Self-Enforced Job Matching*

Ce Liu† Ziwei Wang‡ Hanzhe Zhang§

December 28, 2023

Abstract

The classic two-sided many-to-one job matching model with transfers assumes gross substitutability (firms treat workers as substitutes) and no peer effects (workers ignore colleagues when choosing where to work) to guarantee the existence of stable matching in static settings. Though complementarities and peer effects are common in practice, incorporating these features may lead to nonexistence of static stable matchings. However, matching is often not a static allocation, but an ongoing process with long-lived firms and short-lived workers. We show that stability is always guaranteed dynamically when firms are sufficiently patient, even with complementarities in firm technologies and peer effects in worker preferences. One way to maintain dynamic stability is through no-poaching agreements, which may explain their prevalence in labor markets.
Keywords: matching, repeated games, no-poaching agreements
JEL: C73

*We thank Nageeb Ali, Mariagiovanna Baccara, Vincent Crawford, Henrique de Oliviera, Laura Doval, Jon Eguia, Ed Green, Scott Kominers, Maciej Kotowski, SangMok Lee, Xiao Lin, Elliot Lipnowski, Qingmin Liu, Arijit Mukherjee, Marzena Rostek, Joel Sobel, Quan Wen, Kun Zhang, and audiences at various seminars for their helpful comments. Zhang acknowledges the National Science Foundation.
†Department of Economics, Michigan State University. E-mail: celiu@msu.edu.
‡School of Economics and Management, Wuhan University. E-mail: zwang.econ@gmail.com.
§Department of Economics, Michigan State University. E-mail: hanzhe@msu.edu.

1 Introduction

Two-sided many-to-one matching models with monetary transfers are crucial for understanding a variety of transactional environments, including auctions, labor markets, and housing markets. In their seminal work, Kelso and Crawford (1982) established the existence of stable matchings, provided that firms' production technologies satisfy gross substitutability (no workers serve as complementary inputs in a firm's production) and workers' preferences do not exhibit peer effects (workers ignore their colleagues' identities and characteristics).

However, complementarities and peer effects are present in many important matching markets. In industries such as manufacturing, healthcare, and software development, it is common for tasks to require the collaboration of workers with complementary skills. In labor markets such as academia, startups, and construction engineering, colleagues are an important consideration when choosing where to work. Although recent literature has obtained positive existence results by modeling large markets or considering alternative assumptions on preferences, accommodating arbitrary production technologies and peer effects has been challenging.

In this paper, we propose an approach that does not rely on restrictions on market size, technologies, or preferences. We observe that matching is often an ongoing process in which long-lived players on one side of the market interact with short-lived players or objects on the other side over time. This dynamic is salient in many matching environments: Long-lived sellers can serve a sequence of short-lived buyers, firms can hire a new batch of interns in every recruitment season, and brokers can purchase newly issued securities in every trading cycle. Given this observation, we consider stable matching as an unfolding process rather than a static allocation. We find that in the presence of complementarities and peer effects that destabilize static matchings, no-poaching agreements among firms can maintain stability in the dynamic setting, even if these agreements are not legally binding. The key idea is that dynamic incentives can act as both carrots and sticks to deter firms' deviations. We provide an illustrative example at the end of this section.

From a technical standpoint, our setting differs from standard repeated games since our stage game is a cooperative game. In standard repeated games, the noncooperative stage game is guaranteed to have a Nash equilibrium; playing a Nash equilibrium at every history delivers a subgame perfect Nash equilibrium. In our setting, by contrast, the cooperative stage game may not have a stable matching, so the existence of a stable matching process cannot be guaranteed simply by forming a static stable matching at every history. Instead, we prove the claim by explicitly constructing a dynamic matching process. We do this in three steps. First, we define and characterize the analog of "minmax payoffs" in the cooperative stage game. We then use a folk-theorem like construction to show that payoff profiles above these minmax payoffs can be sustained in a dynamically stable matching process when firms are sufficiently patient. Finally, we show that in a no-poaching agreement implemented through a random serial dictatorship, each firm obtains random payoffs that first-order stochastically dominate their respective minmax payoffs.

No-poaching agreements appear in many labor markets (Krueger and Ashenfelter 2022). They have attracted intense scrutiny in antitrust regulations, although these agreements are often informal and nonbinding.[1] A case in point is the ongoing civil antitrust litigation, *Henry, et al. v. Brown University, et al.*: Despite antitrust exemptions granted to universities under the Improving America's Schools Act of 1994, this case scrutinizes the use of a consensus approach by universities in allocating financial aid to students. Our finding adds a new perspective to this debate: While no-poaching agreements are inherently anti-competitive, they may serve a vital role in sustaining market stability, especially in environments lacking static stable matchings due to significant complementarities and peer effects.

An Example. Two long-lived firms f_1 and f_2 each offer two internship positions every period. Each firm treats workers as complements: It generates a revenue of \$6 only when both of its vacancies are filled and is unproductive otherwise. On the other side of the market, every period, three new identical workers w_1, w_2, and w_3 look for positions. For simplicity, assume that workers' payoffs are equal to the wages they receive. If the matching market in every period is treated as an isolated one-shot interaction, there is no static stable matching: In any static matching, at most one firm can be productive, and for both firms to break even, there must be one worker earning zero and another worker earning no more than \$3. The

[1]In *US v. Adobe Systems Inc., et al.*, these agreements are reached via emails among CEOs and HR officers. In *Seaman v. Duke University* and *Binotti v. Duke University*, the debate revolved around hiring guidelines that have been mutually agreed upon by university administrators.

unproductive firm can then form a blocking coalition with these two workers and split the gain.

Instead of static matchings, consider the following history-dependent matching process (illustrated in Figure 1). In each period, the market is in one of four possible states: Two wage-suppressing collusion states C1 and C2 and two punishment states P1 and P2. Each collusion state has a dominant firm, f_1 in C1 and f_2 in C2, and each punishment state has a punished firm, f_1 in P1 and f_2 in P2. The market starts at and remains in state C1 until a firm deviates. Whenever a firm deviates, the market transitions to the punishment state for that firm and remains there for four periods; if a firm deviates during a punishment phase, then the process transitions to the punishment phase for that firm. After this punishment phase, the market enters the collusion state in which the non-deviating firm assumes the role of the dominant firm and stays there until a firm deviates.

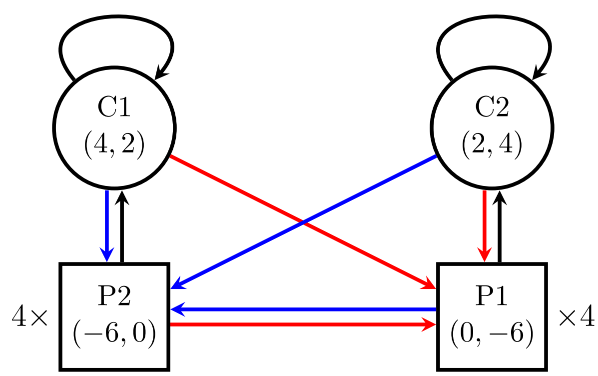

FIGURE 1. A stable matching process. The first number in parentheses is the per-period payoff of f_1, while the second number is the payoff of f_2. Black arrows represent transitions when no firm deviates, red arrows represent transitions after f_1's deviations, and blue arrows represent transitions after f_2's deviations.

At the beginning of every period in collusion states C1 and C2, a biased coin toss determines which firm secures the hiring rights for the current period. The dominant firm wins with a probability of 2/3, and the nondominant firm wins with a probability of 1/3. The firm that wins the coin toss hires two workers at zero wage, while the losing firm temporarily shuts down. Therefore, the dominant firm obtains \$4 on average, while the other obtains \$2. Notice that the collusion states feature a no-poaching agreement: In each period, the non-hiring firm refrains from soliciting workers from the other firm, thus forgoing an immediate gain in the current period. In punishment states P1 and P2, the firm being punished shuts down, while the punishing firm hires two workers each at a wage of $6. Figure 1 depicts firms' payoffs in each state.

To see that this matching process is dynamically stable, we use a one-shot deviation principle (Lemma 3), which establishes dynamic stability by checking two requirements at every history: (1) no worker wishes to unilaterally leave her matched firm to be unemployed and (2) no firm has a profitable one-shot deviation with a group of workers who also find this deviation profitable. Note that the first requirement is satisfied in every state of the matching process, since all workers are weakly better off than being unemployed. We next verify that the second requirement is also met in every state.

In punishment states, the punished firm cannot find a profitable deviation in the stage game. The punishing firm can gain a profit in the current period by deviating, but by doing so it would forgo the advantage of being the dominant firm in future collusive periods. As a result, for high values of δ, the punishing firm does not find one-shot deviations profitable.

In collusion states, regardless of which firm has the hiring rights for the period, the dominant firm finds no profitable one-shot deviations when δ is high: Any deviation would lead to the loss of its position as the dominant firm and thus a lower long-run payoff. The nondominant firm also has no incentive to deviate: Even though it could potentially obtain a current-period gain of $6 by poaching workers from the winning firm, this would be followed by a net loss of $2 for each of the next four periods (it gets $0 from shutting down, compared with $2 in the collusive state). When firms are patient, the loss outweighs the potential gain, so the nondominant firm also has no profitable one-shot deviations.

Some discussion of the importance of various assumptions is in order. The coin tosses in the collusion states are public randomization devices that convexify players' payoffs. Public randomization simplifies our exposition but is not essential to our results, since we can also convexify payoffs using sequence of play. Our results do, however, require firms to have the ability to make transfers to workers, since otherwise firms would have less scope to punish each other. In Section 4.2 we provide an example of matching without transfers in which no dynamically stable matching process exists.

Related Literature. The literature on static job matching has guaranteed the existence of stable matching by (i) imposing restrictions on preferences (Hatfield and Milgrom 2005; Sun and Yang 2006; Echenique and Yenmez 2007; Hatfield and Kojima 2008; Pycia 2012; Hatfield, Kominers, Nichifor, Ostrovsky and Westkamp 2013; Rostek and Yoder 2020; Kojima, Sun and Yu 2020, 2023; Pycia and Yenmez 2023); (ii) considering large markets (Kojima, Pathak and Roth 2013; Ashlagi, Braverman and Hassidim 2014; Azevedo and Hatfield 2018; Che, Kim and Kojima 2019); or (iii) making minimal adjustments to quotas (Nguyen and Vohra 2018).[2] In this paper, we propose a different approach based on firms' dynamic incentives.

Our paper also contributes to the literature on dynamic matching (Corbae, Temzelides and Wright 2003; Damiano and Lam 2005; Du and Livne 2016; Kadam and Kotowski 2018a,b; Altınok 2020; Kotowski 2020; Kurino 2020; Doval 2022; Pęski 2022). The work most closely related to this paper is Liu (2023), who studies repeated matching markets *without* transfers, while imposing assumptions so that static stable matchings exist in the stage game. By contrast, this paper studies matching markets *with* transfers and shows that dynamic stability can be restored even when static stability fails. The example without transfers in Section 4 in which dynamically stable matchings do not exist demonstrates the necessity of transfers for our existence result. In addition, Ali and Liu (2020) more broadly explore when and how dynamic incentives deter coalitional deviations in a general repeated cooperative games framework; however, they focus on settings in which all players are long-lived.

The rest of the paper is organized as follows. Section 2 introduces the model. Section 3 describes the main results and their proofs. Section 4 provides extensions and discussions of crucial assumptions. Section 5 concludes. Appendix A collects omitted proofs.

2 Model

Players. At the beginning of every period $t = 0, 1, 2, \ldots$, a new generation of workers enters the market to match with a fixed set of long-lived firms. Let $\mathcal{F}$ denote the finite set of firms, where $|\mathcal{F}| \geq 2$. Matching is many-to-one: In every period, each firm $f \in \mathcal{F}$ has $q_f > 0$ hiring slots to fill. Workers are short-lived and remain in the market for only one

[2] A related paper by Echenique and Yenmez (2007) considers matching with peer effects and proposes an algorithm that finds all static stable matchings whenever they exist.

period. For expositional convenience, we use the same notation $\mathcal{W}$ to denote the finite set of workers in every period; however, each worker $w \in \mathcal{W}$ can have a different *type* in every period.[3] A finite set Θ_w contains all possible types θ_w of worker w, and $\Theta = \times_{w\in\mathcal{W}} \Theta_w$ is the set of type profiles of all workers. Workers' type profile $\theta \in \Theta$ is randomly drawn from a distribution $\pi \in \Delta(\Theta)$ for each generation. It is important to note that while we can accommodate random θ, we do not rely on this randomness for our results. In fact, all results hold in the special case when π is a degenerate distribution.

Each firm f has a stage-game revenue function $\widetilde{u}_f : 2^{\mathcal{W}} \times \Theta \to \mathbb{R}$ defined over subsets of workers and type profiles.[4] We normalize the revenue of staying unmatched, $\widetilde{u}_f(\emptyset, \theta)$, to 0 for every $\theta \in \Theta$. Note that for each type profile, we do not require firms' revenue functions to satisfy the gross substitutes condition. Firms share a common discount factor δ and evaluate a sequence of flow utilities through exponential discounting.

Each worker cares about both her employer and her colleagues, which we will refer to collectively as her *work environment*. Let $\Phi_w = (\mathcal{F} \times 2^{\mathcal{W}\setminus\{w\}}) \cup \{(\emptyset, \emptyset)\}$ denote the set of possible work environments of worker w, where $(\emptyset, \emptyset)$ represents staying unmatched. Each worker w has utility function $\widetilde{v}_w : \Phi_w \times \Theta \to \mathbb{R}$ over work environments. For every worker $w \in \mathcal{W}$ and type profile $\theta \in \Theta$, $\widetilde{v}_w(\emptyset, \emptyset, \theta)$ is normalized to 0.

Stage Matching. The outcome in every period is a static many-to-one matching among the firms and workers. Formally, a stage matching $m = (\phi, p)$ is described by an *assignment* ϕ and a *wage vector* p. In particular, ϕ is a mapping defined on the set $\mathcal{F} \cup \mathcal{W}$ such that (i) for every $w \in \mathcal{W}$, $\phi(w) \in \Phi_w$; (ii) for every $f \in \mathcal{F}$, $\phi(f) \subseteq \mathcal{W}$ and $|\phi(f)| \leq q_f$; and (iii) for every $w \in \mathcal{W}$ and every $f \in \mathcal{F}$, $w \in \phi(f)$ if and only if $\phi(w) = (f, W')$ for some $W' \in 2^{\mathcal{W}\setminus\{w\}}$. The wage vector $p = (p_{fw}) \in \mathbb{R}^{|\mathcal{F}|\times|\mathcal{W}|}$ describes the transfer from firms to workers. We assume that any nonzero transfer occurs only between a firm and its own employees: $p_{fw} = 0$ for every $w \notin \phi(f)$.

[3]Our model can accommodate fluctuations in the number of workers across periods by introducing undesirable types, so that matching with an undesirable worker renders all of her partners strictly worse off (before transfers) than staying unmatched.

[4]While it may be reasonable to assume that a firm's revenue depends solely on the types of workers it employs, none of our results require this assumption. Hence, for notational simplicity, we allow revenue to depend on the entire profile of worker types. The same applies to workers' payoff functions.

Players have quasilinear utilities: For every stage matching $m = (\phi, p)$ and type profile θ,

$$u_f(m, \theta) \equiv \widetilde{u}_f(\phi(f), \theta) - \sum_{w' \in \mathcal{W}} p_{fw'} \quad \text{and} \quad v_w(m, \theta) \equiv \widetilde{v}_w(\phi(w), \theta) + \sum_{f' \in \mathcal{F}} p_{f'w}$$

are the stage-game payoffs received by firm f and worker w, respectively.[5]

A stable stage matching is immune to three kinds of deviation: (i) a unilateral deviation by a firm f who fires a subset of its employees and leaves those positions unfilled; (ii) a unilateral deviation by a worker w who leaves her employer and remains unmatched; and (iii) a coalitional deviation $(f, W', p'_f) \in \mathcal{F} \times 2^{\mathcal{W}} \times \mathbb{R}^{|\mathcal{W}|}$ with $|W'| \leq q_f$ and $p'_{fw} = 0$ for every $w \notin W'$, where firm f and workers W' match at wages specified in p'_f and abandon any other pre-existing match partners. A deviation is profitable if all participants strictly prefer the deviating matching to the original matching m. A stage matching is *individually rational* if no unilateral deviations are profitable, and is *stable* if none of the three kinds of deviations above is profitable.

When studying the stability of static matchings, there is no need to specify how other players will be matched after a deviation by a coalition. However, in our dynamic setting, players' future behavior is influenced by past histories, so to study the stability of matching processes, we need to specify the realized stage-game outcome after a deviation. To this end, we adopt the following assumption.

Assumption 1. *After coalitional deviation (f, W', p'_f) from matching $m = (\phi, p)$, let matching $m' = [m, (f, W', p'_f)] \in M$ denote the resulting stage matching and let ϕ' denote the assignment in m'. Assume $\phi'(f) = W'$ and $\phi'(f') = \phi(f') \backslash W'$ for every $f' \neq f$.*

Assumption 1 says that in the stage matching that results from a coalitional deviation, the deviators are matched together, players abandoned by the deviators remain unmatched, and those untouched by the deviation remain matched with their original partners. The sole purpose of this assumption is to ensure "perfect monitoring": When a matching m is blocked by a coalition (f, W', p'_f), the firm in the deviating coalition is identifiable. Lemma 4 formally proves this identifiability property from Assumption 1. Any alternative assumption

[5]Our results do not hinge on quasilinearity and hold as long as utilities are monotone and unbounded in wages.

that delivers the same identifiability property will not change our results.[6]

Repeated Matching. To convexify players' stage-game payoffs, we employ a public randomization device on the unit interval $S = [0, 1]$ endowed with the Lebesgue measure. The use of public randomization streamlines our proofs, but our results do not depend on it.[7]

The timing in each period is as follows. First, a new cohort of workers arrives. A type profile $\theta \in \Theta$ is drawn from the distribution π, and a public randomization $s \in S$ is realized. Based on the realized (θ, s), a stage matching is recommended for active players in the market. Players then decide whether to deviate from this recommendation, which determines the outcome of the stage game.

A t-period ex ante history $\overline{h} = (\theta_\tau, s_\tau, m_\tau)_{\tau=0}^{t-1}$ specifies a sequence of past realizations of the type profile, the public correlation device, and the stage matching up to period $t-1$.[8] We write $\overline{\mathcal{H}}_t$ for the set of all t-period ex ante histories, with $\overline{\mathcal{H}}_0 = \{\emptyset\}$ being the singleton set comprising the null history. Let $\overline{\mathcal{H}} \equiv \bigcup_{t=0}^{\infty} \overline{\mathcal{H}}_t$ be the set of all ex ante histories. Moreover, let $\mathcal{H}_t \equiv \overline{\mathcal{H}}_t \times \Theta \times S$ denote the set of t-period ex post histories, and $\mathcal{H} \equiv \overline{\mathcal{H}} \times \Theta \times S$ the set of all ex post histories.

A *matching process* $\mu : \mathcal{H} \to M$ specifies a stage matching for every ex post history. It represents a shared understanding among players regarding how past histories impact future employment. For every ex post history $h \in \mathcal{H}$, we let $\mu(f|h)$ and $\mu(w|h)$ denote the matching partners for firm f and worker w in the stage matching $\mu(h)$, respectively.

Let $\overline{\mathcal{H}}_\infty = (\Theta \times S \times M)^\infty$ be the set of outcomes h_∞ of the repeated cooperative matching game. Let $m_t(h_\infty)$ denote the stage matching in the t-th period of h_∞. Following every t-period (ex ante or ex post) history $\widehat{h} \in \overline{\mathcal{H}} \cup \mathcal{H}$, let

$$U_f(\widehat{h} \,|\, \mu) \equiv (1-\delta)\mathbb{E}_\mu\Big[\sum_{\tau=t}^{\infty} \delta^{\tau-t} u_f(m_\tau(h_\infty), \theta_\tau) \,\Big|\, \widehat{h}\,\Big]$$

[6] A similar assumption that serves the same purpose is seen in Roth and Sotomayor (1992), Mauleon, Vannetelbosch and Vergote (2011), and Liu (2023).

[7] For example, we could follow arguments in Sorin (1986) and Fudenberg and Maskin (1991) to convexify players' payoffs using sequences of play instead.

[8] Note that this assumes observable wages. However, since workers are short-lived and none of the self-enforcing matching processes constructed in this paper make use of information about wages, the existence result continues to hold even if wages are unobservable.

denote the continuation payoff firm f obtains from μ following $\widehat{h}$, where the expectation is taken with respect to the measure over $\overline{\mathcal{H}}_\infty$ induced by μ conditional on $\widehat{h}$.

Deviation Plan. We make two observations that allow us to tractably analyze firms' deviations in the repeated cooperative matching game. Recall that there are three kinds of deviations within a period. We first observe that a unilateral deviation by firm f to fire a subset of its employees is equivalent to a deviation by the coalition consisting of f and the workers that remain with f. It is therefore without loss to focus on two kinds of deviation: (i) unilateral deviations by workers and (ii) deviations by coalitions consisting of a firm and a set of workers in the stage game.

Our second observation is that as a long-lived player, a firm can participate in a sequence of deviations by forming coalitions with workers across periods. Each of these coalitions must be immediately profitable for the participating short-lived workers but not necessarily for the firm, since the firm cares about the profit it collects from the entire sequence.

Motivated by this second observation, we define a *deviation plan* for firm f as a complete contingent plan that specifies, at every ex post history, a set of workers to recruit and their wage offers. Formally, a deviation plan for firm f is a pair $(d : \mathcal{H} \to 2^{\mathcal{W}}, \eta : \mathcal{H} \to \mathbb{R}^{|\mathcal{W}|})$ such that $|d(h)| \leq q_f$ for any h and $\eta_w(h) \neq 0$ only if $w \in d(h)$. Together with the original matching process, a deviation plan generates a distribution over the outcomes of the game $\overline{\mathcal{H}}_\infty$. Given a matching process μ and f's deviation plan (d, η), the *manipulated matching process*, denoted by $[\mu, (f, d, \eta)] : \mathcal{H} \to M$, is a matching process defined by

$$\Big[\mu, (f, d, \eta)\Big](h) \equiv \Big[\mu(h),\ \Big(f, d(h), \eta(h)\Big)\Big] \quad \forall h \in \mathcal{H}.$$

Firm f's deviation plan (d, η) from μ is *feasible* if at every ex post history $h = (\overline{h}, \theta, s)$,

$$v_w\Big(\Big[\mu, (f, d, \eta)\Big](h),\ \theta\Big) \geq v_w\Big(\mu(h), \theta\Big) \quad \forall w \in d(h),$$

and the inequality is strict if $w \notin \mu(f|h)$. That is, workers participate in the deviation only if they find the new stage matching weakly better than the recommendation from μ, and this preference must be strict if the worker is employed by a different firm under μ.

Lastly, the deviation plan (d, η) is *profitable* if there exists an ex post history h such that $U_f\big(h \,\big|\, [\mu, (f, d, \eta)]\big) > U_f(h|\mu)$.

Self-Enforcing Matching Process. The two observations above motivate our notion of dynamic stability.

Definition 1. Matching process μ is *self-enforcing* if (i) $v_w(\mu(h), \theta) \geq 0$ for every $w \in \mathcal{W}$ at every ex post history $h \in \mathcal{H}$ and (ii) no firm has a strictly profitable feasible deviation plan.

The first requirement guards against deviations by a single worker in any generation, and the second guards against deviations by firms. These requirements are imposed on the matching process at *all* ex post histories, including those that are off path. This restriction on all ex post histories embeds a form of sequential rationality in the same way as subgame perfection does in a repeated noncooperative game. It is also worth noting that Definition 1 focuses on coalitions of a single firm and multiple workers: This is stronger than pairwise stability but weaker than group stability. In dynamic environments, these stability notions are not equivalent.[9] Nevertheless, in static settings, Definition 1 reduces to the definition of core allocation in Kelso and Crawford (1982), so our definition can be viewed as its dynamic generalization.

3 Results

Minmax Payoffs. We begin by defining firms' minmax payoffs. For every realized type profile $\theta \in \Theta$, let

$$M^\circ(\theta) \equiv \{m \in M : v_w(m, \theta) \geq 0 \text{ for all } w \in \mathcal{W}\}$$

denote the set of stage matchings that are individually rational for workers. From Definition 1, a self-enforcing matching process can only recommend stage matchings in $M^\circ(\theta)$. Moreover,

[9]Allowing infinite-horizon deviations by multiple long-lived players creates a conceptual difficulty in assessing whether those deviations themselves can be self-enforcing. Ali and Liu (2020) show that if coalitions cannot commit to long-run behavior and are unable to make anonymous transfers, then modeling coalitions with multiple long-run players does not alter the set of sustainable outcomes when players are patient.

for every firm $f \in \mathcal{F}$ and recommended stage matching $m = (\phi, p)$, let

$$D_f(m,\theta) \equiv \left\{ (W', p'_f) : \begin{array}{r} |W'| \leq q_f,\ p'_{fw} = 0 \text{ if } w \notin W', \\ \widetilde{v}_w(f, W', \theta) + p'_{fw} \geq v_w(m,\theta) \text{ for every } w \in W', \text{ and} \\ \widetilde{v}_w(f, W', \theta) + p'_{fw} > v_w(m,\theta) \text{ for every } w \in W' \backslash \phi(f) \end{array} \right\}$$

denote the set of feasible stage-game deviations for f at type profile θ. Similar to the definition of feasible deviation plans, a deviation in the stage matching is feasible if (i) the deviating firm does not exceed its quota in hiring, (ii) all participating workers are weakly better off compared to the recommended stage matching, and (iii) any worker recruited elsewhere should be strictly incentivized compared to the recommended stage matching. Firm f's minmax payoff is its payoff from "best responding" to the worst recommendation.

Definition 2. Firm f's minmax payoff at type profile θ is

$$\underline{u}_f(\theta) \equiv \inf_{m \in M^\circ(\theta)} \sup_{(W', p'_f) \in D_f(m,\theta)} u_f([m, (f, W', p'_f)], \theta).$$

To characterize this minmax payoff, let

$$\chi(f, W, \theta) \equiv \widetilde{u}_f(W, \theta) + \sum_{w \in W} \widetilde{v}_w(f, W \backslash \{w\}, \theta)$$

denote the total surplus of coalition (f, W) at type profile θ. The following lemma characterizes each firm's minmax payoff.

Lemma 1. *Let $Q \equiv \sum_{f' \in \mathcal{F}} q_{f'}$ represent the sum of all firms' hiring quotas. For every firm f and type profile θ, f's minmax payoff satisfies*

$$\underline{u}_f(\theta) = \min_{W' \subseteq \mathcal{W}, |W'| \leq Q} \ \max_{W \subseteq \mathcal{W} \backslash W', |W| \leq q_f} \chi(f, W, \theta). \tag{1}$$

Lemma 1 states that a firm's minmax payoff equals the maximum surplus it can generate after $\sum_{f'} q_{f'}$ workers have been removed from the market in an adversarial manner. The intuition is as follows. To punish f, other firms can offer high wages to $\sum_{f' \neq f} q_{f'}$ workers, which renders them unattractive as potential partners for f. Also, if the employees of f are paid high wages by f and require consent to join a deviation, then f would be better off abandoning them and looking for cheaper workers for its deviation. Altogether, this excludes $\sum_{f'} q_{f'}$ workers from f's potential match pool in case of a deviation. Instead, f will select its partners from the remaining workers and secure all the matching surplus by offering those workers zero payoff.

Characterization. We first introduce some notation. For every $\theta \in \Theta$ and $m \in M^\circ(\theta)$, let $u(m,\theta) \equiv (u_f(m,\theta))_{f\in\mathcal{F}}$ denote firms' payoff profile under m. For every $\theta \in \Theta$, let

$$\mathcal{U}(\theta) \equiv \text{co}\left(\left\{u \in \mathbb{R}^{\mathcal{F}} : u = u(m,\theta) \text{ for some } m \in M^\circ(\theta)\right\}\right)$$

denote the convex hull of these payoff profiles. In addition, define

$$\mathcal{U}^* \equiv \left\{\sum_{\theta\in\Theta} \pi(\theta)u(\theta) : u(\theta) \in \mathcal{U}(\theta) \text{ for every } \theta \in \Theta\right\}.$$

Finally, for every firm f, let

$$\underline{u}_f^* \equiv \mathbb{E}_\pi[\underline{u}_f(\theta)]$$

denote its expected minmax payoff over type profiles.

The next result characterizes the discounted payoffs from self-enforcing matching processes.

Proposition 1. *(i) If $u \in \mathcal{U}^*$ satisfies $u_f > \underline{u}_f^*$ for all $f \in \mathcal{F}$, then there is a $\underline{\delta} \in (0,1)$ such that for every $\delta \in (\underline{\delta},1)$, there exists a self-enforcing matching process with firms' discounted payoffs u. (ii) Suppose μ is a self-enforcing matching process for a given $\delta \in (0,1)$. For every ex ante history $\overline{h} \in \overline{\mathcal{H}}$, firms' discounted payoff profile satisfies $\left(U_f(\overline{h} \mid \mu)\right)_{f\in\mathcal{F}} \in \mathcal{U}^*$ and $U_f(\overline{h} \mid \mu) \geq \underline{u}_f^*$ for every $f \in \mathcal{F}$.*

To prove statement (i), we first show that $\mathcal{U}^*$ satisfies the NEU condition (Abreu, Dutta and Smith 1994). For every payoff profile in $\mathcal{U}^*$ that is strictly above every firm's average minmax payoff, this allows us to construct a set of payoff profiles, each individually tailored for a specific firm, which can be used as carrots and sticks to discipline the firms. Finally, these payoff profiles are integrated into the classic idea of Fudenberg and Maskin (1986) to construct a matching process that is self-enforcing when firms are patient.

The cooperative nature of the stage game gives rise to a technical difficulty. Since stage matchings lack the usual product structure of strategy profiles in normal-form games, when a firm f is being minmaxed, it might be necessary for f to obtain stage payoffs that are even lower than its minmax value. When $\pi \in \Delta(\Theta)$ is not a degenerate distribution, this may create incentives for f to deviate at some ex post histories when it is being minmaxed, even though it is not profitable to do so on average. We tackle this by adjusting wages in the punishment scheme so that the ex post benefits of deviations are equalized across realizations of θ (see Lemma 5(iii)).

To prove statement (ii), note that given realized type profile θ, by the definition of minmax payoffs, every firm f can secure the minmax payoff $\underline{u}_f(\theta)$ by deviating with workers. Taking expectation over the distribution of type profiles delivers the result.

Proposition 1, however, does not guarantee the existence of self-enforcing matching processes. In fact, statement (i) would be vacuously true if there is no feasible payoff profile $u \in \mathcal{U}^*$ that satisfies $u_f > \underline{u}_f^*$ for all $f \in \mathcal{F}$. Note that this is not a concern for noncooperative games, but in two-sided matching environments such scenarios can indeed arise, for example, when the market has no wages or there are matching externalities. We provide detailed illustrations of these possibilities in Section 4.

Proving the existence of a self-enforcing matching process therefore amounts to showing that the existence of $u \in \mathcal{U}^*$ that satisfies $u_f > \underline{u}_f^*$ for all $f \in \mathcal{F}$. We will show that such a payoff profile can arise from a random serial dictatorship among the firms.

No-Poaching Agreement. Let $\Gamma \equiv \{\gamma : \mathcal{F} \to \{1, \ldots, |\mathcal{F}|\}\}$ denote the set of orderings over firms. In the serial dictatorship corresponding to an ordering $\gamma \in \Gamma$, firms sequentially choose workers based on γ while setting all of their matched workers' payoffs to zero, thereby capturing the entire matching surplus.

Definition 3 (Serial Dictatorship). Given $\theta \in \Theta$ and $\gamma \in \Gamma$, stage matching $\widehat{m}(\theta, \gamma)$ is induced by the following procedure. Initialize $W_0^\# \equiv \emptyset$. For every step $i = 1, \ldots, |\mathcal{F}|$, denote $\widehat{f}_i \equiv \gamma^{-1}(i)$, and let

$$\widehat{W}_i \in \underset{W \subseteq \mathcal{W} \setminus W_{i-1}^\#, |W| \leq q_{\widehat{f}_i}}{\arg\max} \chi(\widehat{f}_i, W, \theta);$$

Set $\phi(\widehat{f}_i) = \widehat{W}_i$, $p_{\widehat{f}_i w} = -\widetilde{v}_w(\widehat{f}_i, \widehat{W}_i \backslash \{w\})$ for every $w \in \widehat{W}_i$, and $p_{\widehat{f}_i w} = 0$ for every $w \notin \widehat{W}_i$; Update $W_i^{\#} \equiv W_{i-1}^{\#} \cup \widehat{W}_i$.

In a *random serial dictatorship (RSD)*, firms randomize over Γ and play matching $\widehat{m}(\theta, \gamma)$ based on the realized order γ. Note that RSD can be seen as a form of no-poaching agreement since a firm refrains from soliciting workers already employed by another firm even if it is feasible to do so.

We make an assumption to simplify our existence proof. Let

$$\overline{u}_f(\theta) \equiv \max_{W \subseteq \mathcal{W}, |W| \leq q_f} \chi(f, W, \theta) \tag{2}$$

denote firm f's maximum feasible payoff at type profile θ. It is easy to see that $\overline{u}_f(\theta) \geq \underline{u}_f(\theta)$ for every $\theta \in \Theta$ and $f \in \mathcal{F}$, since removing Q workers in an adversarial manner reduces f's maximum surplus. For now, we assume that for every firm, this inequality is strict with a positive probability.

Assumption 2. *For every firm f, there exists $\theta \in \Theta$ with $\pi(\theta) > 0$ such that $\overline{u}_f(\theta) > \underline{u}_f(\theta)$.*

Assumption 2 holds generically, for example, when players' payoffs given every type profile θ is randomly drawn from a continuous distribution. We make this assumption to simplify the exposition of Proposition 2. In Section 4 we show that our results hold even without this assumption.

The following lemma shows that when firms randomize uniformly over serial dictatorships, they can simultaneously obtain payoffs that strictly exceed their expected minmax payoffs. Figure 2 provides an illustration of this result when the market has only two firms.

Lemma 2. *Under Assumption 2, for every firm f,*

$$\frac{1}{|\Gamma|} \sum_{\gamma \in \Gamma} \mathbb{E}_\pi \Big[u_f\big(\widehat{m}(\theta, \gamma), \theta\big)\Big] > \underline{u}_f^*.$$

Let us explain the intuition for Lemma 2. If firms randomize uniformly over Γ for every realized type profile θ, then each firm f receives the payoff on the left side. Fix f and θ. In this RSD, the worst-case scenario for f arises when f is ranked last in γ (i.e., $\gamma(f) = |\mathcal{F}|$):

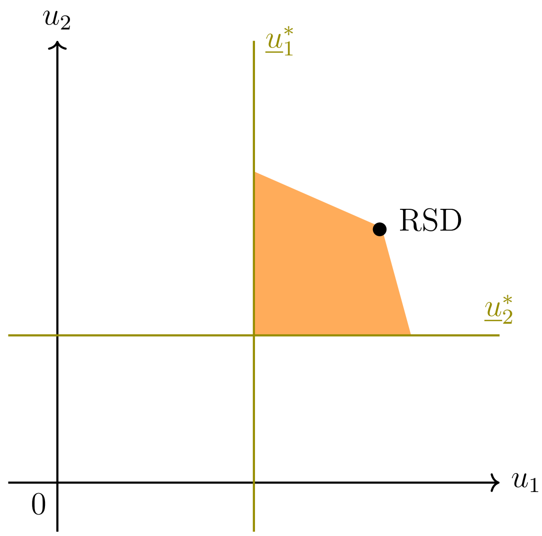

Figure 2. Firms' payoff space. Payoff profiles in the orange region (feasible and strictly above firms' expected minmax payoffs) can be sustained by Proposition 1. Lemma 2 ensures the orange region is nonempty with a random serial dictatorship (RSD).

By the time f selects its employees, $Q_{-f} \equiv \sum_{f' \neq f} q_{f'}$ workers are already off the market. However, for f, this worst case under RSD is still weakly more desirable than being minmaxed, in which case f must choose its employees after $Q = \sum_{f'} q_{f'}$ workers have been eliminated in an *adversarial* manner. Therefore, the distribution of payoffs f obtained under RSD first-order stochastically dominates the distribution of its minmax payoffs, so in expectation, every firm prefers RSD over being minmaxed.

Combined with Proposition 1, Lemma 2 delivers our existence result below.

Proposition 2. *Under Assumption 2, when firms are sufficiently patient, there exists a self-enforcing matching process in which firms play RSD in every period on path.*

This result shows that as a particular form of no-poaching agreement, RSD can always be sustained dynamically. To understand the intuition, consider the uniform RSD. For every θ and γ, matching $\widehat{m}(\theta, \gamma)$ is in $M^{\circ}(\theta)$, so the randomization over Γ places firms' payoff profile in $\mathcal{U}(\theta)$. Taking expectation over θ, we know that firms' payoff profile from the uniform RSD is in $\mathcal{U}^*$. Furthermore, according to Lemma 2, each firm's payoff is strictly higher than its average minmax payoff. Proposition 1(i) then delivers the existence result.

4 Extensions and Discussions

In this section, we explore the extensions and limitations of our framework. We first demonstrate that Assumption 2 is not crucial to our main results. That is, even if some firms can never be harmed by others and thus cannot be incentivized dynamically, we can still apply the analysis in Section 3 to the *remaining* firms to establish existence of self-enforcing matching processes. We then highlight the importance of the presence of transfers and absence of externalities in our existence result by providing examples of nonexistence of stable matching processes under alternative settings without transfers or with externalities.

4.1 Relaxing Assumption 2

We first explain how our results continue to hold when Assumption 2 is not satisfied: There exists some firm f such that for every $\theta \in \Theta$ with $\pi(\theta) > 0$, we have $\overline{u}_f(\theta) = \underline{u}_f(\theta)$. In this case, other players' matching decisions have *no impact* on the maximum static payoff firm f can derive, since it can always turn to the unmatched workers and extract their surpluses. This means that future punishments cannot affect the matching behavior of f via dynamic incentives, nor does f find it beneficial to participate in any punishment scheme of other firms. Therefore, we can treat such a firm as "inactive" in our analysis, assign it the maximum payoff it can receive in every period, and ignore it for the rest of our analysis. An iteration is needed to identify all such firms that cannot be incentivized dynamically.

Formally, for a set of firms $\mathcal{F}' \subseteq \mathcal{F}$, denote by $Q(\mathcal{F}') \equiv \sum_{f \in \mathcal{F}'} q_f$ the total hiring capacity of $\mathcal{F}'$. When all firms in $\mathcal{F} \backslash \mathcal{F}'$ are inactive, the effective minmax payoff of a firm $f \in \mathcal{F}'$ at θ is

$$\underline{u}_f(\mathcal{F}', \theta) \equiv \min_{W' \subseteq \mathcal{W}, |W'| \leq Q(\mathcal{F}')} \max_{W \subseteq \mathcal{W} \backslash W', |W| \leq q_f} \chi(f, W, \theta).$$

Using a similar argument as in Lemma 1, we can show that this is exactly firm f's payoff from "best responding" to the worst punishment by firms in $\mathcal{F}'$, while firms in $\mathcal{F} \backslash \mathcal{F}'$ leave no surplus to their employees. Note that for each θ, the value $\underline{u}_f(\mathcal{F}', \theta)$ weakly increases as $\mathcal{F}'$ becomes smaller.

Definition 4. A hierarchical partition $\mathcal{P} \equiv \{\mathcal{P}_1, \ldots, \mathcal{P}_N, \mathcal{R}\}$ over firms $\mathcal{F}$ is induced by the following procedure. Initialize $\mathcal{P}_0 \equiv \emptyset$. For $n \geq 1$:

- If $\{f \in \mathcal{F} \backslash \bigcup_{k=0}^{n-1} \mathcal{P}_k : \overline{u}_f(\theta) = \underline{u}_f(\mathcal{F} \backslash \bigcup_{k=0}^{n-1} \mathcal{P}_k, \theta)\ \forall \theta\} \neq \emptyset$, let this set be $\mathcal{P}_n$. Assign $n = n+1$ and continue;
- If $\{f \in \mathcal{F} \backslash \bigcup_{k=0}^{n-1} \mathcal{P}_k : \overline{u}_f(\theta) = \underline{u}_f(\mathcal{F} \backslash \bigcup_{k=0}^{n-1} \mathcal{P}_k, \theta)\ \forall \theta\} = \emptyset$, let $\mathcal{R} = \mathcal{F} \backslash \bigcup_{k=0}^{n-1} \mathcal{P}_k$ and stop.

Intuitively, each $\mathcal{P}_n$ consists of firms that cannot be punished in the matching process without cooperation from those in $\bigcup_{k=0}^{n-1} \mathcal{P}_k$. If $\mathcal{R} \neq \emptyset$, by construction, $\overline{u}_f(\theta) > \underline{u}_f(\mathcal{R}, \theta)$ for every $f \in \mathcal{R}$ and $\theta \in \Theta$. Let

$$\overline{u}_f^* \equiv \mathbb{E}_\pi[\overline{u}_f(\theta)] \quad \forall f \in \mathcal{F} \backslash \mathcal{R},$$

and

$$\underline{u}_f^*(\mathcal{R}) \equiv \mathbb{E}_\pi[\underline{u}_f(\mathcal{R}, \theta)] \quad \forall f \in \mathcal{R}.$$

A generalized version of Proposition 1 can be stated as follows.

Proposition 1*. (i) If $u \in \mathcal{U}^*$ satisfies $u_f = \overline{u}_f^*$ for all $f \in \mathcal{F} \backslash \mathcal{R}$ and $u_f > \underline{u}_f^*(\mathcal{R})$ for all $f \in \mathcal{R}$, then there is a $\underline{\delta} \in (0,1)$ such that for every $\delta \in (\underline{\delta}, 1)$, there exists a self-enforcing matching process with firms' discounted payoffs u. (ii) Suppose μ is a self-enforcing matching process for a given $\delta \in (0,1)$. For every ex ante history $\overline{h} \in \overline{\mathcal{H}}$, firms' discounted payoff profile satisfies $\left(U_f(\overline{h} \,|\, \mu)\right)_{f \in \mathcal{F}} \in \mathcal{U}^*$, $U_f(\overline{h} \,|\, \mu) = \overline{u}_f^*$ for every $f \in \mathcal{F} \backslash \mathcal{R}$, and $U_f(\overline{h} \,|\, \mu) \geq \underline{u}_f^*(\mathcal{R})$ for every $f \in \mathcal{R}$.

We next define a random serial dictatorship with respect to the hierarchical partition $\mathcal{P} = \{\mathcal{P}_1, \ldots, \mathcal{P}_N, \mathcal{R}\}$. To do so, we first introduce a subset of orderings $\Gamma_{\mathcal{P}} \subseteq \Gamma$ that contains those that (i) give firms in $\mathcal{R}$ the highest priorities and (ii) rank firms in $\mathcal{P}_n$ higher than those in $\mathcal{P}_k$ if $n > k$. That is,

$$\Gamma_{\mathcal{P}} \equiv \left\{ \gamma \in \Gamma : \begin{array}{c} \gamma(\mathcal{R}) = \{1, 2, \ldots, |\mathcal{R}|\}, \text{ and} \\ \text{if } f \in \mathcal{P}_k,\ f' \in \mathcal{P}_n, \text{ and } k < n, \text{ then } \gamma(f) > \gamma(f') \end{array} \right\}.$$

The stage matching $\widehat{m}(\theta, \gamma)$ induced by a serial dictatorship according to γ is defined as in Definition 3. By using a random serial dictatorship restricted to $\Gamma_{\mathcal{P}}$, the following lemma generalizes Lemma 2 and shows that Proposition 1*(i) is not vacuously true.

Lemma 2*. For every firm $f \in \mathcal{R}$,

$$\frac{1}{|\Gamma_{\mathcal{P}}|} \sum_{\gamma \in \Gamma_{\mathcal{P}}} \mathbb{E}_{\pi} \Big[u_f \big(\widehat{m}(\theta, \gamma), \theta \big) \Big] > \underline{u}^*_f(\mathcal{R}).$$

For every firm $f \in \mathcal{F} \backslash \mathcal{R}$,

$$\frac{1}{|\Gamma_{\mathcal{P}}|} \sum_{\gamma \in \Gamma_{\mathcal{P}}} \mathbb{E}_{\pi} \Big[u_f \big(\widehat{m}(\theta, \gamma), \theta \big) \Big] = \overline{u}^*_f.$$

In view of Lemma 2* and Proposition 1*, we can establish Proposition 2 without Assumption 2.

Proposition 2*. When firms are sufficiently patient, there exists a self-enforcing matching process in which firms in $\mathcal{R}$ play RSD in every period on path.

4.2 Matching without Transfers

The assumption of transferable utility plays a crucial role in our analysis. When punishing a firm f who deviates from a recommended stage matching, large transfers can be used to attract workers that f prefers the most. Therefore, we can simplify the minmax payoff of a firm and express it as the highest attainable payoff when workers are removed in an adversarial manner (see Lemma 1). In contrast, without the flexibility of transfers, firms encounter limitations in punishing deviations. The set of workers that can be lured away during punishment is intricately tied to the preferences of all workers. This dependence raises the minmax payoffs of the firms, which may lead to nonexistence of a self-enforcing matching process, that is, the payoff set characterized by Proposition 1 is empty.

Consider a model in which transfers are not allowed. That is, the wage vector is restricted to $p = 0$, and any deviation plan (d, η) should satisfy $\eta(h) = 0$ for any ex post history h. Suppose there are three workers $\mathcal{W} = \{w_1, w_2, w_3\}$ and two firms $\mathcal{F} = \{f_1, f_2\}$. Each firm has a quota of 2. A firm earns a stage payoff of 2 with two employees, and 1 with only one employee; it does not care about the identities of the employees. On the other hand, workers'

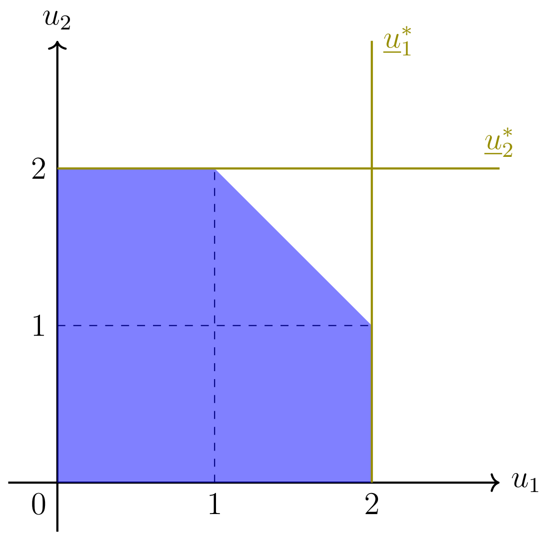

FIGURE 3. Matching without transfers. The blue region includes the set of firm payoffs that can be achieved via public randomization. The green lines represent the firms' minmax payoffs. There is no feasible payoff vector such that both firms are strictly better off than receiving the minmax payoffs.

preferences are determined by their colleagues: w_1 prefers to work with w_2, w_2 prefers w_3, and w_3 prefers w_1. Moreover, having any coworker is strictly better than working alone or being unemployed. Workers may possess strict preferences for firms, but this aspect is subordinate and is dominated by their preferences for coworkers.

In this example, each firm can guarantee a deviating payoff of 2 in the stage matching game by hiring an unemployed worker w together with the worker w' who prefers w. This is true regardless of the other firm's hiring decision; therefore, the firms' minmax payoffs should equal 2. However, there is no feasible way to generate a payoff of 2 for every firm in the stage matching even with a randomizing device. We can conclude that no self-enforcing matching process exists for any value of δ in this nontransferable utility setting. This nonexistence issue is illustrated by Figure 3.

4.3 Matching with Externalities

In our setting, we implicitly assume that firms are exclusively concerned with the identities and types of their own employees, while workers only care about their own work environments. However, in a more general context, agents' payoffs may be contingent on the entire matching

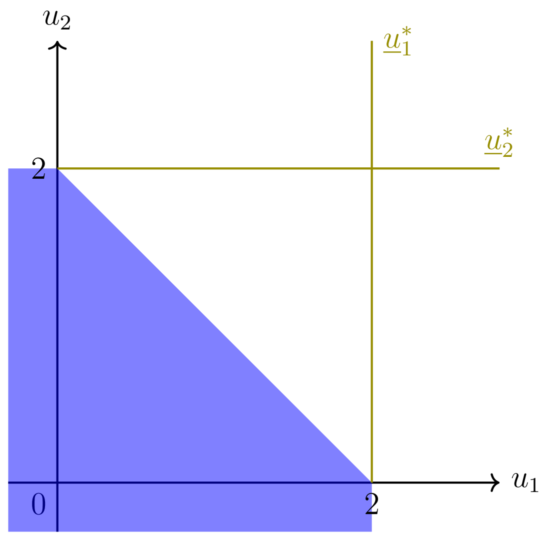

Figure 4. Matching with externalities. The blue region includes the set of firm payoffs that can be achieved via public randomization. The green lines represent the firms' minmax payoffs. There is no feasible payoff vector such that both firms are strictly better off than receiving the minmax payoffs.

assignment, including the hiring decisions of other firms. In this case, even with the availability of transfers, the characterization of minmax payoffs in Lemma 1 does not hold. We provide an example below in which the minmax payoffs are excessively high, and thus a self-enforcing matching process no longer exists.

Consider a matching market with transferable utilities in which each firm's stage payoff depends on the assignments of other firms. Suppose there are three workers $\mathcal{W} = \{w_1, w_2, w_3\}$ and two firms $\mathcal{F} = \{f_1, f_2\}$. Each firm can only hire one worker. The workers are indifferent between being unemployed and working for any firm at a payment of 0. Firm f_1 obtains a stage payoff of 2 if it hires the same number of workers as firm f_2 does and a stage payoff of 0 otherwise. On the other hand, firm f_2 gets 2 if the two firms hire different numbers of workers and 0 otherwise.

In this example, each firm can ensure a deviating payoff arbitrarily close to 2 regardless of the other firm's hiring decisions: For example, firm f_1 can always turn to the unemployed worker and offer an infinitesimal wage if f_2 has an employee, and it can choose not to hire and shut down if f_2 does the same. Therefore, the minmax payoffs of both firms should be equal to 2. However, there is no feasible way to generate a payoff of 2 for both firms in the stage matching, because either they hire the same number of workers or they do not. In other words, the frontier of the feasible payoff set has an intercept of 2 and a slope of -1; see Figure 4 for an illustration. Hence, no self-enforcing matching process exists for any value of δ with the presence of matching externalities.

5 Conclusion

In this paper, we propose a new approach of studying stability in matching markets with complementarities and peer effects, which have traditionally challenged the existence of static stable outcomes. Our approach treats stability as the result of a dynamic process that is self-sustained by expectations; importantly, these expectations must also be themselves consistent with stability.

We leverage the repeated interactions of long-lived firms to discipline these market participants and sustain stability over time. As a result, we can define notions such as minmax payoffs and feasible payoff sets that are analogous to those from standard repeated games. However, unlike standard repeated games in which players play a noncooperative game, in our model firms and workers participate in a cooperative game every period. This difference makes the nonissue of equilibrium existence in noncooperative setting more challenging in our cooperative setting. We prove the existence of a stable matching process by explicitly constructing a feasible payoff profile through random serial dictatorship where firms take turns to receive their more favorable outcome. This can be interpreted as a form of no-poaching agreement, which is prevalent and subject to intense legal debates in real-world contexts.

From a theoretical standpoint, the existence of self-enforcing matching processes reconciles the lack of static stable matching and absence of complete chaos in many matching markets. From a practical point of view, the sustainability of outcomes through RSD adds a new dimension to the ongoing debate regarding the anti-trust implications of no-poaching agreements. While no-poaching agreements are generally regarded as anti-competitive, our analysis suggests they may play a crucial role in preserving market stability.

A Appendix

A.1 Intermediate Results

Proof of Lemma 1. Fix $f \in \mathcal{F}$, $\theta \in \Theta$, and stage matching $m = (\phi, p) \in M^\circ(\theta)$. Define $W_m \equiv \{w \in \mathcal{W} : \phi(w) \neq (\emptyset, \emptyset)\}$, so $|W_m| \leq Q$. Hiring from $\mathcal{W} \backslash W_m$ at infinitesimal wages is always a feasible deviation for f, which means

$$
\begin{aligned}
\sup_{(W',p'_f)\in D_f(m,\theta)} u_f([m,(f,W',p'_f)],\theta) &\geq \max_{W'\subseteq \mathcal{W}\backslash W_m, |W'|\leq q_f} \chi(f,W',\theta) \\
&\geq \min_{W\subseteq\mathcal{W},|W|=|W_m|} \max_{W'\subseteq \mathcal{W}\backslash W, |W'|\leq q_f} \chi(f,W',\theta) \\
&\geq \min_{W\subseteq\mathcal{W},|W|\leq Q} \max_{W'\subseteq \mathcal{W}\backslash W, |W'|\leq q_f} \chi(f,W',\theta).
\end{aligned}
$$

Taking infimum on the LHS yields

$$
\underline{u}_f(\theta) \geq \min_{W\subseteq\mathcal{W},|W|\leq Q} \max_{W'\subseteq \mathcal{W}\backslash W, |W'|\leq q_f} \chi(f,W',\theta).
$$

For the other direction, take any $W \subseteq \mathcal{W}$ with $|W| \leq Q$. We can always construct a stage matching $\widehat{m} = (\widehat{\phi}, \widehat{p}) \in M^\circ(\theta)$ such that (i) $\widehat{\phi}$ assigns all workers in W to firms, and (ii) all workers in W receive sufficiently high wages so that f never finds it profitable to deviate with any workers in W. Therefore,

$$
\begin{aligned}
\max_{W'\subseteq \mathcal{W}\backslash W, |W'|\leq q_f} \chi(f,W',\theta) &= \sup_{(W',p'_f)\in D_f(\widehat{m},\theta)} u_f([\widehat{m},(f,W',p'_f)],\theta) \\
&\geq \inf_{m\in M^\circ(\theta)} \sup_{(W',p'_f)\in D_f(m,\theta)} u_f([m,(f,W',p'_f)],\theta)
\end{aligned}
$$

Minimizing over W on the LHS yields the other direction. $\square$

Lemma 3. *Deviation plan (d, η) is a* one-shot deviation *from matching process μ if there is a unique ex post history $\widehat{h}$ where $[\mu, (f, d, \eta)](\widehat{h}) \neq \mu(\widehat{h})$. Matching process μ is self-enforcing if and only if (i) $v_w(\mu(h), \theta) \geq 0$ for every $w \in \mathcal{W}$ at every $h \in \mathcal{H}$; and (ii) no firm has a* *feasible and profitable* one-shot deviation.

Proof of Lemma 3. Standard arguments (Blackwell 1965) suffice. We include them for completeness. Suppose one-shot deviation plan (d, η) from matching process μ for firm f is feasible and profitable. Since stage payoffs are bounded for firm f and there is discounting, the standard one-shot deviation principle for individual decision-making (Blackwell 1965) implies that there exists an ex post history $\widehat{h} = (\overline{\widehat{h}}, \widehat{\theta}, \widehat{s})$ such that

$$(1-\delta)\left[\widetilde{u}_f(d_j(\widehat{h}), \widehat{\theta}) - \sum_{w \in d(\widehat{h})} \eta_w(\widehat{h})\right] + \delta U_f\big([\mu(\widehat{h}), (f, d(\widehat{h}), \eta(\widehat{h}))] \,\big|\, \mu\big) > U_f(\widehat{h} \,|\, \mu).$$

Consider deviation plan (d_f^o, η_f^o) that satisfies

$$[\mu, (f, d_f^o, \eta_f^o)](h) = \begin{cases} [\mu, (f, d, \eta)](h) & \text{if } h = \widehat{h}, \\ \mu(h) & \text{otherwise.} \end{cases}$$

Then (d_f^o, η_f^o) is a profitable one-shot deviation plan for firm f. □

Lemma 4. *For any stage matching m,*

$$[m, (f_1, W_{f_1}, p_{f_1})] = [m, (f_2, W_{f_2}, p_{f_2})] \neq m \text{ implies } f_1 = f_2.$$

Proof of Lemma 4. Let $m = (\phi, p)$, $[m, (f_1, W'_{f_1}, p'_{f_1})] = (\overline{\phi}, \overline{p})$, and $[m, (f_2, W'_{f_2}, p'_{f_2})] = (\widehat{\phi}, \widehat{p})$. Toward contradiction, suppose $f_1 \neq f_2$, but $[m, (f_1, W'_{f_1}, p'_{f_1})] = [m, (f_2, W'_{f_2}, p'_{f_2})] \neq m$. Then $\overline{\phi} = \widehat{\phi}$ and $\overline{p}_f = \widehat{p}_f$ for every $f \in \mathcal{F}$. Each of the three cases to consider yields a contradiction.

1. Suppose $W'_{f_1} = \phi(f_1)$ and $W'_{f_2} = \phi(f_2)$. Since $[m, (f_1, W'_{f_1}, p'_{f_1})] \neq m$, we have $p'_{f_1} \neq p_{f_1}$. Then $\overline{p}_{f_1} = p'_{f_1} \neq p_{f_1} = \widehat{p}_{f_1}$, a contradiction.

2. Suppose $W'_{f_1} \subseteq \phi(f_1)$ and $W'_{f_2} \subseteq \phi(f_2))$ but, without loss of generality, $W'_{f_1} \neq \phi(f_1)$. We have $\overline{\phi}(f_1) = W'_{f_1} \neq \phi(f_1) = \widehat{\phi}(f_1)$, so $\overline{\phi} \neq \widehat{\phi}$, a contradiction.

3. Suppose, without loss of generality, $W'_{f_1} \nsubseteq \phi(f_1)$. Let $w' \in W'_{f_1} \backslash \phi(f_1)$. We have

$\overline{\phi}(w') = f_1$, whereas $\widehat{\phi}(w') \in \{\phi(w'), f_2\}$. Since $w' \notin \phi(f_1)$ and $f_1 \neq f_2$, $f_1 \notin \{\phi(w'), f_2\}$. Hence, $\overline{\phi} \neq \widehat{\phi}$, a contradiction.

Therefore, $[m, (f_1, W_{f_1}, p_{f_1})] = [m, (f_2, W_{f_2}, p_{f_2})] \neq m$ implies $f_1 = f_2$. □

A.2 Omitted Proofs for Section 3

A.2.1 Proof of Proposition 1

Lemma 5. *For each $\theta \in \Theta$, there exist stage matchings $\{\underline{m}_f(\theta)\}_{f\in\mathcal{F}} \subseteq M^{\circ}(\theta)$ such that $\forall f \in \mathcal{F}$,*

(i) $\sup_{(W',p'_f)\in D_f(\underline{m}_f,\theta)} u_f([\underline{m}_f(\theta), (f, W', p'_f)], \theta) = \underline{u}_f(\theta)$;

(ii) $u_f(\underline{m}_f(\theta), \theta) \leq \underline{u}_f(\theta)$;

(iii) $\underline{u}_f(\theta) - u_f(\underline{m}_f(\theta), \theta) = \underline{u}^*_f - \mathbb{E}_\pi[u_f(\underline{m}_f(\theta), \theta)]$.

Proof of Lemma 5. Fix $\theta \in \Theta$. For every $f \in \mathcal{F}$, let

$$\underline{W}_f(\theta) \in \underset{W\subseteq\mathcal{W},|W|\leq Q}{\arg\min} \max_{W'\subseteq\mathcal{W}\backslash W,|W'|\leq q_f} \chi(f, W', \theta)$$

denote a set of workers to eliminate to minmax firm f, and

$$b_f(\theta) = \overline{u}_f(\theta) - \underline{u}_f(\theta) \geq 0$$

denotes the difference between f's maximum feasible payoff and minmax payoff.

For every $f \in \mathcal{F}$, let $\{\underline{W}^{f'}_f(\theta)\}_{f'\in\mathcal{F}}$ be a partition of $\underline{W}_f(\theta)$ such that $\left|\underline{W}^f_f(\theta)\right| \geq 1$ and $\left|\underline{W}^{f'}_f(\theta)\right| \leq q_{f'}$ for every $f' \in \mathcal{F}$. Define

$$B_f \equiv \max_{\theta\in\Theta} \max_{w\in\underline{W}^f_f(\theta)} \left|\underline{W}^f_f(\theta)\right| \left[b_f(\theta) - \widetilde{v}_w\big(f, \underline{W}^f_f(\theta)\backslash\{w\}, \theta\big)\right] + \underline{u}_f(\theta) - \widetilde{u}_f(\underline{W}^f_f(\theta), \theta).$$

For every $\theta \in \Theta$, define stage matching $\underline{m}_f(\theta) = (\underline{\phi}^f(\theta), \underline{p}^f(\theta))$, where

$$\underline{\phi}^f(\theta)(f') = \underline{W}^{f'}_f(\theta),$$

and

$$\underline{p}^f_{f'w}(\theta) = \begin{cases} \frac{\widetilde{u}_f(\underline{W}^f_f(\theta),\theta) - \underline{u}_f(\theta) + B_f}{\left|\underline{W}^f_f(\theta)\right|} & \text{if } f' = f \text{ and } w \in \underline{W}^f_f(\theta) \\ b_f(\theta) - \widetilde{v}_w\big(f', \underline{W}^{f'}_f(\theta)\backslash\{w\}, \theta\big) & \text{if } f' \neq f \text{ and } w \in \underline{W}^{f'}_f(\theta) \neq \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

For any $\theta \in \Theta$, if $w \in \underline{W}^f_f(\theta)$,

$$v_w(\underline{m}_f(\theta), \theta) = \widetilde{v}_w\big(f, \underline{W}^f_f(\theta)\backslash\{w\}, \theta\big) + \frac{\widetilde{u}_f(\underline{W}^f_f(\theta), \theta) - \underline{u}_f(\theta) + B_f}{\left|\underline{W}^f_f(\theta)\right|} \geq b_f(\theta) \geq 0$$

by the definition of B_f. If $w \in \underline{W}^{f'}_f(\theta)$ and $f' \neq f$,

$$v_w(\underline{m}_f(\theta), \theta) = b_f(\theta) \geq 0.$$

Hence, $\underline{m}_f(\theta) \in M^\circ(\theta)$.

(i) Consider any feasible deviation $(W', p'_f) \in D_f(\underline{m}_f, \theta)$. Suppose $W' \subseteq \mathcal{W}\backslash\underline{W}_f(\theta)$. By feasibility, every $w \in W'$ finds the deviation individually rational, so

$$\widetilde{v}_w(f, W'\backslash\{w\}, \theta) + p'_{fw} \geq 0.$$

This implies

$$\begin{aligned} \widetilde{u}_f(W', \theta) - \sum_{w \in W'} p'_{fw} &\leq \widetilde{u}_f(W', \theta) + \sum_{w \in W'} \widetilde{v}_w(f, W'\backslash\{w\}, \theta) & (3) \\ &= \chi(f, W', \theta) \\ &\leq \max_{W \subseteq \mathcal{W}\backslash\underline{W}_f(\theta), |W| \leq q_f} \chi(f, W, \theta) & (4) \\ &= \underline{u}_f(\theta), \end{aligned}$$

where (3) follows from $W' \subseteq \mathcal{W}\backslash\underline{W}_f(\theta)$, and (4) follows from the definition of $\underline{W}_f(\theta)$. Suppose instead $W' \nsubseteq \mathcal{W}\backslash\underline{W}_f(\theta)$. Fix $w \in W' \cap \underline{W}_f(\theta)$. By the construction of $\underline{m}_f(\theta)$, $v_w(\underline{m}_f(\theta), \theta) \geq b_f(\theta)$. For the deviation to be feasible, w must obtain a payoff weakly higher than in $\underline{m}_f(\theta)$:

$$\widetilde{v}_w(f, W'\backslash\{w\}, \theta) + p'_{fw} \geq b_f(\theta);$$

meanwhile, every $w' \in W'\backslash\{w\}$ needs to find the deviation individually rational:

$$\widetilde{v}_{w'}(f, W'\backslash\{w'\}, \theta) + p'_{fw'} \geq 0.$$

Then

$$\begin{aligned}
\widetilde{u}_f(W', \theta) - \sum_{w' \in W'} p'_{fw'} &= \widetilde{u}_f(W', \theta) - p'_{fw} - \sum_{w' \in W', w' \neq w} p'_{fw'} \\
&\leq \widetilde{u}_f(W', \theta) + \sum_{w \in W'} \widetilde{v}_w(f, W'\backslash\{w\}, \theta) - b_f(\theta) \\
&= \chi(f, W', \theta) - \overline{u}_f(\theta) + \underline{u}_f(\theta) \\
&\leq \underline{u}_f(\theta).
\end{aligned}$$

We have shown that for any W',

$$u_f([\underline{m}_f, (f, W', p'_f)], \theta) \leq \underline{u}_f(\theta).$$

Hence,

$$\sup_{(W', p'_f) \in D_f(\underline{m}_f, \theta)} u_f([\underline{m}_f, (f, W', p'_f)], \theta) \leq \underline{u}_f(\theta).$$

By the definition of $\underline{u}_f(\theta)$, the above holds with equality.

(ii) For every $f \in \mathcal{F}$ and $\theta \in \Theta$,

$$\begin{aligned}
u_f(\underline{m}_f(\theta), \theta) &= \widetilde{u}_f(\underline{W}^f_f(\theta), \theta) - \sum_{w \in \underline{W}^f_f(\theta)} \underline{p}^f_{fw}(\theta) \\
&= \widetilde{u}_f(\underline{W}^f_f(\theta), \theta) - \widetilde{u}_f(\underline{W}^f_f(\theta), \theta) + \underline{u}_f(\theta) - B_f \\
&= \underline{u}_f(\theta) - B_f && (5) \\
&\leq \widetilde{u}_f(\underline{W}^f_f(\theta), \theta) + \sum_{w \in \underline{W}^f_f(\theta)} \widetilde{v}_i(f, \underline{W}^f_f(\theta)\backslash\{w\}, \theta) - \left|\underline{W}^f_f(\theta)\right| b_f(\theta) && (6)
\end{aligned}$$

$$
\begin{aligned}
&\leq \widetilde{u}_f(\underline{W}_f^f(\theta), \theta) + \sum_{w \in \underline{W}_f^f(\theta)} \widetilde{v}_i(f, \underline{W}_f^f(\theta) \backslash \{w\}, \theta) - b_f(\theta) \quad (7) \\
&= \chi(f, \underline{W}_f^f(\theta), \theta) - \overline{u}_f(\theta) + \underline{u}_f(\theta) \\
&\leq \underline{u}_f(\theta),
\end{aligned}
$$

where (6) follows from the definition of B_f, and (7) follows from $\left|\underline{W}_f^f(\theta)\right| \geq 1$.

(iii) By (5), for every $f \in \mathcal{F}$ and $\theta \in \Theta$,

$$u_f(\underline{m}_f(\theta), \theta) = \underline{u}_f(\theta) - B_f.$$

Therefore,

$$\mathbb{E}_\pi[u_f(\underline{m}_f(\theta), \theta)] = \underline{u}_f^* - B_f,$$

which means

$$\underline{u}_f(\theta) - u_f(\underline{m}_f(\theta), \theta) = \underline{u}_f^* - \mathbb{E}_\pi[u_f(\underline{m}_f(\theta), \theta)].$$

□

Lemma 6. *For every $u \in \mathcal{U}^*$, there exist vectors $\{u^f\}_{f \in \mathcal{F}} \subseteq \mathcal{U}^*$ such that for all distinct $f, f' \in \mathcal{F}$,*

$$u_f^f < u_f \quad \textit{and} \quad u_f^f < u_f^{f'}.$$

Proof of Lemma 6. Take any distinct $f, f' \in \mathcal{F}$. If f pays higher wages to its employees, f is worse off while f' is indifferent. $\mathcal{U}^*$ satisfies the NEU condition (Abreu et al. 1994), and $\{u^f\}_{f \in \mathcal{F}}$ can be obtained via a similar construction. The detailed arguments are as follows.

Fix stage matching $m = (\phi, p)$ and type profile θ. For every $f \in \mathcal{F}$, define a matching $m^f = (\phi^f, p^f)$ by $\phi^f = \phi$, and

$$p_{f'w}^f = \begin{cases} p_{f'w} + 1 & \text{if } f' = f \text{ and } w \in \phi^f, \\ p_{f'w} & \text{otherwise.} \end{cases}$$

Clearly, $v_w(m^f, \theta) = v_w(m, \theta)$ for every $w \in \mathcal{W} \backslash \phi^f$, and $v_w(m^f, \theta) > v_w(m, \theta)$ for every $w \in \phi^f$, so m^f is individually rational for all workers whenever m is—i.e., $m \in M^\circ(\theta)$ implies

$m^f \in M^\circ(\theta)$. Moreover, by the construction of $\{m^f\}_{f\in\mathcal{F}}$, we have $u_f(m^f,\theta) \leq u_f(m,\theta)$ for every $f \in \mathcal{F}$, with the inequality being strict if $\phi(f) \neq \emptyset$. In the following, we write $\zeta^f : m \mapsto m^f$.

For an arbitrary vector $u \in \mathcal{U}^*$, there exists a collection $(\lambda(\theta))_{\theta\in\Theta}$, where $\lambda(\theta) \in \Delta(M^\circ(\theta))$ for every $\theta \in \Theta$, such that $u_f = \mathbb{E}_\pi[\mathbb{E}_{\lambda(\theta)}[u_f(m,\theta)]] > \underline{u}_f^*$ for every f. For every firm $f \in \mathcal{F}$, define

$$u^f = \epsilon\mathbb{E}_\pi[\mathbb{E}_{\lambda(\theta)}[u(\zeta^f(m),\theta)]] + (1-\epsilon)u.$$

First observe that because $u_f > \underline{u}_f^* \geq 0$, there exist $\theta \in \Theta$ and $m = (\phi, p) \in M^\circ(\theta)$ such that $\lambda(\theta)[m] > 0$ and $\phi(f) \neq \emptyset$; by the construction of $\{\zeta^f(m)\}_{f\in\mathcal{F}}$, $u_f(\zeta^f(m),\theta) < u_f(m,\theta) = u_f(\zeta^{f'}(m),\theta)$, so $u_f^f < u_f$ and $u_f^f < u_f^{f'}$ if $\epsilon > 0$. Second, every u^f can be written as

$$u^f = \mathbb{E}_\pi[\mathbb{E}_{\lambda^f(\theta)}[u_f(m,\theta)]],$$

where

$$\lambda^f(\theta) = \epsilon\lambda(\theta) \circ (\zeta^f)^{-1} + (1-\epsilon)\lambda(\theta) \quad f \in \mathcal{F}. \tag{8}$$

Note that the support of every $\lambda^f(\theta)$ is also bounded. Finally, for small enough $\epsilon > 0$, $u_f^f > \underline{u}_f^*$. Therefore, there must exist $\{u^f\}_{f\in\mathcal{F}} \subseteq \mathcal{U}^*$ such that for all distinct $f, f' \in \mathcal{F}$,

$$u_f^f < u_f \quad \text{and} \quad u_f^f < u_f^{f'}.$$

□

Proof of Proposition 1. (i) Fix $u \in \mathcal{U}^*$. Let $(\lambda(\theta))_{\theta\in\Theta}$ be a tuple of lotteries such that $\lambda(\theta) \in \Delta(M^\circ(\theta))\ \forall\theta \in \Theta$, and $u = \mathbb{E}_\pi[\mathbb{E}_{\lambda(\theta)}[u(m,\theta)]]$. Let $\{\underline{m}_f(\theta)\}_{\theta\in\Theta,f\in\mathcal{F}}$ be the minmax stage matchings constructed in Lemma 5. By Lemma 5(ii),

$$u_f(\underline{m}_f(\theta),\theta) \leq \underline{u}_f(\theta) \quad \forall(f,\theta) \in \mathcal{F}\times\Theta. \tag{9}$$

By Lemma 6, there exist $\{u^f\}_{f\in\mathcal{F}} \subseteq \mathcal{U}^*$ such that for all distinct $f, f' \in \mathcal{F}$,

$$u_f^f < u_f \quad \text{and} \quad u_f^f < u_f^{f'}.$$

For every $f \in \mathcal{F}$, let $(\lambda^f(\theta))_{\theta\in\Theta}$ be the tuple of lotteries that satisfies $\lambda^f(\theta) \in \Delta(M^\circ(\theta))\ \forall\theta \in \Theta$, and $u^f = \mathbb{E}_\pi[\mathbb{E}_{\lambda^f(\theta)}[u(m,\theta)]]$. By Carathéodory's theorem, it is without loss to assume every $\lambda^f(\theta)$ has bounded support. Define

$$Z_f(\theta) \equiv \inf_{m\in\text{supp}(\lambda^f(\theta))} u_f(m,\theta). \tag{10}$$

At any θ, let $\overline{u}_f(\theta)$ denote a firm f's maximum payoff from any matching in $M^\circ(\theta)$ as defined in (2). Note that $\overline{u}_f(\theta)$ is an upper bound for the payoff f can obtain in any feasible deviation from matchings in $M^\circ(\theta)$. Let L be an integer greater than

$$\max_{f\in\mathcal{F},\theta\in\Theta} \frac{\overline{u}_f(\theta) - Z_f(\theta)}{u_f^f - \underline{u}_f^*}.$$

Consider a matching process with the following three phases:

(I) If past realizations from $\lambda(\cdot)$ were followed: Match according to $\lambda(\cdot)$;

(II) If f deviates: For the next L periods, match according to $\underline{m}_f(\cdot)$;

(III) If $\underline{m}_f(\cdot)$ was followed for L periods: Match according to the realization from $\lambda^f(\cdot)$ until a firm deviates.

Note that if firm f' deviates from (II) or (III), the process restarts (II) with f replaced by f'. All deviations by individual workers are ignored.

By Lemma 4, for any stage matching that can result from a firm's deviation, we can uniquely identify the firm, so the transition above is well-defined. All stage matchings in this matching process are individually rational for workers. It remains to check that no firm has profitable one-shot deviations.

<u>(I)</u> f has no profitable one-shot deviation if

$$(1-\delta)u_f(m,\theta) + \delta u_f \geq (1-\delta)\overline{u}_f(\theta) + \delta(1-\delta^L)\sum_{\theta'\in\Theta}\pi[\theta']u_f(\underline{m}_f(\theta'),\theta') + \delta^{L+1}u_f^f.$$

Since $u_f > u_f^f$ by construction, f has no profitable one-shot deviation for δ high enough.

<u>(II)</u> Consider two cases.

Case a: $f' \neq f$. Without deviation, f' gets

$$(1-\delta)u_{f'}(\underline{m}_f(\theta),\theta)+\delta(1-\delta^{L-\tau-1})\sum_{\theta'\in\Theta}\pi[\theta']u_{f'}(\underline{m}_f(\theta'),\theta')+\delta^{L-\tau}u^f_{f'}. \tag{11}$$

By deviating, f' gets less than

$$(1-\delta)\overline{u}_{f'}(\theta)+\delta(1-\delta^{L})\sum_{\theta'\in\Theta}\pi[\theta']u_{f'}(\underline{m}_{f'}(\theta'),\theta')+\delta^{L+1}u^{f'}_{f'}. \tag{12}$$

As $\delta \to 1$, (11) converges to $u^f_{f'}$ and (12) to $u^{f'}_{f'}$. Because $u^f_{f'} > u^{f'}_{f'}$ by construction, no one-shot deviation is profitable for high δ.

Case b: $f' = f$. Without deviation, f gets

$$(1-\delta)u_{f}(\underline{m}_f(\theta),\theta)+\delta(1-\delta^{L-\tau-1})\sum_{\theta'\in\Theta}\pi[\theta']u_{f}(\underline{m}_f(\theta'),\theta')+\delta^{L-\tau}u^f_{f}. \tag{13}$$

With deviation, f gets less than

$$(1-\delta)\underline{u}_{f}(\theta)+\delta(1-\delta^{L})\sum_{\theta'\in\Theta}\pi[\theta']u_{f}(\underline{m}_{f}(\theta'),\theta')+\delta^{L+1}u^{f}_{f}. \tag{14}$$

Since $\sum_{\theta'\in\Theta}\pi[\theta']u_f(\underline{m}_f(\theta'),\theta') \leq \underline{u}^*_f < u^f_f$, expression (13) is increasing in τ. Hence it suffices to prove the case when $\tau = 0$. When $\tau = 0$,

$$\left[u_f(\underline{m}_f(\theta),\theta)-\underline{u}_f(\theta)\right]-\delta^L\sum_{\theta'\in\Theta}\pi[\theta']u_f(\underline{m}_f(\theta'),\theta')+\delta^L u^f_f \geq 0. \tag{15}$$

By Lemma 5(iii),

$$u_f(\underline{m}_f(\theta),\theta)-\underline{u}_f(\theta)=\sum_{\theta'\in\Theta}\pi[\theta']u_f(\underline{m}_f(\theta'),\theta')-\underline{u}^*_f. \tag{16}$$

Substituting (16) into (15) yields

$$(1-\delta^L)\sum_{\theta'\in\Theta}\pi[\theta']u_f(\underline{m}_f(\theta'),\theta')+\delta^L u^f_f \geq \underline{u}^*_f.$$

By construction, $u_f^f > \underline{u}_f^*$, so the inequality holds for δ high enough.

(III) Consider two cases.

Case a: $f' \neq f$. The argument for (I) holds analogously here, once we replace f with f' and u_f with $u_{f'}^f$.

Case b: $f' = f$. There is no profitable one-shot deviation for f' if

$$(1-\delta)u_{f'}(m,\theta) + \delta u_{f'}^{f'} \geq (1-\delta)\overline{u}_{f'}(\theta) + \delta(1-\delta^L)\sum_{\theta'\in\Theta}\pi[\theta']u_{f'}(\underline{m}_{f'}(\theta'),\theta') + \delta^{L+1}u_{f'}^{f'}.$$

This is equivalent to

$$\delta\frac{1-\delta^L}{1-\delta}\left[u_{f'}^{f'} - \sum_{\theta'\in\Theta}\pi[\theta']u_{f'}(\underline{m}_{f'}(\theta'),\theta')\right] \geq \overline{u}_{f'}(\theta) - u_{f'}(m,\theta).$$

Because $\frac{1-\delta^L}{1-\delta} \to L$ as $\delta \to 1$, the LHS converges to

$$\begin{aligned} & L\left[u_{f'}^{f'} - \sum_{\theta'\in\Theta}\pi[\theta']u_{f'}(\underline{m}_{f'}(\theta'),\theta')\right] \\ > \; & \frac{\overline{u}_{f'}(\theta) - Z_{f'}(\theta)}{u_{f'}^{f'} - \underline{u}_{f'}^*}\left[u_{f'}^{f'} - \sum_{\theta'\in\Theta}\pi[\theta']u_{f'}(\underline{m}_{f'}(\theta'),\theta')\right] \qquad (17) \\ \geq \; & \overline{u}_{f'}(\theta) - u_{f'}(m,\theta), \qquad (18) \end{aligned}$$

where (17) follows from the definition of L, and (18) follows from (9), (10), and $m \in \text{supp}(\lambda^f(\theta))$. Thus, no deviation is profitable for δ high enough. □

Proof of Proposition 1(ii). By definition, every firm f can secure $\underline{u}_f(\theta)$ by deviating with workers. Taking expectation over θ delivers the result. It suffices to show that for any self-enforcing matching process μ,

$$U_f(h\,|\,\mu) \geq (1-\delta)\underline{u}_f(\theta) + \delta\underline{u}_f^* \quad \forall f \in \mathcal{F}, h \in \mathcal{H}.$$

Suppose by contradiction that $U_f(h\,|\,\mu) < (1-\delta)\underline{u}_f(\theta) + \delta\underline{u}_f^*$ for some $f \in \mathcal{F}$ and $h \in \mathcal{H}$. We will show that firm f has a feasible profitable deviation plan from μ, so μ cannot be self-enforcing.

Consider the following deviation plan (d, η) from μ: For all ex post histories that precede h, the deviation plan is consistent with the matching process μ. For any $h' = (\overline{h}', \theta', s') \in \mathcal{H}$ that succeeds h (including h itself),

$$d(h') = \underset{A \subseteq \mathcal{W} \setminus \bigcup_{f' \in \mathcal{F}} \mu(f'|h')}{\arg\max} \chi(f, A, \theta').$$

If $d(h') \neq \emptyset$, then

$$\eta_w(h') = \begin{cases} -\widetilde{v}_w(f, d(h')) + \frac{1}{2|d(h')|} \left[(1-\delta)\underline{u}_f(\theta) + \delta \underline{u}_f^* - U_f(h \,|\, \mu)\right] & \text{if } w \in d(h'); \\ 0 & \text{otherwise.} \end{cases}$$

If $d(h') = \emptyset$, define

$$\eta_w(h') = 0 \quad \forall w \in \mathcal{W}.$$

Note that by construction, $\chi(f, d(h'), \theta') \geq \underline{u}_f(\theta')$ because $\bigcup_{f' \in \mathcal{F}} \mu(f'|h) \leq Q$.

We first verify that deviation plan (d, η) is feasible. At any history h' that succeeds h, any $w \in d(h')$ is unmatched under $\mu(h')$ by construction, which means $v_w(\mu(h'), \theta') = 0$. Meanwhile, according to the deviation plan,

$$\begin{aligned} & \widetilde{v}_w(f, d(h')) + \eta_w(h') \\ = & \widetilde{v}_w(f, d(h')) - \widetilde{v}_w(f, d(h')) + \frac{1}{2\,|d(h')|} \left[(1-\delta)\underline{u}_f(\theta) + \delta \underline{u}_f^* - U_f(h \,|\, \mu)\right] \\ = & \frac{1}{2\,|d(h')|} \left[(1-\delta)\underline{u}_f(\theta) + \delta \underline{u}_f^* - U_f(h \,|\, \mu)\right] \\ > & \, 0. \end{aligned}$$

So at every possible history h', every worker in $d(h')$ finds herself strictly better off by joining the deviation, which ensures the feasibility of (d, η).

To see that (d, η) is profitable, observe that at every history h' after h, firm f's stage-game payoff from the manipulated static matching $[\mu(h'), (f, d(h'), \eta(h'))]$ is

$$\begin{aligned}
&\widetilde{u}_f(d(h'), \theta') - \sum_{w \in d(h')} \eta_w(h') \\
&= \widetilde{u}_f(d(h'), \theta') + \sum_{w \in d(h'))} \widetilde{v}_w(f, d(h')) - \frac{1}{2}\left[(1-\delta)\underline{u}_f(\theta) + \delta \underline{u}_f^* - U_f(h \,|\, \mu)\right] \\
&\geq \underline{u}_f(\theta') - \frac{1}{2}\left[(1-\delta)\underline{u}_f(\theta) + \delta \underline{u}_f^* - U_f(h \,|\, \mu)\right].
\end{aligned}$$

Since this is true for every history h' after h, f's total discounted payoff at ex post history h from the deviation plan is

$$\begin{aligned}
U_f(h \,|\, [\mu, (f, d, \eta)]) &= (1-\delta)\underline{u}_f(\theta) + \delta \underline{u}_f^* - \frac{1}{2}\left[(1-\delta)\underline{u}_f(\theta) + \delta \underline{u}_f^* - U_f(h \,|\, \mu)\right] \\
&= \frac{1}{2}\left[(1-\delta)\underline{u}_f(\theta) + \delta \underline{u}_f^* + U_f(h \,|\, \mu)\right] \\
&> \frac{1}{2}[(U_f(h \,|\, \mu) + U_f(h \,|\, \mu)] \\
&= U_f(h \,|\, \mu).
\end{aligned}$$

Hence, deviation plan (d, η) is both feasible and profitable for firm f, which contradicts the self-enforcement of μ. Hence, $U_f(h \,|\, \mu) \geq (1-\delta)\underline{u}_f(\theta) + \delta \underline{u}_f^* \quad \forall f \in \mathcal{F}, h \in \mathcal{H}$. □

A.2.2 Proof of Lemma 2

For every $f \in \mathcal{F}$, let $\Gamma(f) = \{\gamma \in \Gamma : \gamma(f) = 1\}$ denote the set of orderings in which f is ranked first. Recall that for every $\gamma \in \Gamma$ and $\theta \in \Theta$, matching $\widehat{m}(\theta, \gamma)$ is produced by serial dictatorship (Definition 3). It is straightforward that

$$u_f(\widehat{m}(\theta, \gamma)) = \overline{u}_f(\theta). \tag{19}$$

Additionally, in serial dictatorship, for each $\widehat{m}(\theta, \gamma)$ and $f \in \mathcal{F}$,

$$u_f(\widehat{m}(\theta, \gamma)) = u_f(\widehat{W}_{\gamma(f)}) = \max_{W' \subseteq \mathcal{W} \setminus \cup_{1 \leq i < \gamma(f)} \widehat{W}_i, |W'| \leq q_f} \chi(f, W', \theta).$$

Since every $\gamma \in \Gamma$ satisfies $\left|\cup_{1\le i<\gamma(f)} \widehat{W}_i\right| \le Q$,

$$u_f(\widehat{m}(\theta,\gamma)) \ge \min_{W\subseteq\mathcal{W},|W|\le Q} \ \max_{W'\subseteq\mathcal{W}\setminus W,|W'|\le q_f} \chi(f,W',\theta) = \underline{u}_f(\theta) \quad \forall\theta\in\Theta, f\in\mathcal{F}, \gamma\in\Gamma. \quad (20)$$

For every $f \in \mathcal{F}$,

$$\begin{aligned}\frac{1}{|\Gamma|}\sum_{\gamma\in\Gamma}\mathbb{E}_\pi\Big[u_f\big(\,\widehat{m}(\theta,\gamma),\theta\,\big)\Big] &= \sum_{\theta\in\Theta}\pi(\theta)\sum_{\gamma\in\Gamma}\frac{1}{|\Gamma|}\Big[u_f\big(\,\widehat{m}(\theta,\gamma),\theta\,\big)\Big]\\ &> \sum_{\theta\in\Theta}\pi(\theta)\,\underline{u}_f(\theta) = \underline{u}^*_{f'},\end{aligned}$$

where the strict inequality follows from (19), (20), and Assumption 2.

A.3 Omitted Proofs for Section 4.1

A.3.1 Proof of Proposition 1*

For part (i), let $(\lambda(\theta))_{\theta\in\Theta}$ be the tuple of lotteries such that $\lambda(\theta) \in \Delta(M^\circ(\theta))$ for every $\theta \in \Theta$, and $u = \mathbb{E}_\pi[\mathbb{E}_{\lambda(\theta)}[u(m,\theta)]]$. There are three cases to consider.

Case 1: $\mathcal{R} = \emptyset$. If a firm f receives $u_f = \overline{u}^*_f$ on average, it is necessary that this firm receives the highest possible payoff $\overline{u}_f(\theta)$ at every realization of θ. This in turn implies that, for each θ, $\lambda(\theta)$ only assigns positive probability to stage matchings that are *stable* in a static sense. Therefore, a matching process that recommends according to $(\lambda(\theta))_{\theta\in\Theta}$ in every period is self-enforcing.

Case 2: $|\mathcal{R}| = 1$. Let f denote the single firm in $\mathcal{R}$, and let $(\underline{m}_f(\theta))_{\theta\in\Theta}$ be the matchings that give f the most severe punishment by f itself, while all other firms receive the highest possible payoff $\overline{u}_f(\theta)$. Consider the following matching process:

(I) Match according to $\lambda(\cdot)$ if $\lambda(\cdot)$ was followed in the last period or $\underline{m}_f(\cdot)$ was followed for L periods;

(II) If firm f deviates from (I), match according to $\underline{m}_f(\cdot)$ for L periods.

If firm f deviates from (II), restart (II).

It is easy to check that when L is sufficiently large, firm f has no incentive to deviate in either phase, since $\delta \to 1$. All other firms have no incentive to deviate, since they already receive maximum stage payoff in every period.

Case 3: $|\mathcal{R}| \geq 2$. The proof for this case essentially follows the one for Proposition 1 with proper adjustments.

For part (ii), by construction, any firm $f \in \mathcal{P}_1$ can secure a stage payoff $\overline{u}_f(\theta)$ by deviation at every θ regardless of the stage matching. Taking expectation yields $U_f(\overline{h} \,|\, \mu) = \overline{u}^*_f$ for every $f \in \mathcal{P}_1$.

Suppose $U_{f'}(\overline{h} \,|\, \mu) = \overline{u}^*_{f'}$ for every $f' \in \bigcup_{k=1}^{n} \mathcal{P}_k$ with $n < N$. Then all firms in $\bigcup_{k=1}^{n} \mathcal{P}_k$ offer zero wages to their employees. By construction, each firm $f \in \mathcal{P}_{n+1}$ can secure a stage payoff $\overline{u}_f(\theta)$ by deviation at every θ regardless of how firms in $\mathcal{F} \backslash \bigcup_{k=1}^{n} \mathcal{P}_k$ are matched in each period. Taking expectation yields $U_f(\overline{h} \,|\, \mu) = \overline{u}^*_f$ for every $f \in \mathcal{P}_{n+1}$. By induction, the equality holds for all $f \in \mathcal{F} \backslash \mathcal{R}$.

By definition of the effective minmax payoff, every firm $f \in \mathcal{R}$ can secure $\underline{u}_f(\mathcal{R}, \theta)$ in each period by deviating with workers. Taking expectation over θ yields $U_f(\overline{h} \,|\, \mu) \geq \underline{u}^*_f(\mathcal{R})$ for these firms. Rigorous proof can be adapted from that of Proposition 1(ii) in Appendix A.2.1.

A.3.2 Proof of Lemma 2*

The proof of the first statement follows that of Lemma 2.

For the second statement, take any $f \in \mathcal{P}_n$, $n = 1, 2, \ldots, N$. By definition, for every $\gamma \in \Gamma_{\mathcal{P}}$, we have $\left| W^{\#}_{\gamma(f)} \right| < Q(\mathcal{F} \backslash \bigcup_{k=0}^{n-1} \mathcal{P}_k)$. This means

$$
\begin{aligned}
u_f\big(\,\widehat{m}(\theta, \gamma), \theta\,\big) &= \max_{W \subseteq \mathcal{W} \backslash W^{\#}_{\gamma(f)}, |W| \leq q_f} \chi(f, W, \theta) \\
&\geq \min_{W' \subseteq \mathcal{W}, |W'| \leq Q(\mathcal{F} \backslash \bigcup_{k=0}^{n-1} \mathcal{P}_k)} \; \max_{W \subseteq \mathcal{W} \backslash W', |W| \leq q_f} \chi(f, W, \theta) \\
&= \overline{u}_f(\theta), \quad \forall \theta \in \Theta,
\end{aligned}
$$

where the last equality comes from the definition of $\mathcal{P}_n$. Taking expectation over Θ gives $\mathbb{E}_\pi\big[u_f\big(\,\widehat{m}(\theta, \gamma), \theta\,\big)\big] = \overline{u}^*_f$, which suffices for the second statement to hold.

References

Abreu, Dilip, Prajit K Dutta, and Lones Smith, "The folk theorem for repeated games: A NEU condition," *Econometrica*, 1994, *62* (4), 939–948.

Ali, S Nageeb and Ce Liu, "Coalitions in repeated games," 2020. Working Paper.

Altınok, Ahmet, "Dynamic many-to-one matching," 2020. Working Paper.

Ashlagi, Itai, Mark Braverman, and Avinatan Hassidim, "Stability in large matching markets with complementarities," *Operations Research*, 2014, *62* (4), 713–732.

Azevedo, Eduardo M and John William Hatfield, "Existence of equilibrium in large matching markets with complementarities," 2018. Working Paper.

Blackwell, David, "Discounted dynamic programming," *Annals of Mathematical Statistics*, 1965, *36* (1), 226–235.

Che, Yeon-Koo, Jinwoo Kim, and Fuhito Kojima, "Stable matching in large economies," *Econometrica*, 2019, *87* (1), 65–110.

Corbae, Dean, Ted Temzelides, and Randall Wright, "Directed matching and monetary exchange," *Econometrica*, 2003, *71* (3), 731–756.

Damiano, Ettore and Ricky Lam, "Stability in dynamic matching markets," *Games and Economic Behavior*, 2005, *52* (1), 34–53.

Doval, Laura, "Dynamically stable matching," *Theoretical Economics*, 2022, *17* (2), 687–724.

Du, Songzi and Yair Livne, "Rigidity of transfers and unraveling in matching markets," 2016. Working Paper.

Echenique, Federico and M Bumin Yenmez, "A solution to matching with preferences over colleagues," *Games and Economic Behavior*, 2007, *59* (1), 46–71.

Fudenberg, Drew and Eric Maskin, "The folk theorem in repeated games with discounting or with incomplete information," *Econometrica*, 1986, pp. 533–554.

_ **and** _ , "On the dispensability of public randomization in discounted repeated games," *Journal of Economic Theory*, 1991, *53* (2), 428—438.

Hatfield, John William and Fuhito Kojima, "Matching with contracts: Comment," *American Economic Review*, 2008, *98* (3), 1189–1194.

_ **and Paul R Milgrom**, "Matching with contracts," *American Economic Review*, 2005, *95* (4), 913–935.

_ **, Scott Duke Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp**, "Stability and competitive equilibrium in trading networks," *Journal of Political Economy*, 2013, *121* (5), 966–1005.

Kadam, Sangram V and Maciej H Kotowski, "Multiperiod matching," *International Economic Review*, 2018, *59* (4), 1927–1947.

_ **and** _ , "Time horizons, lattice structures, and welfare in multi-period matching markets," *Games and Economic Behavior*, 2018, *112*, 1–20.

Kelso, Alexander S and Vincent P Crawford, "Job matching, coalition formation, and gross substitutes," *Econometrica*, 1982, *50* (6), 1483–1504.

Kojima, Fuhito, Ning Sun, and Ning Neil Yu, "Job matching under constraints," *American Economic Review*, 2020, *110* (9), 2935–2947.

_ **,** _ **, and** _ , "Job matching with subsidy and taxation," *Review of Economic Studies*, 2023, p. rdad032.

_ **, Parag A Pathak, and Alvin E Roth**, "Matching with couples: Stability and incentives in large markets," *Quarterly Journal of Economics*, 2013, *128* (4), 1585–1632.

Kotowski, Maciej H, "A perfectly robust approach to multiperiod matching problems," 2020. Working Paper.

Krueger, Alan B and Orley Ashenfelter, "Theory and evidence on employer collusion in the franchise sector," *Journal of Human Resources*, 2022, *57* (S), S324–S348.

Kurino, Morimitsu, "Credibility, efficiency, and stability: A theory of dynamic matching markets," *Japanese Economic Review*, 2020, *71*, 135–165.

Liu, Ce, "Stability in repeated matching markets," *Theoretical Economics*, 2023, *18*, 1711–1757.

Mauleon, Ana, Vincent J Vannetelbosch, and Wouter Vergote, "Von Neumann–Morgenstern farsightedly stable sets in two-sided matching," *Theoretical Economics*, 2011, *6* (3), 499–521.

Nguyen, Thanh and Rakesh Vohra, "Near-feasible stable matchings with couples," *American Economic Review*, 2018, *108* (11), 3154–3169.

Pęski, Marcin, "Tractable model of dynamic many-to-many matching," *American Economic Journal: Microeconomics*, 2022, *14* (2), 1–43.

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

_ **and M Bumin Yenmez**, "Matching with externalities," *Review of Economic Studies*, 2023, *90* (2), 948–974.

Rostek, Marzena and Nathan Yoder, "Matching with complementary contracts," *Econometrica*, 2020, *88* (5), 1793–1827.

Roth, Alvin E and Marilda Sotomayor, "Two-sided matching," *Handbook of Game Theory with Economic Applications*, 1992, *1*, 485–541.

Sorin, Sylvain, "On repeated games with complete information," *Mathematics of Operations Research*, 1986, *11* (1), 147–160.

Sun, Ning and Zaifu Yang, "Equilibria and indivisibilities: Gross substitutes and complements," *Econometrica*, 2006, *74* (5), 1385–1402.