

Optimal Incentive Design for Decentralized Dynamic Matching Markets

Chen Chen¹, Pengyu Qian², and Jingwei Zhang³

¹New York University Shanghai

²Questrom School of Business, Boston University

³School of Data Science, The Chinese University of Hong Kong, Shenzhen (CUHK-SZ)
cc8029@nyu.edu, pqian@bu.edu, zhangjingwei@cuhk.edu.cn

November 25, 2024

Abstract

In decentralized dynamic matching markets, agents can benefit from sharing their resources, but they will collaborate only if properly incentivized. Motivating applications include multi-hospital kidney exchanges and job-hunting markets. This paper explores the design of non-monetary mechanisms to incentivize collaboration in decentralized matching markets. We study a model with multiple self-interested agents, each managing a local multi-way dynamic matching problem, where jobs of different types arrive stochastically at each agent and remain available to match for a limited time. Each agent holds private information about her job amounts and actions and aims to maximize her long-run average matching reward.

We design a simple, non-monetary mechanism that incentivizes agents to submit all their jobs to a shared pool upon arrival, enabling centralized matching at the shared pool. The mechanism achieves this by randomly selecting an agent (with a specified probability) to perform a matching and collect the associated matching reward. We demonstrate that under the mechanism, it is in each agent's best interest to submit all her jobs to a shared pool—resulting in system efficiency that matches that of a centralized setting—in a large market regime with many agents.

Subject classifications: Decentralized matching, market design, multi-way matching, impatient jobs, stochastic optimization, fluid approximation.

1 Introduction

Matching markets play a crucial role in connecting diverse items to form mutually beneficial relationships, with applications spanning both for-profit and non-profit platforms. A well-studied setting is centralized matching, where a central planner controls all items, and the focus is on designing algorithms to maximize match value as different types of items enter and exit the marketplace. In this paper, we shift our attention to the decentralized setting, which introduces additional challenges beyond algorithm design. Specifically, we explore how to incentivize self-interested entities, each controlling their own sources of items, to contribute to the marketplace in a way that maximizes overall efficiency. To do so, we need to handle strategic behavior and prevent free-riding. We begin by presenting two motivating examples.

Example 1: Multi-Hospital Kidney Exchange. Kidney exchanges enable (two or more) patients with willing but incompatible donors to swap kidneys, ensuring each patient receives a compatible organ (Roth et al., 2005). Each kidney transplant saves the expensive cost of dialysis and results in a welfare gain exceeding \$1 million (Held et al., 2016). In 2019, more than 1,500 kidney exchange transplants were performed in the US, and the number is rapidly increasing (Agarwal et al., 2021). However, donor-patient pairs are typically registered at individual hospitals or transplant centers, leading to a decentralized matching process where hospitals may prioritize their own interests. Despite efforts to create a national-level exchange to facilitate more compatible matches, hospitals often exhibit *free-riding* behavior by withholding easy-to-match pairs for internal match and submitting only hard-to-match pairs to the national exchange (Ashlagi and Roth, 2014). A critical research challenge in this application is determining how to incentivize hospitals to enroll all donor-patient pairs in the national exchange, thereby maximizing overall system efficiency.

Example 2: Collaboration among Matching Intermediaries. Two-sided matching intermediaries, such as real estate and head-hunting agents, facilitate the matching of supply with demand. A distinctive feature of these markets is the network effect, where the value of an intermediary increases with its user base (Banton, 2024). This suggests that even competing intermediaries could benefit from granting each other access to their customers and resources, paving the way for collaboration among them. For example, in job-hunting markets, both job openings and candidates are available on the market for only a limited time, due to evolving outside options or changing market conditions. Headhunters can share job opportunities and candidates among themselves to facilitate matching. However, it is not immediately clear that sharing resources benefits all parties, and the

central research question is how to design incentives that encourage intermediaries to collaborate fully and willingly.

These applications share several common features. The market is fragmented, and greater value can be achieved in a thicker market. Decisions are dynamic—items such as donor-patient pairs and job candidates and opportunities arrive over time and can expire if left unmatched for too long. Participants, such as hospitals and other matching intermediaries, are strategic and possess private information about their items and will only participate (e.g., by sharing items) if it aligns with their interests to do so. In this paper, we develop and analyze a model that incorporates these key features. We focus on a decentralized dynamic matching market with multiple strategic agents. Each agent faces a dynamic multi-way matching¹ problem (the local problem), where items (or jobs) of different types arrive over time and remain available for matching for a limited amount of time. Information regarding the arrival of jobs is private to each agent. Different types of matches require different combinations of jobs and yield different rewards, although the reward structure is the same for all agents. In addition to making local matches, an agent has the option to submit jobs to a shared entity, with the reward for this action being determined by a designed mechanism. The goal of each agent is to maximize her long-run average payoff from matching and submitting. We address the mechanism design problem aimed at maximizing social welfare, defined as the total matching reward accumulated across all agents.

As nicely summarized by Ashlagi and Roth (2021), the decentralized matching problem comprises two key components: submission and matching. Extensive literature examines each of these aspects in isolation. Studies on the submission component, primarily from the market design literature, focuses on incentivizing agents (such as hospitals) to submit all items (see, e.g., Sönmez and Ünver 2013, Ashlagi and Roth 2014, and Agarwal et al. 2019). However, these studies often rely on static models, which do not capture market frictions related to limited item lifespans (e.g., donor-patient pairs in kidney exchange) and market dynamics. On the other hand, works on (centralized) dynamic matching (e.g., Aouad and Sarıtaç 2022 and Aveklouris et al. 2024) develop algorithms with theoretical performance guarantees but generally do not account for the incentives of market participants. Our work advances the literature by jointly considering both submission and matching in a decentralized dynamic setting. Tackling this market design problem is challenging yet important. The primary challenge lies in analyzing a dynamic game with a large strategy space and incomplete information. This problem is important because, as highlighted by Ashlagi

¹That is, a match can connect two or more jobs.

and Roth (2021), mechanisms that offer good incentives in static models may prove inefficient in practical settings. Hence, there is a need for a deeper understanding of dynamic mechanisms that can achieve high efficiency.

1.1 Overview of Results and Contributions

Our main contribution includes proposing both monetary and non-monetary mechanisms for the decentralized dynamic matching problem and rigorously analyzing their performance in the resulting equilibrium. These mechanisms are simple and intuitive. Below, we elaborate on our main results.

The key observation underlying all our proposed mechanisms is as follows (Section 4.1): If each submitted job earns the submitting agent a reimbursement equal to the job’s marginal value in the aggregate (centralized) matching problem, it becomes a dominant strategy for all agents to submit all jobs upon arrival, regardless of the number of jobs they have or the strategies of other agents. Notably, obtaining the jobs’ marginal values only requires solving the fluid relaxation of the centralized matching problem, which involves only the aggregate job arrival rates and the matching rewards, thus requiring minimal information about each individual agent. This observation leads to a monetary mechanism that may be well-suited for for-profit platforms, such as head-hunting and ride-sharing. Under this mechanism, the agents’ incentives align with social welfare as characterized by the centralized matching problem, although the agents may have different job arrival rates.

We next demonstrate that the above form of reimbursement can be implemented equivalently without relying on monetary transfers, using a randomized matching allocation (Section 4.2). This mechanism may be more suitable for non-profit platforms, such as kidney exchanges. Specifically, we propose a randomized matching allocation where, when a match is about to be performed, it is assigned to one of the participant jobs with a probability proportional to the job’s marginal value. The submitting agent of that job then executes the match (e.g., performs the transplant in the kidney exchange setting) and receives the matching reward. As we demonstrate, such a randomized allocation ensures that the expected value of a submitted job equals its marginal value, conditional on it being matched, regardless of the specific match it participates in. We then run a matching algorithm at the shared pool, which is asymptotically optimal as the job arrival rates scale to infinity; as demonstrated in Section 3, a simple periodic matching policy can achieve this.

We rigorously prove that the proposed mechanism incentivizes all agents to submit all their jobs upon arrival in a large market regime with many agents (Section 4.3). Intuitively, when all agents submit their jobs, the job arrival rates at the shared pool grow to infinity in the large

market regime. Consequently, the shared pool becomes congested, and the fluid relaxation of the centralized matching problem approximates the dynamics of the shared pool well. As long as the shared pool operates under an asymptotically optimal matching policy, the probability that a job with a positive marginal value is matched before departing approaches one. As a result, if an agent submits a job, it yields an expected payoff equal to its marginal value by the mechanism’s design; this incentivizes an agent to submit all her jobs. In contrast, if the agent withholds a job, it may expire before being matched with a positive probability, resulting in a tangible loss. We rigorously validate this intuition in Section 4.3 and show that all agents submitting jobs fully constitutes an approximate Nash equilibrium in the original problem, meaning that the benefit of unilaterally deviating from full submission becomes negligible as the number of agents increases.

Finally, in Section 5, we numerically evaluate the performance of our non-monetary mechanism on both a simple synthetic example (Section 5.1) and a more realistic example using kidney exchange data (Section 5.2). We demonstrate that full submission is an approximate equilibrium within our mechanism in practical settings.

The rest of the paper is organized as follows. Section 1.2 reviews some related work. Section 2 formulates the problem. Section 3 investigates the centralized setting as a benchmark. Section 4 examines the decentralized setting and introduces our monetary and non-monetary mechanisms. In Section 4.3, we analyze our mechanism and show that all agents submitting their jobs upon arrival is an approximate equilibrium in the original problem when the number of agents is large. Section 5 numerically evaluates the performance of our mechanism. Section 6 concludes.

1.2 Related Literature

Our paper relates to several strands of fast-growing literature spanning operations research, economics, and computer science.

Incentive Issues in Multi-Hospital Kidney Exchange Since the early days of kidney exchange, researchers have observed that hospitals participating in multi-hospital kidney exchanges may benefit from strategically reserving easy-to-match donor-patient pairs for internal matches (Sönmez and Ünver 2013 and Ashlagi and Roth 2014). Subsequent literature propose various mechanisms to encourage hospitals to submit all their pairs and use a range of kidney exchange models to provide theoretical guarantees on the effectiveness of these mechanisms. The pioneering works in this literature employ static random graphs to model kidney exchanges and design incentivizing mechanisms (e.g., Ashlagi and Roth 2014, Toulis and Parkes 2015, and Ashlagi et al. 2015). Leveraging

neoclassical producer theory, Agarwal et al. (2019) argue through a static optimization problem that hospitals should be rewarded based on the marginal contribution of the pairs they submit. Biró et al. (2019) utilize a cooperative game model on a graph and propose mechanisms based on the Shapley value. This stream of literature predominantly considers a static modeling approach, where donor-patient pairs do not arrive or depart over time. However, as noted in the survey of Ashlagi and Roth (2021), matching in a dynamic setting is a critical challenge in kidney exchange, and mechanisms that offer strong incentives in static models might be inefficient in practice. Our contribution to this literature lies in proposing and rigorously analyzing a general dynamic model, which uncovers strategic hospital behaviors absent in static models. It is worth noting that our proposed mechanism resembles the one introduced in Agarwal et al. (2019), where rewards are based on the marginal value of submitted pairs. However, there are some key differences. First, we propose a randomized matching allocation to reimburse submissions, whereas Agarwal et al. (2019) use a point system. Additionally, we rigorously characterize the performance of the marginal value-based mechanism in a more realistic dynamic model with limited information, whereas Agarwal et al. (2019) analyze it in a static model.

Centralized Dynamic Matching While our focus is on incentivizing agents in a decentralized dynamic matching market to submit all items, a related question is how to design algorithms to match these submitted items. Our mechanism is agnostic to the specific matching algorithm used in the shared pool (Remark 3.2), although the choice of the the algorithm does influence our performance guarantees. Many studies study the dynamic matching problem where items exit the system if they remain unmatched for a certain duration, somewhat aligning with our setting (see, e.g., Aouad and Sarıtaç 2022, Avelouris et al. 2024, Akbarpour et al. 2020, Ashlagi et al. 2023). More recently, a stream of literature has focused on any-time low-regret dynamic matching algorithms within a finite-horizon setting (Kerimov et al. 2023, Kerimov et al. 2024, Gupta 2024, Wei et al. 2023). Although the context differs, these studies highlight the connection between algorithm performance (i.e., regret) and the structural properties of the fluid approximation of the matching problem, such as dual degeneracy, which we also observe in our problem setting.

Cooperation Mechanism in Other Setups Some researchers have explored ways to incentivize cooperation among companies and platforms in various applications beyond dynamic matching. For example, studies examine cooperation among airline companies, i.e., airline alliances (Netessine and Shumsky 2005, Wright et al. 2010, Hu et al. 2013), as well as among ride-sharing companies (Cohen

and Zhang, 2022) and manufacturing firms (Roels and Tang, 2017).

Finally, some papers apply cooperative game theory to allocate the total costs or profits among agents to sustain cooperation (see, e.g., Anily and Haviv 2010, Karsten et al. 2015, Liu and Yu 2022). However, the implementation of this approach to our setting would require agents to be able to monitor the actions and job counts of all other agents to ensure that all pairs are indeed submitted. Given our focus on a private information setting, we adopt a non-cooperative game approach in this paper. In a similar manner, Chen et al. (2023) explore a non-cooperative game approach to incentivize collaboration in a different decentralized queueing setting.

1.3 Notation and Terminology

We let \mathbb{N} denote the set of nonnegative integers and \mathbb{N}_+ the set of strictly positive integers. For any two integers $a, b \in \mathbb{N}$ with $a \leq b$, we let $[a : b] = \{a, a + 1, \dots, b - 1, b\}$ denote a sequence of integers starting from a and ending with b , and we denote $[n] = [1 : n]$ for any $n \in \mathbb{N}_+$. For any real number $x \in \mathbb{R}$, we let $(x)^+ \triangleq \max\{x, 0\}$ denote the maximum of x and 0. Given a vector $\mathbf{x} = (x_i)_{i \in [n]} \in \mathbb{R}^n$, we let $\mathbf{x} \geq \mathbf{0}$ denote $x_i \geq 0$ for any entry $i \in [n]$.

2 Model Formulation

We consider a continuous-time model with N strategic agents. Each agent $i \in [N]$ manages a local dynamic multi-way matching. Specifically, there are J types of jobs in the system. Jobs of each type $j \in [J]$ arrive at agent i following an independent Poisson process with a rate of $\lambda_{ij} \geq 0$. Let $\boldsymbol{\lambda}_i = (\lambda_{ij})_{j \in [J]} \in \mathbb{R}_+^J$ denote the concatenation of arrival rates at agent i . Jobs expire if not matched within a certain amount of time. The waiting time of a type- j job is exponentially distributed with a rate parameter of $\theta_j > 0$ for any $j \in [J]$.

There are K types of matchings. A matching $m \in [K]$ requires $M_{jm} \in \mathbb{N}$ units of type- j jobs and creates a reward of $r_m > 0$ for any agent. Each match requires at least two jobs; that is, $\sum_{j \in [J]} M_{jm} \geq 2$ for any $m \in [K]$. Matched jobs leave the system immediately. We let $\mathbf{r} = (r_m)_{m \in [K]} \in \mathbb{R}_+^K$ represent the concatenation of matching rewards, $r_{\max} \triangleq \max_{m \in [K]} r_m < \infty$ the maximum matching reward, and $\mathbf{M} = (M_{jm})_{j \in [J], m \in [K]} \in \mathbb{N}^{J \times K}$ the matching matrix.

We shall consider a large market regime with many agents, and we make the following regularity assumption.

Assumption 2.1. The total job arrival rate $\sum_{i \in [N]} \boldsymbol{\lambda}_i$ satisfies $\sum_{i \in [N]} \boldsymbol{\lambda}_i = N \cdot \boldsymbol{\lambda} \in \mathbb{R}_+^J$ for some vector $\boldsymbol{\lambda} = (\lambda_j)_{j \in [J]}$, where $\lambda_j > 0$ for all $j \in [J]$. In addition, there exists some constants $C_j > 0$

such that for any $i \in [N]$ and $j \in [J]$, we have $\lambda_{ij}/\lambda_j \leq C_j$.

Assumption 2.1 requires that the agents' job arrival rates are not too different from one another. This assumption holds, for example, for a high multiplicity model in which agents are partitioned into a fixed number L of types, and the number of each type- ℓ agents is a fixed proportion α_ℓ of the total agent population. In addition, agents of the same type ℓ have the same arrival rate vector $\lambda_\ell = (\lambda_{\ell j})_{j \in [J]} \in \mathbb{R}_+^J$, leading to $\lambda \triangleq \sum_{\ell \in [L]} \alpha_\ell \lambda_\ell$ and $C_j \triangleq \max_{\ell \in [L]} \lambda_{\ell j}/\lambda_j$.

Local Dynamics For each agent $i \in [N]$ and job type $j \in [J]$, we let $D_{ij}(t)$ denote the number of type- j jobs that have arrived at agent i by time t , which follows a Poisson process with rate λ_{ij} by assumption. Let $N_{im}(t)$ represent the number of type- m matchings that agent i performs and $T_{ih,j}(t)$ the number of type- j jobs that agent i transfers to another agent $h \neq i$, both by time t . Both $N_{im}(t)$ and $T_{ih,j}(t)$ depend on the strategies of agent i . Finally, let $X_{ij}(t)$ denote the number of type- j jobs agent i possesses at time t , and $A_{ij}(t)$ the number of type- j jobs that have expired at agent i by time t . The process $A_{ij}(t)$ depends on $X_{ij}(t)$, as each of the $X_{ij}(t)$ type- j jobs expires within the next Δt units of time with a probability of $\theta_j \Delta t$, provided it has not been matched and assuming that Δt is sufficiently small. The dynamics of agent i 's jobs can be described as follows:

$$X_{ij}(t) = X_{ij}(0) + D_{ij}(t) + \sum_{h \neq i} \left(T_{hi,j}(t) - T_{ih,j}(t) \right) - \sum_{m \in [K]} M_{jm} N_{im}(t) - A_{ij}(t), \forall j \in [J].$$

Agents' Problem Agents operate in self-interest, and each maximizes her own expected long-run average payoff from matching over an infinite time horizon, as follows:

$$\liminf_{t \rightarrow \infty} \frac{1}{t} V_i(t), \quad \text{where } V_i(t) \triangleq \mathbb{E} \left[\sum_{m \in [K]} r_m N_{im}(t) \right].$$

We assume that each agent holds private information about her job arrival process and actions. As a result, the number of jobs of different types she currently holds and the internal matching she performs are unobservable to others.

Designer's Problem On the other hand, the designer would like to design a mechanism to incentivize agent cooperation and maximize the social welfare—that is, the total long-run average payoff from matching:

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \sum_{i \in [N]} V_i(t).$$

The designer has limited information about individual agents and knows only the aggregate job arrival rate $\boldsymbol{\lambda} \in \mathbb{R}_+^J$,² the matching reward vector $\mathbf{r} \in \mathbb{R}_+^K$, and the matching matrix $\mathbf{M} \in \mathbb{N}^{J \times K}$.

3 The Centralized Matching

In this section, we study the centralized, full-information setting as a benchmark. In this setting, a central planner has full control over the agents, observes the job arrivals and departures at each agent, and maximizes the expected long-run average payoff of the entire system. Since there is only one agent (the central planner), we adopt the notation defined for the local dynamics and drop the dependence on the index i throughout the section. Let V^* denote the long-run average payoff from centralized matching. Recalling that $N\boldsymbol{\lambda}$ is the aggregate job arrival rates, problem (1) provides a fluid relaxation for the centralized control problem.

$$\begin{aligned} V^{\text{F}}(\boldsymbol{\lambda}) &= \max_{\mathbf{x} \in \mathbb{R}_+^K} \mathbf{r}^{\text{T}} \mathbf{x} \\ \text{s.t.} \quad & \mathbf{M} \mathbf{x} \leq \boldsymbol{\lambda}. \end{aligned} \tag{1}$$

Lemma 3.1. *We have that $V^* \leq V^{\text{F}}(N\boldsymbol{\lambda}) = N \cdot V^{\text{F}}(\boldsymbol{\lambda})$.*

We prove Lemma 3.1 in Appendix A.1. Let $V^{\text{F}} \triangleq V^{\text{F}}(\boldsymbol{\lambda})$ and $\bar{\mathbf{x}} = (\bar{x}_m)_{m \in [K]} \in \mathbb{R}_+^K$ denote an optimal solution to (1). The dual problem of (1) is given by (2).

$$\begin{aligned} \min_{\mathbf{p} \in \mathbb{R}_+^J} \quad & \boldsymbol{\lambda}^{\text{T}} \mathbf{p} \\ \text{s.t.} \quad & \mathbf{M}^{\text{T}} \mathbf{p} \geq \mathbf{r}. \end{aligned} \tag{2}$$

Let $\bar{\mathbf{p}} = (\bar{p}_j)_{j \in [J]} \in \mathbb{R}_+^J$ denote an optimal solution to (2). We can interpret the value of \bar{p}_j as the marginal value of a type- j job to the payoff of the centralized problem.

3.1 Asymptotically Optimal Matching Policy

In this section, we examine a simple periodic matching policy, denoted by π^{P} , which is asymptotically optimal in the large market regime (i.e., when $N \rightarrow \infty$).

Let $\bar{\mathbf{p}} = (\bar{p}_j)_{j \in [J]}$ be an optimal solution to (2). Define $\mathcal{M}_0 \triangleq \{m \in [K] : \sum_{j \in [J]} \bar{p}_j M_{jm} > r_m\}$ as the set of suboptimal matchings in the fluid relaxation (1). To interpret this definition, note that for any optimal solution $\bar{\mathbf{x}} = (\bar{x}_m)_{m \in [K]}$ to (1), complementary slackness implies that $\bar{x}_m = 0$

²Knowing only the relative values of the aggregate job arrival rate $(\lambda_j)_{j \in [J]}$ will be sufficient for our mechanism.

for all $m \in \mathcal{M}_0$. Therefore, no match in \mathcal{M}_0 is used by any optimal solution to (1). Let $\mathcal{M}_+ \triangleq [K] \setminus \mathcal{M}_0 = \{m \in [K] : \sum_{j \in [J]} \bar{p}_j M_{jm} = r_m\}$ represents the set of matchings that may be used in the fluid relaxation (1). In what follows, we impose the constraint $x_m = 0$ for all $m \in \mathcal{M}_0$ into (1) without loss of optimality.

Definition 3.1 (Periodic Matching Policy). At each time point $t = k\Delta$, where $\Delta \in \mathbb{R}$ is a fixed time interval and $k \in \mathbb{N}_+$, given the system state $\mathbf{X}(t) \in \mathbb{N}^J$ (representing the number of jobs of each type at time t), the periodic matching policy π^P first solves the fluid relaxation $V^F(\mathbf{X}(t))$ (imposing the constraint $x_m = 0, \forall m \in \mathcal{M}_0$) to obtain an optimal solution $\mathbf{x}^* \in \mathbb{R}_+^K$, and then performs $\lfloor x_m^* \rfloor$ type- m matchings for every m .

In words, the central planner solves an optimal matching in every Δ units of time and implements the matching. In the following, we show that this policy is asymptotically optimal with an appropriately selected value of Δ and we analyze the dynamics of the system under this policy.

We first demonstrate in Lemma 3.2 that the policy π^P is asymptotically optimal in the large market regime in the centralized setting.

Lemma 3.2 (Asymptotic Optimality). *Let the interval length be $\Delta = N^{-\frac{1}{3}}$. The performance of the periodic matching policy, denoted by V^P , satisfies*

$$\frac{V^* - V^P}{V^*} \leq \frac{NV^F - V^P}{NV^F} \leq C_1 \cdot N^{-\frac{1}{3}}$$

for some constant $C_1 > 0$.

We prove Lemma 3.2 in Appendix A.2. Lemma 3.2 indicates that the fluid relaxation (1) is asymptotically tight and the policy π^P is asymptotically optimal. In the proof, we decompose the loss of policy π^P relative to the fluid relaxation upper bound over a time interval of length Δ into three components: the expiration loss, concavity loss, and rounding loss. First, jobs that arrive depart at a certain rate, leading to the expiration loss. To mitigate this, the planner would prefer a small Δ to conduct matches frequently. However, even in the absence of job departures, there remains a loss relative to the fluid relaxation due to fluctuations in job arrivals, as the function $V^F(\boldsymbol{\lambda})$ is concave (concavity loss).³ To mitigate this, the planner would prefer a large Δ to accumulate sufficient jobs before matching. Finally, implementing a rounded solution incurs a bounded loss per period. To minimize this loss relative to the payoff during the time interval, the planner would again prefer a large Δ to reduce the frequency of rounding. In the proof, we bound

³The concavity of $V^F(\boldsymbol{\lambda})$ follows from (2) and strong duality.

each of these three losses. The concavity loss dominates the rounding loss, and we determine the optimal value of Δ , which turns to be $\Theta\left(N^{-\frac{1}{3}}\right)$, to balance the concavity loss and expiration loss.

We next demonstrate that the performance guarantee of policy π^P can be strengthened if the fluid relaxation (1) satisfies some regularity condition termed non-degeneracy, which requires that the optimal dual variable $\bar{\mathbf{p}}$ is stable under small perturbations of the aggregate job arrival rates $\boldsymbol{\lambda}$. We first define the non-degeneracy condition in Definition 3.2.

Definition 3.2 (Non-Degeneracy Condition). Problem (1) is non-degenerate if there exists a positive constant $\delta > 0$ and a dual variable $\bar{\mathbf{p}} \in \mathbb{R}_+^J$ such that $\bar{\mathbf{p}}$ is the unique optimal dual variable of $V^F(\boldsymbol{\lambda}')$ for any $\boldsymbol{\lambda}'$ such that $\|\boldsymbol{\lambda}' - \boldsymbol{\lambda}\|_\infty \leq \delta$.

Remark 3.1 presents necessary and sufficient conditions for the non-degenerate condition. These conditions follow from the fact that $V^F(\boldsymbol{\lambda})$ is piecewise linear and concave in $\boldsymbol{\lambda}$, and that $\bar{\mathbf{p}} \in \mathbb{R}_+^J$ is an optimal dual variable of $V^F(\boldsymbol{\lambda})$ if and only if it is a subgradient of $V^F(\cdot)$ at the point $\boldsymbol{\lambda}$. We provide more details in Appendix A.3.

Remark 3.1 (Necessary and Sufficient Conditions for Definition 3.2). We have the following:

1. (*Necessary and Sufficient Condition*) (1) is non-degenerate if and only if it has a unique optimal dual solution $\bar{\mathbf{p}}$.
2. (*Sufficient Condition*) (1) is non-degenerate if it has a non-degenerate primal optimal basic feasible solution.⁴

Lemma 3.3 strengthens the performance of the policy π^P under the non-degeneracy condition.

Lemma 3.3. *Suppose that (1) is non-degenerate (i.e., Definition 3.2 holds), and denote by V^P the performance of the periodic matching policy π^P . The following hold:*

1. *Let the interval length be $\Delta = N^{-\frac{1}{2}}$; then*

$$\frac{V^* - V^P}{V^*} \leq \frac{NV^F - V^P}{NV^F} \leq C_2 \cdot N^{-\frac{1}{2}}$$

for some constant $C_2 > 0$.

2. *If, in addition, the matching matrix \mathbf{M} is totally unimodular,⁵ then by choosing $\Delta = c \cdot \frac{\ln N}{N}$ for some constant $c > 0$ (which depends only on the values of aggregate job arrival rates $\boldsymbol{\lambda}$*

⁴See Section 2.4 of Bertsimas and Tsitsiklis (1997) for the definition of a non-degenerate basic solution.

⁵For instance, the matching matrix \mathbf{M} is totally unimodular when every match requires two jobs of different types.

and non-degeneracy parameter δ), we have

$$\frac{V^* - V^P}{V^*} \leq \frac{NV^F - V^P}{NV^F} \leq C_3 \frac{\ln N}{N}$$

for some constant $C_3 > 0$.

We prove Lemma 3.3 in Appendix A.4. Intuitively, when the non-degeneracy condition holds, $V^F(\boldsymbol{\lambda})$ is linear in $\boldsymbol{\lambda}$ in the vicinity of $\boldsymbol{\lambda}$. As a result, the concavity loss is negligible, and we only need to balance the rounding and expiration losses. Furthermore, if the matching matrix \mathbf{M} is totally unimodular, the rounding loss is eliminated, as the optimal solution to the matching problem solved at each time point $k\Delta$ is always integral. In this case, we can simply select a small interval length Δ to minimize the expiration loss.

3.1.1 Convergence of Dynamics

In this section, we analyze the system dynamics under policy π^P . We first show in Lemma 3.4 that the long-run average matching rates under policy π^P converge to an optimal solution of (1) in the large market regime.

Lemma 3.4 (Convergence of Dynamics). *Let $x_m^P \triangleq \lim_{t \rightarrow \infty} \mathbb{E}[N_m(t)]/t$ denote the long-run average matching rate of match m under policy π^P , and let $\mathbf{x}^P = (x_m^P)_{m \in [K]} \in \mathbb{R}_+^K$. There exists a constant $C_4 > 0$ such that, for any number of agents N , there exists an optimal solution $\bar{\mathbf{x}}$ of (1) satisfying:*

$$\left\| \frac{\mathbf{x}^P}{N} - \bar{\mathbf{x}} \right\|_{\infty} \leq C_4 \cdot \frac{NV^F - V^P}{N}.$$

This, combined with Lemmas 3.2 and 3.3, implies the existence of a constant $C_5 > 0$ such that, for any number of agents N , there exists an optimal solution $\bar{\mathbf{x}}$ of (1) that satisfies the following:

1. If the interval length is $\Delta = N^{-\frac{1}{3}}$, then $\left\| \frac{\mathbf{x}^P}{N} - \bar{\mathbf{x}} \right\|_{\infty} \leq C_5 \cdot N^{-\frac{1}{3}}$;
2. If (1) is non-degenerate and $\Delta = N^{-\frac{1}{2}}$, then $\left\| \frac{\mathbf{x}^P}{N} - \bar{\mathbf{x}} \right\|_{\infty} \leq C_5 \cdot N^{-\frac{1}{2}}$;
3. If (1) is non-degenerate, the matching matrix M is totally unimodular, and $\Delta = c \cdot \frac{\ln N}{N}$, where c is the constant specified in Lemma 3.3, then $\left\| \frac{\mathbf{x}^P}{N} - \bar{\mathbf{x}} \right\|_{\infty} \leq C_5 \cdot \frac{\ln N}{N}$.

We prove Lemma 3.4 in Appendix A.5 by observing that the normalized matching rates \mathbf{x}^P/N are feasible to (1) and by leveraging the Lipschitz continuity of the LP optimal solutions with respect to changes in the constraints' right-hand side.

Remark 3.2 (Agnostic to Specific Matching Algorithm). We note that Lemma 3.4 applies to any asymptotically optimal matching policy besides policy π^P , where the convergence rate of the dynamics aligns with the respective convergence rate of performance, following the same proof.

Lemma 3.4 implies that the long-run fraction of type- j jobs that participate in the type- m match is approximately $M_{jm}\bar{x}_m/\lambda_j$ in the large market regime, as we formally state in Corollary 3.5 and prove in Appendix A.6.

Corollary 3.5. *Let $D_j(t)$ denote the number of type- j jobs that have arrived by time t , and $D_{jm}(t)$ the number of type- j jobs that have participated in type- m matches by time t under policy π . We have:*

$$q_{jm} \triangleq \lim_{t \rightarrow \infty} \mathbb{E} \left[\frac{D_{jm}(t)}{D_j(t)} \right] = \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{D_{jm}(t)}{D_j(t)} \right],$$

where q_{jm} represents the long-run average fraction of type- j jobs that participate in type- m matches, and the second equality follows from the bounded convergence theorem because $D_{jm}(t) \leq D_j(t)$. Moreover, there exists a constant $C_6 > 0$ such that, for any number of agents N , there exists an optimal solution $\bar{\mathbf{x}}$ of (1) satisfying:

$$\left| q_{jm} - M_{jm}\bar{x}_m/\lambda_j \right| \leq C_6 \cdot \frac{NV^F - V^P}{N}$$

for any $j \in [J]$ and $m \in [K]$.

Let $\bar{\mathbf{p}} = (\bar{p}_j)_{j \in [J]}$ be an optimal dual variable of (1). In the following, we define a job type j as *over-demanded* if $\bar{p}_j > 0$, and *under-demanded* if $\bar{p}_j = 0$. To interpret this definition, let \mathbf{m}_j^T denote the j -th row of the matching matrix \mathbf{M} and $\bar{\mathbf{x}}$ any optimal solution of (1). By complementary slackness, $\mathbf{m}_j^T \bar{\mathbf{x}} = \lambda_j$ for any job type j with $\bar{p}_j > 0$. Therefore, the capacity constraint in (1) is binding for all over-demanded job types. We denote the set of over-demanded job types by $\mathcal{N}_+ \triangleq \{j \in [J] : \bar{p}_j > 0\}$ and the set of under-demanded job types by $\mathcal{N}_0 \triangleq [J] \setminus \mathcal{N}_+$. According to Corollary 3.5, an over-demanded job is matched before departure with probability one in the large market regime, as we illustrate in Remark 3.3.

Remark 3.3 (Over-Demanded Jobs Matched with Probability One in the Large Market Regime). In the long run, a type- j job is matched before departure with probability $\sum_{m \in [K]} q_{jm}$, which converges to $\mathbf{m}_j^T \bar{\mathbf{x}}/\lambda_j$ in the large market regime (i.e., as $N \rightarrow \infty$) according to Corollary 3.5. For an over-demanded type- j job, since $\mathbf{m}_j^T \bar{\mathbf{x}} = \lambda_j$, it is matched with probability one before expiration. In contrast, an under-demanded type- j job remains unmatched before expiration with probability $1 - \mathbf{m}_j^T \bar{\mathbf{x}}/\lambda_j$.

4 Mechanism for Decentralized Setting

In this section, we describe our mechanism for the decentralized, limited information setting. We first introduce a mechanism with monetary transfer in Section 4.1, and we describe our non-monetary mechanism in Section 4.2. In Section 4.3, we show that our mechanism ensures the achievement of the first-best performance in the large market regime.

4.1 Warm-Up: Mechanism with Monetary Transfer

We first present a mechanism that uses monetary transfer as a preparation for our non-monetary mechanism in Section 4.2. Let $\bar{\mathbf{p}} = (\bar{p}_j)_{j \in [J]}$ be an optimal dual variable to (1). The mechanism rewards an agent with a value of \bar{p}_j whenever she submits a type- j job to the shared pool. All the submissions are irrevocable. The mechanism then performs an asymptotically optimal matching in the shared pool (e.g., implementing the periodic matching policy π^P described in Section 3.1). In the following, we show that this mechanism incentivizes all agents to submit their jobs fully.

To do so, we consider the fluid relaxation to the problem of any agent i under the above mechanism, which is given in (3).

$$\begin{aligned} \max_{\mathbf{x} \in \mathbb{R}_+^K, \mathbf{s} \in \mathbb{R}_+^J} \quad & \mathbf{r}^\top \mathbf{x} + \bar{\mathbf{p}}^\top \mathbf{s} \\ \text{s.t.} \quad & \mathbf{M}\mathbf{x} + \mathbf{s} \leq \boldsymbol{\lambda}_i. \end{aligned} \tag{3}$$

In (3), the decision variable x_m represents the rate of performing match m locally, and s_j represents the rate of submitting type- j jobs to the shared pool. The constraint requires that, for each job type, the combined rate of using jobs for local matching and submitting jobs to the shared pool can not exceed the job arrival rate. The dual problem of (3) is given in (4).

$$\begin{aligned} \min_{\mathbf{p} \in \mathbb{R}_+^J} \quad & \boldsymbol{\lambda}_i^\top \mathbf{p} \\ \text{s.t.} \quad & \mathbf{M}^\top \mathbf{p} \geq \mathbf{r}, \\ & \mathbf{p} \geq \bar{\mathbf{p}}. \end{aligned} \tag{4}$$

Proposition 4.1 gives optimal solutions to the primal and dual problems.

Proposition 4.1. $\mathbf{x} = \mathbf{0}$ and $\mathbf{s} = \boldsymbol{\lambda}_i$ is an optimal primal solution and $\bar{\mathbf{p}}$ is an optimal dual solution to (3), with the optimal value being $\boldsymbol{\lambda}_i^\top \bar{\mathbf{p}}$.

Proof. Note that $\mathbf{x} = \mathbf{0}$ and $\mathbf{s} = \boldsymbol{\lambda}_i$ is feasible to (3) and $\mathbf{p} = \bar{\mathbf{p}}$ is feasible to (4) because $\mathbf{p} = \bar{\mathbf{p}}$ is

feasible to (2). Moreover, these solutions achieve the same objective value. By weak duality, they must be optimal solutions to the primal and dual problems, respectively. \square

Proposition 4.1 shows that submitting all arriving jobs to the shared pool is optimal to an agent’s fluid relaxation problem (3). Moreover, under the mechanism, the marginal value of a type- j job to an agent’s problem equals the job’s marginal value in the centralized problem (1), which is \bar{p}_j , regardless of the agent’s individual job arrival rates. Finally, Proposition 4.1 also implies that it is optimal for an agent to submit all jobs to the shared pool in the original problem, as we state in Proposition 4.2.

Proposition 4.2. *For any agent $i \in [N]$, it is optimal to submit all the jobs to the shared pool upon arrival, and this yields an expected long-run average payoff of $\boldsymbol{\lambda}_i^T \bar{\mathbf{p}}$.*

Proof. This is because full submission achieves the fluid relaxation upper bound (3). \square

Note that the budget is balanced under the monetary mechanism in the large market regime, as we illustrate in Remark 4.1.

Remark 4.1 (Budget Balance in the Large Market Regime). According to Proposition 4.2, each agent $i \in [N]$ submits her jobs fully and receives a payoff of $\boldsymbol{\lambda}_i^T \bar{\mathbf{p}}$ under the mechanism. Consequently, the mechanism pays the agents $\sum_{i \in [N]} \boldsymbol{\lambda}_i^T \bar{\mathbf{p}} = N \boldsymbol{\lambda}^T \bar{\mathbf{p}} = NV^F(\boldsymbol{\lambda})$ in total, where the last equality follows from the strong duality between (1) and (2). On the other hand, since the shared pool implements an asymptotically optimal matching policy, it collects a matching reward of $N(V^F(\boldsymbol{\lambda}) - o(1))$ from the submitted jobs according to Lemmas 3.2 and 3.3. As a result, the mechanism is budget-balanced in the large market regime, and it distributes the payoff from centralized matching performed at the shared pool to agents fully and proportionally to each agent’s contribution to centralized matching.

Finally, Proposition 4.3 shows that submitting all jobs to the shared pool is an agent’s unique optimal strategy if each match requires at least two over-demanded jobs.

Proposition 4.3. *Submitting all jobs to the shared pool is the unique optimal strategy for an agent if every match m requires at least two over-demanded jobs (which may belong to the same type).*

We prove Proposition 4.3 in Appendix A.7. Intuitively, when an agent performs some matches locally, some over-demanded jobs will inevitably expire without being matched and thus be wasted. As a result, the long-run average payoff will be strictly less than the payoff $\boldsymbol{\lambda}_i^T \bar{\mathbf{p}}$ achieved by full submission.

4.2 Mechanism without Money Transfer

The key design principle in Section 4.1 is to ensure that submitting a type- j job to the shared pool is rewarded with its marginal value, \bar{p}_j , in the centralized setting. In this section, we show that we can achieve this without relying on monetary transfer, by allocating a match to an agent with a specified probability and letting the agent perform the match and collect the corresponding reward.

Let $\bar{\mathbf{p}} = (\bar{p}_j)_{j \in [J]}$ be an optimal solution to (2). Recall from Section 3.1 that we define $\mathcal{M}_+ \triangleq \{m \in [K] : \sum_{j \in [J]} \bar{p}_j M_{jm} = r_m\}$ as the set of matches that might be used in the fluid relaxation (1). Notably, for any optimal solution $\bar{\mathbf{x}} = (\bar{x}_m)_{m \in [K]}$ to (1), complementary slackness implies that $\bar{x}_m > 0$ only if $m \in \mathcal{M}_+$. Therefore, matches outside the set \mathcal{M}_+ , denoted by $\mathcal{M}_0 = [K] \setminus \mathcal{M}_+$, are suboptimal to (1) and can be discarded. Additionally, recall that we impose the constraint $x_m = 0$ for any $m \in \mathcal{M}_0$ whenever the periodic matching policy π^P solves a matching problem (Definition 3.1). Therefore, only matches in the set \mathcal{M}_+ are performed by policy π^P .

For any match $m \in \mathcal{M}_+$, its reward r_m equals the sum of the optimal dual variables \bar{p}_j associated with the participant jobs. Leveraging this fact, instead of reimbursing a job with its marginal value \bar{p}_j as done in Section 4.1, the non-monetary mechanism (approximately) achieves the same outcome through a random matching allocation when there are many participants. We formally describe the non-monetary mechanism in Definition 4.1.

Definition 4.1 (Non-Monetary Mechanism). The non-monetary mechanism proceeds as follows:

1. Implement an asymptotically optimal matching policy at the shared pool (e.g., the periodic matching policy π^P described in Definition 3.1).
2. Whenever a type- m match is being performed at the shared pool, assign the match to a participant job with a probability proportionally to its marginal value (i.e., with a probability of \bar{p}_j/r_m) and let the agent of that job perform the match and collect the matching reward r_m . This can be implemented by an automated algorithm such as a smart contract.

Finally, the mechanism permits agents to withdraw any submitted job, provided it has not yet been matched at the shared pool.

Note that the allocation probabilities in Step 2 are well-defined because $\sum_{j \in [J]} M_{jm} \cdot \bar{p}_j / r_m = 1$ by the definition of set \mathcal{M}_+ . Furthermore, a type- j job submitted to the shared pool yields an expected payoff of \bar{p}_j upon being matched, regardless of the specific match type. As a result, an agent is concerned only with the probability that a job will be matched, rather than the specific match type, when submitting a job.

Finally, our mechanism permits agents to withdraw any submitted job as long as it has not been matched at the shared pool, making it more practical in real-life applications. For example, in the context of kidney exchange, when a hospital “submits” a donor-patient pair to the national-level exchange (a virtual shared pool), the pair still remains physically at the hospital. As a result, it would be difficult for the national-level exchange to prohibit the hospital from using the pair, even after it has been submitted to the shared pool.

In Section 4.3, we analyze this non-monetary mechanism and demonstrate that all agents submitting their jobs fully is an approximate equilibrium in the large market regime. Intuitively, when all agents fully submit their jobs, the probability of an over-demanded job being matched before expiration approaches one as the number of agents increases (Remark 3.3). Consequently, an expected payoff of \bar{p}_j is ensured by our randomized matching allocation. This, in turn, incentivizes all agents to submit their jobs fully according to Proposition 4.2.

4.3 Performance Analysis

In this section, we analyze the non-monetary mechanism described in Section 4.2. Specifically, we show that under the mechanism, all agents submitting their jobs fully to the shared pool constitutes an approximate Nash equilibrium, meaning that the benefit from unilaterally deviating to other strategies becomes negligible as the number of agents grows large.

First, we show in Lemma 4.4 that the optimal value of (3) is an upper bound on the long-run average payoff of an agent under the non-monetary mechanism, regardless of the strategies of other agents and even when the agent has complete information about the system.

Lemma 4.4. *Under the non-monetary mechanism, the long-run average payoff of any agent i is at most the optimal value of (3), which is $\lambda_i^T \bar{\mathbf{p}}$, regardless of the strategies chosen by other agents, and even when agent i has complete information about the system.*

Proof. Let x_m represent the long-run average rate at which agent i performs match m locally, and s_j the long-run average rate at which type- j jobs submitted by agent i to the shared pool are successfully matched. The vectors $\mathbf{x} = (x_m)$ and $\mathbf{s} = (s_j)$ are feasible to (3) because, for each job type, the combined rate of jobs used for local matching and jobs matched at the shared pool can not exceed the job arrival rate. Moreover, since every type- j job matched at the shared pool yields agent i an expected payoff of \bar{p}_j , irrespective of the match type, the agent’s long-run average payoff is $\mathbf{r}^T \mathbf{x} + \bar{\mathbf{p}}^T \mathbf{s}$, which is no larger than the optimal value of (3). \square

Next, suppose that all agents submit their jobs fully to the shared pool. Since the dynamics of the shared pool converge to the fluid relaxation (1) for the centralized setting (Lemma 3.4), the long-run average payoff of agent i is at least $\lambda_i^T \bar{\mathbf{p}}$ minus a diminishing term; we state this formally in Lemma 4.5.

Lemma 4.5. *Suppose all agents submit their jobs fully to the shared pool. Then, under the non-monetary mechanism, the long-run average payoff of agent i is at least $\lambda_i^T \bar{\mathbf{p}} - C_6 \cdot \frac{NV^F - V^P}{N}$, where $C_6 \triangleq r_{\max} \cdot C_4 \cdot \sum_{j \in [J]} C_j \sum_{m \in [K]} M_{jm} > 0$ is a constant (with C_j and C_4 positive constants specified in Assumption 2.1 and Lemma 3.4, respectively), V^P denotes the long-run average matching rewards collected at the shared pool, and the term $\frac{NV^F - V^P}{N}$ diminishes to zero at a rate characterized by Lemmas 3.2 and 3.3.*

We prove Lemma 4.5 in Appendix A.8. According to Lemmas 4.4 and 4.5, the benefit of unilaterally deviating from full submission is only a negligible term that is $O(N^{-\frac{1}{3}})$ in the general case and $O(N^{-\frac{1}{2}})$ or $O(\frac{\ln N}{N})$ when (1) is non-degenerate, even when agent i has complete information about the system.

5 Numerical Experiments

In Section 4.3, we rigorously justified that submitting jobs fully when all other agents do the same is a near-optimal best response when the number of agents is large. In this section, we numerically evaluate the performance of our non-monetary mechanism in practical settings and demonstrate that full submission is approximately an equilibrium under our mechanism even with a moderate number of agents. We consider a simple synthetic example in Section 5.1 and a more realistic example using kidney exchange data in Section 5.2. All experiments are implemented in Matlab on a personal computer.

5.1 A Simple Example

We first consider a simple example with N statistically identical agents having the same job arrival rates λ_i . Specifically, there are $J = 3$ types of jobs, with arrival rates $\lambda_i = [7.5 \ 5 \ 2.5]^T$ for each agent $i \in [N]$ and departure rates $\theta = [1 \ 1 \ 1]^T$. Thus, on average, jobs arrive five times faster than they depart. Jobs can form $K = 5$ different types of matchings. The matching rewards are

$\mathbf{r} = [1 \ 1 \ 1 \ 2 \ 4]^T$, and the matching matrix is given by

$$\mathbf{M} = \begin{bmatrix} 2 & 0 & 0 & 1 & 1 \\ 0 & 2 & 0 & 1 & 1 \\ 0 & 0 & 2 & 0 & 1 \end{bmatrix}.$$

Specifically, each of the first three matching types requires two jobs of the same type and provides a reward of 1. The fourth matching type requires one job of type one and one job of type two and leads to a reward of 2. Finally, the fifth matching type requires one job from each type and yields the highest reward of 4.

Intuitively, without job departures, it is preferable to perform the fourth and final matchings and match the remaining type-one jobs with each other. This is because each job generates a higher reward from participating in the last two matching types than from matching within its own type. Notably, the fluid relaxation (1), which disregards job expiration, has a unique optimal solution, $\bar{\mathbf{x}} = [1.25 \ 0 \ 0 \ 2.5 \ 2.5]^T$, and this solution is non-degenerate because all basic variables are positive. In the original problem, however, job expirations introduce a trade-off: each agent must decide whether to wait for jobs to form a more desired match with the risk that on-hand jobs may expire. Therefore, although our example setup is simple, it captures the key trade-off agents face.

Numerical Results Suppose agents two to N submit all their jobs to the shared pool. We compare the long-run average payoff of agent one from doing the same to the optimal value of (3), which serves as an upper bound on agent one’s payoff under any strategy, even when she holds complete information about the system (Lemma 4.4).

Specifically, we take $\Delta = \frac{1}{2} \cdot N^{-\frac{1}{2}}$ for the periodic matching policy π^P based on Lemma 3.3, and number of agents N increasing linearly from 20 to 200 with a step size of 20. We present our numerical results in Figure 1. Figure 1(a) illustrates the long-run average payoff of agent one from full submission alongside the corresponding fluid upper bound (3). As the number of agents increases, the payoff that agent one obtains from full submission converges to her fluid relaxation upper bound, consistent with the theoretical result (Lemmas 4.4 and 4.5). Figure 1(b) illustrates the relative suboptimal gap $\left(= \frac{\text{fluid relaxation (3)} - \text{payoff from full submission}}{\text{payoff from full submission}} \right)$. As the number of agents increases, the sub-optimality of full submission decreases fast. In particular, the sub-optimality gap is smaller than 5% when there are 40 agents (4.78%) and drops below 2% when there are 180 agents (1.99%). Note that the sub-optimality gap we evaluate is conservative in that the fluid relaxation (3) for an agent might be a bit loose when the number of agents N is small. This is because, with




Figure 1: Simulation results for the simple example.

fewer agents, the shared pool is less congested, and the probability of an over-demanded job being matched is lower. In contrast, in the fluid relaxation (3), an over-demanded job always secures the expected payoff from matching when submitted to the shared pool. Despite this, our numerical results indicate that an agent has a limited ability to strategize when all other servers submit fully and the number of agents is not very small.

5.2 A Kidney Exchange Example

In this section, we consider a more realistic multi-hospital kidney exchange example based on real data. In this example, jobs represent incompatible patient-donor pairs, and agents represent hospitals that can perform kidney exchanges. We restrict attention to bilateral exchanges.

An exchange between two patient-donor pairs can occur if each patient is both blood-type (ABO) compatible and tissue-type compatible with the other pair's donor. ABO compatibility requires that a patient cannot receive a kidney from a donor who has a blood antigen (A or B) that the patient does not have. Figure 2 (provided in Ashlagi and Roth 2021) illustrates the ABO compatibility structure. In addition to ABO compatibility, the patient must also be tissue-type compatible with the donor, meaning that the patient cannot have antibodies against the donor's human leukocyte antigens. A commonly used measure of a patient's difficulty in finding a tissue-type compatible donor, among those who are ABO compatible, is the panel-reactive antibody (PRA). PRA quantifies the likelihood that a patient is tissue-type incompatible with a random

donor in the population based on the patient’s antibodies. Patients with a high PRA are more likely to have difficulty finding a tissue-type compatible donor. For further details, see Section 2 of Ashlagi and Roth (2021).




Figure 2: Illustration of ABO compatibility (Figure 1 of Ashlagi and Roth 2021). A directed arc of X to Y means that a donor with blood type X is compatible with a recipient with blood type Y.

Parameter Setup For simplicity, we assume that all hospitals are statistically identical and, therefore, have the same arrival rates of patient-donor pairs. We set the total arrival rate at each hospital to be one patient-donor pair every two weeks, or approximately 26 pairs per year.

Patient-donor pairs are classified based on the patient’s and donor’s blood types and the patient’s PRA score. For example, an (A–B, 65) patient-donor pair consists of a donor with blood type B and a patient with blood type A and a PRA score of 65 (indicating a 0.65 probability of incompatibility with an ABO-compatible donor).

To determine the frequency of different patient-donor types, we use statistics about the APKD (a major U.S. kidney exchange program) historical pool composition from 2010 to 2019, as provided in Table A.1 of Ashlagi and Roth (2021) and reproduced in Table 1 below. Table 1 categorizes the PRA scores into 7 intervals. Consequently, we set the total number of pair types to be the 16 ABO blood type combinations multiplied by the 7 PRA intervals, resulting in 112 types. For each PRA interval, we use the median value as the representative PRA score.⁶ In Table 1, the second column presents the empirical distribution of the patient-donor ABO pairs, and columns three to nine present the frequencies of PRA intervals, conditional on each ABO pair. As a result, the arrival rate for each type can be calculated by multiplying the total arrival rate by the percentage of the corresponding ABO pair and the conditional PRA frequency.

Finally, following the empirical setup in Ashlagi et al. (2023), we assume an average waiting time

⁶The seven intervals are $[0, 1]$, $[1, 10]$, $[10, 50]$, $[50, 80]$, $[80, 95]$, $[95, 99]$, and $[99, 100]$, and their median values are 0.5, 5.5, 30, 65, 87.5, 97, and 99.5, respectively.

of 360 days, or an expiration rate of $1/360$ per day, for any pair type.⁷ We consider each hospital’s objective to be maximizing the expected number of exchanges. Therefore, two pairs can be matched only if they are ABO compatible, and the expected payoff from matching is the probability of tissue-type compatibility, which is $(1 - \text{PRA of the first patient}) \times (1 - \text{PRA of the second patient})$.

Patient–donor ABO	% of pairs	Marginal frequencies (PRA intervals)						
		0 – 1	1 – 10	10 – 50	50 – 80	80 – 95	95 – 99	99 – 100
AB–AB	0.2	0.0	0.0	0.0	50.0	0.0	25.0	25.0
AB–B	0.4	0.0	0.0	0.0	16.7	16.7	0.0	66.7
AB–A	0.7	0.0	8.3	0.0	8.3	50.0	0.0	33.3
AB–O	0.6	10.0	0.0	20.0	10.0	0.0	20.0	40.0
B–AB	0.9	37.5	6.2	18.8	6.2	12.5	0.0	18.8
B–B	2.4	0.0	4.9	12.2	12.2	31.7	9.8	29.3
B–A	5.8	46.5	8.1	13.1	9.1	12.1	1.0	10.1
B–O	4.2	9.9	1.4	4.2	16.9	19.7	15.5	32.4
A–AB	1.0	41.2	5.9	5.9	11.8	17.6	0.0	17.6
A–B	3.6	30.6	9.7	6.5	14.5	9.7	1.6	27.4
A–A	9.7	4.2	1.8	16.9	19.3	18.1	10.8	28.9
A–O	8.8	12.7	4.7	9.3	19.3	15.3	18.0	20.7
O–AB	2.3	46.2	10.3	23.1	5.1	12.8	0.0	2.6
O–B	9.2	47.1	10.8	14.0	7.6	8.3	4.5	7.6
O–A	29.4	49.9	10.0	12.8	8.8	6.4	3.6	8.6
O–O	20.7	4.5	2.8	13.9	17.3	23.9	16.2	21.3

Table 1: APKD pool composition (2010-2019) (Table A.1 of Ashlagi and Roth 2021).

Numerical Results We again assume that hospitals two to N submit all their jobs to the shared pool, and compare the long-run average payoff of hospital one from full submission to the fluid upper bound (3). In the simulation, we vary the number of hospitals N linearly from 50 to 500 with a step size of 50, and set the matching interval length Δ in the periodic matching policy π^P to be $\Delta \in \{1, 2, 4, 7, 30\}$, corresponding to 1 day, 2 days, 4 days, 1 week, and 1 month, respectively. Our numerical results are presented in Figure 3. Figure 3(a) illustrates the long-run average payoff of hospital one from full submission under the non-monetary mechanism for different matching interval lengths, and Figure 3(b) illustrates the relative suboptimal gap. From Figure 3(b), even with only 50 hospitals participating in the kidney exchange program, the suboptimal gap of full submission remains around 5% (specifically, 5.17% when $\Delta = 7$) assuming all other hospitals fully submit their patient-donor pairs. The suboptimal gap drops further to 2% when 250 hospitals

⁷See footnote 27 in Section 5 of Ashlagi et al. (2023).

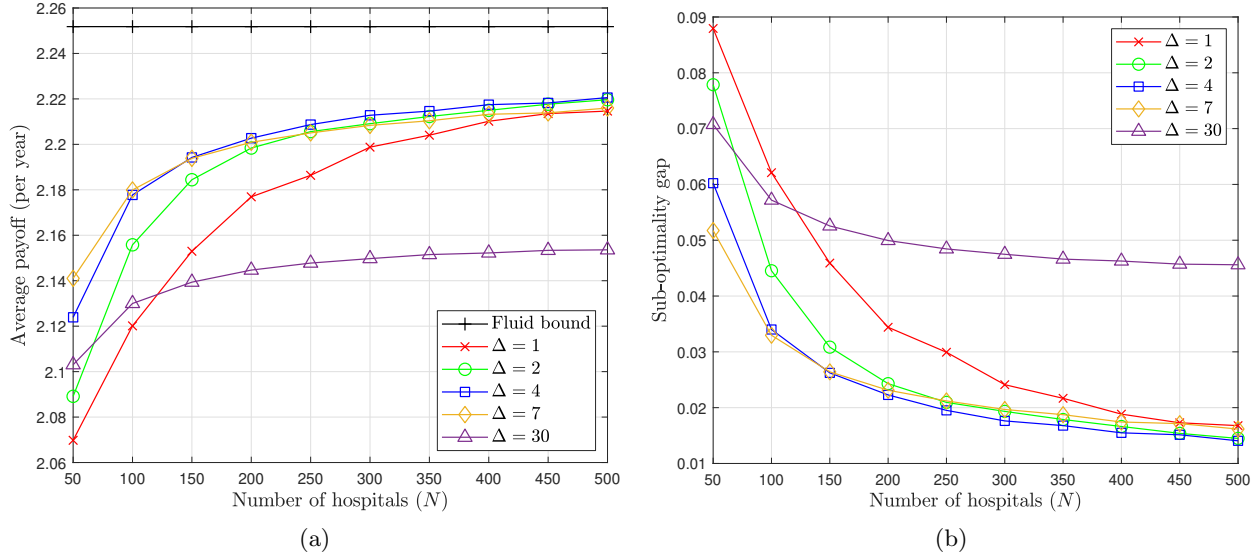


Figure 3: Simulation results for the kidney exchange example.

participate (1.95% when $\Delta = 4$). It is worth noting that there are currently 256 kidney transplant centers in the U.S. (Wang and Hart 2021). When there are 250 or more hospitals, shorter matching periods (e.g., daily or every few days) outperform less frequent matching (e.g., monthly), which aligns with current practices in the U.S. (Section 2.3 of Ashlagi and Roth 2021).

6 Conclusions

Motivated by applications in multi-hospital kidney exchanges and collaboration among intermediaries in other matching markets, we studied the problem of incentivizing participation among N strategic agents, each managing a local multi-way matching problem. We focus on a limited information setup, where an agent’s job arrivals and actions are unobservable to others, and the designer knows only the aggregate job arrival rates and matching rewards. As our main contribution, we develop a simple non-monetary mechanism (Definition 4.1) that (i) implements an asymptotically optimal matching algorithm at the shared pool, and (ii) properly distributes matching rewards through a well-designed random matching allocation. We demonstrate that this mechanism incentivizes all agents to fully submit their jobs when the number of agents is large. Consequently, the shared pool effectively operates in a centralized setting, achieving the same dynamics and performance as under centralized control. Numerical experiments based on synthetic examples and kidney exchange data demonstrate that full submission is approximately an equilibrium within our mechanism in practical settings. Ongoing work involves extending our mechanism further to

more general settings, such as when agents have heterogeneous matching rewards or hold private information about those rewards.

References

- Agarwal, N., I. Ashlagi, E. Azevedo, C. R. Featherstone, and Ö. Karaduman (2019). Market failure in kidney exchange. *American Economic Review* 109(11), 4026–70.
- Agarwal, N., I. Ashlagi, M. Rees, and A. Roth (2021). Expanding kidney exchange. Available at <http://onepercentsteps.com/policy-briefs/expanding-kidney-exchange/>. Accessed: 2024-10-03.
- Akbarpour, M., S. Li, and S. O. Gharan (2020). Thickness and information in dynamic matching markets. *Journal of Political Economy* 128(3), 783–815.
- Anily, S. and M. Haviv (2010). Cooperation in service systems. *Operations Research* 58(3), 660–673.
- Aouad, A. and O. Saritaç (2022). Dynamic stochastic matching under limited time. *Operations Research* 70(4), 2349–2383.
- Ashlagi, I., M. Burq, C. Dutta, P. Jaillet, A. Saberi, and C. Sholley (2023). Edge-weighted online windowed matching. *Mathematics of Operations Research* 48(2), 999–1016.
- Ashlagi, I., F. Fischer, I. A. Kash, and A. D. Procaccia (2015). Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior* 91, 284–296.
- Ashlagi, I., A. Nikzad, and P. Strack (2023). Matching in dynamic imbalanced markets. *The Review of Economic Studies* 90(3), 1084–1124.
- Ashlagi, I. and A. E. Roth (2014). Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics* 9(3), 817–863.
- Ashlagi, I. and A. E. Roth (2021). Kidney exchange: An operations perspective. *Management Science* 67(9), 5455–5478.
- Aveklouris, A., L. DeValve, M. Stock, and A. Ward (2024). Matching impatient and heterogeneous demand and supply. *Operations Research*.
- Banton, C. (2024). What is the network effect? Available at <https://www.investopedia.com/terms/n/network-effect.asp>. Accessed: 2024-10-03.
- Bertsimas, D. and J. N. Tsitsiklis (1997). *Introduction to linear optimization*, Volume 6. Athena scientific Belmont, MA.
- Biró, P., W. Kern, D. Pálvölgyi, and D. Paulusma (2019). Generalized matching games for international kidney exchange. In *18th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2019*, pp. 413–421. Association for Computing Machinery.
- Chen, C., Y. Chen, and P. Qian (2023). Incentivizing resource pooling. Available at SSRN 4586771.
- Cohen, M. C. and R. Zhang (2022). Competition and cooperation for two-sided platforms. *Production and Operations Management* 31(5), 1997–2014.
- Gupta, V. (2024). Greedy algorithm for multiway matching with bounded regret. *Operations Research* 72(3), 1139–1155.
- Held, P. J., F. McCormick, A. Ojo, and J. P. Roberts (2016). A cost-benefit analysis of government compensation of kidney donors. *American Journal of Transplantation* 16(3), 877–885.
- Hu, X., R. Caldentey, and G. Vulcano (2013). Revenue sharing in airline alliances. *Management Science* 59(5), 1177–1195.

- Karsten, F., M. Slikker, and G.-J. Van Houtum (2015). Resource pooling and cost allocation among independent service providers. *Operations Research* 63(2), 476–488.
- Kerimov, S., I. Ashlagi, and I. Gurvich (2023). On the optimality of greedy policies in dynamic matching. *Operations Research*.
- Kerimov, S., I. Ashlagi, and I. Gurvich (2024). Dynamic matching: Characterizing and achieving constant regret. *Management Science* 70(5), 2799–2822.
- Liu, H. and Y. Yu (2022). Incentives for shared services: multiserver queueing systems with priorities. *Manufacturing & Service Operations Management* 24(3), 1751–1759.
- Mangasarian, O. L. and T.-H. Shiau (1987). Lipschitz continuity of solutions of linear inequalities, programs and complementarity problems. *SIAM Journal on Control and Optimization* 25(3), 583–595.
- Netessine, S. and R. A. Shumsky (2005). Revenue management games: Horizontal and vertical competition. *Management Science* 51(5), 813–831.
- Roels, G. and C. S. Tang (2017). Win-win capacity allocation contracts in coproduction and codistribution alliances. *Management science* 63(3), 861–881.
- Ross, S. M. (1995). *Stochastic processes* (2 ed.). John Wiley & Sons.
- Ross, S. M. and E. A. Peköz (2023). *A Second Course in Probability* (2 ed.). Cambridge University Press.
- Roth, A. E., T. Sönmez, and M. U. Ünver (2005). Pairwise kidney exchange. *Journal of Economic theory* 125(2), 151–188.
- Royden, H. L. and P. Fitzpatrick (2010). *Real analysis* (4 ed.). Pearson.
- Sönmez, T. and M. U. Ünver (2013). Market design for kidney exchange. In *The Handbook of Market Design*, Chapter 4. Oxford University Press.
- Toulis, P. and D. C. Parkes (2015). Design and analysis of multi-hospital kidney exchange mechanisms using random graphs. *Games and Economic Behavior* 91, 360–382.
- Wang, J. H. and A. Hart (2021). Global perspective on kidney transplantation: United states. *Kidney360* 2(11), 1836–1839.
- Wei, Y., J. Xu, and S. H. Yu (2023). Constant regret primal-dual policy for multi-way dynamic matching. *Available at SSRN 4357216*.
- Wright, C. P., H. Groenevelt, and R. A. Shumsky (2010). Dynamic revenue management in airline alliances. *Transportation Science* 44(1), 15–37.

A Proofs

A.1 Proof of Lemma 3.1

Recall that

$$V^* = \max_{\pi \in \Pi} \liminf_{t \rightarrow \infty} \frac{1}{t} V^\pi(t), \quad \text{where } V^\pi(t) = \mathbb{E} \left[\sum_{m \in [K]} r_m N_m^\pi(t) \right],$$

and Π denotes the set of all feasible policies.

For any feasible policy $\pi \in \Pi$, the number of type- j jobs that have been matched by time t is $\sum_{m \in [K]} M_{jm} N_m^\pi(t)$, which must be less than $X_j(0) + D_j(t)$, the number of type- j jobs that have arrived by time t . Therefore, for any sample path, the total matching rewards by time t under a policy π is bounded from above by the following:

$$\begin{aligned} V(t)[\mathbf{X}(0) + \mathbf{D}(t)] &\triangleq \max_{\mathbf{x} \in \mathbb{N}^K} \mathbf{r}^\top \mathbf{x} \\ \text{s.t.} \quad &\mathbf{M}\mathbf{x} \leq \mathbf{X}(0) + \mathbf{D}(t). \end{aligned} \tag{5}$$

Let $\mathbf{x}[\mathbf{X}(0) + \mathbf{D}(t)]$ denote an optimal solution to (5); it holds that

$$\mathbf{M} \mathbb{E}[\mathbf{x}[\mathbf{X}(0) + \mathbf{D}(t)]] \leq \mathbf{X}(0) + \mathbb{E}[\mathbf{D}(t)] = \mathbf{X}(0) + N\boldsymbol{\lambda}t.$$

Therefore, $\mathbb{E}[\mathbf{x}[\mathbf{X}(0) + \mathbf{D}(t)]]$ is feasible to the problem of $V^F(\mathbf{X}(0) + N\boldsymbol{\lambda}t)$, and it follows that

$$V^\pi(t) \leq \mathbb{E}[V(t)[\mathbf{X}(0) + \mathbf{D}(t)]] \leq V^F(\mathbf{X}(0) + N\boldsymbol{\lambda}t).$$

According to the definition of $V^F(\boldsymbol{\lambda})$ and its dual problem (2), the function $V^F(\boldsymbol{\lambda})$ is concave and piecewise linear, and thus is continuous. As a result, we have

$$\lim_{t \rightarrow \infty} \frac{1}{t} V^F(\mathbf{X}(0) + N\boldsymbol{\lambda}t) = \lim_{t \rightarrow \infty} V^F \left(N\boldsymbol{\lambda} + \frac{1}{t} \mathbf{X}(0) \right) = V^F(N\boldsymbol{\lambda}) = NV^F(\boldsymbol{\lambda}).$$

In conclusion, we have $V^* \leq V^F(N\boldsymbol{\lambda}) = NV^F(\boldsymbol{\lambda})$.

A.2 Proof of Lemma 3.2

A.2.1 Preparations

To start, we provide some useful lemmas as preparation for the proof of Lemma 3.2. First, Lemma A.1 bounds the optimal dual variable $\bar{\mathbf{p}}$ from above.

Lemma A.1. *Let $\bar{\mathbf{p}} = (\bar{p}_j)_{j \in [J]}$ be an optimal solution to (2). It follows that $\bar{p}_j \leq r_{\max}$ for any $j \in [J]$.*

Proof. We prove by contradiction. Suppose, instead, that $\bar{p}_j > r_{\max}$ for some index $j \in [J]$. If we set $\bar{p}_j = r_{\max}$ while keeping all other components of $\bar{\mathbf{p}}$ unchanged, $\bar{\mathbf{p}}$ remains feasible to (2), and this adjustment strictly improves the objective value, contradicting the optimality of $\bar{\mathbf{p}}$. \square

Second, Lemma A.2 shows that the function $V^F(\boldsymbol{\lambda})$ is Lipschitz continuous in $\boldsymbol{\lambda}$.

Lemma A.2. *For any two arrival rate vectors $\boldsymbol{\lambda}, \boldsymbol{\lambda}' \geq 0$, we have*

$$V^F(\boldsymbol{\lambda}') - V^F(\boldsymbol{\lambda}) \leq r_{\max} \cdot \sum_{j \in [J]} (\lambda'_j - \lambda_j)^+.$$

Proof. By the strong duality, $V^F(\boldsymbol{\lambda})$ equals the optimal value of (2). Let $\bar{\mathbf{p}} = (\bar{p}_j)_{j \in [J]}$ denote an optimal solution of (2). Therefore, we have $V^F(\boldsymbol{\lambda}) = \boldsymbol{\lambda}^T \bar{\mathbf{p}}$. Moreover, Since $\bar{\mathbf{p}}$ satisfies $\mathbf{M}^T \bar{\mathbf{p}} \geq \mathbf{r}$, it holds that $V^F(\boldsymbol{\lambda}') \leq \boldsymbol{\lambda}'^T \bar{\mathbf{p}}$ by weak duality. As a result, we have:

$$V^F(\boldsymbol{\lambda}') - V^F(\boldsymbol{\lambda}) \leq \sum_{j \in [J]} \bar{p}_j (\lambda'_j - \lambda_j) \leq \sum_{j \in [J]} \bar{p}_j (\lambda'_j - \lambda_j)^+ \leq r_{\max} \cdot \sum_{j \in [J]} (\lambda'_j - \lambda_j)^+,$$

where the last inequality follows from Lemma A.1. \square

Finally, Lemma A.3 shows that the number of jobs that arrive within a time interval and remain till the end of the interval follows a specific Poisson distribution.

Lemma A.3. *For any time interval of length Δ , let Z_j denote the number of type- j jobs that arrive within the time interval and remain till the end of the interval. Then, Z_j follows a Poisson distribution with a mean value of $\frac{N\lambda_j}{\theta_j} (1 - \exp(-\theta_j \Delta))$.*

Proof. Without loss of generality, we consider the time interval $(0, \Delta]$. We first compute the ex-ante probability $\mathbb{P}_j(\Delta)$ that a type- j job which arrived during $(0, \Delta]$ is still in the system at time Δ . Note that a type- j job will still be in the system if and only if its arrival time + waiting time $> \Delta$. Therefore, it holds that

$$\begin{aligned} \mathbb{P}_j(\Delta) &= \mathbb{P}(\text{arrival time} + \text{waiting time} > \Delta) = \mathbb{P}(\text{waiting time} > \Delta - \text{arrival time}) \\ &= \frac{1}{\Delta} \int_0^\Delta \mathbb{P}(\text{waiting time} > x) dx = \frac{1}{\Delta} \int_0^\Delta \exp(-\theta_j x) dx \\ &= \frac{1}{\Delta \theta_j} (1 - \exp(-\theta_j \Delta)), \end{aligned}$$

where the third equality follows from the fact that the arrival time is uniformly distributed over $(0, \Delta]$, conditional on the job arriving within this time interval. (see Section 2.3 of Ross 1995).

Let $D_j(\Delta)$ denote the number of type- j jobs that have arrived by time Δ , which follows a Poisson distribution with a mean value of $\Delta N \lambda_j$. Since each of the $D_j(\Delta)$ jobs remains in the system at time Δ with a probability of $\mathbb{P}_j(\Delta)$, Z_j follows a Poisson distribution with a mean value of

$$\Delta N \lambda_j \cdot \mathbb{P}_j(\Delta) = \frac{N \lambda_j}{\theta_j} (1 - \exp(-\theta_j \Delta)). \quad \square$$

A.2.2 Proof of Lemma 3.2

Consider a time interval of length Δ . Let Z_j denote the number of type- j jobs that arrive within the time interval and remain till the end of the interval. According to Lemma A.3, we have that $Z_j \sim \text{Poisson}\left(\frac{N \lambda_j}{\theta_j} (1 - \exp(-\theta_j \Delta))\right)$.

Let $\mathbf{Z} = (Z_j)_{j \in [J]} \in \mathbb{R}_+^J$ be the concatenation of job arrivals, $V^P(\Delta)$ the expected payoff of the periodic matching policy π^P in a period of length Δ , and $\mathbf{X} \in \mathbb{N}^J$ the number of jobs in the system by the end of the time interval and before matches are conducted. It holds that $\mathbf{X} \geq \mathbf{Z}$ for any sample path, because there could be jobs that arrived before the time interval and still remain in the system. By the definition of policy π^P , we solve the fluid relaxation $V^F(\mathbf{X})$, obtain an optimal solution $\mathbf{x}^* \in \mathbb{R}_+^K$, and perform $\lfloor x_m^* \rfloor$ type- m matchings for every $m \in \mathcal{M}_+$.

We first bound the rounding error. Note that:

$$\sum_{m \in \mathcal{M}_+} r_m \lfloor x_m^* \rfloor \geq \sum_{m \in \mathcal{M}_+} r_m x_m^* - \min \{K, J\} \cdot r_{\max},$$

where the inequality follows from the fact that the fluid relaxation (1) has at most $\min \{K, J\}$ basic variables. Therefore, the following holds:

$$V^{\text{P}}(\Delta) \geq V^{\text{F}}(\mathbf{X}) - \min \{K, J\} \cdot r_{\max} \geq V^{\text{F}}(\mathbf{Z}) - \min \{K, J\} \cdot r_{\max}, \quad (6)$$

where the second inequality follows from the fact that $X \geq Z$ for every sample path. Consequently,

$$\begin{aligned} V^{\text{F}}(\Delta N \boldsymbol{\lambda}) - \mathbb{E}[V^{\text{P}}(\Delta)] &\leq V^{\text{F}}(\Delta N \boldsymbol{\lambda}) - \mathbb{E}[V^{\text{F}}(\mathbf{Z})] + \min \{K, J\} \cdot r_{\max} \\ &\leq r_{\max} \cdot \sum_{j \in [J]} \mathbb{E} \left(\Delta N \lambda_j - Z_j \right)^+ + \min \{K, J\} \cdot r_{\max} \\ &\leq r_{\max} \cdot \sum_{j \in [J]} \left\{ \mathbb{E} \left(\Delta N \lambda_j - \mathbb{E}[Z_j] \right)^+ + \mathbb{E} \left(\mathbb{E}[Z_j] - Z_j \right)^+ \right\} + \min \{K, J\} \cdot r_{\max} \\ &\leq r_{\max} \cdot \sum_{j \in [J]} \left\{ \frac{1}{2} \theta_j N \lambda_j \Delta^2 + \sqrt{N \lambda_j \Delta} \right\} + \min \{K, J\} \cdot r_{\max}. \end{aligned}$$

where the first inequality follows from (6), the second inequality from Lemma A.2, and the fourth inequality from the facts that

$$\mathbb{E}[Z_j] = N \lambda_j \int_0^\Delta \exp(-\theta_j x) dx \geq N \lambda_j \int_0^\Delta (1 - \theta_j x) dx = N \lambda_j \Delta \left(1 - \frac{1}{2} \theta_j \Delta \right), \quad (7)$$

and that

$$\begin{aligned} \mathbb{E} \left(\mathbb{E}[Z_j] - Z_j \right)^+ &\leq \mathbb{E} \left| \mathbb{E}[Z_j] - Z_j \right| \leq \left[\mathbb{E} \left(\mathbb{E}[Z_j] - Z_j \right)^2 \right]^{\frac{1}{2}} = \sqrt{\text{Var}(Z_j)} \\ &= \sqrt{\frac{N \lambda_j}{\theta_j} (1 - \exp(-\theta_j \Delta))} \leq \sqrt{N \lambda_j \Delta}. \end{aligned}$$

By taking $\Delta = N^{-\frac{1}{3}}$, we have

$$\begin{aligned} \frac{N V^{\text{F}} - V^{\text{P}}}{N V^{\text{F}}} &= \frac{V^{\text{F}}(\Delta N \boldsymbol{\lambda}) - \mathbb{E}[V^{\text{P}}(\Delta)]}{V^{\text{F}}(\Delta N \boldsymbol{\lambda})} \\ &\leq r_{\max} \cdot \frac{\sum_{j \in [J]} \left\{ \sqrt{N \lambda_j \Delta} + \frac{1}{2} \theta_j N \lambda_j \Delta^2 \right\} + \min \{K, J\}}{\Delta N \cdot V^{\text{F}}} \\ &\leq \frac{r_{\max}}{V^{\text{F}}} \cdot \sum_{j \in [J]} \left(\sqrt{\lambda_j} \cdot \frac{1}{\sqrt{\Delta N}} + \frac{\lambda_j \theta_j}{2} \cdot \Delta + \frac{1}{\Delta N} \right) \\ &\leq \frac{r_{\max}}{V^{\text{F}}} \cdot \sum_{j \in [J]} \left(\sqrt{\lambda_j} + \frac{\lambda_j \theta_j}{2} + 1 \right) \cdot N^{-\frac{1}{3}}. \end{aligned}$$

A.3 Details of Remark 3.1

Proof of Bullet One We note that function $V^F(\boldsymbol{\lambda})$ is piecewise linear and concave in $\boldsymbol{\lambda}$, and that $\bar{\mathbf{p}} \in \mathbb{R}_+^J$ is an optimal dual variable of $V^F(\boldsymbol{\lambda})$ if and only if it is a subgradient of $V^F(\cdot)$ at the point $\boldsymbol{\lambda}$ (see Sections 5.2 and 5.3 of Bertsimas and Tsitsiklis 1997). Therefore, $\bar{\mathbf{p}}$ is the unique optimal dual solution if and only if $\boldsymbol{\lambda}$ is not a kink of the function $V^F(\cdot)$, and since $V^F(\boldsymbol{\lambda})$ is piecewise linear, this condition holds if and only if $\bar{\mathbf{p}}$ is the unique optimal dual solution in the neighborhood of $\boldsymbol{\lambda}$.

Proof of Bullet Two Let (8) be the standard form of the fluid relaxation (1) by introducing the slack variable $\mathbf{s} \in \mathbb{R}_+^J$.

$$\begin{aligned} V^F(\boldsymbol{\lambda}) &= \max_{\mathbf{x} \in \mathbb{R}_+^K, \mathbf{s} \in \mathbb{R}_+^J} \mathbf{r}^T \mathbf{x} \\ \text{s.t.} \quad & \mathbf{M}\mathbf{x} + \mathbf{s} = \boldsymbol{\lambda}. \end{aligned} \quad (8)$$

We show that $\mathbf{x} \in \mathbb{R}^K$ is a non-degenerate basic feasible solution of (1) if and only if \mathbf{x} together with $\mathbf{s} = \boldsymbol{\lambda} - \mathbf{M}\mathbf{x}$ is a non-degenerate basic feasible solution of (8). For a matrix $\mathbf{M} \in \mathbb{R}^{J \times K}$ and subsets $A \subseteq [J]$ and $B \subseteq [K]$, let $\mathbf{M}_{A,B}$ denote a sub-matrix of \mathbf{M} containing only the rows in set A and columns in set B . Additionally, for a feasible solution \mathbf{x} to (1), define $I \triangleq \{i \in [K] : x_i = 0\}$ and $H \triangleq \{j \in [J] : \mathbf{m}_j^T \mathbf{x} = \lambda_j\}$, where \mathbf{m}_j^T represents the j -th row of the matching matrix \mathbf{M} . On one hand, \mathbf{x} is a non-degenerate basic feasible solution of (1) if and only if \mathbf{x} is feasible to (1) and the following conditions hold:

$$\begin{aligned} |I| + |H| &= K, \\ \text{The sub-matrix } \mathbf{M}_{H, [K] \setminus I} &\text{ has full rank.} \end{aligned} \quad (9)$$

On the other hand, \mathbf{x} together with $\mathbf{s} = \boldsymbol{\lambda} - \mathbf{M}\mathbf{x}$ is a non-degenerate basic solution of (8) if and only if \mathbf{x} is feasible to (1) and the following conditions hold:

$$\begin{aligned} |I| + |H| &= K, \\ \text{The matrix } (\mathbf{M}_{[J], [K] \setminus I}, \mathbf{I}_{[J], [J] \setminus [H]}) &\text{ has full rank,} \end{aligned} \quad (10)$$

where $\mathbf{I}_{[J], [J] \setminus [H]}$ is a sub-matrix of the $J \times J$ identity matrix. Since $(\mathbf{M}_{[J], [K] \setminus I}, \mathbf{I}_{[J], [J] \setminus [H]})$ has full rank if and only if $\mathbf{M}_{H, [K] \setminus I}$ has full rank, (9) and (10) are equivalent.

Note that (2) is also the dual of (8). Therefore, if (1) has a non-degenerate optimal basic feasible solution, so does (8), and this implies that (2) has a unique optimal solution according to Bertsimas and Tsitsiklis (1997) (see, for example, the discussion following Example 4.6 or preceding Theorem 5.1).

A.4 Proof of Lemma 3.3

Let $\bar{\mathbf{p}}$ be the unique optimal dual variable of $V^F(\hat{\boldsymbol{\lambda}})$ for any $\hat{\boldsymbol{\lambda}}$ such that $\|\hat{\boldsymbol{\lambda}} - \boldsymbol{\lambda}\|_\infty \leq \delta$. We have:

$$V^F(\hat{\boldsymbol{\lambda}}) = \hat{\boldsymbol{\lambda}}^T \bar{\mathbf{p}}, \quad \forall \hat{\boldsymbol{\lambda}} \text{ such that } \|\hat{\boldsymbol{\lambda}} - \boldsymbol{\lambda}\|_\infty \leq \delta. \quad (11)$$

Analogous to Appendix A.2, we focus on a time interval of length Δ and define random variable Z_j to be the number of type- j jobs that arrive within the time interval and remain till the end of the interval, which follows a Poisson distribution with a mean value of $\frac{N\lambda_j}{\theta_j} (1 - \exp(-\theta_j \Delta))$ by

Lemma A.3. Note that the following holds:

$$(1 - \theta_j \Delta) \Delta N \lambda_j \leq \left(1 - \frac{1}{2} \theta_j \Delta\right) \Delta N \lambda_j \leq \frac{N \lambda_j}{\theta_j} (1 - \exp(-\theta_j \Delta)) \leq \Delta N \lambda_j,$$

where the second inequality follows from (7) and the third inequality follows from the fact that $e^x \geq 1 + x$ for any $x \in \mathbb{R}$. We can construct two Poisson random variables $\bar{Z}_j \sim \text{Poisson}(\Delta N \lambda_j)$ and $\underline{Z}_j \sim \text{Poisson}((1 - \theta_j \Delta) \Delta N \lambda_j)$, such that with proper coupling,

$$\underline{Z}_j \leq Z_j \leq \bar{Z}_j$$

holds for every sample path.

A.4.1 Proof of Bullet One

Within a time interval of length Δ , the expected payoff of the periodic matching policy $V^P(\Delta)$ is at least $V^F(\mathbf{Z}) - \min\{K, J\} r_{\max}$ according to (6), and the payoff of the fluid relaxation is $V^F(\Delta N \boldsymbol{\lambda})$.

Take the interval length to be $\Delta = N^{-\frac{1}{2}}$, and define the following events:

$$\begin{aligned} \bar{B}_j &= \{Z_j > \lfloor \Delta N(\lambda_j + \delta) \rfloor\}, \forall j \in [J], \\ \underline{B}_j &= \{Z_j < \lceil \Delta N(\lambda_j - \delta) \rceil\}, \forall j \in [J], \\ A &= (\cup_{j \in [J]} (\bar{B}_j \cup \underline{B}_j))^c \subseteq \{\|\mathbf{Z} - \Delta N \boldsymbol{\lambda}\|_\infty \leq \Delta N \delta\}. \end{aligned}$$

By the concentration inequality for Poisson random variables (Lemma A.6), there exists positive constants $c_1, c_2 > 0$, such that the following hold:

$$\begin{aligned} \mathbb{P}(\bar{B}_j) &\leq \mathbb{P}(\bar{Z}_j \geq \lfloor \Delta N(\lambda_j + \delta) \rfloor) \leq c_1 \cdot \exp\left(-c_2 \cdot N^{\frac{1}{2}}\right), \\ \mathbb{P}(\underline{B}_j) &\leq \mathbb{P}(\underline{Z}_j \leq \lceil \Delta N(\lambda_j - \delta) \rceil) \leq c_1 \cdot \exp\left(-c_2 \cdot N^{\frac{1}{2}}\right). \end{aligned} \tag{12}$$

In the first line, the first inequality follows from the fact that $Z_j \leq \bar{Z}_j$, and the second inequality from taking $\Delta = N^{-\frac{1}{2}}$ and Lemma A.6. The second line can be justified analogously.

Using the union bound, we obtain:

$$\mathbb{P}(A^c) \leq \sum_{j \in [J]} (\mathbb{P}(\bar{B}_j) + \mathbb{P}(\underline{B}_j)) \leq 2Jc_1 \cdot \exp\left(-c_2 \cdot N^{\frac{1}{2}}\right). \tag{13}$$

In addition, we have $\|\frac{\mathbf{Z}}{\Delta N} - \boldsymbol{\lambda}\|_\infty \leq \delta$ conditioning on event A . Therefore, from (11) we have:

$$V^F(\Delta N \boldsymbol{\lambda}) - V^F(\mathbf{Z}) \leq \mathbf{1}[A] \cdot \bar{\mathbf{p}}^T \cdot (\Delta N \boldsymbol{\lambda} - \mathbf{Z}) + \mathbf{1}[A^c] \cdot V^F(\Delta N \boldsymbol{\lambda}).$$

Consequently,

$$\begin{aligned} \frac{NV^F - V^P}{NV^F} &= \frac{V^F(\Delta N \boldsymbol{\lambda}) - \mathbb{E}[V^P(\Delta)]}{V^F(\Delta N \boldsymbol{\lambda})} \\ &\leq \frac{V^F(\Delta N \boldsymbol{\lambda}) - \mathbb{E}[V^F(\mathbf{Z})] + \min\{K, J\} \cdot r_{\max}}{V^F(\Delta N \boldsymbol{\lambda})} \\ &\leq \frac{\mathbb{E}[\mathbf{1}[A] \cdot \bar{\mathbf{p}}^T \cdot (\Delta N \boldsymbol{\lambda} - \mathbf{Z})] + \min\{K, J\} \cdot r_{\max}}{\Delta N \cdot V^F} + \mathbb{P}(A^c), \end{aligned} \tag{14}$$

where the first inequality follows from (6). Lemma A.4 serves as preparation to bound the first term in (14) from above.

Lemma A.4. *For any job type $j \in [J]$, we have the following:*

$$\mathbb{E}[\mathbf{1}[A] \cdot Z_j] \geq \Delta N \lambda_j \cdot (1 - \theta_j \Delta) - \Delta N \cdot (\lambda_j + 2J(\lambda_j + \delta)) \cdot c_1 \cdot \exp\left(-c_2 \cdot N^{\frac{1}{2}}\right).$$

Proof. First, note that:

$$\mathbb{E}[\mathbf{1}[A] \cdot Z_j] = \mathbb{E}[Z_j] - \mathbb{E}[\mathbf{1}[A^c] \cdot Z_j] \geq \Delta N \lambda_j \cdot (1 - \theta_j \Delta) - \mathbb{E}[\mathbf{1}[A^c] \cdot Z_j],$$

where the inequality follows from $Z_j \geq \underline{Z}_j$, and hence $\mathbb{E}[Z_j] \geq \mathbb{E}[\underline{Z}_j] = \Delta N \lambda_j \cdot (1 - \theta_j \Delta)$.

We now bound the second term from above, as follows:

$$\begin{aligned} \mathbb{E}[\mathbf{1}[A^c] \cdot Z_j] &= \mathbb{E}[\mathbf{1}[A^c \cap \bar{B}_j] \cdot Z_j] + \mathbb{E}[\mathbf{1}[A^c \cap \bar{B}_j^c] \cdot Z_j] \\ &\leq \mathbb{E}[\mathbf{1}[\bar{B}_j] \cdot Z_j] + \mathbb{P}[A^c \cap \bar{B}_j^c] \cdot \Delta N (\lambda_j + \delta) \\ &\leq \mathbb{E}[\mathbf{1}[\bar{Z}_j > \lfloor \Delta N (\lambda_j + \delta) \rfloor] \cdot \bar{Z}_j] + \mathbb{P}[A^c] \cdot \Delta N (\lambda_j + \delta) \\ &\leq \Delta N \lambda_j \cdot \mathbb{P}(\bar{Z}_j \geq \lfloor \Delta N (\lambda_j + \delta) \rfloor) + 2J \cdot \Delta N (\lambda_j + \delta) \cdot c_1 \cdot \exp\left(-c_2 \cdot N^{\frac{1}{2}}\right) \\ &\leq \Delta N \cdot (\lambda_j + 2J(\lambda_j + \delta)) \cdot c_1 \cdot \exp\left(-c_2 \cdot N^{\frac{1}{2}}\right), \end{aligned}$$

where the first inequality follows from the facts that $\bar{B}_j \subseteq A^c$ and that $Z_j \leq \Delta N (\lambda_j + \delta)$ conditional on event \bar{B}_j^c , the second inequality from $Z_j \leq \bar{Z}_j$, the third inequality from (13) and Lemma A.7, and the fourth inequality from (12). \square

From Lemma A.4, we have the following:

$$\begin{aligned} \mathbb{E}[\mathbf{1}[A] \cdot \bar{\mathbf{p}}^\top \cdot (\Delta N \boldsymbol{\lambda} - \mathbf{Z})] &\leq \sum_{j \in [J]} \bar{p}_j \cdot \left(\Delta N \lambda_j - \mathbb{E}[\mathbf{1}[A] \cdot Z_j] \right) \\ &\leq r_{\max} \cdot \sum_{j \in [J]} \left(\Delta N \lambda_j \cdot \theta_j \Delta + \Delta N \cdot (\lambda_j + 2J(\lambda_j + \delta)) \cdot c_1 \cdot \exp\left(-c_2 \cdot N^{\frac{1}{2}}\right) \right), \end{aligned} \tag{15}$$

where the second inequality follows from Lemma A.4 and the fact that $\bar{p}_j \leq r_{\max}$ for any $j \in [J]$ (Lemma A.1). Combining (13) – (15), we obtain:

$$\begin{aligned} \frac{NV^F - V^P}{NV^F} &\leq \frac{V^F(\Delta N \boldsymbol{\lambda}) - \mathbb{E}[V^F(\mathbf{Z})] + \min\{K, J\} \cdot r_{\max}}{V^F(\Delta N \boldsymbol{\lambda})} \\ &\leq \frac{r_{\max}}{V^F} \cdot \sum_{j \in [J]} \left(\lambda_j \theta_j \Delta + \frac{1}{\Delta N} + (\lambda_j + 2J(\lambda_j + \delta)) \cdot c_1 \cdot \exp\left(-c_2 \cdot N^{\frac{1}{2}}\right) \right) + 2Jc_1 \cdot \exp\left(-c_2 \cdot N^{\frac{1}{2}}\right) \\ &= O\left(N^{-\frac{1}{2}}\right) \end{aligned}$$

by using $\Delta = N^{-\frac{1}{2}}$.

A.4.2 Proof of Bullet Two

Assume that the matching matrix \mathbf{M} is totally unimodular. Then, the optimal solution to the matching problem solved at any time point $k\Delta$ is always integral. Therefore, rounding is not

needed to perform the matching, and we have the following:

$$V^{\mathbf{P}}(\Delta) = V^{\mathbf{F}}(\mathbf{X}) \geq V^{\mathbf{F}}(\mathbf{Z}), \quad (16)$$

following the same notation as in (6).

Define a constant $c \triangleq \max_{j \in [J]} \left(\frac{1}{\lambda_j g(\delta/\lambda_j)} \vee \frac{2}{\lambda_j g(-\delta/\lambda_j)} \right)$, where $g(u) \triangleq (1+u) \ln(1+u) - u \geq 0$ is defined in Lemma A.6. Note that $g(0) = 0$ and $g(u) > 0$ for any $u \neq 0$. Take the interval length to be $\Delta = 2c \cdot \frac{\ln N}{N}$. Moreover, assume that N is sufficiently large so that: (i) $\Delta \leq \min_{j \in [J]} \frac{\delta}{\lambda_j \theta_j}$ and (ii) $\frac{1}{2} \cdot g(-\delta/\lambda_j) \leq (1 - \theta_j \Delta) \cdot g\left(\frac{\lambda_j \theta_j \Delta - \delta}{\lambda_j(1 - \theta_j \Delta)}\right)$ for all $j \in [J]$. Finally, define the following events:

$$\begin{aligned} \bar{B}_j &= \{Z_j > \Delta N(\lambda_j + \delta)\}, \forall j \in [J], \\ \underline{B}_j &= \{Z_j < \Delta N(\lambda_j - \delta)\}, \forall j \in [J], \\ A &= \left(\cup_{j \in [J]} (\bar{B}_j \cup \underline{B}_j) \right)^c = \{\|\mathbf{Z} - \Delta N \boldsymbol{\lambda}\|_\infty \leq \Delta N \delta\}. \end{aligned}$$

By the concentration inequality for Poisson random variables (Lemma A.6), we have:

$$\begin{aligned} \mathbb{P}(\bar{B}_j) &\leq \mathbb{P}(\bar{Z}_j \geq \Delta N(\lambda_j + \delta)) \leq \exp(-\Delta N \lambda_j \cdot g(\delta/\lambda_j)) \leq \frac{1}{N^2}, \\ \mathbb{P}(\underline{B}_j) &\leq \mathbb{P}(\underline{Z}_j \leq \Delta N(\lambda_j - \delta)) \leq \exp\left(-\Delta N \lambda_j \cdot (1 - \theta_j \Delta) \cdot g\left(\frac{\lambda_j \theta_j \Delta - \delta}{\lambda_j(1 - \theta_j \Delta)}\right)\right) \leq \frac{1}{N^2}. \end{aligned}$$

In the first line, the first inequality follows from the fact that $Z_j \leq \bar{Z}_j$, the second inequality from Lemma A.6, and the third inequality from the definitions of the interval length Δ and constant c . Analogously, in the second line, the first inequality follows from the fact that $\underline{Z}_j \leq Z_j$, the second inequality from Lemma A.6, and the third inequality from the definitions of the interval length Δ and constant c and the aforementioned properties of Δ given that N is sufficiently large.

Using the union bound, we obtain:

$$\mathbb{P}(A^c) \leq \sum_{j \in [J]} (\mathbb{P}(\bar{B}_j) + \mathbb{P}(\underline{B}_j)) \leq \frac{2J}{N^2}. \quad (17)$$

In addition, we have $\|\frac{\mathbf{Z}}{\Delta N} - \boldsymbol{\lambda}\|_\infty \leq \delta$ conditioning on event A . Therefore, from (11) we get:

$$V^{\mathbf{F}}(\Delta N \boldsymbol{\lambda}) - V^{\mathbf{F}}(\mathbf{Z}) \leq \mathbf{1}[A] \cdot \bar{\mathbf{p}}^{\mathbf{T}} \cdot (\Delta N \boldsymbol{\lambda} - \mathbf{Z}) + \mathbf{1}[A^c] \cdot V^{\mathbf{F}}(\Delta N \boldsymbol{\lambda}).$$

Consequently,

$$\begin{aligned} \frac{NV^{\mathbf{F}} - V^{\mathbf{P}}}{NV^{\mathbf{F}}} &= \frac{V^{\mathbf{F}}(\Delta N \boldsymbol{\lambda}) - \mathbb{E}[V^{\mathbf{P}}(\Delta)]}{V^{\mathbf{F}}(\Delta N \boldsymbol{\lambda})} \\ &\leq \frac{V^{\mathbf{F}}(\Delta N \boldsymbol{\lambda}) - \mathbb{E}[V^{\mathbf{F}}(\mathbf{Z})]}{V^{\mathbf{F}}(\Delta N \boldsymbol{\lambda})} \\ &\leq \frac{\mathbb{E}[\mathbf{1}[A] \cdot \bar{\mathbf{p}}^{\mathbf{T}} \cdot (\Delta N \boldsymbol{\lambda} - \mathbf{Z})]}{\Delta N \cdot V^{\mathbf{F}}} + \mathbb{P}(A^c), \end{aligned} \quad (18)$$

where the first inequality follows from (16). Analogous to Lemma A.4, Lemma A.5 helps bound the first term in (18) from above.

Lemma A.5. For any job type $j \in [J]$, we have the following:

$$\mathbb{E}[\mathbf{1}[A] \cdot Z_j] \geq \Delta N \lambda_j \cdot (1 - \theta_j \Delta) - \Delta \lambda_j - 2J \cdot (\lambda_j + \delta) \cdot \frac{\Delta}{N}$$

when N is larger than a certain constant.

Proof. First, note that:

$$\mathbb{E}[\mathbf{1}[A] \cdot Z_j] = \mathbb{E}[Z_j] - \mathbb{E}[\mathbf{1}[A^c] \cdot Z_j] \geq \Delta N \lambda_j \cdot (1 - \theta_j \Delta) - \mathbb{E}[\mathbf{1}[A^c] \cdot Z_j],$$

where the inequality follows from $Z_j \geq \underline{Z}_j$, and hence $\mathbb{E}[Z_j] \geq \mathbb{E}[\underline{Z}_j] = \Delta N \lambda_j \cdot (1 - \theta_j \Delta)$.

We now bound the second term from above, as follows:

$$\begin{aligned} \mathbb{E}[\mathbf{1}[A^c] \cdot Z_j] &= \mathbb{E}[\mathbf{1}[A^c \cap \bar{B}_j] \cdot Z_j] + \mathbb{E}[\mathbf{1}[A^c \cap \bar{B}_j^c] \cdot Z_j] \\ &\leq \mathbb{E}[\mathbf{1}[\bar{B}_j] \cdot Z_j] + \mathbb{P}[A^c \cap \bar{B}_j^c] \cdot \Delta N (\lambda_j + \delta) \\ &\leq \mathbb{E}[\mathbf{1}[\bar{Z}_j > \Delta N (\lambda_j + \delta)] \cdot \bar{Z}_j] + \mathbb{P}[A^c] \cdot \Delta N (\lambda_j + \delta) \\ &\leq \Delta N \lambda_j \cdot \mathbb{P}[\bar{Z}_j \geq \lfloor \Delta N (\lambda_j + \delta) \rfloor] + 2J \cdot (\lambda_j + \delta) \cdot \frac{\Delta}{N} \\ &\leq \Delta \lambda_j + 2J \cdot (\lambda_j + \delta) \cdot \frac{\Delta}{N}. \end{aligned}$$

In the above, the first inequality holds because $\bar{B}_j \subseteq A^c$ and $Z_j \leq \Delta N (\lambda_j + \delta)$ conditional on the event \bar{B}_j^c . The second inequality follows from the fact that $Z_j \leq \bar{Z}_j$. The third inequality follows from (17) and Lemma A.7. Finally, let $\tilde{\delta}$ be a positive constant satisfying that $\tilde{\delta} < \delta$ and that $c \geq \max_{j \in [J]} \frac{1}{2 \cdot \lambda_j g(\tilde{\delta}/\lambda_j)}$. Such a constant $\tilde{\delta}$ exists by the definition of c . The fourth inequality follows from the fact that $\Delta N (\lambda_j + \tilde{\delta}) \leq \lfloor \Delta N (\lambda_j + \delta) \rfloor$ when N is sufficiently large (specifically, when $\Delta N \geq 1/(\delta - \tilde{\delta})$) and that

$$\mathbb{P}(\bar{Z}_j \geq \Delta N (\lambda_j + \tilde{\delta})) \leq \exp(-\Delta N \lambda_j \cdot g(\tilde{\delta}/\lambda_j)) \leq \frac{1}{N}$$

by the definition of $\tilde{\delta}$ and Lemma A.6. □

From Lemma A.5, we have the following:

$$\begin{aligned} \mathbb{E}[\mathbf{1}[A] \cdot \bar{\mathbf{p}}^T \cdot (\Delta N \boldsymbol{\lambda} - \mathbf{Z})] &\leq \sum_{j \in [J]} \bar{p}_j \cdot (\Delta N \lambda_j - \mathbb{E}[\mathbf{1}[A] \cdot Z_j]) \\ &\leq r_{\max} \cdot \sum_{j \in [J]} \left(\Delta N \lambda_j \cdot \theta_j \Delta + \Delta \lambda_j + 2J \cdot (\lambda_j + \delta) \cdot \frac{\Delta}{N} \right), \end{aligned} \tag{19}$$

where the second inequality follows from Lemma A.5 and the fact that $\bar{p}_j \leq r_{\max}$ for any $j \in [J]$

(Lemma A.1). Combining (17) - (19), we obtain:

$$\begin{aligned} \frac{NV^F - V^P}{NV^F} &\leq \frac{V^F(\Delta N\boldsymbol{\lambda}) - \mathbb{E}[V^F(\mathbf{Z})]}{V^F(\Delta N\boldsymbol{\lambda})} \\ &\leq \frac{r_{\max}}{V^F} \cdot \sum_{j \in [J]} \left(\lambda_j \theta_j \Delta + \frac{\lambda_j}{N} + 2J \cdot (\lambda_j + \delta) \cdot \frac{1}{N^2} \right) + \frac{2J}{N^2} \\ &= O\left(\frac{\ln N}{N}\right) \end{aligned}$$

by using $\Delta = \Theta\left(\frac{\ln N}{N}\right)$.

Lemma A.6 (Concentration Inequality for Poisson Random Variables). *Let $X \sim \text{Poisson}(\lambda)$ be a Poisson random variable with a mean value of $\lambda > 0$. We have*

$$\mathbb{P}(X \geq \lambda(1 + u)) \leq \exp(-\lambda g(u))$$

for any $u \geq 0$, and

$$\mathbb{P}(X \leq \lambda(1 - u)) \leq \exp(-\lambda g(-u))$$

for any $0 \leq u < 1$, where $g(u) = (1 + u) \ln(1 + u) - u$.

Proof. We first prove the first inequality. For any $t \geq 0$, we have

$$\begin{aligned} \mathbb{P}(X \geq \lambda(1 + u)) &= \mathbb{P}\left[\exp(tX) \geq \exp(t\lambda(1 + u))\right] \\ &\leq \mathbb{E}\left[\exp(tX)\right] \exp(-t\lambda(1 + u)) \\ &= \exp(\lambda(e^t - 1) - t\lambda(1 + u)) \end{aligned}$$

where the inequality follows from Markov's inequality and the second equality from the moment generating function of the Poisson distribution. It turns out that the right-hand side of the second equality is minimized by setting $t = \ln(1 + u)$, in which case the right-hand side simplifies to $\exp(-\lambda g(u))$. The second inequality can be proven similarly. \square

Lemma A.7. *Let $X \sim \text{Poisson}(\lambda)$ be a Poisson random variable with a mean value of $\lambda > 0$ and $z \in \mathbb{N}$ be a nonnegative integer. The following holds:*

$$\mathbb{E}[X \cdot \mathbb{1}[X > z]] = \sum_{k=z+1}^{\infty} k \cdot \frac{\lambda^k e^{-\lambda}}{k!} = \lambda \cdot \sum_{k=z}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!} = \lambda \cdot \mathbb{P}(X \geq z).$$

A.5 Proof of Lemma 3.4

Let x_m^P represent the long-run average matching rate of match m under policy π^P , and let $\mathbf{x}^P = (x_m^P)_{m \in [K]} \in \mathbb{R}_+^K$. It follows that $\mathbf{M}\mathbf{x}^P \leq N\boldsymbol{\lambda}$ because the rate of jobs matched must be less than the rate of job arrivals. Therefore, $\frac{\mathbf{x}^P}{N}$ is an optimal solution to the following LP (20):

$$\begin{aligned} &\max_{\mathbf{x} \in \mathbb{R}_+^K} \mathbf{r}^T \mathbf{x} \\ \text{s.t.} \quad &\mathbf{M}\mathbf{x} \leq \boldsymbol{\lambda}, \\ &\mathbf{r}^T \mathbf{x} \leq \frac{\mathbf{r}^T \mathbf{x}^P}{N}, \end{aligned} \tag{20}$$

because $\frac{\mathbf{x}^P}{N}$ is feasible and obtains the highest possible objective value. On the other hand, since the optimal value of the fluid relaxation (1) is V^F , a solution is optimal to (1) if and only if it is optimal to the following LP (21):

$$\begin{aligned} \max_{\mathbf{x} \in \mathbb{R}_+^K} \quad & \mathbf{r}^T \mathbf{x} \\ \text{s.t.} \quad & \mathbf{M} \mathbf{x} \leq \boldsymbol{\lambda}, \\ & \mathbf{r}^T \mathbf{x} \leq V^F. \end{aligned} \tag{21}$$

By Lipschitz continuity of optimal solutions of LPs with respect to constraints' right-hand side (i.e., Lemma A.8 below), there exists an optimal solution $\bar{\mathbf{x}}$ to the fluid relaxation (1) such that

$$\left\| \frac{\mathbf{x}^P}{N} - \bar{\mathbf{x}} \right\|_{\infty} \leq C_4 \left| V^F - \frac{\mathbf{r}^T \mathbf{x}^P}{N} \right| = C_4 \cdot \frac{NV^F - V^P}{N}.$$

for some constant $C_4 > 0$. The rest of Lemma 3.4 follows from the fact that policy π^P is asymptotically optimal (i.e., Lemmas 3.2 and 3.3 hold).

Lemma A.8 (Theorem 2.4 of Mangasarian and Shiao 1987). *Consider the following LP:*

$$P(\mathbf{b}) = \max \{ \mathbf{r}^T \mathbf{x} : \mathbf{A} \mathbf{x} \leq \mathbf{b} \}.$$

There exists a constant κ such that, for any vectors $\mathbf{b}, \hat{\mathbf{b}} \in \mathbb{R}_+^n$ and any optimal solution \mathbf{x} to $P(\mathbf{b})$, there exists an optimal solution $\hat{\mathbf{x}}$ to $P(\hat{\mathbf{b}})$ such that $\|\mathbf{x} - \hat{\mathbf{x}}\|_{\infty} \leq \kappa \|\mathbf{b} - \hat{\mathbf{b}}\|_{\infty}$.

A.6 Proof of Corollary 3.5

Let $N_m(t)$ denote the number of type- m matches completed by time t ; we have $D_{jm}(t) = M_{jm}N_m(t)$. Therefore, the following holds:

$$\mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{D_{jm}(t)}{D_j(t)} \right] = \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{D_{jm}(t)}{t} \frac{t}{D_j(t)} \right] = \frac{M_{jm}}{N\lambda_j} \cdot \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{N_m(t)}{t} \right] = \frac{M_{jm}}{N\lambda_j} \cdot \lim_{t \rightarrow \infty} \mathbb{E} \left[\frac{N_m(t)}{t} \right],$$

where the second equality follows from $D_{jm}(t) = M_{jm}N_m(t)$ and the fact that $\lim_{t \rightarrow \infty} D_j(t)/t = N\lambda_j$ almost surely by the strong law for renewal processes (Proposition 6.2 in Ross and Peköz 2023), and the third equality follows from the generalized dominated convergence theorem (Theorem 19 in Royden and Fitzpatrick 2010) because $N_m(t) \leq D_{j'}(t)$ and $\lim_{t \rightarrow \infty} \mathbb{E} \left[\frac{D_{j'}(t)}{t} \right] = \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{D_{j'}(t)}{t} \right]$, where j' represents a job type included in match m .

This equality, together with Lemma 3.4, yields Corollary 3.5 by setting $C_6 \triangleq C_4 \cdot \max_{j \in [J], m \in [K]} \frac{M_{jm}}{\lambda_j}$.

A.7 Proof of Proposition 4.3

Let x_m denote the long-run average rate at which agent i performs match m locally, and s_j denote the rate of submitting type- j jobs to the shared pool; we have $\mathbf{x} = (x_m)$ and $\mathbf{s} = (s_j)$ are feasible to (3). We will show that if $x_m > 0$ for some type- m match—that is, the agent performs some matches locally—the long-run average payoff is strictly less than $\boldsymbol{\lambda}_i^T \bar{\mathbf{p}}$ and is therefore suboptimal by Proposition 4.2.

Assume $x_m > 0$ for some $m \in [K]$. Since the jobs' interarrival and sojourn time follow exponential distributions, by the memoryless property, it is optimal for the agent to perform match m only when its component jobs arrive. We apply the approach of Aouad and Saritaç (2022) to bound the expiration rate of some over-demanded job type involved in match m from below. Specifically, since

a match requires two or more jobs and, at any time, at most one job arrives due to Poisson arrival processes, there exist two component job types j and j' for match m , such that the followings hold: (i) match m requires at least one type- j job and one type- j' job, (ii) with a certain probability, match m is performed when a job of type j arrives (referred to as a passive job), which implies a type- j' job is actively waiting for the match (referred to as an active job), and (iii) job type j' is over-demanded; that is, $\bar{p}_{j'} > 0$.⁸

Let $x_m^j > 0$ denote the rate of performing match m when a type- j job arrives, and let $y_{j'}$ denote the expiration rate of type- j' jobs. Based on the proof of Lemma 1 from Aouad and Saritaç (2022), the following inequality holds:

$$y_{j'} \geq \frac{\theta_{j'}}{\lambda_j} \cdot x_m^j.$$

Intuitively, an (active) type- j' job expires before the next type- j job arrives with probability $\frac{\theta_{j'}}{\lambda_j + \theta_{j'}}$. Therefore, type- j' jobs expire and are wasted at a positive rate. Consequently, the long-run average payoff for agent i is no greater than $\lambda_i^T \bar{\mathbf{p}} - y_{j'} \bar{p}_{j'}$, which is less than the payoff $\lambda_i^T \bar{\mathbf{p}}$ achieved by submitting all jobs to the shared pool.

A.8 Proof of Lemma 4.5

Suppose all agents fully submit their jobs to the shared pool. Without loss of generality, assume that neither the shared pool nor any agent holds any job at time zero. Let $N_m(t)$ denote the number of type- m matches formed by the shared pool by time t , $x_m^P \triangleq \lim_{t \rightarrow \infty} \mathbb{E}[N_m(t)]/t$ the long-run average matching rate of match m at the shared pool, and $S_{ijm}(t)$ the number of type- j jobs submitted by agent i that participate in match m by time t . We then have the following equation:

$$\mathbb{E}[S_{ijm}(t)] = \frac{\lambda_{ij}}{N\lambda_j} \cdot M_{jm} \cdot \mathbb{E}[N_m(t)], \quad (22)$$

where $\lambda_j = \sum_{i \in [N]} \lambda_{ij}/N$ denote the normalized aggregate arrival rate of type- j jobs. To interpret (22), note that a total of $M_{jm}N_m(t)$ type j jobs participate in match m by time t , and each of these jobs comes from agent i with a probability of $\frac{\lambda_{ij}}{N\lambda_j}$.

Let $\bar{x} = (\bar{x}_m)_{m \in [K]} \in \mathbb{R}_+^K$ be an optimal solution to (1). The long-run average payoff for agent i , denoted by V_i , satisfies the following:

$$\begin{aligned} V_i &= \lim_{t \rightarrow \infty} \mathbb{E} \left[\frac{\sum_{j \in [J]} \bar{p}_j \cdot \sum_{m \in [K]} S_{ijm}(t)}{t} \right] \\ &= \sum_{j \in [J]} \bar{p}_j \cdot \frac{\lambda_{ij}}{\lambda_j} \sum_{m \in [K]} M_{jm} \cdot \frac{x_m^P}{N} \\ &\geq \sum_{j \in [J]} \bar{p}_j \cdot \frac{\lambda_{ij}}{\lambda_j} \sum_{m \in [K]} M_{jm} \cdot \left(\bar{x}_m - C_4 \cdot \frac{NV^F - V^P}{N} \right) \\ &\geq \sum_{j \in [J]} \bar{p}_j \cdot \frac{\lambda_{ij}}{\lambda_j} \sum_{m \in [K]} M_{jm} \cdot \bar{x}_m - C_6 \cdot \frac{NV^F - V^P}{N} \\ &= \lambda_i^T \bar{\mathbf{p}} - C_6 \cdot \frac{NV^F - V^P}{N}. \end{aligned}$$

⁸We allows for $j = j'$, in which case match m requires at least two type- j jobs.

In the above, the first line follows from the fact that every matched type- j job yields an expected value of \bar{p}_j . The second line follows from (22) and the definition of x_m^P . The third line follows from Lemma 3.4. The fourth line follows from the definition of the constant C_6 , Lemma A.1, and Assumption 2.1. Finally, the fifth line follows from the complimentary slackness condition of (1), that is, $\bar{p}_j \cdot (\mathbf{m}_j^T \bar{x} - \lambda_j) = 0$ for all $j \in [J]$, where \mathbf{m}_j^T denotes the j -th row of matrix \mathbf{M} .