Phi_t \right\rangle \right], \\
&\stackrel{(b)}{\geq} \frac{\sum_{i \in V_+ \cup V_-} \lambda_i}{\sum_{i \in V} \lambda_i} - 2\mathbb{E} \left[\left\langle \mathbf{1}^{(V_+, V_-)}, Q_t \right\rangle \left\langle \mathbf{1}^{(V_+, V_-)}, \Phi_t \right\rangle \right], \\
&\geq \frac{\sum_{i \in V_+ \cup V_-} \lambda_i}{\sum_{i \in V} \lambda_i} - 2\sqrt{\mathbb{E} \left[\left\langle \mathbf{1}^{(V_+, V_-)}, Q_t \right\rangle^2 \right] \mathbb{E} \left[\left\langle \mathbf{1}^{(V_+, V_-)}, \Phi_t \right\rangle^2 \right]} \quad (\text{by Cauchy-Schwarz}), \\
&\stackrel{(c)}{\geq} \frac{\sum_{i \in V_+ \cup V_-} \lambda_i}{\sum_{i \in V} \lambda_i} - 4\sqrt{\mathbb{E} \left[\left(\sum_{i \in V} Q_{t,i} \right)^2 \right] \mathbb{E} \left[\left(\sum_{k \in E \setminus E^*} \delta M_{t,k} \right)^2 \right]}, \\
&\stackrel{(d)}{\geq} \frac{\sum_{i \in V_+ \cup V_-} \lambda_i}{\sum_{i \in V} \lambda_i} - 4\sqrt{\mathcal{M} \mathbb{E} \left[\left(\sum_{i \in V} Q_{t,i} \right)^2 \right] \sum_{k \in E \setminus E^*} \mathbb{E} [\delta M_{t,k}]}, \\
&\stackrel{(e)}{\geq} \frac{\sum_{i \in V_+ \cup V_-} \lambda_i}{\sum_{i \in V} \lambda_i} - 4\sqrt{\mathcal{M} \epsilon \mathbb{E} \left[\left(\sum_{i \in V} Q_{t,i} \right)^2 \right]}. \tag{27}
\end{aligned}$$

Here, (a) follows by noting that $\mathbb{E} [\langle \mathbf{1}^{(V_+, V_-)}, A_t \rangle] = \mathbb{E} [\langle \mathbf{1}^{(V_+, V_-)}, \lambda \rangle] = 0$, as μ is a feasible solution to the conservation equations (CE) and $G^\circ = (V_+ \cup V_-, E^\circ)$ is a connected component of the subgraph of G induced by support of μ . Note that we also use the independence of A_t and Q_t in this step. Next, (b) follows by noting that $\langle \mathbf{1}^{(V_+, V_-)}, A_t - \Phi_t \rangle^2 \geq 1$ whenever $\langle \mathbf{1}^{(V_+, V_-)}, A_t \rangle = \pm 1$ (because $\langle \mathbf{1}^{(V_+, V_-)}, \Phi_t \rangle$ is an even number, as each match is counted twice), and that the event $\langle \mathbf{1}^{(V_+, V_-)}, A_t \rangle = \pm 1$ happens with probability $\frac{\sum_{i \in V_+ \cup V_-} \lambda_i}{\sum_{i \in V} \lambda_i}$. Further, (c) follows by noting that

$$\begin{aligned}
\left| \left\langle \mathbf{1}^{(V_+, V_-)}, \Phi_t \right\rangle \right| &= \left| \sum_{i \in V_+} \Phi_{t,i} - \sum_{i \in V_-} \Phi_{t,i} \right| = \left| \sum_{i \in V_+} \sum_{j \in V_i} \delta M_{t,\{i,j\}} - \sum_{i \in V_-} \sum_{j \in V_i} \delta M_{t,\{i,j\}} \right|, \\
&\leq \sum_{i \in V_+} \sum_{j \in V_i \setminus V_-} \delta M_{t,\{i,j\}} + \sum_{i \in V_-} \sum_{j \in V_i \setminus V_+} \delta M_{t,\{i,j\}}, \\
&\quad + \left| \sum_{i \in V_+} \sum_{j \in V_i \cap V_-} \delta M_{t,\{i,j\}} - \sum_{i \in V_-} \sum_{j \in V_i \cap V_+} \delta M_{t,\{i,j\}} \right|, \quad (\text{Triangle Inequality}) \\
&\stackrel{(*)}{=} \sum_{i \in V_+} \sum_{j \in V_i \setminus V_-} \delta M_{t,\{i,j\}} + \sum_{i \in V_-} \sum_{j \in V_i \setminus V_+} \delta M_{t,\{i,j\}}, \\
&\stackrel{(**)}{\leq} 2 \sum_{k \in E \setminus E^*} \delta M_{t,k},
\end{aligned}$$

where (*) follows by noting that $\sum_{i \in V_+} \sum_{j \in V_i \cap V_-} \delta M_{t,\{i,j\}} = \sum_{i \in V_-} \sum_{j \in V_i \cap V_+} \delta M_{t,\{i,j\}}$ as the set of edges that we are summing over are the same in the left and right hand side. In

particular, we are adding over all the edges between V_+ and V_- . Next, (**) follows by noting that all edges connecting $i \in V_+$ to $j \in V \setminus V_-$ are not in E^* . Similarly, all edges connecting $i \in V_-$ to $j \in V \setminus V_+$ are also not in E^* by the definition of G° . Now, (d) follows by the assumption that at most \mathcal{M} pairs of items can be matched in a single time step. Lastly, (e) follows as $\mathbb{E}[\delta M_t] = \mu(\Phi)$ in the steady state and

$$\|\mu(\Phi) - \mu\|_1 \leq \epsilon \implies \sum_{k \in E \setminus E^*} \mu(\Phi)_k \leq \epsilon.$$

Now, to complete the proof, we use (27) which implies

$$\mathbb{E} \left[\sum_{i \in V} Q_{\infty, i}^2 \right] \geq \frac{1}{|V|} \mathbb{E} \left[\left(\sum_{i \in V} Q_{\infty, i} \right)^2 \right] \stackrel{(27)}{\geq} \frac{\left(\sum_{i \in V_+ \cup V_-} \lambda_i \right)^2}{16\mathcal{M}|V| \left(\sum_{i \in V} \lambda_i \right)^2} \cdot \frac{1}{\epsilon} = \Omega \left(\frac{1}{\epsilon} \right).$$

Note that the first inequality follows as $\|q\|_2^2 \geq \|q\|_1^2/|V|$ for all $q \in \mathbb{R}^{|V|}$. This completes the proof. \blacktriangleleft

D.2 Proof of Proposition 19

We first prove two intermediary results, Lemmas 31 and 32, that will be instrumental in proving the upper bound for the mean queue length in Proposition 19.

► **Lemma 31.** *Consider a matching problem (G, λ) . If the CRP gap $\delta(G, \lambda)$ is positive, then the ML policy stabilizes (G, λ) , and*

$$\mathbb{E} \left[\sum_{i \in V} Q_i \right] \leq \frac{n}{2\delta(G, \lambda)} \sum_{i \in V} \lambda_i.$$

Proof. Consider the quadratic test function W defined on \mathcal{Q}_G by

$$W(q) = \sum_{i \in V} q_i^2, \quad q \in \mathcal{Q}_G.$$

We define the drift of this test function by $\Delta W(q) = \mathbb{E}[W(Q_{t+1}) - W(Q_t) | Q_t = q]$ for each $q \in \mathcal{Q}_G$. Lastly, we let $\mathcal{I}_q = \{i \in V : q_i > 0\}$ denote the support of q , for each $q \in \mathcal{Q}_G$. Now, we simplify the drift under the ML policy using the queue evolution recursion given by (2). For all $q \in \mathcal{Q}_G \setminus \{0\}$, we have

$$\begin{aligned} \Delta W(q) &= \mathbb{E} \left[\sum_{i \in V} (Q_{t,i} + \mathbf{1}\{i = I_t\})^2 \mathbf{1}\{J_t = \perp\} + (Q_{t,i} - \mathbf{1}\{i = J_t\})^2 \mathbf{1}\{J_t \neq \perp\} - Q_{t,i}^2 \mid Q_t = q \right], \\ &= \mathbb{E} \left[\sum_{i \in V} \mathbf{1}\{i = I_t, J_t = \perp\} + \mathbf{1}\{i = J_t\} + 2(\mathbf{1}\{i = I_t, J_t = \perp\} - \mathbf{1}\{i = J_t\}) Q_{t,i} \mid Q_t = q \right], \\ &\stackrel{(a)}{=} 1 + 2 \sum_{i \in V} q_i \mathbb{E} \left[\mathbf{1}\{i = I_t, J_t = \perp\} - \mathbf{1}\{i = J_t\} \mid Q_t = q \right], \\ &\stackrel{(b)}{=} 1 + \frac{2}{\sum_{i \in V} \lambda_i} \sum_{i \in V} \lambda_i q_i - 2 \sum_{i \in V} q_i \mathbb{E} \left[\mathbf{1}\{i = J_t\} \mid Q_t = q \right], \\ &\stackrel{(c)}{=} 1 + \frac{2}{\sum_{i \in V} \lambda_i} \sum_{i \in V} \lambda_i q_i - \frac{2}{\sum_{i \in V} \lambda_i} \sum_{j \in V} \lambda_j \max_{i \in V_j} q_i, \end{aligned}$$

$$\begin{aligned}
&\stackrel{(d)}{=} 1 + \frac{2}{\sum_{i \in V} \lambda_i} \sum_{i \in \mathcal{I}_q} \lambda_i q_i - \frac{2}{\sum_{i \in V} \lambda_i} \sum_{j \in V(\mathcal{I}_q)} \lambda_j \max_{i \in V_j} q_i, \\
&\stackrel{(e)}{\leq} 1 - \frac{2\delta(G, \lambda)}{|\mathcal{I}_q| \sum_{i \in V} \lambda_i} \sum_{i \in \mathcal{I}_q} q_i, \\
&= 1 - \frac{2\delta(G, \lambda)}{|\mathcal{I}_q| \sum_{i \in V} \lambda_i} \sum_{i \in V} q_i, \\
&\leq 1 - \frac{2\delta(G, \lambda)}{n \sum_{i \in V} \lambda_i} \sum_{i \in V} q_i. \tag{28}
\end{aligned}$$

Here, (a) follows as

$$\sum_{i \in V} \mathbb{1}\{i = I_t, J_t = \perp\} + \sum_{i \in V} \mathbb{1}\{i = J_t\} = \mathbb{1}\{J_t = \perp\} + \mathbb{1}\{J_t \neq \perp\} = 1.$$

Next, (b) follows as for all $i \in V$, we have

$$\begin{aligned}
q_i \mathbb{E}[\mathbb{1}\{i = I_t, J_t = \perp\} | Q_t = q] &= q_i \mathbb{P}\{i = I_t, J_t = \perp | Q_t = q\}, \\
&= q_i \mathbb{P}\{J_t = \perp | Q_t = q, i = I_t\} \mathbb{P}\{i = I_t | Q_t = q\}, \\
&= q_i \frac{\lambda_i}{\sum_{i \in V} \lambda_i} \mathbb{1}\left\{\max_{j \in V_i} q_j = 0\right\} = q_i \frac{\lambda_i}{\sum_{i \in V} \lambda_i}.
\end{aligned}$$

The last equality follows from the fact that the ML policy is greedy, which implies that, if $q_i > 0$, then $q_j = 0$ for all $j \in V_i$. Now, (c) follows by using the definition of ML policy. In particular, we have

$$\begin{aligned}
\sum_{i \in V} q_i \mathbb{E}\left[\mathbb{1}\{i = J_t\} \mid Q_t = q\right] &= \frac{1}{\sum_{j \in V} \lambda_j} \sum_{j \in V} \lambda_j \sum_{i \in V} q_i \mathbb{E}\left[\mathbb{1}\{i = J_t\} \mid Q_t = q, j = I_t\right], \\
&= \frac{1}{\sum_{j \in V} \lambda_j} \sum_{j \in V} \lambda_j \max_{i \in V_j} q_i.
\end{aligned}$$

The equality (d) follows by noting that \mathcal{I}_q is equal to the support of q , i.e., we have $\mathcal{I}_q = \{i \in V : q_i > 0\}$. The rest of the paragraph is dedicated to proving (e). Consider the bipartite graph

$$G_q = G(\mathcal{I}_q \cup V(\mathcal{I}_q), E_q) \quad \text{with} \quad E_q = \{\{i, j\} \in E : i \in \mathcal{I}_q, j \in V(\mathcal{I}_q)\}.$$

By definition of $\delta(G, \lambda)$, we have, for each $\mathcal{I} \subseteq \mathcal{I}_q$,

$$\sum_{i \in \mathcal{I}} \lambda_i \leq \sum_{j \in V(\mathcal{I})} \lambda_j - \delta(G, \lambda) = \sum_{\substack{j \in V(\mathcal{I}): \\ \exists i \in \mathcal{I}, \{i, j\} \in E_q}} \lambda_j - \delta(G, \lambda),$$

where the equality follows as $V(\mathcal{I}) \subseteq V(\mathcal{I}_q)$, and so, $V(\mathcal{I}) = V(\mathcal{I}) \cap \{j : \exists i \in \mathcal{I}, \{i, j\} \in E_q\}$ by the definition of G_q . Now, the above inequality implies that

$$\sum_{i \in \mathcal{I}} \left(\lambda_i + \frac{\delta(G, \lambda)}{|\mathcal{I}_q|} \right) \leq \sum_{\substack{j \in V(\mathcal{I}): \\ \exists i \in \mathcal{I}, \{i, j\} \in E_q}} \lambda_j \quad \forall \mathcal{I} \subseteq \mathcal{I}_q.$$

So, by the weighted version of Hall's marriage theorem [21, Lemma 2.5], there exists $\tilde{\mu} \in \mathbb{R}_{\geq 0}^{|E_q|}$ such that

$$\lambda_i + \frac{\delta(G, \lambda)}{|\mathcal{I}_q|} = \sum_{\substack{j \in V(\mathcal{I}_q): \\ \{i, j\} \in E_q}} \tilde{\mu}_{i, j}, \quad \forall i \in \mathcal{I}_q, \quad \lambda_j \geq \sum_{\substack{i \in \mathcal{I}_q: \\ \{i, j\} \in E_q}} \tilde{\mu}_{i, j}, \quad \forall j \in V(\mathcal{I}_q).$$

Using the above, we get

$$\sum_{i \in \mathcal{I}_q} \lambda_i q_i - \sum_{j \in V(\mathcal{I}_q)} \lambda_j \max_{i' \in V_j} q_{i'} \leq \sum_{\{i,j\} \in E_q} \tilde{\mu}_{i,j} \left(q_i - \max_{i' \in V_j} q_{i'} \right) - \frac{\delta(G, \lambda)}{|\mathcal{I}_q|} \sum_{i \in \mathcal{I}_q} q_i \leq -\frac{\delta(G, \lambda)}{|\mathcal{I}_q|} \sum_{i \in \mathcal{I}_q} q_i,$$

where the last inequality follows as $\max_{i' \in V_j} q_{i'} \geq q_i$ for all $\{i, j\} \in E_q$. This completes the proof of (e).

Now, by (28) we immediately conclude that

$$\Delta W(q) \leq -1 \quad \text{for each } q \in \mathcal{Q}_G \text{ such that } \sum_{i \in V} q_i \geq \frac{n \sum_{i \in V} \lambda_i}{\delta(G, \lambda)}.$$

Thus, by the Foster-Lyapunov theorem [39], we conclude from (28) that the Markov chain under the ML policy is positive recurrent. Now, by the moment bound theorem [24, Proposition 6.14], we further conclude that in steady state, we have

$$\mathbb{E} \left[\sum_{i \in V} Q_i \right] \leq \frac{n}{2\delta(G, \lambda)} \sum_{i \in V} \lambda_i.$$

This completes the proof of the lemma. \blacktriangleleft

Now we turn our attention to studying the CRP gap of (G', λ'_ϵ) . We start by defining additional notations for convenience. Define the set of “positive” vertices by $V_+ = \{i^+ : i \in V\}$ and the set of “negative” vertices by $V_- = \{i^- : i \in V\}$. Let $\mathbb{I}_{G'}$ be the set of all independent sets in G' . Note that, each independent set $\mathcal{I} \in \mathbb{I}_{G'}$ can be uniquely written as $\mathcal{I} = J \cup K_+$ where $J = \mathcal{I} \cap V_-$ and $K_+ = \mathcal{I} \cap V_+$ (with at least one of them nonempty). Observe we have $J \in \{\emptyset\} \cup \mathbb{I}_G$ and $K_+ \in \{\emptyset\} \cup \mathbb{I}_{G^*}$. Now, we lower bound the CRP gap of the augmented graph in the following lemma. In particular, combining the following lemma with Proposition 8 implies that (G', λ'_ϵ) is stabilizable whenever (G, λ) is.

► **Lemma 32.** *For the augmented problem (G', λ'_ϵ) , we have*

$$\delta(G', \lambda'_\epsilon) \geq \min \left\{ \epsilon \min_{i \in V} \lambda_i, (1 - \epsilon) \min_{i \in V} \lambda_i, \delta(G, \lambda) \right\}.$$

Proof. Consider $\mathcal{I} \in \mathbb{I}_{G'}$ and let $J \in \{\emptyset\} \cup \mathbb{I}_G$ and $K_+ \in \{\emptyset\} \cup \mathbb{I}_{G^*}$ such that $J = \mathcal{I} \cap V_-$ and $K_+ = \mathcal{I} \cap V_+$, so that $\mathcal{I} = J \cup K_+$. Using the definition of G' , we can verify that, for each $0 < \epsilon < 1$, we have

$$\begin{aligned} \sum_{i \in V_{G'}(\mathcal{I})} \lambda'_{\epsilon, i} - \sum_{i \in \mathcal{I}} \lambda'_{\epsilon, i} &= \epsilon \left(\sum_{i \in V_G(J)} \lambda_i - \sum_{i \in J} \lambda_i \right) + (1 - \epsilon) \sum_{i \in V_G(J) \setminus V_{G^*}(K_+)} \lambda_i \\ &\quad + (1 - \epsilon) \left(\sum_{i \in V_{G^*}(K_+)} \lambda_i - \sum_{i \in K_+} \lambda_i \right) + \epsilon \sum_{i \in V_{G^*}(K_+) \setminus V_G(J)} \lambda_i. \end{aligned} \quad (29)$$

Now, we consider the following two cases.

Case I: Either $V_G(K_+) \not\subseteq V_G(J)$ or $V_G(J) \not\subseteq V_{G^*}(K_+)$. In this case, either the second term or the fourth term in (29) is non-zero. So, we have

$$\sum_{i \in V_{G'}(\mathcal{I})} \lambda'_{\epsilon, i} - \sum_{i \in \mathcal{I}} \lambda'_{\epsilon, i} \geq \min \{ \epsilon, 1 - \epsilon \} \min_{i \in V} \lambda_i \quad \forall \mathcal{I} \in \mathbb{I}_{G'}. \quad (30)$$

Case II: We have $V_G(K) \subseteq V_G(J)$ and $V_G(J) \subseteq V_{G^*}(K)$. In this case, we must have $J \neq \emptyset$ and $K \neq \emptyset$. Thus, we have

$$V_{G^*}(K) \subseteq V_G(K) \subseteq V_G(J) \subseteq V_{G^*}(K) \implies V_{G^*}(K) = V_G(K).$$

Now, as $K \in \mathbb{I}_{G^*}$, we have $K \cap V_{G^*}(K) = \emptyset$. Thus, we have $K \cap V_G(K) = \emptyset$ which implies that $K \in \mathbb{I}_G$. Now, using the above, we get

$$\sum_{i \in V_{G^*}(K)} \lambda_i - \sum_{i \in K} \lambda_i = \sum_{i \in V_G(K)} \lambda_i - \sum_{i \in K} \lambda_i \geq \delta(G, \lambda), \quad \text{and} \quad \sum_{i \in V_G(J)} \lambda_i - \sum_{i \in J} \lambda_i \geq \delta(G, \lambda),$$

where the inequalities above follow by the definition of $\delta(G, \lambda)$ and noting that $J, K \in \mathbb{I}_G$. Now, using (29), we get

$$\sum_{i \in V_{G'}(\mathcal{I})} \lambda'_{\epsilon, i} - \sum_{i \in \mathcal{I}} \lambda'_{\epsilon, i} \geq \delta(G, \lambda). \quad (31)$$

Now, by combining (30) and (31), for all $\mathcal{I} \in \mathbb{I}_{G'}$ we get

$$\sum_{i \in V_{G'}(\mathcal{I})} \lambda'_{\epsilon, i} - \sum_{i \in \mathcal{I}} \lambda'_{\epsilon, i} \geq \min \left\{ \epsilon \min_{i \in V} \lambda_i, (1 - \epsilon) \min_{i \in V} \lambda_i, \delta(G, \lambda) \right\},$$

which completes the proof by definition of $\delta(G', \lambda'_\epsilon)$. \blacktriangleleft

Now that we have stated and proved Lemmas 31 and 32, we are ready to prove Proposition 19 below.

Proof of Proposition 19. We divide the proof into two parts. In the first part, we show that the expected queue length is $O(1/\epsilon)$. Then, the second part deals with the matching rates.

Controlling the expected queue length: By Lemma 32, we know that $\delta(G', \lambda'_\epsilon) > 0$. The corresponding queueing system is stable by Proposition 8. Furthermore, we have

$$\mathbb{E} \left[\sum_{i \in V} Q_i \right] = \mathbb{E} \left[\sum_{i \in V^+ \cup V^-} Q_i \right] \leq \frac{2n}{2\delta(G', \lambda'_\epsilon)} \sum_{i \in V} \lambda_i \leq \frac{n}{\epsilon \min_{i \in V} \lambda_i} \sum_{i \in V} \lambda_i,$$

where the first inequality follows by applying Lemma 31 to the extended matching problem (G', λ'_ϵ) , and the second inequality follows by applying Lemma 32 and noting there exists $\epsilon_0 \in (0, 1)$ such that, for all $\epsilon \in (0, \epsilon_0)$, we have $\epsilon \leq 1 - \epsilon$ and $\epsilon \leq \delta(G, \lambda) / \min_{i \in V} \lambda_i$. Now, by defining $C_1 = \frac{n}{\min_{i \in V} \lambda_i} \sum_{i \in V} \lambda_i$, the proof of the first part is complete.

Controlling the matching rates: Now, we show the second part of the proposition. Let Φ'_ϵ denote the ML policy adapted to (G', λ'_ϵ) . As explained at the beginning of Section 6.2.2, the policies Φ_ϵ and Φ'_ϵ are practically equivalent, except that Φ_ϵ is seen as a (non-greedy ϵ -dependent) policy adapted to the original problem (G, λ) , while Φ'_ϵ is seen as a (greedy) policy adapted to the (ϵ -dependent) extended problem (G', λ'_ϵ) . Consistently, $\mu(\Phi_\epsilon)$ is the matching-rate vector (indexed by E) under the ϵ -filtering ML policy adapted to (G, λ) with filter E^* , while $\mu(\Phi'_\epsilon)$ denotes the matching-rate vector (indexed by E') under the ML policy adapted to (G', λ'_ϵ) . The rest of the proof is divided mainly in two parts: first we prove the following bound on our objective $\|\mu - \mu(\Phi_\epsilon)\|_1$ in terms of $\mu(\Phi'_\epsilon)$:

$$\|\mu - \mu(\Phi_\epsilon)\|_1 \leq \sum_{\{i, j\} \in E^*} |\mu_{i, j} - \mu(\Phi'_\epsilon)_{i^+, j^+}| + 2\epsilon \sum_{j \in V} \lambda_j. \quad (32)$$

Then, we show that $\sum_{\{i,j\} \in E^*} |\mu_{i,j} - \mu(\Phi'_\epsilon)_{i^+,j^+}| = O(\epsilon)$ using properties of ML policy to complete the proof. Intuitively, λ'_ϵ is designed in such a way that under any stable policy, all but $O(\epsilon)$ arrivals are matched using edges in E^+ . So, $\mu(\Phi'_\epsilon)_{i^+,j^+}$ must be close to $\mu_{i,j}$ for all $\{i,j\} \in E^*$.

We start by proving (32). As the ϵ -filtering ML policy mimics the ML policy adapted to (G', λ'_ϵ) , we have

$$\mu(\Phi_\epsilon(\mu))_{i,j} = \begin{cases} \mu(\Phi'_\epsilon)_{i^+,j^-} + \mu(\Phi'_\epsilon)_{i^-,j^+} + \mu(\Phi'_\epsilon)_{i^-,j^-} + \mu(\Phi'_\epsilon)_{i^+,j^+} & \text{if } \{i,j\} \in E^* \\ \mu(\Phi'_\epsilon)_{i^+,j^-} + \mu(\Phi'_\epsilon)_{i^-,j^+} + \mu(\Phi'_\epsilon)_{i^-,j^-} & \text{if } \{i,j\} \in E \setminus E^*, \end{cases}$$

Using the above, we can upper bound our objective $\|\mu - \mu(\Phi_\epsilon)\|_1$ in terms of $\mu(\Phi'_\epsilon)$ as follows:

$$\begin{aligned} \|\mu - \mu(\Phi_\epsilon)\|_1 &\stackrel{(a)}{=} \sum_{\{i,j\} \in E^*} |\mu_{i,j} - \mu(\Phi_\epsilon)_{i,j}| + \sum_{\{i,j\} \in E \setminus E^*} \mu(\Phi_\epsilon)_{i,j}, \\ &\stackrel{(b)}{\leq} \sum_{\{i,j\} \in E^*} |\mu_{i,j} - \mu(\Phi'_\epsilon)_{i^+,j^+}| + \sum_{\{i,j\} \in E} \left(\mu(\Phi'_\epsilon)_{i^+,j^-} + \mu(\Phi'_\epsilon)_{i^-,j^+} + \mu(\Phi'_\epsilon)_{i^-,j^-} \right), \\ &\stackrel{(c)}{\leq} \sum_{\{i,j\} \in E^*} |\mu_{i,j} - \mu(\Phi'_\epsilon)_{i^+,j^+}| + 2\epsilon \sum_{j \in V} \lambda_j. \end{aligned}$$

Here, (a) follows by recalling that E^* is defined as the support of μ . Next, (b) follows as, for all $\{i,j\} \in E^*$,

$$\begin{aligned} |\mu_{i,j} - \mu(\Phi_\epsilon)_{i,j}| &= |\mu_{i,j} - \mu(\Phi'_\epsilon)_{i^+,j^+} - \mu(\Phi'_\epsilon)_{i^+,j^-} - \mu(\Phi'_\epsilon)_{i^-,j^+} - \mu(\Phi'_\epsilon)_{i^-,j^-}|, \\ &\leq |\mu_{i,j} - \mu(\Phi'_\epsilon)_{i^+,j^+}| + \mu(\Phi'_\epsilon)_{i^+,j^-} + \mu(\Phi'_\epsilon)_{i^-,j^+} + \mu(\Phi'_\epsilon)_{i^-,j^-}. \end{aligned}$$

Lastly, (c) follows by noting that

$$\begin{aligned} \sum_{\{i,j\} \in E} \left(\mu(\Phi'_\epsilon)_{i^+,j^-} + \mu(\Phi'_\epsilon)_{i^-,j^+} \right) &\leq \sum_{j \in V} \sum_{\substack{i \in V: \\ \{i,j\} \in E}} \left(\mu(\Phi'_\epsilon)_{j^-,i^+} + \mu(\Phi'_\epsilon)_{j^-,i^-} \right) \stackrel{(*)}{=} \sum_{j \in V} \lambda'_{\epsilon,j^-} = \epsilon \sum_{j \in V} \lambda_j, \\ \sum_{\{i,j\} \in E} \mu(\Phi'_\epsilon)_{i^-,j^+} &\leq \sum_{i \in V} \sum_{\substack{j \in V: \\ \{i,j\} \in E}} \mu(\Phi'_\epsilon)_{i^-,j^+} \stackrel{(*)}{\leq} \sum_{i \in V} \lambda'_{\epsilon,i^-} = \epsilon \sum_{i \in V} \lambda_i, \end{aligned}$$

where (*) follows by applying the conservation equations (CE) to the extended matching model $(G', \lambda'_\epsilon, \Phi'_\epsilon)$. This completes the proof of (32).

Now, we proceed to show that $\sum_{\{i,j\} \in E^*} |\mu_{i,j} - \mu(\Phi'_\epsilon)_{i^+,j^+}| = O(\epsilon)$. Recall that, by Proposition 13, the rows of matrix A^* are linearly independent. This observation has multiple consequences:

- The system of linear equations $A^*y = \lambda$ of unknown y has a unique solution, which we know to be equal to μ .
- The system of linear equations $A^*y = \tilde{\lambda}$ of unknown y has a unique solution, which we know to be equal to $\tilde{\mu} = \{\mu(\Phi'_\epsilon)_{i^+,j^+}\}_{\{i,j\} \in E^*}$, where $\tilde{\lambda}$ is given by

$$\tilde{\lambda}_i = \sum_{\substack{j \in V: \\ \{i,j\} \in E^*}} \mu(\Phi'_\epsilon)_{i^+,j^+} \quad i \in V.$$

- The matrix $(A^*)^T A^*$ is invertible, and the Moore-Penrose inverse of A^* is $((A^*)^T A^*)^{-1} (A^*)^T$.

Thus, we have $\mu - \tilde{\mu} = ((A^*)^T A^*)^{-1} (A^*)^T (\lambda - \tilde{\lambda})$ which implies

$$\sum_{\{i,j\} \in E^*} |\mu_{i,j} - \mu(\Phi'_\epsilon)_{i^+,j^+}| \leq \|((A^*)^T A^*)^{-1} (A^*)^T\|_1 \|\lambda - \tilde{\lambda}\|_1. \quad (33)$$

Now, we upperbound the term $\|\lambda - \tilde{\lambda}\|_1$ as follows:

$$\begin{aligned} \|\lambda - \tilde{\lambda}\|_1 &= \sum_{i \in V} \left| \lambda_i - \sum_{\substack{j \in V: \\ \{i,j\} \in E^*}} \mu(\Phi'_\epsilon)_{i^+,j^+} \right|, \\ &\leq \sum_{i \in V} \left| (1 - \epsilon) \lambda_i - \sum_{\substack{j \in V: \\ \{i,j\} \in E^*}} \mu(\Phi'_\epsilon)_{i^+,j^+} \right| + \epsilon \sum_{i \in V} \lambda_i, \\ &\stackrel{(a)}{=} \sum_{i \in V} \sum_{\substack{j \in V: \\ \{i,j\} \in E}} \mu(\Phi'_\epsilon)_{i^+,j^-} + \epsilon \sum_{i \in V} \lambda_i, \\ &\stackrel{(b)}{\leq} 2\epsilon \sum_{i \in V} \lambda_i, \end{aligned} \quad (34)$$

where (a) again follows by applying (CE) to the extended matching model $(G', \lambda'_\epsilon, \Phi'_\epsilon)$, and (b) is justified as follows:

$$\sum_{i \in V} \sum_{\substack{j \in V: \\ \{i,j\} \in E}} \mu(\Phi'_\epsilon)_{i^+,j^-} \stackrel{(i)}{=} \sum_{j \in V} \sum_{\substack{i \in V: \\ \{i,j\} \in E}} \mu(\Phi'_\epsilon)_{i^+,j^-} \stackrel{(ii)}{\leq} \sum_{j \in V} \lambda'_{\epsilon,j^-} = \epsilon \sum_{j \in V} \lambda_j.$$

Here, (i) follows by switching the order of summation, and (ii) follows by noting that $\{i^+ \in V^+ : \{i^+, j^-\} \in E^\pm\} \subseteq \{i \in V' : \{i, j\} \in E'\}$ and again applying (CE) to the matching model $(G', \lambda'_\epsilon, \Phi'_\epsilon)$. Now, by combining (34) with (33), we get

$$\sum_{\{i,j\} \in E^*} |\mu_{i,j} - \mu(\Phi'_\epsilon)_{i^+,j^+}| \leq 2 \|((A^*)^T A^*)^{-1} (A^*)^T\|_1 \sum_{i \in V} \lambda_i \epsilon.$$

Lastly, by injecting the above inequality into (32) and defining $C_2 = 2(1 + \|((A^*)^T A^*)^{-1} (A^*)^T\|_1) \sum_{i \in V} \lambda_i$, the proof is complete. \blacktriangleleft

E

 Supplementary material of Section 7 (Practical performance).

E.1 k -filtering policies – Section 7.1 (Considered policies)

We conjecture that k -filtering policies, introduced in Section 7.1, yield a sequence of matching rate vectors that are arbitrarily close to a vertex of the polytope $\Pi_{\geq 0}$, even if this vertex is injective-only. This conjecture is supported by numerical results shown in Sections 7 and 8.

► **Conjecture 33.** Consider a vertex μ of $\Pi_{\geq 0}$ and let $G^* = (V, E^*)$ denote its support graph. For each $k \in \mathbb{N}$, consider the following semi-filtering policy, denoted by $\Phi_k(\mu)$:

- If the size of the longest queue is less than k , apply the filtering ML policy adapted to G with filter E^* ;
- Otherwise, apply the greedy ML policy adapted to G .

$\Phi_k(\mu)$ is stable for each $k \in \mathbb{N}$ and $\lim_{k \rightarrow \infty} \mu(\Phi_k(\mu)) = \mu$.

Intuition of the proof. In essence, the approach is as follows: we take a possibly unstable policy that achieves the desired goal (a matching rate vector equal to μ), and we make it stable by reverting to a stable policy when the queue sizes become too large. If the threshold is large enough, most matches will be made under the unstable policy that achieves the desired goal.

We present here a sketch of proof for the case where G^* is injective-only with a unique connected component, i.e. when G^* is a tree. If G^* is made of multiple connected components that are connected in G , the proof needs to compare the model to a virtual model where each component acts independently and accounts for the interactions between the components. For example, when one the queue is longer than k , all components may interact when decisions follow the ML policy on the graph G .

The stability of the model $(G, \lambda, \Phi_k(\mu))$ for each $k \in \mathbb{N}$ can be proved using the Lyapunov-Foster theorem, see [8, Theorem 1.1 in Chapter 5], using the fact that, outside a finite set of states (those where all queues are shorter than k), the policy $\Phi_k(\mu)$ behaves like the stable ML policy on the graph G .

Now, consider a random vector $Q = (Q_1, Q_2, \dots, Q_n)$ distributed like the vector of queue sizes in the matching model $(G, \lambda, \Phi_k(\mu))$ in stationary regime, and let $Q_+ = \sum_{i \in V_+} Q_i$ and $Q_- = \sum_{i \in V_-} Q_i$ denote the total queue sizes in the parts V_+ and V_- of the (bipartite) graph G^* . Also let $L_+ = \max(Q_i, i \in V_+)$, $L_- = \max(Q_i, i \in V_-)$, and $L = L_+$ if $L_+ \geq L_-$ and $L = -L_-$ if $L_+ < L_-$, so that $|L|$ is the size of the longest queue and the sign of L indicates the part (V_+ or V_-) to which this queue belongs. The random variable L will play a crucial role in the proof. Lastly, we let $p_\ell = \mathbb{P}(L = \ell)$, $\ell \in \{\dots, -2, -1, 0, 1, 2, \dots\}$, denote its distribution. The dependency of these random variables on k (via the policy $\Phi_k(\mu)$) is left implicit to simplify notation.

To show that $\lim_{k \rightarrow \infty} \mu(\Phi_k(\mu)) = \mu$, we need to upperbound the probability that a match is performed along an edge that does not belong to the support of the vertex μ . To do this, we upperbound the probability of applying the greedy ML policy on the whole graph G , i.e., the probability $\sum_{|\ell| \geq k} p_\ell$ that the size of the longest queue is k or more.

Intuitively, we expect that the stability of the ML policy implies a uniform drift of the size of the longest queue towards the origin whenever this size is larger than k , provided that k is large enough. Formally, we conjecture that there exist $K \in \mathbb{N}$ and $0 < \rho < 1$ such that, for each $k \geq K$, we have under Φ_k that $p_{\ell+1} \leq \rho p_\ell$ for each $\ell \geq k$ and $p_{\ell-1} \leq \rho p_\ell$ for each $\ell \leq -k$. In particular, $\sum_{\ell \geq k} p_\ell \leq \frac{1}{1-\rho} p_k$ and $\sum_{\ell \leq -k} p_\ell \leq \frac{1}{1-\rho} p_{-k}$. This implies that, if p_k and p_{-k} go to 0 as k goes to infinity, the probability that the filter on E^* is disabled goes to 0 as well.

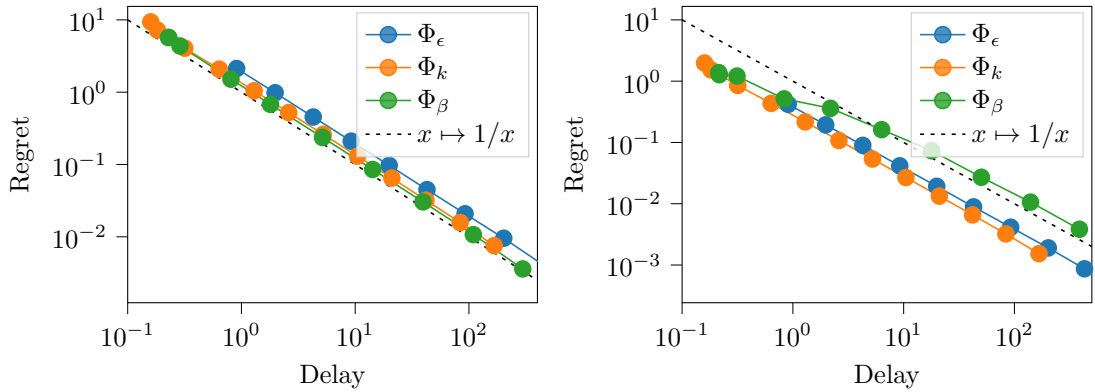
To control p_k and p_{-k} , we now need to consider the cases where the size of the longest queue is at most k . We argue that, as long as $|L| \leq k$, the dynamics of L is mainly controlled by an unbiased random walk between the two parts of the bipartite graph. More precisely, as long as $|L| \leq k$, the difference $Q_+ - Q_-$ behaves like an unbiased random walk because any arrival of a class in V_+ increases the difference by one, any arrival of a class in V_- decreases the difference by one, and the arrival rates of V_+ and V_- are equal (see Equation (26)). We conjecture that, when $|Q_+ - Q_-|$ is large enough (while still satisfying $|L| \leq k$), a distinctive structure emerges due to the filtered ML policy: the queues in one part of the graph have approximately equal sizes, while the queues in the other part are empty. In particular, we conjecture that the typical states of the system (when $|L| \leq k$) verify (in some sense) $L \approx \frac{Q_+ - Q_-}{|V_+|}$ if $Q_+ > Q_-$ and $L \approx \frac{Q_+ - Q_-}{|V_-|}$ if $Q_+ < Q_-$. In a true unbiased random walk over a bounded range of integers, all states have equal probabilities. By comparing the evolutions of L and $Q_+ - Q_-$ and taking into account the border cases ($|L| = k$), we conjecture that there exist $K' \in \mathbb{N}$ and $c > 0$ such that, if $k \geq K'$, then $p_\ell \geq c \max(p_k, p_{-k})$ for all $\ell \in \{-k, \dots, -1, 0, 1, \dots, k\}$. In particular, both p_k and p_{-k} should be upperbounded by $(c(2k + 1))^{-1}$ when k is large enough.

By combining the two cases $|L| \geq k$ and $|L| \leq k$, we get $\sum_{|\ell| \geq k} p_\ell = O(1/k)$. In other words, the probability that the size of the longest queue is greater than k tends to 0 as k goes to infinity. As matchings outside E^* occur only when the size of the longest queue is greater than k , we conclude that the matching rate of an edge outside E^* goes to 0, which by continuity of the conservation principle and injectivity of G^* implies that $\lim_{k \rightarrow \infty} \mu(\Phi_k(\mu)) = \mu$. ◀

E.2 Additional examples

Section 7 evaluated the performance of matching policies on the diamond graph. We present here simulations on other graphs.

E.2.1 Co-domino graph



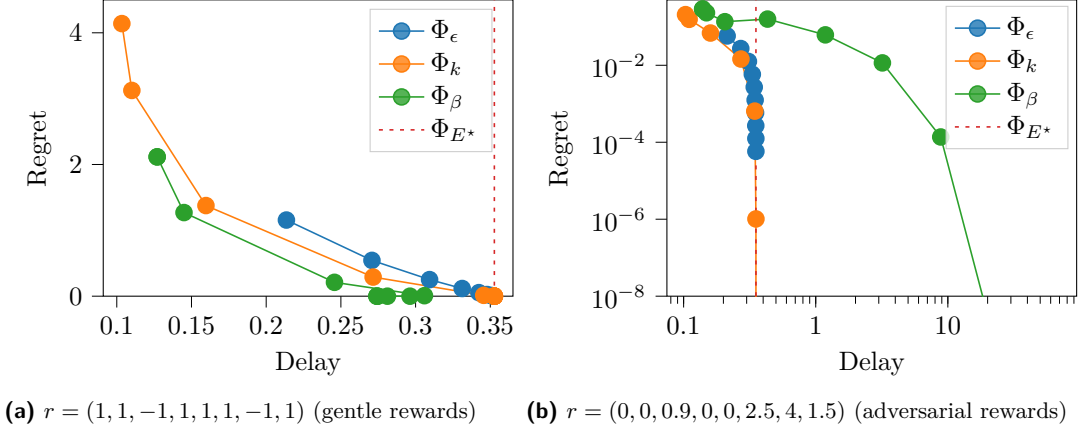
(a) $r = (-1, 1, 1, -1, -1, -1, 1, -1)$ (gentle rewards) (b) $r = (-1, 1, 1, -1, -1, 4.9, 1, 4.9)$ (adversarial rewards)

■ **Figure 21** Matching problem from Figure 19. Targeting the injective-only vertex from Figure 19f.

We first consider the matching problem depicted in Figure 19, with the goal of approaching the injective-only vertex shown in Figure 19f. This vertex disables 5 out of 8 edges: $\{1, 2\}$,

$\{2, 6\}$, $\{3, 4\}$, $\{3, 5\}$, and $\{5, 6\}$. The results are presented in Figure 21.

As observed in Section 7.3, both Φ_β and Φ_k perform best when rewards are gentle. However, the performance of Φ_β significantly deteriorates under adversarial rewards.



■ **Figure 22** Matching problem from Figure 18. Targeting the bijective vertex from Figure 18e.

Still focusing on the codomino graph, we now consider the matching problem from Figure 18. Our objective is to approach the bijective vertex illustrated in Figure 18e, which disables the edges $\{2, 3\}$ and $\{4, 5\}$. The results are presented in Figure 22, with a dotted line representing the delay of Φ_{E^*} , which is reward-optimal and stable for bijective vertices.

As seen in Section 7.4, both Φ_ϵ and Φ_k steadily converge to Φ_{E^*} . When gentle rewards are used, Φ_β manages to optimize the reward with a delay even lower than Φ_{E^*} . However, its performance deteriorates considerably under adversarial rewards.


E.2.2 Larger graph

To observe the performance of our methods on a larger graph, we generated a simple Erdős-Rényi graph with $n = 100$ nodes and an edge probability of $20/(n - 1)$, resulting in $m = 1032$ edges. The edge rewards were drawn uniformly and independently.

We aimed to study both an injective-only vertex and a bijective vertex. On graphs of this size, computing the exhaustive list of vertices is unrealistic, but once the reward and arrival-rate vectors are given, calculating the corresponding optimal vertex is straightforward.

To obtain an injective-only vertex, we used degree-proportional arrival rates, following the discussion in Section 5.2.2. This ensured the problem was stabilizable and, in this case, resulted in an optimal vertex with 96 positive edges (injective-only). For the bijective vertex, we added a noisy perturbation to the degree-proportional arrival rates, yielding a bijective vertex, as hinted by Proposition 14.

The performance results for both cases are shown in Figure 23. We observe that the performance of Φ_ϵ , Φ_k , and Φ_β are often comparable. If some level of regret is acceptable, Φ_β can achieve lower delays than the other policies. However, Φ_k tends to be more effective at minimizing regret, particularly converging to Φ_{E^*} in the bijective case.



(a) Injective-only vertex (degree-proportional arrivals). (b) Bijective vertex (perturbed degree-proportional arrivals)

■ **Figure 23** Reaching a vertex on a larger graph ($n = 100, m = 1032$).

F Supplementary material of Section 8 (Extensions of our results).

F.1 Convexity proofs for $\Pi_{\mathcal{P}}$ and $\Pi_{\mathcal{G}}$ (Propositions 20 and 21).

We consider the extended definition of matching policies, as discussed in Appendix A. Consider a stabilizable matching problem (G, λ) and let $\Pi_{\mathcal{P}}$ (resp. $\Pi_{\mathcal{G}}$) denote the set of vectors of matching rates achievable by stable policies (resp. by stable greedy policies) adapted to the compatibility graph G .

Convexity of $\Pi_{\mathcal{P}}$. Consider two (extended) policies Φ_1 and Φ_2 that stabilize the matching problem (G, λ) . The state-space, queue-size function, and empty state of Φ_1 (resp. Φ_2) are denoted by $\mathcal{S}_1, |\cdot|_1$, and \emptyset_1 (resp. $\mathcal{S}_2, |\cdot|_2$, and \emptyset_2). Given $0 < \beta < 1$, our goal is to build a matching policy Φ_β that also stabilizes the matching problem (G, λ) and satisfies $\mu(\Phi_\beta) = \beta\mu(\Phi_1) + (1 - \beta)\mu(\Phi_2)$.

Intuitively, the policy Φ_β will consistently follow the decisions of either Φ_1 or Φ_2 as long as the system is non-empty, and switch between these two policies each time the system becomes empty (therefore, at renewal times). The probability of selecting each policy after visiting the empty state will be set to achieve the desired matching rate vector in the long run. More formally, the policy Φ_β will have the following characteristics:

- State space $\mathcal{S}_\beta = (\mathcal{S}_1 \times \{\emptyset_2\}) \cup (\{\emptyset_1\} \times \mathcal{S}_2)$.
- Queue-size function $|\cdot|_\beta$ defined by $|(s_1, s_2)|_\beta = |s_1|_1 + |s_2|_2$ for each $(s_1, s_2) \in \mathcal{S}_\beta$.
- Empty state: $\emptyset_\beta = (\emptyset_1, \emptyset_2)$.

For some $0 < \gamma_\beta < 1$ that will be specified later, the matching policy Φ_β is defined as follows:

$$\begin{aligned} \Phi_\beta((\emptyset_1, \emptyset_2), i, j, (s_1, \emptyset_2)) &= \gamma_\beta \Phi_1(\emptyset_1, i, j, s_1), & s_1 \in \mathcal{S}_1, \quad i \in V, \quad j \in V \cup \{\perp\}, \\ \Phi_\beta((\emptyset_1, \emptyset_2), i, j, (\emptyset_1, s_2)) &= (1 - \gamma_\beta) \Phi_2(\emptyset_2, i, j, s_2), & s_2 \in \mathcal{S}_2, \quad i \in V, \quad j \in V \cup \{\perp\}, \\ \Phi_\beta((s_1, \emptyset_2), i, j, (s'_1, \emptyset_2)) &= \Phi_1(s_1, i, j, s'_1), & s_1 \in \mathcal{S}_1 \setminus \{\emptyset_1\}, \quad s'_1 \in \mathcal{S}_1, \\ \Phi_\beta((\emptyset_1, s_2), i, j, (\emptyset_1, s'_2)) &= \Phi_2(s_2, i, j, s'_2), & s_2 \in \mathcal{S}_2 \setminus \{\emptyset_2\}, \quad s'_2 \in \mathcal{S}_2. \end{aligned}$$

The first two equations say that, when the system is empty, we next apply Φ_1 with probability γ_β and Φ_2 with probability $1 - \gamma_\beta$. The third (resp. fourth) equation says that, once we start applying policy Φ_1 (resp. Φ_2), we keep applying this policy until we re-visit the empty state.

The stability of the policies Φ_1 and Φ_2 implies that of the policy Φ_β . According to the elementary renewal theorem for renewal reward processes, the vector giving the long-run

average matching rates under the policy Φ_β is given by

$$\begin{aligned}\mu(\Phi_\beta) &= \frac{\gamma_\beta T_1 \mu(\Phi_1) + (1 - \gamma_\beta) T_2 \mu(\Phi_2)}{\gamma_\beta T_1 + (1 - \gamma_\beta) T_2}, \\ &= \frac{\gamma_\beta T_1}{\gamma_\beta T_1 + (1 - \gamma_\beta) T_2} \mu(\Phi_1) + \frac{(1 - \gamma_\beta) T_2}{\gamma_\beta T_1 + (1 - \gamma_\beta) T_2} \mu(\Phi_2).\end{aligned}$$

where T_1 (resp. T_2) denotes the mean number of matches between two successive visits to the empty state \emptyset_1 (resp. \emptyset_2) in the matching model (G, λ, Φ_1) (resp. (G, λ, Φ_2)). We let the reader verify that the following value of γ_β yields $\mu(\Phi_\beta) = \beta \mu(\Phi_1) + (1 - \beta) \mu(\Phi_2)$:

$$\gamma_\beta = \frac{\beta T_2}{(1 - \beta) T_1 + \beta T_2}.$$

Convexity of Π_G . It suffices to observe that, if the policies Φ_1 and Φ_2 are greedy, so is the policy Φ_β .

F.2 Supplementary material of Section 8.2 (Greedy policies).

F.2.1 Proof of Proposition 23 (complete graphs) and discussion.

Consider a matching problem (K_n, λ) , where K_n is the complete graph with $n \geq 3$ nodes. According to Proposition 8, this matching problem is stabilizable if and only if $\lambda_i < \frac{1}{2} \sum_{j \in V} \lambda_j$ for each $i \in V$. We prove Proposition 23 by first demonstrating that all greedy policies are “equivalent”, in the sense that they always (deterministically) make the same matching decisions. The fact that we can formally define many greedy policies in this context is merely an artifact of our extended definition of a matching policy in Appendix A (see in particular Appendix A.2).

Notice that the only independent sets of a complete graph are the singletons, hence by (7) the state space of the queue-size process under greedy policies is

$$\mathcal{Q}_G(K_n) = \{0\} \cup \left(\bigcup_{i \in V} \{\ell \mathbb{1}_i : \ell \in \mathbb{N}_{>0}\} \right), \quad (35)$$

where $\mathbb{1}_i$ is the n -dimensional vector with one in coordinate i and zero elsewhere, for each $i \in V$. This observation implies that greediness leaves no freedom in the matching decisions: all unmatched items belong to the same class, and an incoming item must be matched with one of them if its class differs. This is formalized in Proposition 34.

► **Proposition 34.** *Consider the complete graph K_n with $n \geq 3$ nodes.*

- (i) *There is a unique size-based greedy policy that is adapted to the compatibility graph K_n . This policy, called the natural greedy policy and denoted by $\Phi_G(K_n)$, is the deterministic size-based policy defined on $\mathcal{Q}_G(K_n) \times V$ by*

$$\Phi_G(K_n)(q, i) = \begin{cases} j & \text{if } q_j \geq 1 \text{ with } j \neq i, \\ \perp & \text{otherwise.} \end{cases} \quad (36)$$

- (ii) *Consider a greedy policy Φ adapted to the graph K_n , and let $(\mathcal{S}, |\cdot|)$ denote its state space. The policy Φ makes the same decisions as the natural greedy policy Φ_G in the sense that*

$$\sum_{s' \in \mathcal{S}} \Phi(s, i, \Phi_G(K_n)(|s|, i), s') = 1, \quad s \in \mathcal{S}, \quad i \in V.$$

In other words, all greedy policies are equivalent in the sense that they can all be reduced to $\Phi_{\mathcal{G}}(K_n)$ (see “Equivalent policies” in Appendix A).

- (iii) For each greedy policy Φ adapted to the graph K_n and each sequence $I = (I_t, t \in \mathbb{N})$ of item classes, we have $(Q_{\mathcal{G}})_t = Q_t$ for each $t \in \mathbb{N}$, where $Q_{\mathcal{G}}$ and Q are the queue-size processes of the models $(K_n, I, \Phi_{\mathcal{G}}(K_n))$ and (K_n, I, Φ) , respectively.

Proof. We prove each statement one by one.

Proposition 34i. Consider a (possibly random) size-based greedy policy $\Phi : \mathcal{Q}_{\mathcal{G}}(K_n) \times V \times (V \cup \{\perp\}) \rightarrow [0, 1]$, as defined in Section 2.3.4, where $\mathcal{Q}_{\mathcal{G}}(K_n)$ is given in (35). Given $(q, i) \in \mathcal{Q}_{\mathcal{G}}(K_n) \times V$, the definition of $\mathcal{Q}_{\mathcal{G}}(K_n)$ implies that $\{j \in V_i : q_j \geq 1\}$ is either a singleton or the empty set. In the former case, letting j denote the unique element of the singleton, we have $\Phi(q, i, \perp) = 1$ if $i = j$, while the greediness of Φ implies that $\Phi(q, i, j) = 1$ if $i \neq j$. In the latter case, we have directly $\Phi(q, i, \perp) = 1$. In all cases, the matching decision is deterministic and makes the same decisions as in (36).

Proposition 34ii. The same argument can be repeated for an arbitrary greedy policy Φ with state space $(\mathcal{S}, |\cdot|)$. Given $(s, i) \in \mathcal{S} \times V$, we know that $|s| \in \mathcal{Q}_{\mathcal{G}}(K_n)$, so that $\{j \in V_i : |s|_j \geq 1\}$ is either a singleton or the empty set. In the former case, letting j denote the unique element of the singleton, we have $\sum_{s' \in \mathcal{S}} \Phi(s, i, \perp, s') = 1$ if $i = j$, while the greediness of Φ implies that $\sum_{s' \in \mathcal{S}} \Phi(s, i, j, s') = 1$ if $i \neq j$. In the latter case, we have directly $\sum_{s' \in \mathcal{S}} \Phi(s, i, \perp, s') = 1$. In all cases, we have $\sum_{s' \in \mathcal{S}} \Phi(s, i, \Phi_{\mathcal{G}}(|s|, i), s') = 1$.

Proposition 34iii. This statement follows directly from Proposition 34ii and from the fact that, under any matching policy, the dynamics of the queue-size process is fully specified by the sequence of incoming item classes and the matching decisions via (2). ◀

Proposition 35 below, which encompasses Proposition 23, builds on this equivalence to conclude that the set $\Pi_{\mathcal{G}}$ collapses to a single point and directly express the unique vector of matching rates achieved by (stable) greedy policies. In contrast, the polytope $\Pi_{\geq 0}$ exhibits a dimension $d = m - n = \frac{n(n-3)}{2}$, where $n = |V|$ is the number of classes and $m = |E| = \binom{n}{2} = \frac{n(n-1)}{2}$ is the number of possible matches. In particular, for $n \geq 4$, the dimension $d = m - n$ of $\Pi_{\mathcal{G}}$ is at least 2. Therefore, limiting ourselves to greedy policies dramatically narrows the range of achievable solutions on complete graphs larger than a triangle.

► **Proposition 35.** Consider a stabilizable matching problem (K_n, λ) , where K_n is the complete graph with $n \geq 3$ nodes, and let Φ denote a greedy policy adapted to K_n . Then:

- (i) The model (K_n, λ, Φ) is stable.
(ii) The matching rate vector $\mu_{\mathcal{G}}$ in this model satisfies

$$(\mu_{\mathcal{G}})_k = \lambda_i p_j + \lambda_j p_i, \quad e_k = \{i, j\} \in E, \quad (37)$$

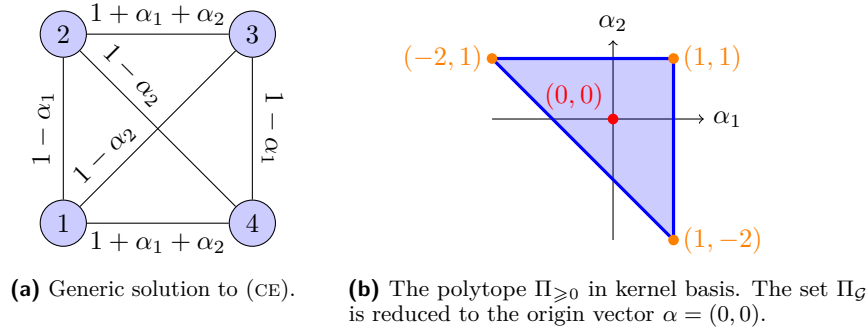
where, for each $i \in V$, p_i is the stationary probability that queue i is non-empty, given by

$$p_i = \frac{\lambda_i}{(\sum_{j \neq i} \lambda_j) - \lambda_i} p_{\emptyset}, \quad (38)$$

and p_{\emptyset} is the stationary probability that the system is empty, given by

$$p_{\emptyset} = \left(1 + \sum_{i \in V} \frac{\lambda_i}{(\sum_{j \neq i} \lambda_j) - \lambda_i} \right)^{-1}. \quad (39)$$

In particular, we have $\Pi_{\mathcal{G}} = \{\mu_{\mathcal{G}}\} \subsetneq \Pi_{>0}$ whenever $n \geq 4$.



■ **Figure 24** Matching problem (K_4, λ) with $\lambda = (3, 3, 3, 3)$.

Proof. We prove each statement one by one.

Proposition 35i. Proposition 35i is a consequence of Corollary 27i (Appendix B.3).

Proposition 35ii. We can focus without loss of generality on the FCFM policy, as Proposition 34iii implies that all greedy matching policies yield the same vector of matching rates. Equation (37) follows by observing that, for each edge $e_k = \{i, j\} \in E$, a match between classes i and j happens if:

- a class- i item arrives while queue j is non-empty, which happens at rate $\lambda_i p_j$;
- a class- j item arrives while queue i is non-empty, which happens at rate $\lambda_j p_i$.

Equations (38) and (39) follow from [15, Proposition 5]. Indeed, for each $i \in V$, applying [15, Equation (10)] to the independent set $\{i\}$ yields (38), and the value of p_\emptyset given in (39) follows from the normalizing equation. This result may also be obtained more directly by observing that, for each $i \in V$, the restriction of the transition diagram of the Markov chain (K_n, λ, Φ_G) to the states where all queues but queue i are empty is similar to a (discrete-time) birth-and-death process with birth probability λ_i and death probability $\sum_{j \neq i} \lambda_j$. ◀

Figure 24 illustrates this result on a complete graph K_4 in which all arrival rates are equal to 3. The particular solution μ° and the basis $\{b_1, b_2\}$ of $\ker(A)$ that we use are shown on Figure 24a. In the corresponding kernel basis, the polytope $\Pi_{\geq 0}$ is defined by the inequalities $\alpha_1 \leq 1$, $\alpha_2 \leq 1$, and $\alpha_1 + \alpha_2 \geq -1$, that is, it is the triangle with vertices $(-2, 1)$, $(1, -2)$, and $(1, 1)$. Yet, Proposition 35 shows that only the solution $\alpha = (0, 0)$ can be achieved by a greedy policy.

F.2.2 Proof of Proposition 24 (diamond graph) and discussion.

Consider the diamond (double fan) graph introduced in Example 4, for which we will show that the set Π_G of matching rate vectors achieved by stable greedy policies is a strict subset of the set $\Pi_{>0}$ of positive solution to (CE). More specifically, we consider the diamond matching problem (D, λ) defined as follows:

$$\begin{aligned}
 D &= (V, E) \text{ with } V = \{1, 2, 3, 4\} \text{ and } E = \{\{1, 2\}, \{1, 3\}, \{2, 3\}, \{2, 4\}, \{3, 4\}\}, \\
 \lambda &= (\lambda_1, \bar{\lambda}_2 + \beta, \bar{\lambda}_3 + \beta, \lambda_4) \\
 &\quad \text{with } \lambda_1 > 0, \bar{\lambda}_2 > 0, \bar{\lambda}_3 > 0, \lambda_4 > 0, \beta > 0, \lambda_1 + \lambda_4 = \bar{\lambda}_2 + \bar{\lambda}_3 = \frac{1}{2}, \\
 \mathcal{B} &= \{(1, -1, 0, -1, 1)\}, \\
 \mu^\circ &= (2\lambda_1 \bar{\lambda}_2, 2\lambda_1 \bar{\lambda}_3, \beta, 2\bar{\lambda}_2 \lambda_4, 2\bar{\lambda}_3 \lambda_4), \\
 \mu &= (2\lambda_1 \bar{\lambda}_2 + \alpha, 2\lambda_1 \bar{\lambda}_3 - \alpha, \beta, 2\bar{\lambda}_2 \lambda_4 - \alpha, 2\bar{\lambda}_3 \lambda_4 + \alpha), \quad \alpha \in \mathbb{R}.
 \end{aligned} \tag{40}$$

Parameterizing λ by β ($= \mu_{2,3}$ for each $\mu \in \Pi$) is a notational convenience that does not lead to any loss of generality: as already observed in Example 4, the stabilizability condition ii writes $\bar{\lambda}_2 > 0$, $\bar{\lambda}_3 > 0$, and $\beta > 0$. We leave it to the reader to verify that our choices for \mathcal{B} and μ° are correct. Equation (40) implies in particular that the sets $\Pi_{\geq 0}$ and $\Pi_{> 0}$, are real intervals, and so are $\Pi_{\mathcal{P}}$ and $\Pi_{\mathcal{G}}$ by convexity.

As in Appendix F.2.1, we first provide an equivalence result that will be useful to describe $\Pi_{\mathcal{G}}$ in details.

Equivalence of greedy policies

As in Appendix F.2.1, we first identify a common behavior shared by all greedy policies and then we exploit this behavior to characterize the matching rates achievable by stable greedy policies.

Since $\{1, 4\}$ is the only independent set of the diamond graph D that is not a singleton, the possible queue states under any greedy policy can be partitioned as follows: either all queues are empty, or exactly one class has a non-empty queue, or (only) classes 1 and 4 have non-empty queues. In other words, the state space of the queue-size process under greedy policies adapted to the graph D is

$$\mathcal{Q}_{\mathcal{G}}(D) = \{0\} \cup \left(\bigcup_{i \in V} \{\ell \mathbf{1}_i, \ell \in \mathbb{N}_{>0}\} \right) \cup \{\ell_1 \mathbf{1}_1 + \ell_4 \mathbf{1}_4 : \ell_1, \ell_4 \in \mathbb{N}_{>0}\}. \quad (41)$$

Therefore, greediness entirely determines the decisions made by greedy policies, except if an item of class 2 or 3 enters while there are unmatched items of classes 1 and 4. Using the fact that classes 2 and 3 are both compatible with classes 1 and 4, we prove formally in Proposition 36 that all greedy policies adapted to the diamond graph make the same matching decisions as the natural greedy policy adapted to the complete graph K_3 obtained by “merging” classes 1 and 4 in the diamond graph.

► **Proposition 36.** *Given the diamond graph D , we introduce the following notation:*

- *Queue-size projection: For each $q = (q_1, q_2, q_3, q_4) \in \mathcal{Q}_{\mathcal{G}}(D)$, we let $\langle q \rangle = (q_1 + q_4, q_2, q_3)$.*
- *Class projection: We let $\langle i \rangle = i$ for each $i \in \{1, 2, 3, \perp\}$ and $\langle 4 \rangle = 1$.*

Every greedy policy Φ adapted to the diamond graph D satisfies the following properties:

- (i) *If Φ is a deterministic size-based policy, then*

$$\langle \Phi(q, i) \rangle = \Phi_{\mathcal{G}}(K_3)(\langle q \rangle, \langle i \rangle), \quad (q, i) \in \mathcal{Q}_{\mathcal{G}}(D) \times V.$$

where $\Phi_{\mathcal{G}}(K_3)$ is the natural greedy policy (36) adapted to the complete graph K_3 .

- (ii) *In general, if Φ is a policy with state space $(\mathcal{S}, |\cdot|)$, then*

$$\sum_{s' \in \mathcal{S}} \sum_{\substack{j \in V \cup \{\perp\}: \\ \langle j \rangle = \Phi_{\mathcal{G}}(K_3)(\langle |s| \rangle, \langle i \rangle)}} \Phi(s, i, j, s') = 1, \quad s \in \mathcal{S}, \quad i \in V.$$

- (iii) *For each sequence I of item classes, we have $\langle Q_t \rangle = (Q_{\mathcal{G}})_t$ for each $t \in \mathbb{N}$, where Q and $Q_{\mathcal{G}}$ are the queue-size processes of the models (D, I, Φ) and $(K_3, \langle I \rangle, \Phi_{\mathcal{G}})$, respectively, with $\langle I \rangle = (\langle I_t \rangle, t \in \mathbb{N})$.*

Loosely speaking, the main take-away of Proposition 36iii is that, in the matching model (D, I, Φ) , the process $((Q_{t,1} + Q_{t,4}, Q_{t,2}, Q_{t,3}), t \in \mathbb{N})$ does not depend on the specific greedy policy Φ that is applied, and it is in fact equal to the queue size process built by applying the natural greedy policy $\Phi_{\mathcal{G}}(K_3)$ in the complete graph K_3 obtained by merging classes 1 and 4 in the diamond graph D .

Proof of Proposition 36. We prove each statement one by one.

Proposition 36i. We leave it to the reader to verify that $\mathcal{Q}_{\mathcal{G}}(K_3)$ is the image of $\mathcal{Q}_{\mathcal{G}}(D)$ by the application $q = (q_1, q_2, q_3, q_4) \mapsto \langle q \rangle = (q_1 + q_4, q_2, q_3)$. This means in particular that $q_1 + q_4$, q_2 , and q_3 cannot be positive simultaneously if $q \in \mathcal{Q}_{\mathcal{G}}(D)$. Since the diamond graph D has only four nodes, we can then conclude by enumerating all relevant cases, depending on the support of $\langle q \rangle = (q_1 + q_4, q_2, q_3)$. For example, if $q_1 + q_4 \geq 1$ and $q_2 = q_3 = 0$, then we have immediately $\Phi(q, i) = \perp$ if $i \in \{1, 4\}$, and the greediness of the policy Φ implies that $\Phi(q, i) \in \{1, 4\}$ if $i \in \{2, 3\}$; in other words, we have $\langle \Phi(q, i) \rangle = \perp$ if $\langle i \rangle = 1$ and $\langle \Phi(q, i) \rangle = 1$ if $\langle i \rangle \in \{2, 3\}$. Similarly, if $q_1 + q_4 = q_3 = 0$ and $q_2 \geq 1$, then $\Phi(q, i) = \perp$ if $i = 2$ and $\Phi(q, i) = 2$ if $i \in \{1, 3, 4\}$, that is, $\langle \Phi(q, i) \rangle = \perp$ if $\langle i \rangle = 2$ and $\langle \Phi(q, i) \rangle = 2$ if $\langle i \rangle \in \{1, 3\}$. In all cases, we can verify that $\langle \Phi(q, i) \rangle$ is equal to $\Phi_{\mathcal{G}}(K_3)(\langle q \rangle, \langle i \rangle)$.

Alternatively, by taking a step back, we can prove Proposition 36i more directly by observing that, for each $(q, i) \in \mathcal{Q}_{\mathcal{G}}(D) \times V$, (i) the support of $\langle q \rangle$ is a singleton $\{j\}$ whenever $q \neq 0$, and (ii) whether $j \in V_i$ depends on i only via $\langle i \rangle$.

Proposition 36ii. The same argument can be repeated for an arbitrary greedy policy Φ with state space $(\mathcal{S}, |\cdot|)$. As in the proof of Proposition 34ii, the key argument consists of observing that we still have $q = |s| \in \mathcal{Q}_{\mathcal{G}}(D)$ for each $(s, i) \in \mathcal{S} \times V$, so that $q_1 + q_4$, q_2 , and q_3 cannot be positive simultaneously.

Proposition 36iii. The conclusion follows in much the same way as in the proof of Proposition 34iii, by injecting statements i and ii from Proposition 36 into (2). ◀

Matching rates under stable greedy policies.

Propositions 37 and 38, which encompass Proposition 24, use the equivalence result of Proposition 36 to characterize the matching rate vector under stable greedy policies. Proposition 37 bounds the coordinates of the matching-rate vectors achievable by a stable greedy policy, while Proposition 38 details the bounds α_+ and α_- of $\Pi_{\mathcal{G}}$. These propositions illustrate the complementarity of edge and kernel coordinates: the results of Proposition 37 are easier to state using edge coordinates, while those of Proposition 38 are easier to state using kernel coordinates. The lower bound (42) is obtained by following a similar approach as in Proposition 35ii.

► **Proposition 37 (Edge coordinates).** *Consider the matching model (D, λ, Φ) , where (D, λ) is the diamond problem (40) and Φ is a greedy policy adapted to the graph D .*

(i) *This matching model is stable.*

(ii) *The matching rate vector $\mu = \mu(D, \lambda, \Phi)$ satisfies $\mu_{2,3} = \beta = \frac{1}{2}(\lambda_2 + \lambda_3 - \lambda_1 - \lambda_4)$, and*

$$\begin{aligned} \mu_{1,2} &\geq \lambda_1 p_2 + \lambda_2 p_1, & \mu_{1,3} &\geq \lambda_1 p_3 + \lambda_3 p_1, \\ \mu_{2,4} &\geq \lambda_2 p_4 + \lambda_4 p_2, & \mu_{3,4} &\geq \lambda_3 p_4 + \lambda_4 p_3, \end{aligned} \quad (42)$$

where p_i is the stationary probability that the system contains unmatched items that belong exclusively to class i , for each $i \in V$.

(iii) *Let $p_{1,4}$ denote the stationary probability that the system contains unmatched items that belong to class 1 or 4 (or both). We have*

$$\begin{aligned} p_2 &= \frac{\lambda_2}{\lambda_1 + \lambda_3 + \lambda_4 - \lambda_2} p_{\emptyset}, & p_3 &= \frac{\lambda_3}{\lambda_1 + \lambda_2 + \lambda_4 - \lambda_3} p_{\emptyset}, \\ p_{1,4} &= \frac{\lambda_1 + \lambda_4}{\lambda_2 + \lambda_3 - \lambda_1 - \lambda_4} p_{\emptyset}, \end{aligned} \quad (43)$$

where p_\emptyset is the stationary probability that the system is empty, whose value follows from the normalization equation $p_\emptyset + p_{1,4} + p_2 + p_3 = 1$. We also have

$$p_1 > \frac{\lambda_1}{\lambda_2 + \lambda_3 + \lambda_4} p_\emptyset, \quad p_4 > \frac{\lambda_4}{\lambda_1 + \lambda_2 + \lambda_3} p_\emptyset. \quad (44)$$

Proof of Proposition 37. We prove each statement one after another.

Proposition 37i. This is a consequence of Corollary 27ii (Appendix B.3).

Proposition 37ii. The equation $\mu_{2,3} = \beta = \frac{1}{2}(\lambda_2 + \lambda_3 - \lambda_1 - \lambda_4)$ is a direct consequence of (CE). The inequalities (42) for $\mu_{1,2}$, $\mu_{1,3}$, $\mu_{2,4}$, and $\mu_{3,4}$ follow by observing that, for each edge $\{i, j\} \in \{\{1, 2\}, \{1, 4\}, \{2, 3\}, \{3, 4\}\}$, a match between classes i and j happens at least in one of the following cases:

- a class- i item arrives while the system contains unmatched items that all belong to class j ,
 - a class- j item arrives while the system contains unmatched items that all belong to class i .
- These events occur at rates $\lambda_i p_j$ and $\lambda_j p_i$, respectively. Equations (42) are not equalities in general because the above list is not exhaustive. For example, depending on the greedy policy, a match between classes 1 and 2 may happen if a class-2 item arrives while the system contains unmatched items of class 1 and unmatched items of class 4.

Proposition 37iii. The expressions for p_\emptyset , p_2 , p_3 , and $p_{1,4}$ follow directly by combining Proposition 36iii with Equations (38) and (39) in Proposition 35. We now derive the lower bound (44) for p_1 . The one for p_4 follows by symmetry.

First assume that the greedy policy Φ is deterministic and size-based, so that it satisfies Proposition 36i. For each $q \in \mathcal{Q}_G(D)$, let π_q denote the probability that the Markov chain (D, λ, Φ) is in state q in stationary regime. The probability that we want to lower-bound is

$$p_1 = \sum_{\ell=1}^{+\infty} \pi_{\ell \mathbb{1}_1}. \quad (45)$$

Now let $\ell \in \mathbb{N}_{>0}$ and consider the balance equation for state $\ell \mathbb{1}_1$, given by

$$(\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4) \pi_{\ell \mathbb{1}_1} = \lambda_1 \pi_{(\ell-1) \mathbb{1}_1} + (\lambda_2 + \lambda_3) \pi_{(\ell+1) \mathbb{1}_1} + C, \quad (46)$$

where C is a non-negative real that depends on the model parameters, the integer ℓ , and the policy Φ , and that accounts for the flow to state $\ell \mathbb{1}_1$ from state $\ell \mathbb{1}_1 + \mathbb{1}_4$, if any. It follows that $(\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4) \pi_{\ell \mathbb{1}_1} > \lambda_1 \pi_{(\ell-1) \mathbb{1}_1}$. An inductive argument allows us to conclude that

$$\pi_{\ell \mathbb{1}_1} > \left(\frac{\lambda_1}{\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4} \right)^\ell p_\emptyset, \quad \ell \in \mathbb{N}_{>0}. \quad (47)$$

Injecting this inequality into (45) allows us to conclude:

$$p_1 > \sum_{\ell=1}^{+\infty} \left(\frac{\lambda_1}{\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4} \right)^\ell p_\emptyset = \frac{\lambda_1}{\lambda_2 + \lambda_3 + \lambda_4} p_\emptyset. \quad (48)$$

If Φ is an extended policy with state space $(\mathcal{S}, |\cdot|)$, we can still write (45)–(48) and reach the same conclusion. The only difference is that π can no longer be defined as the stationary distribution of a Markov chain: instead, we define $\pi_q = \sum_{s \in \mathcal{S}: |s|=q} \varpi_s$ for each $q \in \mathcal{Q}_G(D)$, where ϖ_s is the probability that the Markov chain (D, λ, Φ) is in state s in stationary regime, for each $s \in \mathcal{S}$. ◀

► **Proposition 38** (Kernel coordinates). *Consider the diamond matching problem (D, λ) of (40).*

(i) The intervals $\Pi_{\geq 0}$, $\Pi_{> 0}$, and Π_G are defined as follows in the kernel basis:

$$\begin{aligned}\Pi_{\geq 0} &= [-2 \min(\lambda_1 \bar{\lambda}_2, \bar{\lambda}_3 \lambda_4), 2 \min(\lambda_1 \bar{\lambda}_3, \bar{\lambda}_2 \lambda_4)], \\ \Pi_{> 0} &= (-2 \min(\lambda_1 \bar{\lambda}_2, \bar{\lambda}_3 \lambda_4), 2 \min(\lambda_1 \bar{\lambda}_3, \bar{\lambda}_2 \lambda_4)), \\ \Pi_G &= [\alpha_-, \alpha_+],\end{aligned}$$

with $-2 \min(\lambda_1 \bar{\lambda}_2, \bar{\lambda}_3 \lambda_4) < \alpha_- \leq \alpha_+ < 2 \min(\lambda_1 \bar{\lambda}_3, \bar{\lambda}_2 \lambda_4)$. Hence, $\Pi_G \subsetneq \Pi_{> 0} \subsetneq \Pi_{\geq 0}$.

(ii) The coordinates α_+ and α_- satisfy the following properties:

- a. $\alpha_+ = \alpha(\Phi_+)$, where Φ_+ is the Highest-Reward-First (HRF) policy adapted to the graph D whereby edges $\{1, 2\}$ and $\{3, 4\}$ have the highest reward.
- b. $\alpha_- = \alpha(\Phi_-)$, where Φ_- is the HRF policy adapted to the graph D whereby edges $\{1, 3\}$ and $\{2, 4\}$ have the highest reward.
- c. If $\beta \rightarrow +\infty$ while λ_1 , $\bar{\lambda}_2$, $\bar{\lambda}_3$, and λ_4 remain fixed, we have $\alpha_+ \rightarrow 0$ and $\alpha_- \rightarrow 0$.

Proof of Proposition 38. We prove each statement one after another.

Proposition 38i–iia–iib. The diamond graph D has $n = 4$ nodes and $m = 5$ edges. Therefore, according to Propositions 10 and 21, the sets $\Pi_{\geq 0}$, $\Pi_{> 0}$, and Π_G have dimension $d = m - n = 1$, meaning that they are intervals in \mathbb{R} . The equations for the intervals $\Pi_{\geq 0}$ and $\Pi_{> 0}$ follow directly from the change-of-basis equation $\mu = (2\lambda_1 \bar{\lambda}_2 + \alpha, 2\lambda_1 \bar{\lambda}_3 - \alpha, \beta, 2\bar{\lambda}_2 \lambda_4 - \alpha, 2\bar{\lambda}_3 \lambda_4 + \alpha)$ from (40). That Π_G is also an interval is a consequence of its convexity (Proposition 21). The (non-strict) inequality $-2 \min(\lambda_1 \bar{\lambda}_2, \bar{\lambda}_3 \lambda_4) \leq \alpha_- \leq \alpha_+ \leq 2 \min(\lambda_1 \bar{\lambda}_3, \bar{\lambda}_2 \lambda_4)$ is a consequence of Proposition 21, which states that $\Pi_G \subseteq \Pi_{> 0}$. The first and third inequalities are also strict because α_+ and α_- belong to Π_G (see below), while $2 \min(\lambda_1 \bar{\lambda}_3, \bar{\lambda}_2 \lambda_4)$ and $-2 \min(\lambda_1 \bar{\lambda}_2, \bar{\lambda}_3 \lambda_4)$ do not belong to $\Pi_{> 0}$. That Π_G is a closed interval of the form $\Pi_G = [\alpha_-, \alpha_+]$ and that α_+ and α_- are as given in Proposition 38iia–iib are consequences of Lemma 39 below, which will be proved by a coupling argument later in this appendix¹⁵.

► **Lemma 39.** Consider the HRF policy Φ_+ adapted to the diamond graph D whereby edges $\{1, 2\}$ and $\{3, 4\}$ have the highest reward. For each greedy policy Φ adapted to the compatibility graph D , we have

$$\begin{aligned}\mu_{1,2}(\Phi) &\leq \mu_{1,2}(\Phi_+), & \mu_{3,4}(\Phi) &\leq \mu_{3,4}(\Phi_+), \\ \mu_{1,3}(\Phi) &\geq \mu_{1,3}(\Phi_+), & \mu_{2,4}(\Phi) &\geq \mu_{2,4}(\Phi_+),\end{aligned}\tag{49}$$

Equivalently, in kernel coordinates, we have $\alpha(\Phi) \leq \alpha(\Phi_+)$.

Proposition 38iic. Most of the quantities we consider in this proof are functions of β , but this dependency is left implicit to simplify notation. In particular, we let μ denote (the edge coordinates of) the matching rate vector in the matching model (D, λ, Φ_+) , where $\lambda = (\lambda_1, \bar{\lambda}_2 + \beta, \bar{\lambda}_3 + \beta, \lambda_4)$, with $\lambda_1 + \lambda_4 = \bar{\lambda}_2 + \bar{\lambda}_3 = \frac{1}{2}$, and Φ_+ is the greedy HRF policy defined in Lemma 39. By combining (42) with the conservation equation $\mu_{1,2} + \mu_{1,3} = \lambda_1$, we obtain the following lower and upper bounds for $\mu_{1,2}$:

$$\lambda_1 p_2 + \lambda_2 p_1 \leq \mu_{1,2} \leq \lambda_1 - (\lambda_1 p_3 + \lambda_3 p_1).\tag{50}$$

¹⁵ Lemma 39 focuses on Proposition 38iia. The case Proposition 38iib is trivially deduced by switching the class labels 2 and 3.

In addition, injecting the definition (40) of λ into (43) and (44) shows that, in the model (D, λ, Φ_+) , we have

$$\begin{aligned} p_2 &= p_\emptyset \frac{\bar{\lambda}_2 + \beta}{2\bar{\lambda}_3}, & p_3 &= p_\emptyset \frac{\bar{\lambda}_3 + \beta}{2\bar{\lambda}_2}, & p_{1,4} &= p_\emptyset \frac{1}{3\beta}, \\ p_1 &> p_\emptyset \frac{\lambda_1}{\frac{1}{2} + \lambda_4 + 2\beta}, & p_4 &> p_\emptyset \frac{\lambda_4}{\frac{1}{2} + \lambda_1 + 2\beta}, \end{aligned} \quad (51)$$

with, by the normalization equation $p_\emptyset + p_2 + p_3 + p_{1,4} = 1$,

$$p_\emptyset = \left(1 + \frac{1}{3\beta} + \frac{\bar{\lambda}_2 + \beta}{2\bar{\lambda}_3} + \frac{\bar{\lambda}_3 + \beta}{2\bar{\lambda}_2} \right)^{-1}. \quad (52)$$

Taking the limit of (51) and (52) as $\beta \rightarrow +\infty$, we conclude that both the lower-bound and the upper-bound in (50) tend to $2\lambda_1\bar{\lambda}_2$ as $\beta \rightarrow +\infty$. Then combining (50) with the squeeze theorem allows us to conclude that $\mu_{1,2}$ also tends to $2\lambda_1\bar{\lambda}_2$ as $\beta \rightarrow +\infty$. By symmetry, we obtain directly $\mu_{1,3} \rightarrow 2\lambda_1\bar{\lambda}_3$, $\mu_{2,4} \rightarrow 2\bar{\lambda}_2\lambda_4$, and $\mu_{3,4} \rightarrow 2\bar{\lambda}_3\lambda_4$ as $\beta \rightarrow +\infty$. According to the change-of-basis equation $\mu = (2\lambda_1\bar{\lambda}_2 + \alpha, 2\lambda_1\bar{\lambda}_3 - \alpha, \beta, 2\bar{\lambda}_2\lambda_4 - \alpha, 2\bar{\lambda}_3\lambda_4 + \alpha)$ from (40), this means that $\alpha_+ \rightarrow 0$ as $\beta \rightarrow +\infty$. ◀

Proof of Lemma 39. Consider a greedy policy Φ adapted to the graph D . We will prove the inequality relations in (49) using a coupling argument. More specifically, we will compare the matching models (D, I, Φ) and (D, I, Φ_+) , where $I = (I_t, t \in \mathbb{N})$ is a sequence of i.i.d. classes, such that $I_t = i$ with probability $\lambda_i/(\lambda_1 + \lambda_2 + \lambda_4 + \lambda_3)$ for each $t \in \mathbb{N}$ and $i \in V$.

Considering the model (D, I, Φ) , we let Q_t denote the vector of queue sizes at time t , $L_{t,i}$ the number of class- i items among the first t arrivals, and $M_{t,\{i,j\}}$ (or $M_{t,i,j}$ for short) the number of times that classes i and j are matched over the first t arrivals, for each $t \in \mathbb{N}$ and $i, j \in V$, as defined in (2)–(5). We introduce similar notation for the model (D, I, Φ^+) , the only difference being that all quantities have superscript $+$. Since both models have the same sequence of incoming items, we have $L_{t,i} = L_{t,i}^+$ for each $t \in \mathbb{N}$ and $i \in V$. As usual, we also assume that $Q_0 = Q_0^+ = 0$. Neither $(Q_t, t \in \mathbb{N})$ nor $(Q_t^+, t \in \mathbb{N})$ need to be Markov chains for our argument to hold.

Our end goal is to prove that the following inequalities are satisfied at each time $t \in \mathbb{N}$:

$$M_{t,1,2} \leq M_{t,1,2}^+, \quad (53-1,2)$$

$$M_{t,3,4} \leq M_{t,3,4}^+, \quad (53-3,4)$$

$$M_{t,1,3} \geq M_{t,1,3}^+, \quad (53-1,3)$$

$$M_{t,2,4} \geq M_{t,2,4}^+. \quad (53-2,4)$$

Injecting these inequalities into the definition (8) of the matching rates yields the inequalities (49). We will prove (53) by induction over time $t \in \mathbb{N}$. The following equations will be instrumental to prove the induction step:

$$Q_{t,1} + Q_{t,4} = Q_{t,1}^+ + Q_{t,4}^+, \quad Q_{t,2} = Q_{t,2}^+, \quad Q_{t,3} = Q_{t,3}^+, \quad (54)$$

$$L_{t,1} = Q_{t,1} + M_{t,1,2} + M_{t,1,3} = Q_{t,1}^+ + M_{t,1,2}^+ + M_{t,1,3}^+, \quad (55)$$

$$L_{4,t} = Q_{t,4} + M_{t,2,4} + M_{t,3,4} = Q_{t,4}^+ + M_{t,2,4}^+ + M_{t,3,4}^+. \quad (56)$$

Equation (54) follows from Proposition 36iii, while (55) and (56) follow from our assumption that the arrivals in both models are coupled. Furthermore, given the definition of $Q_{\mathcal{G}}(D)$

in (41), we know that only one integer among $Q_{t,1} + Q_{t,4}$, $Q_{t,2}$, and $Q_{t,3}$ can be positive, for each $t \in \mathbb{N}$.

We now proceed to the induction step. Let $t \in \mathbb{N}$. We now prove that, assuming that the inequalities (53) are satisfied at time t , these inequalities are also satisfied at time $t + 1$. We distinguish several cases depending on the value of I_t :

Case $I_t = 1$: We have directly $M_{t+1,i,j} = M_{t,i,j}$ and $M_{t+1,i,j}^+ = M_{t,i,j}^+$ for $(i,j) \in \{(2,4), (3,4)\}$, hence the induction assumption implies that (53-2,4) and (53-3,4) hold at time $t + 1$.

Since the policy Φ is greedy, we only have three mutually-exclusive cases:

Case $Q_{t,2} \geq 1$: The class-1 item is matched with a class-2 item already present, and we obtain $M_{t+1,1,2} = M_{t,1,2} + 1$ and $M_{t+1,1,3} = M_{t,1,3}$.

Case $Q_{t,3} \geq 1$: The class-1 item is matched with a class-3 item already present, and we obtain $M_{t+1,1,2} = M_{t,1,2}$ and $M_{t+1,1,3} = M_{t,1,3} + 1$.

Case $Q_{t,2} = Q_{t,3} = 0$: The class-1 item is left unmatched, and we obtain $M_{t+1,1,2} = M_{t,1,2}$ and $M_{t+1,1,3} = M_{t,1,3}$.

Since the policy Φ^+ is also greedy, we can repeat the same argument for the quantities associated with Φ^+ . Combining this observation with (54) yields $M_{t+1,1,2} - M_{t+1,1,2}^+ = M_{t,1,2} - M_{t,1,2}^+$ and $M_{t+1,1,3} - M_{t+1,1,3}^+ = M_{t,1,3} - M_{t,1,3}^+$. Hence, the induction assumption implies directly that (53-1,2) and (53-1,3) are satisfied at time $t + 1$.

Case $I_t = 2$: We have directly $M_{t+1,i,j} = M_{t,i,j}$ and $M_{t+1,i,j}^+ = M_{t,i,j}^+$ for $(i,j) \in \{(1,3), (3,4)\}$, hence the induction assumption implies that (53-1,3) and (53-3,4) hold at time $t + 1$. Proving that (53-1,2) and (53-2,4) also hold at time $t + 1$ is more intricate, and we will distinguish three mutually-exclusive cases depending on the values of $Q_{t,1}$, $Q_{t,4}$, $Q_{t,1}^+$, and $Q_{t,4}^+$:

Case $Q_{t,1}^+ + Q_{t,4}^+ = 0$: Under both policies, the class-2 item is either matched with a class-3 item or added to the queue. In particular, we obtain $M_{t+1,i,j} = M_{t,i,j}$ and $M_{t+1,i,j}^+ = M_{t,i,j}^+$ for $(i,j) \in \{(1,2), (2,4)\}$, so that (53-1,2) and (53-2,4) are again satisfied at time $t + 1$ thanks to the induction assumption.

Case $Q_{t,1}^+ + Q_{t,4}^+ \geq 1$: We again subdivide this case into three mutually-exclusive cases:

Case $Q_{t,1}^+ \geq 1$ and $Q_{t,4}^+ \geq 1$: We have $M_{t+1,1,2}^+ = M_{t,1,2}^+ + 1$ and $M_{t+1,2,4}^+ = M_{t,2,4}^+$ by definition of the policy Φ^+ , while for the policy Φ we only know that $M_{t+1,1,2} \in \{M_{t,1,2}, M_{t,1,2} + 1\}$ and $M_{t+1,2,4} \in \{M_{t,2,4}, M_{t,2,4} + 1\}$. We can verify that (53-1,2) and (53-2,4) hold at time $t + 1$ thanks to the induction assumption.

Case $Q_{t,1}^+ = 0$ and $Q_{t,4}^+ \geq 1$: By greediness, the policy Φ^+ matches the incoming class-2 item with a class-4 item, and we obtain $M_{t+1,1,2}^+ = M_{t,1,2}^+$ and $M_{t+1,2,4}^+ = M_{t,2,4}^+ + 1$. If the policy Φ makes the same decision, then we also have $M_{t+1,1,2} = M_{t,1,2}$ and $M_{t+1,2,4} = M_{t,2,4} + 1$, hence (53-1,2) and (53-2,4) hold at time $t + 1$ thanks to the induction assumption. Otherwise, the policy Φ matches the class-2 item with a class-1 item, meaning that $M_{t+1,1,2} = M_{t,1,2} + 1$ and $M_{t+1,2,4} = M_{t,2,4}$. Importantly, this is only possible if $Q_{t,1} \geq 1$. We now prove (53-1,2) and (53-2,4) as follows:

– Proving (53-1,2) boils down to proving $M_{t+1,1,2}^+ \geq M_{t,1,2} + 1$. We have successively:

$$M_{t+1,1,2}^+ - M_{t,1,2} = (Q_{t,1} - Q_{t,1}^+) + (M_{t,1,3} - M_{t,1,3}^+) \geq 1 + 0 = 1,$$

where the equality follows from (55) and the inequality follows from the induction assumption and the fact that $Q_{t,1} \geq 1$ and $Q_{t,1}^+ = 0$.

– Proving (53-2,4) boils down to proving $M_{t+1,2,4}^+ + 1 \leq M_{t,2,4}$. We have successively

$$M_{t,2,4} - M_{t+1,2,4}^+ = (Q_{t,4}^+ - Q_{t,4}) + (M_{t,3,4}^+ - M_{t,3,4}) \geq 1 + 0 = 1,$$

where the equality follows from (56) and the inequality follows from the induction assumption and the observation that $Q_{t,4}^+ - Q_{t,4} = Q_{t,1} - Q_{t,1}^+ \geq 1$.

Intuitively, the only way that an incoming class-2 item is matched at time t with a class-1 item under the policy Φ and with a class-4 item under the policy Φ^+ is if, in the past, the policy Φ^+ had made one more match along edge $\{1, 2\}$ and one less match along edge $\{2, 4\}$ compared to the policy Φ .

Case $Q_{t,1}^+ \geq 1$ and $Q_{t,4}^+ = 0$: This case is symmetrical to the previous case.

Case $I_t = 3$: This case is symmetrical to the case $I_t = 2$.

Case $I_t = 4$: This case is symmetrical to the case $I_t = 1$.

◀

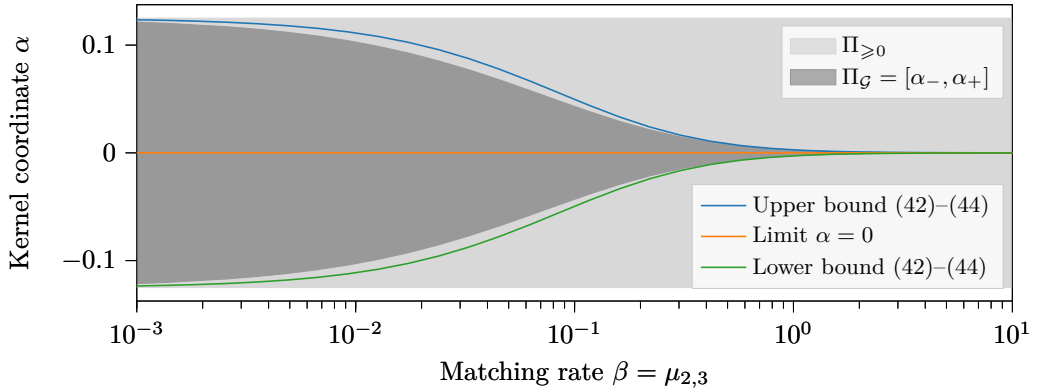
Proposition 38 can be interpreted as follows. Intuitively, the kernel coordinate α given in (40) (also shown in Figure 3d) acts like a slider that determines how much edges $\{1, 2\}$ and $\{3, 4\}$ are (dis)avored compared to edges $\{1, 3\}$ and $\{2, 4\}$ on the long run. Remarkably, Proposition 38 shows that the greedy policy that favors edges $\{1, 2\}$ and $\{3, 4\}$ (resp. $\{1, 3\}$ and $\{2, 4\}$) the most *in the long run* is also the policy that favors these edges the most *in the short run*. This result is proved by a coupling argument. In addition, Proposition 38 uses the lower bound (42) in Proposition 37 to prove that, in the limit as $\beta \rightarrow +\infty$, all greedy policies yield the same matching rate vector.

Taken together, Items i and iic in Proposition 38 show that, as $\beta \rightarrow +\infty$, the interval $\Pi_{\mathcal{G}}$ becomes reduced to a single point $\alpha = 0$, meaning that all greedy policies yield the same vector of matching rates, with edge coordinates $\mu = (2\lambda_1\bar{\lambda}_2, 2\lambda_1\bar{\lambda}_3, \beta, 2\bar{\lambda}_2\lambda_4, 2\bar{\lambda}_3\lambda_4)$. The rationale behind this result is the following. In the regime where $\beta = \mu_{2,3} \rightarrow +\infty$, we see in (40) that the arrival rates of classes 2 and 3 become large compared to those of classes 1 and 4. As a result, items of classes 1 and 4 are matched (almost) always immediately, and unmatched items belong to either class 2 or class 3 (but not both at the same time). In the proof of Proposition 38iic, this intuition is formalized by taking the limit of (43) as $\beta \rightarrow +\infty$, which yields $p_{\emptyset} \rightarrow 0$, $p_2 \rightarrow 2\bar{\lambda}_2$, $p_3 \rightarrow 2\bar{\lambda}_3$, and $p_{1,4} \rightarrow 0$. In this regime, the greediness of the policy allows no degree of flexibility in choosing the class of the item to which an incoming item is matched, so the matching rates are unique.

► **Remark 12.** *Some quantities in Proposition 38 are functions of the matching rate $\beta = \mu_{2,3}$, but this dependency is kept implicit to simplify notation. In particular, the intervals $\Pi_{\geq 0}$ and $\Pi_{> 0}$ do not depend on β , but the interval $\Pi_{\mathcal{G}}$ and the coordinates α_+ and α_- do.*

Numerical results.

To illustrate Proposition 38, Figure 25 shows a symmetric example with $\lambda_1 = \bar{\lambda}_2 = \bar{\lambda}_3 = \lambda_4 = \frac{1}{4}$. The figure compares $\Pi_{\geq 0}$ and $\Pi_{\mathcal{G}}$ with the bounds (42)–(44) (converted in the kernel coordinates) and the limit $\alpha = 0$. Each point forming the shape of $\Pi_{\mathcal{G}} = [\alpha_-, \alpha_+]$ is obtained by running a simulation consisting of 10^{10} steps (as specified in Section 7). As announced by Proposition 38, $\Pi_{\mathcal{G}}$ becomes reduced to a single point $\alpha = 0$ when $\beta \rightarrow +\infty$. We also notice that the bounds (42)–(44) are not tight when β is small (in the sense that the difference between α_+ and the upper-bound is no longer negligible compared to the difference between α_+ and the boundary $2 \min(\lambda_1\bar{\lambda}_3, \bar{\lambda}_2, \lambda_4)$ of $\Pi_{\geq 0}$), while at the same time α_+ and α_- become arbitrarily close to the borders of $\Pi_{\geq 0}$ when β tends to zero (which does not contradict the fact that $\Pi_{\mathcal{G}}$ is a strict subset of $\Pi_{\geq 0}$ as long as $\beta > 0$). The gap between $\Pi_{\mathcal{G}}$ and the bounds (42)–(44) comes from the fact that, to obtain the lower bounds for p_1 and p_4 in (44), we neglected the case where there are both class-1 and class-4 unmatched items, and this case is not negligible when β is small.




■ **Figure 25** Evolution of $\Pi_{\mathcal{G}}$ as a function of β in the diamond problem (40) with $\lambda_1 = \bar{\lambda}_2 = \bar{\lambda}_3 = \lambda_4 = \frac{1}{4}$. For each β , $\Pi_{\mathcal{G}}$ is estimated by simulating the model (D, λ, Φ_+) and leveraging the symmetry of the problem. $\Pi_{\geq 0}$ and other bounds are displayed for comparison. All results are expressed in kernel coordinates.

F.2.3 Intuition of Conjecture 25 (Fish matching problem) and discussion.

We consider a matching problem, which we call the Fish problem, where we conjecture that $\Pi_{\mathcal{G}} = \Pi_{>0}$. The Fish matching problem is defined as follows:

$$\begin{aligned}
 V &= \{1, 2, 3, 4, 5, 6\}, \\
 E &= \{\{1, 2\}, \{1, 3\}, \{2, 3\}, \{3, 4\}, \{3, 6\}, \{4, 5\}, \{5, 6\}\}, \\
 \lambda &= (4, 4, 3, 2, 3, 2), \\
 \mathcal{B} &= \{(0, 0, 0, 1, -1, -1)\}, \\
 \mu^\circ &= (3, 1, 1, \frac{1}{2}, \frac{1}{2}, \frac{3}{2}, \frac{3}{2}), \\
 \mu &= (3, 1, 1, \frac{1}{2} + \alpha, \frac{1}{2} - \alpha, \frac{3}{2} - \alpha, \frac{3}{2} + \alpha), \quad \alpha \in \mathbb{R}.
 \end{aligned} \tag{57}$$

The general solution μ given in (57) is shown in Figure 12. We can verify by a direct inspection that this matching problem is stabilizable and that $\Pi_{\geq 0} = [-\frac{1}{2}, \frac{1}{2}]$ in kernel coordinates. The kernel coordinate α acts like a slider that is positive (resp. negative) if matches along edges $\{3, 4\}$ and $\{5, 6\}$ are more (resp. less) frequent than matches along edges $\{3, 6\}$ and $\{4, 5\}$.




■ **Figure 12 (repeated)** Generic solution to (CE) in the Fish matching problem of (57) with $\lambda = (4, 4, 3, 2, 3, 2)$.


Intuition of proof. To support Conjecture 25, we build two families of stable greedy policies, denoted by $(\Phi_k^+)_{k \in \mathbb{N}}$ and $(\Phi_k^-)_{k \in \mathbb{N}}$, such that $\lim_{k \rightarrow +\infty} \alpha(\Phi_k^+) = \frac{1}{2}$ and $\lim_{k \rightarrow +\infty} \alpha(\Phi_k^-) = -\frac{1}{2}$. The conclusion then follows from the convexity of the set $\Pi_{\mathcal{G}}$ (Proposition 21). We focus on

the family $(\Phi_k^+)_{k \in \mathbb{N}}$, as the family $(\Phi_k^-)_{k \in \mathbb{N}}$ is symmetrical (in the sense that it suffices to exchange classes 4 and 6).

The family $(\Phi_k^+)_{k \in \mathbb{N}}$ is defined as follows. Let Φ_∞^+ denote the HRF policy where edges have the following decreasing reward order: $\{1, 3\}$, $\{2, 3\}$, $\{3, 4\}$, $\{5, 6\}$, followed by all other edges in an arbitrary order. Let Φ_0^+ denote the HRF policy where edges have the



(a) Policy Φ_∞^+ , followed by policy Φ_k^+ when the size of queue 4 is at most $k - 1$



(b) Policy Φ_0^+ , followed by policy Φ_k^+ when the size of queue 4 is at least k

■ **Figure 26** HRF (greedy) policies considered in the Fish matching problem of (57).

following decreasing reward order: $\{3, 4\}$, $\{2, 3\}$, $\{1, 3\}$, $\{5, 6\}$, followed by all other edges in an arbitrary order. The important point is that both policies prioritize edges $\{3, 4\}$ and $\{5, 6\}$ over edges $\{3, 6\}$ and $\{4, 5\}$ (with the hope that this lead to a high α), but that Φ_∞^+ gives higher priority to the “tail” of the fish, while Φ_0^+ gives higher priority to the “trunk”. Now, for each $k \in \mathbb{N}$, Φ_k^+ is the deterministic size-based policy that follows Φ_∞^+ when the queue size of class 4 is at most $k - 1$ and Φ_0^+ when the queue size of class 4 is at least k (that is, $\Phi_k^+(q, i) = \Phi_\infty^+(q, i)$ if $q_4 \leq k - 1$ and $\Phi_k^+(q, i) = \Phi_0^+(q, i)$ if $q_4 \geq k$).

The rationale behind this definition is as follows. According to (57), α is maximal (i.e., equal to $\frac{1}{2}$) when $\mu_{3,4}$ is equal to 1 and $\mu_{3,6}$ is equal to 0. If we allow *non-greedy* policies, $\alpha = \frac{1}{2}$ can be achieved by merely applying the edge-filtering variant of ML (or any maximally-stable policy) on the bijective subgraph of G obtained by eliminating edge $\{3, 6\}$, that is, by *never* performing a match between classes 3 and 6. As we will see, the family $(\Phi_k^+)_{k \in \mathbb{N}}$ of (stable greedy) policies emulates this (non-greedy) edge-filtering policy by favoring edge $\{3, 4\}$ over edge $\{3, 6\}$ while making the probability that $q_4 = 0$ arbitrary small as k increases. Roughly speaking, mixing the policies Φ_∞^+ and Φ_0^+ allows us to control q_4 : Φ_∞^+ is not stable and makes q_4 go to infinity, while Φ_0^+ is stable and reduces q_4 when it is large. All in all, Φ_k^+ keeps the value of q_4 around k , so that the probability that $q_4 = 0$ is low when k is large. In the limit, q_4 is always positive, so that the class-3 items that are not matched with classes 1 or 2 are drained by class 4 (while class 6 is matched only with class 5), so that $\mu_{3,4}(\Phi_k^+)$ tends to 1 and $\mu_{3,6}(\Phi_k^+)$ tends to 0 as $k \rightarrow \infty$.

We expect that a rigorous proof that $\mu_{3,6}(\Phi_k^+)$ tends to 0 as $k \rightarrow +\infty$ will involve the following steps, some of which may require further investigation or analysis:

1. For each $k \in \mathbb{N}$, the matching model (G, λ, Φ_k^+) is stable. Stability can be proved by applying a fluid-limit argument that generalizes the framework of [18]¹⁶. Therefore, for each $k \in \mathbb{N}$, we can consider a random vector $Q_k = (Q_{k,1}, Q_{k,2}, \dots, Q_{k,n})$ distributed like the vector of queue sizes in the matching model (G, λ, Φ_k^+) in stationary regime, and we know that the vector $\mu(\Phi_k^+)$ is well-defined.
2. For each $k \in \mathbb{N}$, by definition of the policy Φ_k^+ , we have

$$\begin{aligned} \mu(\Phi_k^+)_{3,6} &= \lambda_3 P(Q_{k,4} = 0, Q_{k,6} > 0) + \lambda_6 P(Q_{k,3} > 0, Q_{k,5} = 0), \\ &\leq (\lambda_3 + \lambda_6) P(Q_{k,4} = 0), \end{aligned}$$

¹⁶ Formalizing this generalization is outside the scope of this paper and could be the topic of a future work.

where the inequality arises because the events $(Q_{k,4} = 0, Q_{k,6} > 0)$ and $(Q_{k,3} > 0, Q_{k,5} = 0)$ are both included into the event $(Q_{k,4} = 0)$ (in the latter case, this is because the greediness of the policy prevents the system from containing unmatched items of classes 3 and 4 at the same time). Consequently, to prove that $\mu(\Phi_k^+)_{3,6} \rightarrow 0$, it suffices to prove that $P(Q_{k,4} = 0) \rightarrow 0$ as $k \rightarrow +\infty$. This argument rigorously supports the intuition that, if class 4 becomes “unstable”, then this class “drains” all class-3 items that are not matched with classes 1 and 2.

3. To prove that $P(Q_{k,4} = 0) \rightarrow 0$ as $k \rightarrow +\infty$, we believe that we can reason as follows:
- We can prove that, for each $k \in \mathbb{N}$, the conditional average rate at which class 3 is matched with either class 4 or class 6, *given that* $Q_{k,4} \leq k - 1$, is upperbounded by

$$\lambda_3 \left(1 + \frac{\lambda_1}{\lambda_2 + \lambda_3 - \lambda_1} + \frac{\lambda_2}{\lambda_1 + \lambda_3 - \lambda_2} \right)^{-1} = \frac{9}{11}.$$

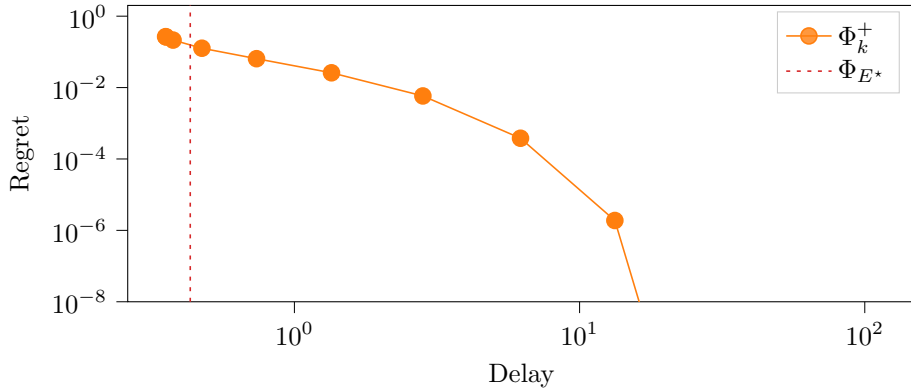
This upper bound follows by coupling the fish problem with a “worst-case” matching problem in which there is an infinite backlog of items of classes 4 and 6 to be matched with class 3. Roughly speaking, this worst-case matching problem can be obtained by taking $\lambda_4 = \lambda_6 = +\infty$ in the fish problem; then, focusing on nodes 1, 2, and 3, the system behaves like a paw graph with an unstable pendant node obtained by merging classes 4 and 6.

- Using this upper bound, we prove that at least one class among classes 4 and 6 becomes unstable, in the sense that $\lim_{k \rightarrow +\infty} P(Q_{k,4} = 0) = 0$ or $\lim_{k \rightarrow +\infty} P(Q_{k,6} = 0) = 0$. Since we can verify otherwise that class 6 does not become unstable, it follows that $\lim_{k \rightarrow +\infty} P(Q_{k,4} = 0) = 0$.

◀

The complete proof is left as an open question for future works.

Numerical results.



■ **Figure 27** Approaching the vertex $\alpha = 1/2$ in the Fish graph from Figure 12 with greedy policies ($r = (2, 2, 2, 1, -1, 0, 1)$).

The performance of Φ_k^+ is shown in Figure 27. Since the vertex at $\alpha = 1/2$ is bijective, we include the delay of the reward-optimal policy Φ_{E^*} for comparison. The figure corroborates Conjecture 25, demonstrating that the regret of Φ_k^+ converges to zero. However, as mentioned earlier, this reduction in regret comes at the cost of increasing the size of the class-4 queue, thus increasing the delay of the policy.