Journal of Dynamics and Games

doi:10.3934/jdg.2025039

DYNAMIC MATCHING GAMES: STATIONARY EQUILIBRIA UNDER VARYING COMMITMENTS

Nadia Guiñazú, Pablo Neme* and Jorge Oviedo

Instituto de Matemática Aplicada San Luis,
Universidad Nacional de San Luis and CONICET, Argentina

(Communicated by Elvio Accinelli)

Abstract. This paper examines equilibria in dynamic two-sided matching games, extending Gale and Shapley's foundational model to a non-cooperative, decentralized, and dynamic framework. We focus on markets where agents have utility functions and commitments vary. Specifically, we analyze a dynamic matching game in which firms make offers to workers in each period, considering three types of commitment: (i) no commitment from either side, (ii) firm commitment, and (iii) worker commitment.

Our main contributions are threefold: (i) we show that stable matchings can be supported as stationary equilibria under different commitment scenarios, depending on the strategies adopted by firms and workers; (ii) we characterize the conditions under which agents are willing to switch partners, highlighting the role of discount factors in shaping equilibrium outcomes; and (iii) we provide a unified framework that connects dynamic incentives with classical stability, bridging the gap between cooperative and non-cooperative approaches to matching.

1. Introduction.

Since the influential work of [9], two-sided matching theory has focused on the idea of *stability*: a matching is stable if every agent is paired with an acceptable partner and no firm–worker pair would prefer to match with each other rather than stay with their current partners [17]. Most of the literature studies *centralized* settings, where a mechanism—such as the deferred-acceptance algorithm used by the National Resident Matching Program (NRMP)—collects participants' preferences and computes a stable matching.

However, many real-world markets are *decentralized*: firms and workers interact directly, make offers, and negotiate without a central coordinator. This raises two key questions: (i) can decentralized interactions lead to stable outcomes? and (ii) how do frictions like limited commitment or impatience affect the long-run result?

To study these questions, we extend the classical static model to an infinite-horizon, non-cooperative game. In the static setting, agents make one-time decisions

---

2020 *Mathematics Subject Classification.* Primary: 90C05, 91B68.

Key words and phrases. Dynamic matching game, stationary equilibria, commitment, stable matchings.

We would like to thank the anonymous referees for their valuable comments. We acknowledge the financial support from UNSL through grants 032016 and 030320, from Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET) through the grant PIP 112-200801-00655, and from Agencia Nacional de Promoción Científica y Tecnológica through grant PICT 2017-2355.

*Corresponding author: Pablo Neme.

and cannot respond to earlier outcomes. In contrast, the *dynamic matching game* takes place over discrete periods: in each period, firms make offers to workers, who decide whether to accept or reject them. Matches can be formed and later dissolved, so current decisions affect the set of options available in the future. We consider different commitment structures and study how agents' *patience*—that is, how much they value future opportunities—impacts whether the matchings that emerge in equilibrium are stable.

Because the same strategic situation can arise repeatedly over time, we focus on subgame-perfect Nash equilibria in *stationary* strategies. A stationary strategy depends only on the current situation—who is matched and what offers are available—not on past events or the period number. This assumption is reasonable in anonymous markets and simplifies the analysis, while still capturing the key dynamic aspects. Stationary equilibria thus offer a natural framework to study stability in decentralized matching over time.

A crucial aspect of dynamic models is the degree of *commitment* that agents can maintain with their partners over time. In many real-world applications, commitment is not symmetric across the two sides of the market. For example, in school systems, teachers often hold tenured positions, suggesting that schools are bound by long-term commitments. In contrast, in professional sports, athletes typically sign binding contracts with teams, shifting the burden of commitment to the workers. These examples illustrate how the nature of contractual commitment can vary across markets and highlight the importance of understanding how such asymmetries shape agents' behavior in dynamic environments.

Before delineating the game, it is essential to define what it means for an agent to be *active* or *inactive* in a given period. A worker is considered active in a period if she does not hold a commitment with the agent she was matched with within the previous period or if she was unmatched in the previous period. Similarly, we define when a firm is active or inactive in a given period.

The dynamic game analyzed in this paper is structured as a *two-stage game* for each period. In the first stage, firms simultaneously make offers to at most one worker. An active firm can extend an offer to any worker, while an inactive firm must retain its current employee. Firms remain unaware of the offers made by other firms in that period, but they are cognizant of the matching history from previous periods. Inactive firms are permitted to make offers to their current employees as well.

In the second stage, each worker privately observes the offers received in the first stage, including the renewal offer from their current employer, if applicable. Workers do not know the offers made to others, but they are aware of the matchings from prior periods. Each worker can accept at most one offer and may choose to reject all offers. An active worker can accept any offer, while an inactive worker must accept the renewal offer from their current employer.

Given the dynamic nature of the game and the assumption that agents do not base their decisions on *beliefs* about the history of play—that is, agents do not hold beliefs about the past actions of others—we adopt the concept of stationary equilibrium as our solution concept. Recall that a strategy profile (i.e., one strategy for each agent) constitutes a stationary equilibrium if no agent can profitably deviate, considering only the subgame involving the currently active agents.

In dynamic games, the *commitment* of agents plays a critical role. In this paper, we consider three scenarios: (i) when neither side of the market holds commitment,

(ii) when only firms hold commitment, and (iii) when only workers hold commitment.

When both sides of the market possess commitment, the dynamic game reduces to a two-stage game: firms make offers, and workers either accept or reject them—either because the offer is unacceptable or because they receive multiple offers—thus ending the game. This simplification relies on the assumption that agents do not form beliefs about the actions of others. In particular, workers cannot anticipate or strategize based on potential future offers. As a result, they have no incentive to reject an acceptable offer in the hope of receiving a better one later. This assumption is consistent with environments where agents face informational constraints or cannot coordinate expectations, and it is further reinforced when agents discount the future—making current acceptable offers more attractive than uncertain future opportunities. Under this setting, several studies have shown that stable matchings can be supported as equilibrium outcomes [1, 2].

In the case where only one side of the market holds commitment, we present two stationary equilibria arising from two different scenarios, depending on the firms' strategies: one in which firms behave in a strategically constant manner, and another in which firms adopt responsive strategies. In both cases, workers follow the same strategy: in each period, they accept their most preferred available option as the best possible response. We refer to the resulting equilibria as the *stationary equilibrium in constant strategies* and the *stationary equilibrium in responsive strategies*, respectively. In the scenario where firms hold commitment (i.e., they offer tenured positions to workers), given a stable matching, if firms make offers to workers assigned under the stable matching and each worker accepts the offer yielding the highest utility, this (constant) strategy constitutes a stationary equilibrium that supports the stable matching (Theorem 3.3).

Now, consider the first stage in which firms make offers to the workers assigned to them under the stable matching, and each worker accepts the offer. In this scenario, firms behave strategically responsive: they are allowed to make new offers in subsequent periods if they remain active. Given a stable matching, if a worker seeks to improve her situation, she must resign from her current position with the firm she is assigned to under the stable matching and wait for a new offer. Due to the nature of firms' commitment, this worker must wait for a new offer, which incurs a cost referred to as the *discount factor*. [5] present a re-stabilization process in which a better offer eventually arrives, in expectation, for any worker who resigns. They also compute the expected time until such an offer is received. Once a worker decides to improve her situation and resigns from a firm, that firm becomes active. Since each active firm aims to maximize its utility, workers will only resign if the new offer yields higher utility, considering their discount factors; that is, their final utility is positive despite the discount factor. Thus, we can establish a threshold for discount factors such that if these factors exceed this threshold, it will not be worthwhile for workers to resign and await new offers. In this case, we prove that the proposed (responsive) strategy constitutes a stationary equilibrium of the dynamic game (Theorem 3.4).

In the scenario where workers possess commitment (i.e., workers sign binding contracts with firms), we demonstrate that, given a stable matching, two stationary equilibria support it, depending on the patience of firms in making offers. In this case, we also analyze two distinct firms' strategies: one that remains strategically constant and another that is strategically responsive. In the strategically constant case, similar to the scenario where firms hold commitment, we assume that firms' strategies are always to offer to the worker employed under the stable matching. Workers accept the offer that provides them with the highest utility . In the strategically responsive case, firms initially offer to the worker assigned under the stable matching. In subsequent stages, if firms observe that the stable matching was maintained in the previous period, they renew their offers to the same workers. Conversely, if firms observe that the stable matching was not formed in the prior period, they make offers to the workers employed under the firm-optimal stable matching.[1] Both possible firm strategies—constant and responsive—along with the workers' strategy of always accepting the offer that yields the highest utility, result in two stationary equilibria that support the same stable matching (Theorem 3.5 and Theorem 3.6, respectively). Note that the number of equilibria in all possible scenarios strictly depends on the number of stable matchings in the game.

Related Literature. The closest paper to ours is that of [6], which analyzes decentralized matching markets—such as labor markets or school admissions—modeled as infinitely repeated dynamic games, where firms with vacancies make offers and workers choose which ones to accept. Unlike classical centralized models involving a planner or one-shot Gale–Shapley mechanisms, their analysis focuses on stationary equilibria in settings where agents may or may not be able to commit to long-term relationships. A central contribution of the paper is to show that, even in the absence of commitment, stable matchings can be supported as outcomes of stationary equilibria, thereby implementing the core of the market. A key feature of their analysis is that equilibrium outcomes depend critically on agents' beliefs about others' future strategies and behavior. For instance, a firm might choose not to make an offer to its top candidate—not because it does not value the match, but because it believes the worker will reject it in anticipation of a better offer in future rounds. These strategic expectations shape both the dynamics of the offer process and the resulting equilibria. The paper also distinguishes between unilateral and bilateral commitment, showing that both can lead to inefficient or unstable matchings, depending on how beliefs affect incentives. In contrast, we assume that agents do not form beliefs about others' future actions. As a result, behavior in our model is not shaped by expectations or anticipatory reasoning. This assumption allows us to guarantee that a stable matching supports each equilibrium.

Similarly, [10] examine a dynamic game where payments depend solely on the final matching. In contrast to their approach, this paper considers a model in which matchings are formed in each period, and agents accumulate payoffs as the game progresses, highlighting the significance of both the duration of matchings and the final matching.

Other key works include those by [11] and [14], who also study dynamic models by introducing factors such as wages and the duration of offers. However, these studies differ in that agents only observe the final matching, whereas our approach considers a continuous process of matching formation over time.

[4] develop an algorithm that identifies stable matchings when some agents are already paired. This provides an interesting perspective on how prior matchings can affect the dynamics of offers in repeated markets.

[1] [9] define the **firm-optimal stable matching** as the one that gives the highest utility for all firms among every other stable matching.

When considering dynamic matching problems with variable populations, a related paper is [15]. Using the teacher-proposing deferred acceptance mechanism shows the existence of a dynamically stable matching that (i) minimizes unjustified claims and (ii) Pareto-dominates any other dynamically stable matching.

More recently, [3] build on the framework of [15] by analyzing a dynamic many-to-many school choice model where the population of teachers varies across periods. Teachers may enter or leave the market, hold tenured positions (school commitment), and are allowed to work in multiple schools simultaneously. The authors introduce a new notion of dynamic stability tailored to this setting and propose the TRDA mechanism, which always produces a dynamically stable matching that is constrained-efficient and minimizes unjustified claims. To improve efficiency beyond this class, they develop the TREADA mechanism, which yields efficient outcomes under teacher consent. They also analyze manipulability and show that both mechanisms are immune to obvious dynamic manipulations under suitable assumptions on school priorities.

Following the line of research on dynamic markets with changing populations, several recent papers explore alternative formulations where agents arrive over time. [7] introduces a different notion of dynamic stability in two-sided, one-to-one markets that addresses timing frictions and externalities from unmatched agents, ensuring timely participation and stable outcomes. [13] propose the Dynamic Core for dynamic one-sided markets and design the Intertemporal Top Trading Cycle algorithm, which finds Pareto-efficient and group strategy-proof allocations. Building on these contributions, [18] refine dynamic stability further by incorporating agents' rationalizable conjectures, defining a new notion that strengthens Doval's and is always non-empty.

Finally, [8] introduces a revelation principle for single-agent dynamic mechanism design in Markov environments, where the agent's private information evolves over time and the designer can only commit to short-term mechanisms. The paper shows that any equilibrium payoff can be implemented using flow-direct Blackwell mechanisms, which align truthful reporting with equilibrium beliefs. This result streamlines the design of optimal mechanisms in dynamic environments, including dynamic Mirrlees and social insurance models.

The paper is structured as follows: In Section 2, we present the preliminaries, the two-stage game, and the solution concept. Section 3 analyzes various commitment scenarios and their corresponding stationary equilibria. Finally, Section 4 offers concluding remarks.

2. **Preliminaries.**

We consider two disjoint finite sets of agents, the set of firms F and the set of workers W. An agent refers to either a firm or a worker. Each agent i has a utility function u_i such that

$$u_f : W \cup \{f\} \to \mathbb{R} \text{ for each } f \in F,$$

$$u_w : F \cup \{w\} \to \mathbb{R} \text{ for each } w \in W.$$

We denote by $u_i(j)$ to agent i's utility of being matched to agent j. Here, $u_i(i)$ is the agent's utility of being unmatched, and we ask for utilities to fulfill that $u_i(i) = 0$ for each i. If $u_i(j) \geq 0$ we say j is acceptable to i. We assume "strict" utilities: $u_i(j) = u_i(k)$ only if $j = k$.

A matching is then a mapping $\mu : F \cup W \to F \cup W$ such that: (i) $\mu(f) \in W \cup \{f\}$ for each $f \in F$, (ii) $\mu(w) \in F \cup \{w\}$ for each $w \in W$, and (iii) $\mu(i) = j$ if and only if

$\mu(j) = i$ for each $i, j \in F \cup W$. Here, $\mu(i)$ denotes the agent with whom i is matched and $\mu(i) = i$ means that agent i is unmatched. Let $\mu_\emptyset$ denote the matching in which no one is matched. A matching is individually rational if it asings to each agent i a non-negative utility $(u_i(\mu(i)) \geq 0)$. We say that a matching μ is blocked by a pair $(f, w) \in F \times W$ if

$$u_f(w) > u_f(\mu(f)), \text{ and}$$
$$u_w(f) > u_w(\mu(w)).$$

This means that f and w obtain higher utility from being matched with each other than from being matched with their respective partners under μ. We say that a matching μ is stable if it is individually rational and has no blocking pair. The stable matching set is always non-empty. Moreover, there is always a stable matching that all firms agree to be the one that provides the highest utility, called the firm-optimal stable matching and denoted by μ_F. Conversely, there is always a stable matching that all workers agree to be the one that provides the highest utility, called the worker-optimal stable matching, and denoted by μ_W, [9].

In this paper, we consider a "decentralized" dynamic matching game. To formally define a dynamic matching game, we first define the payoff in each period.

We consider discrete periods: $t = 1, 2, 3, \ldots$. In each period, agents derive a payoff from the realized matching. The period-payoff function for agent i is the utility function u_i, which is invariant across periods. Each agent i maximizes the discounted sum of period payoffs,

$$\hat{u}_i = \sum_{t=1}^{\infty} \delta_i^{t-1} u_i(\mu^t(i)),$$

where $\delta_i \in (0, 1)$ is the discount factor, and μ^t is the realized matching in period t. The discount factor δ_i captures the patience of agent i: the agent is considered patient if δ_i is close to 1, and impatient if δ_i is close to 0.

Now, we describe how μ^t is defined. In order to do so, note that the matching μ^{t-1} observed by each agent will determine, at the beginning of each period t, the set of firms and workers who cannot move in period t depending on the type of commitment considered. We assume $\mu^0 = \mu_\emptyset$, i.e., no one is matched before the initial period.

We denote by $F_c(\mu^{t-1}) \subseteq F$ to the set of inactive firms at period t. These firms have committed themselves to their employees in μ^{t-1}. This means that at period t, the firms in $F_c(\mu^{t-1})$ can neither fire their employees nor hire new ones. That is, their current employees have tenure and their jobs are protected. Conversely, we refer to the firms in $F \setminus F_c(\mu^{t-1})$ as active firms. These firms do not hold a commitment to their current employees, so they can fire their employees if they have any. Similarly, we denote by $W_c(\mu^{t-1}) \subseteq W$ the set of inactive workers at period t. This means that at period t, the workers in $W_c(\mu^{t-1})$ cannot switch their employers. Thus, we refer to the workers in $W \setminus W_c(\mu^{t-1})$ as active workers. These workers do not hold a commitment to their employers and can leave their positions if employed.

Note that, depending on the characteristics of the sets F_c and W_c at a period t, there are three different scenarios to consider:

Case 1: Two-sided commitment. All matched agents are inactive:

$$F_c(\mu^{t-1}) = \{f \in F : \mu^{t-1}(f) \neq f\}, \text{ and}$$
$$W_c(\mu^{t-1}) = \{w \in W : \mu^{t-1}(w) \neq w\}.$$

This means that once a firm and a worker are matched, they stay so permanently.

Case 2: No commitment. All firms and workers are active regardless of the previous matching: $F_c(\mu^{t-1}) = W_c(\mu^{t-1}) = \emptyset$. This means that firms can fire their employees, and workers can leave their current employers.

Case 3: One-sided commitment. In this case, despite the market being symmetric, since only firms are making offers, we must analyze two subcases: when only firms hold commitment, and when only workers hold commitment. If only firms hold commitment, all matched firms are inactive, while all workers remain active:

$$F_c(\mu^{t-1}) = \{f \in F : \mu^{t-1}(f) \neq f\}, \text{ and}$$
$$W_c(\mu^{t-1}) = \emptyset.$$

Workers cannot be fired but may switch to other firms when receiving new offers. When only workers hold commitment, all matched workers are inactive, while all firms remain active:

$$F_c(\mu^{t-1}) = \emptyset, \text{ and}$$
$$W_c(\mu^{t-1}) = \{w \in W : \mu^{t-1}(w) \neq w\}.$$

Workers must honor their employment contracts, even when receiving better offers. However, firms are allowed to lay off their workers in order to hire new ones, eventually.

Now, we are in a position to describe each period of the dynamic matching game. Each period will be decomposed into a two-stage game.

First stage : Each firm simultaneously makes an offer to at most one worker. An active firm can make an offer to any worker while an inactive firm has no option but to keep its employee under μ^{t-1}. Firms do not observe any offer made by another firm in the current period, but each firm observes the matching realized in previous periods. We consider that inactive firms make new offers to their current employees. The action of firm f, denoted by o_f, must fulfill that (i) $o_f \in W \cup \{f\}$ if $f \in F \setminus F_c(\mu^{t-1})$ and (ii) $o_f = \mu^{t-1}(f)$ if $f \in F_c(\mu^{t-1})$, where $o_f = f$ means that f makes no offer to any worker. Furthermore, we denote by $O_f(\mu^{t-1})$ to the set of all possible offers that firm f can make in period t given μ^{t-1}.

Second stage : Each worker w privately observes the offers made to her in the first stage, denoted $O_w = \{f \in F : o_f = w\}$. Recall that O_w includes the renewal offer from the current employer if w has tenure. Workers do not observe any offer made to other workers in the current period, but each worker observes the entire matching realized in previous periods. Thus, each worker simultaneously accepts at most one offer. An active worker w can accept any offer or reject all offers. Inactive workers have no choice but to accept the renewal offers from their current employers. Thus, worker w's response, denoted by r_w, must fulfill that (i) $r_w \in O_w \cup \{w\}$ if $w \in W \setminus W_c(\mu^{t-1})$ and (ii) $r_w = \mu^{t-1}(w)$ if $w \in W_c(\mu^{t-1})$. Furthermore, we denote by $R_w(\mu^{t-1}, O_w)$ to the set of admissible responses for w.

Then, given the actions of firms and workers, the matching in period t, denoted by μ^t, is defined by $\mu^t(w) = r_w$ for each $w \in W$, $\mu^t(f) = w$ if $f = r_w$, and $\mu^t(f) = f$ if $f \neq r_w$. Thus, by definition, μ^t is individually rational at each period t. Therefore, a dynamic matching game is given by $(F, W, (u_i, \delta_i)_{i \in F \cup W}, F_c, W_c)$.

Another two important concepts that we need to present used in our results are the histories and the strategies of the agents. A history at the beginning of period t is an ordered list of past actions, given by

$$h_t = \left((o_f^\tau)_{f\in F}, (r_w^\tau)_{w\in W}\right)_{\tau=1}^{t-1},$$

where o_f^τ is the offer made by firm f in period $\tau = 1,\ldots,t-1$ and r_w^τ is the reply of worker w in period τ. After the first stage of period t, a history is given by $(h^t, (o_f^t)_{f\in F})$, where h^t is a history at the beginning of this period and $(o_f^t)_{f\in F}$ is the profile of offers made in this period.

The profile of replies $(r_w^\tau)_{w\in W}$ in h^t contains the same information as the realized matching μ^t, which becomes public information. Since offers are private information, players do not have complete information about the history. Each player observes only his private history, defined as follows. A private history for firm f in period t is an ordered list $h_f^t = (\mu^0 = \mu_\emptyset, o_f^1, \mu^1, o_f^2, \mu^2, \ldots, o_f^{t-1}, \mu^{t-1})$, where μ^τ is the matching realized in period τ. While μ^τ is public information, o_f^τ is private information. Let H_f^t denote the set of private histories for f in period t. Let

$$H_f = \bigcup_{t=1}^{\infty} H_f^t$$

denote the set of all private histories for f.

A (pure) strategy of firm f is a function $\sigma_f : H_f \to W \cup \{f\}$ such that for each $h_f^t \in H_f, \sigma_f(h_f^t) \in O_f(\mu^{t-1})$, where μ^{t-1} is the last entry of h_f^t. Similarly, a private history for worker w in the middle of period t (when she makes a decision) is an ordered list $h_w^t = (\mu^0 = \mu_\emptyset, O_w^1, r_w^1, \mu^1, O_w^2, r_w^2, \mu^2, \ldots, O_w^{t-1}, r_w^{t-1}, \mu^{t-1}, O_w^t)$, where O_w^τ is the set of offers made to w in period τ (including a renewal offer if any) and r_w^τ is her reply in that period. Let H_w^t denote the set of all private histories for w in period t. Let

$$H_w = \bigcup_{t=1}^{\infty} H_w^t$$

denote the set of all private histories for w. A strategy of worker w is then a function $\sigma_w : H_w \to F \cup \{w\}$ such that, for each $h_w^t \in H_w$,

$$\sigma_w(h_w^t) \in R_w(\mu^{t-1}, O_w^t),$$

where μ^{t-1} and O_w^t are the last two entries of h_w^t.[2] A strategy profile $\sigma = (\sigma_i)_{i\in F\cup W}$ determines the payoff for each agent in the dynamic game. Given a dynamic matching game $(F, W, (u_i, \delta_i)_{i\in F\cup W}, F_c, W_c)$ and a strategy profile σ, μ_σ is the matching resulting of playing the dynamic matching game when agents declare the strategy profile σ, and $\hat{u}_i(\mu_\sigma)$ denotes the utility of each agent i. Moreover, the utility for agent i can be computed as

$$\hat{u}_i(\mu_\sigma) = \sum_{t=1}^{\infty} \delta_i^{t-1} u_i(\mu^t(i)).$$

Given the nature of the game, it is necessary to incorporate information from previous stages. For this reason, we consider the notion of subgame-perfect equilibrium. In our setting, the relevant history includes the matching from the previous period, the offers made and received in earlier stages, and, in the case of workers,

[2] Note that firms' strategies correspond to the offers they make, while workers' strategies consist of acceptance or rejection decisions in response to the offers they receive.

the offers currently available. Based on this information, the game at stage t, along with its corresponding history, can be viewed as a game in itself. Payoffs depend on the sequence of matchings over time, and we introduce a discount factor to capture the dynamic nature of the environment. Strategies are defined as mappings from histories to actions—offers or responses—depending on the agent's role.

Accordingly, we adopt the concept of subgame-perfect Nash equilibrium, which requires that agents' strategies form a Nash equilibrium in every subgame that follows any possible history.

Moreover, since we assume that agents hold no beliefs about the actions of others, we restrict attention to subgame-perfect equilibria in *stationary strategies*, that is, strategies that depend only on the payoff-relevant state of the game and not on agents' beliefs about past behavior. By a *belief*, we refer to an agent's subjective assessment of which actions may have been taken by others at earlier stages of the game, particularly in settings where such actions are not directly observed.

As a preliminary step toward defining a subgame-perfect Nash equilibrium, we first recall the formal definition of a Nash equilibrium.

Definition 2.1. A strategy profile $\sigma^\star$ is a **Nash equilibrium** if $u_i(\mu_{\sigma^\star}) \geq u_i(\mu_{\sigma_i,\sigma^\star_{-i}})$ for each agent i and each strategy σ_i .[3]

Given a dynamic matching game, a subgame that starts in time $\widetilde{t}$ is given by

$$(F, W, (u_i, \delta_i)_{i\in F\cup W}, F_c, W_c)_{t=\widetilde{t}}^{\infty}.$$

Note that there are as many subgames starting at time $\widetilde{t}$ as there are histories at time $\widetilde{t}-1$. A subgame is a part of the (original) dynamic matching game that starts at a point where the game's history up to that point is public knowledge among all agents and includes all subsequent decisions in the original game from that point onward. Now, we are in a position to present the definition of a subgame-perfect Nash equilibrium formally.

Definition 2.2. A strategy $\sigma^\star$ is a **subgame-perfect Nash equilibrium** if for each $\widetilde{t}$, $\sigma^\star$ is a Nash equilibrium for the subgame $(F, W, (u_i, \delta_i)_{i\in F\cup W}, F_c, W_c)_{t=\widetilde{t}}^{\infty}$.

Since we do not rely on the full history of the game to analyze its continuation, we employ stationary strategies to address the dynamic problem more effectively. A crucial element in defining stationary strategies is the concept of continuation-equivalent games.

For an agent to decide which strategy to play at a given time, she will only consider the matching generated in the previous period. Therefore, if two different strategies result in equivalent matchings, the same strategy will be employed in the subsequent period in both cases. Different matchings can induce distinct continuation-equivalent games depending on what commitment types agents possess. Given $\mu, \mu' \in \mathcal{M}$, the equivalence relation $\sim$ depends on the commitment structure of the game, and is defined as follows:

- In the case of bilateral commitment, two matchings are continuation-equivalent if and only if the set of unmatched agents is identical. Formally, $\mu \sim \mu'$ if and only if $\{j \in F \cup W : \mu(j) = j\} = \{j \in F \cup W : \mu'(j) = j\}$. Agents who have been matched cannot change their partner for the rest of the game.

[3] Here, $\mu_{\sigma_i,\sigma^\star_{-i}}$ denotes the matching resulting of playing the dynamic matching game when agent i declare strategy σ_i and the rest of agent j declare strategy $\sigma^\star_j$.

- In the no-commitment case, all matchings are continuation-equivalent. Formally, $\mu \sim \mu'$ for any two matchings μ and μ'. In terms of commitment, the continuation of the game is the same regardless of what happened in the previous period.
- In the case of unilateral commitment, no two different matchings are continuation-equivalent. Formally, $\mu \sim \mu'$ if and only if $\mu = \mu'$. Even if the set of matched agents is the same, the continuation of the game depends on how the agents are currently matched.

Given a preference relation, we can define stationary strategies as follows. The strategy σ_f of firm f is stationary if for any pair of private histories $h_f = (\ldots, \mu)$ and $h'_f = (\ldots, \mu')$ (possibly with different lengths) we have that $\mu \sim \mu'$, then $\sigma_f(h_f) = \sigma_f(h'_f)$. For workers' strategies, there is an additional requirement that the set of offers received in the current period is also identical. That is, the strategy σ_w of worker w is stationary if for any pair of private histories $h_w = (\ldots, \mu, O_w)$ and $h'_w = (\ldots, \mu', O'_w)$, if $\mu \sim \mu'$ and $O_w = O'_w$, then $\sigma_w(h_w) = \sigma_w(h'_w)$.

Definition 2.3. A **stationary equilibrium** is a subgame-perfect Nash equilibrium in which all strategies are stationary.

3. **Stationary equilibria in dynamic games.** In this section, we present the main results of this paper. This section consists of three subsections. In Subsection 3.1, we consider a game where neither firms nor workers hold commitment. In this case, we show that any stable matching is the result of a stationary equilibrium. In Subsection 3.2, we consider a market in which only firms hold commitment. In this case, we show that, given a stable matching, two different equilibria yield this matching as the outcome of the game: one in which firms are strategically constant and another in which they are strategically responsive based on their observations, and workers' discount factors. In Subsection 3.3, we analyze a market in which only workers hold commitment. Here, we similarly show that a stable matching arises from two distinct equilibria, following the same firms' strategic patterns as in the previous case.

Note that when both sides of the market hold commitment, the dynamic game reduces to a one-period static game in which firms make offers, workers either accept or reject, and the game then ends. In this setting, stable matchings are also supported by equilibria [1, 2, for more details].

For the remainder of the paper, we focus on two distinct types of strategic behavior that firms may adopt over time. These differ in how firms make decisions across periods, particularly in their responsiveness to past outcomes.

In the first type, which we refer to as *strategically constant*, firms follow a fixed strategy across all periods. That is, a firm's behavior remains unchanged from one period to the next, regardless of the outcome in the previous period. This type of strategy models firms that are either rigid in their decision-making or that do not update their behavior in response to market dynamics.

In the second type, which we call *strategically responsive*, firms adapt their behavior based on the outcome observed in the previous period. This formulation allows firms to adjust their strategy over time in response to the evolution of the market, for instance by reacting to the match they obtained, the offers made, or the behavior of the workers.

Accordingly, when firms behave strategically constant, any stationary equilibrium that arises will be referred to as a *stationary equilibrium in constant strategies*.

In contrast, when firms behave strategically responsive, the resulting stationary equilibrium will be referred to as a *stationary equilibrium in responsive strategies.*

3.1. **The dynamic game without commitments.** In this subsection, we study the case where no one has any commitment. When there is no commitment, the history leading to the current period does not change the continuation of the game and is therefore ignored by the agents in stationary equilibria. In other words, what happens in the current period does not affect future outcomes. Due to this independence, agents ignore the future and behave as in the static model.

The following result indicates that, in the absence of commitment, the static notion of stability captures the outcome of stationary equilibrium.

Theorem 3.1. *Given a dynamic matching game without commitment, the matching achieved in any stationary equilibrium is stable in all periods. Conversely, for any stable matching, there exists a stationary equilibrium that achieves this matching in every period.*

Proof. Consider any stationary equilibrium σ. The strategies of the firms $(\sigma_f)_{f\in F}$ in the equilibrium are stationary, i.e., firms always make the same offer ($o^t_f = o^{t'}_f$ for any periods t, t'). The action of the workers, that is, their response in any period, does not affect the offers they will receive in subsequent periods because they do not commit. This implies that the best decision each worker w can make is to accept the offer that provides the highest utility in each period. Thus, if the resulting matching of the equilibrium is μ, then it is the same in each period, i.e., $\mu = \mu^t$ for each t.

Suppose that equilibrium σ results in an unstable matching μ. Since μ is individually rational, let (f, w) be a blocking pair for μ. Since w blocks μ together with f, it follows that $u_w(f) > u_w(\mu(w))$. Note that, w's strategy is always to accept the best offer. This means she does not receive an offer from f; otherwise, she would accept it.

Suppose that in some period t, firm f deviates from the equilibrium σ and makes an offer to w; that is, $\hat{o}^t_f = \{w\}$. By the previous observation, the worker w will accept the offer and be matched with f. Then, the firm obtains a higher profit by deviating and making an offer to its blocking pair, contradicting that the strategy σ is an equilibrium.

Conversely, let us choose a stable matching μ and consider the following profile of strategies σ such that:

(i) Each firm f makes an offer to $\mu(f)$; that is, $o^t_f = \{\mu(f)\}$ for each t.
(ii) Each worker w accepts the offer that provides the highest utility.

Note that the strategies are stationary since firms always make the same offer and workers accept the best among the received offers.

The workers' strategies are optimal because without commitment, accepting the offer that provides the highest utility will not affect future choices. The firms' strategies are also optimal because if a firm f makes an offer to a worker $w \neq \mu(f)$ it is because $u_f(w) > u_f(\mu(f))$. Then, for worker w it must hold that $u_w(\mu(w)) > u_w(f)$. Otherwise, (f, w) will block μ contradicting its stability. Furthermore, since w follows the equilibrium, the offer from f will be rejected. If the firm makes an offer to w such that $u_f(\mu(f)) > u_f(w)$, the offer may be accepted and thus the firm's utility decreases implying that f has no incentive to deviate. Therefore, σ is a stationary equilibrium. □

The following Corollary is a consequence of the previous theorem and the Single Agent Theorem.[4]

Corollary 3.2. *Workers and firms who do not have a match in a stationary equilibrium remain unmatched in all stationary equilibria.*

3.2. **The dynamic game when firms hold commitment.** In this subsection, we study the dynamic matching game in which firms, as the offer-makers, are committed to their decisions. Workers are offered permanent positions: they may choose to resign, but they cannot be dismissed.

In the following theorem, we show that for every stable matching, there is a stationary equilibrium in constant strategies that sustains it.

Theorem 3.3. *Given a dynamic matching game where firms hold commitment, for every stable matching μ, there is a stationary equilibrium in constant strategies that yields μ as the outcome.*

Proof. Let μ be a stable matching, and consider the following strategy profile σ at each period t:

(i) Each firm f makes an offer to $\mu(f)$; that is, $o_f^t = \mu(f)$.
(ii) Each worker w accepts the offer that provides the highest utility.

We will prove that the strategy σ is a stationary equilibrium in constant strategies.

Assume that f, an active firm in the first period, deviates from σ by making an offer such that $w \neq \mu(f)$ $(\hat{o}_f^1 = w)$, and all other firms $\hat{f} \neq f$ follow the strategy σ $(o_{\hat{f}} = \mu(\hat{f}))$. We analyze the possible deviations for firm f:

(i) $\hat{o}_f^1 = w$ and $u_f(w) > u_f(\mu(f))$:: Since μ is stable, for the worker w it must hold that $u_w(\mu(w)) > u_w(f)$. Otherwise, (f, w) would block μ contradicting its stability. As workers play σ, meaning they choose the offer that provides the highest utility, the offer from f will be rejected.

(ii) $\hat{o}_f^1 = w$ and $u_f(\mu(f)) > u_f(w)$:: Then, w might accept the offer. Since firms hold commitment, they cannot layoff any worker, and therefore, the utility of firm f will be:

$$\sum_{t=1}^{\infty} \delta_f^{t-1} u_f(w) < \sum_{t=1}^{\infty} \delta_f^{t-1} u_f(\mu(f))$$

Hence, firm f will not benefit by making offer $\hat{o}_f^1$.

(iii) $\hat{o}_f^1 = f$, meaning that f deviates by making no offer:: In the next period, the best outcome for f is to obtain $\mu(f)$, thus

$$\sum_{t=2}^{\infty} \delta_f^{t-1} u_f(\mu(f)) < \sum_{t=1}^{\infty} \delta_f^{t-1} u_f(\mu(f)).$$

Therefore, the firm does not benefit by not making offers in period 1.

Now consider the strategies of workers. Firms do not deviate, meaning $o_f = \mu(f)$ for all $f \in F$. For each $w \in W$, in every period t, we have $\mu(w) \in O_w^t$. Since μ is stable, the offer that provides the highest utility to each w is $\mu(w)$. Therefore, the

[4]The Single Agent Theorem (also known as the "Lone Wolf Theorem" or as the "Rural Hospital Theorem" in more general environments) states that if an agent is single in a stable matching, it is single in all stable matchings [12, 16].

workers' strategies are optimal, as without commitment, they accept the offer that maximizes their utility.

Thus, the strategy σ is a stationary equilibrium in constant strategies, whose outcome is μ. □

Note that in the strategy σ considered in Theorem 3.3, firms act in a strategically constant manner. This raises the following question: can stable matchings be achieved through a different type of strategy, allowing firms to adjust their behavior based on their observations, i.e., by acting strategically responsive? Fortunately, the answer is affirmative. However, before presenting the result that addresses this, we must discuss the re-stabilization process introduced by [5] and adapted to our context.

Assume that the market clears at a stable matching other than the worker-optimal one, leaving room for improvement from the workers' perspective. We can interpret each period in a dynamic matching game where firms are committed to their offers as an iteration of the re-stabilization process presented in [5]. In the first stage, each firm makes an offer to the worker assigned to them under the initial matching μ, and each worker privately observes the offers she receives. Suppose that worker w wishes to improve her labor situation and decides to adopt a strategy of resigning and waiting for a better offer. Consequently, worker w rejects all offers, while the other workers accept the offer that provides them with the highest utility. The firm left unmatched in the previous stage, and then active, makes an offer to the next worker with a higher utility who has not previously rejected it and is willing to accept it, while the other firms repeat their previous offers. This process continues until worker w receives a better offer. They observe that the stable matching in which worker w improves her outcome is the closest to μ in the lattice of stable matchings–namely, among all stable matchings that give every worker weakly higher utility than μ and assign w to a different partner than under μ, it is the one that yields the lowest utility to the workers. Under this situation, the main result presented by [5] establishes how many steps are required by the algorithm to re-stabilize the market. The length of this process reflects the time the worker must wait to secure a new position. Understanding this timeline is crucial, as workers will be unemployed during this adjustment period, which affects their decision to resign or remain in their current position. In this way, we can determine whether the strategy of resigning and waiting for a better offer constitutes a stationary equilibrium.

Let $w \in W$ and μ, ν be two stable matchings where μ is the initial matching, and ν is the re-stabilized matching when w decides to improve her labor situation, provided by the re-stabilization process presented in [5]. Denote by $k(w)$ the number of periods necessary for the firm $\nu(w)$ to make an offer to worker w provided by the results in [5]. Let c^w defined as follows:

$$c^w = \left(\frac{u_w(\mu(w))}{u_w(\nu(w))} \right)^{\frac{1}{k(w)}}.$$

The following theorem guarantees that if the discount factor of each worker w is bounded by c^w, then the stable matching μ results from a stationary equilibrium, where firms act strategically responsive based on what they observe.

Let μ be a stable matching, and consider the following strategy profile σ at each period t in the dynamic game, where each firm follows a responsive strategy, and each worker uses the same strategy as before.

(i) For $t = 1$, each firm f makes an offer a $\mu(f)$, i.e. $o_f^1 = \mu(f)$. For $t > 1$, if $\mu^{t-1}(f) = \mu(f)$, then $o_f^t = \mu(f)$. Otherwise ($\mu^{t-1}(f) = \emptyset$), then $o_f^t = \overline{w}$ where $\overline{w} \in W$ is such that $u_{\overline{w}}(f) > u_{\overline{w}}(\mu^{t-1}(\overline{w}))$ and $u_f(\overline{w}) > u_f(w')$ for each $w' \in W \setminus \{\overline{w}, \mu(f)\}$.

(ii) Each worker w accepts the offer that provides the highest utility.

Note that condition (i) for $t > 1$ states that, if firm f is rejected by its partner under μ, it makes a new offer to a worker w who provides the highest utility among those workers who would not reject it.

Theorem 3.4. *Given a dynamic matching game where firms hold commitment, let μ be a stable matching. If $\delta_w \in (0, c^w)$ for each $w \in W$, then strategy σ is a stationary equilibrium in responsive strategies that yields μ as the outcome.*

Proof. First, we prove that the strategy σ is a stationary equilibrium in responsive strategies. To do this, assume that f, an active firm in the first period, deviates from σ by making an offer to some $\widehat{w} \in W$ such that $\widehat{w} \neq \mu(f)$ ($\hat{o}_f^1 = \widehat{w}$), and all other firms $\hat{f} \neq f$ follow the strategy σ ($o_{\hat{f}} = \mu(\hat{f})$). We analyze the possible deviations for firm f:

(i) $\hat{o}_f^1 = \widehat{w}$ and $u_f(\widehat{w}) > u_f(\mu(f))$:: Since μ is stable, for the worker $\widehat{w}$ it must hold that $u_{\widehat{w}}(\mu(\widehat{w})) > u_{\widehat{w}}(f)$. Otherwise, $(f, \widehat{w})$ would block μ contradicting its stability. As workers play σ, meaning they choose the offer that provides the highest utility, the offer from f will be rejected.

(ii) $\hat{o}_f^1 = \widehat{w}$ and $u_f(\mu(f)) > u_f(\widehat{w})$:: Then, $\widehat{w}$ might accept the offer. Since firms hold commitment, they cannot layoff any worker and, therefore, the utility of firm f will be:

$$\sum_{t=1}^{\infty} \delta_f^{t-1} u_f(\widehat{w}) < \sum_{t=1}^{\infty} \delta_f^{t-1} u_f(\mu(f)).$$

Hence, firm f will not benefit by making offer $\hat{o}_f^1$.

(iii) $\hat{o}_f^1 = f$, meaning that f deviates by making no offer:: In the next period, the best outcome for f is to obtain $\mu(f)$, thus

$$\sum_{t=2}^{\infty} \delta_f^{t-1} u_f(\mu(f)) < \sum_{t=1}^{\infty} \delta_f^{t-1} u_f(\mu(f)).$$

Then, the firm does not benefit by not making offers in period 1.

Now, w.l.o.g. consider that firm f inactive in the first period but rejected at the end of that period, i.e., $\mu^1(f) = \emptyset$. Assume f deviates from strategy σ, i.e., $\hat{o}_f^2 \neq \overline{w}$. Thus, there are three possible deviations for firm f to analyze:

(i) $\hat{o}_f^2 = \widehat{w}$ and $u_f(\widehat{w}) > u_f(\overline{w})$:: for the worker $\widehat{w}$, it must hold that $u_{\widehat{w}}$ $(\mu^1(\widehat{w})) > u_{\widehat{w}}(f)$. Otherwise, worker w is not the one who provides the highest utility among those who would accept an offer from f, contradicting the definition of the strategy σ. As the workers play σ, meaning they choose the offer that provides the highest utility, the offer from f will be rejected.

(ii) $\hat{o}_f^2 = \widehat{w}$ and $u_f(\overline{w}) > u_f(\widehat{w})$:: since workers have no commitment, then $\widehat{w}$ accept the offer $\hat{o}_f^2$. Since firms hold commitment, they cannot layoff any worker, and therefore, the utility of firm f is:

$$\sum_{t=2}^{\infty} \delta_f^{t-1} u_f(\widehat{w}) < \sum_{t=2}^{\infty} \delta_f^{t-1} u_f(\overline{w})$$

Hence, firm f will not benefit by making offer $\hat{o}_f^2$.

(iii) $\hat{o}_f^2 = f$, meaning that f deviates by making no offer:: In the next period, the best outcome for f is to make offer $\hat{o}_f^3 = \overline{w}$ and worker $\overline{w}$ accept such an offer. Thus,

$$\sum_{t=3}^{\infty} \delta_f^{t-1} u_f(\overline{w}) < \sum_{t=2}^{\infty} \delta_f^{t-1} u_f(\overline{w}).$$

Hence, the firm does not benefit by not making offers in period 2.

Then, from cases (i)–(iii), firm f does not benefit by deviating in the second period and, therefore, follows the strategy σ making offer $o_f^2 = \overline{w}$, where $\overline{w}$ is the next worker on the list who will not reject it.

Now consider the strategies of the workers, i.e., each worker w accepts the offer that provides the highest utility. In this case, workers are the only ones that can deviate. For each $w \in W$, in every period $t > 1$ where $\mu^{t-1} = \mu$, we have $\mu(w) \in O_w$. Since μ is stable, the offer that provides the highest utility to each w is $\mu(w)$. Assume that at stage $t = 1$, worker w decides to reject all offers ($r_w^1 = w$). Following the re-stabilization process of [5], starting from matching μ when worker w wants to improve her labor situation, there is a stable matching ν such that $u_w(\nu(w)) > u_w(\mu(w))$. Let $k(w)$ be the time that takes firm $\nu(w)$ to make an offer to worker w, i.e., $\nu(w) \in O_w$. Worker w will benefit from deviating by rejecting all offers in period 1 and waiting for the offer from $\nu(w)$ if

$$\sum_{t=1}^{\infty} \delta_w^{t-1} u_w(\mu(w)) < \sum_{t=1+k(w)}^{\infty} \delta_w^{t-1} u_w(\nu(w)). \tag{1}$$

By the resolution of the Geometric Series, we have[5]

$$u_w(\mu(w)) \frac{1}{1-\delta_w} < u_w(\nu(w)) \frac{\delta_w^{k(w)}}{1-\delta_w},$$

where $k(w)$ is the number of periods necessary for the firm $\nu(w)$ to make an offer to worker w. Now, operating we obtain

$$\frac{u_w(\mu(w))}{u_w(\nu(w))} < \delta_w^{k(w)},$$

and, thus

$$\left(\frac{u_w(\mu(w))}{u_w(\nu(w))}\right)^{\frac{1}{k(w)}} < \delta_w.$$

That is, worker w will satisfy (1) if her discount factor meets the condition $\delta_w > \left(\frac{u_w(\mu(w))}{u_w(\nu(w))}\right)^{\frac{1}{k(w)}}$. Since by hypothesis we have $\delta_w \in (0, c^w)$, where $c^w = \left(\frac{u_w(\mu(w))}{u_w(\nu(w))}\right)^{\frac{1}{k(w)}}$, worker w does not benefit from deviating from σ.

Therefore, σ is a stationary equilibrium in responsive strategies, whose outcome is μ. □

[5]For the resolution of Geometric Series see [19].

3.3. **The dynamic game when workers hold commitment.** In this subsection, we examine the dynamic matching game where workers hold commitment, but firms do not. In this case, workers have job offers that are not permanent, i.e. although they cannot quit, they may be fired. Assuming that firms act strategically constantly, we demonstrate that any stable matching results from a stationary equilibrium. Moreover, we prove that given a stable matching, and assuming that firms act strategically responsive depending on their patience, such a matching is the outcome of a stationary equilibrium.

Theorem 3.5. *Given a dynamic matching game where workers hold commitment, any stable matching μ can be supported as a stationary equilibrium in constant strategies.*

Proof. Let μ be a stable matching, and consider the following strategy profile σ at each period t:

(i) Each firm f makes an offer to $\mu(f)$; that is, $o_f^t = \mu(f)$.
(ii) Each worker w accepts the offer that provides the highest utility.

We will prove that the strategy σ is a stationary equilibrium in constant strategies.

Assume that until period $\widetilde{t}-1 > 2$, all agents play the strategy profile σ. Also, assume that there is a firm that deviates at stage $\widetilde{t}$, i.e., $\hat{o}_f^{\widetilde{t}} \neq \mu(f)$. Now, we have three cases to consider:

(i) $\hat{o}_f^{\widetilde{t}} = w$, $u_f(w) > u_f(\mu(f))$, and $\mu^{\widetilde{t}-1}(w) \neq w$:: Since workers hold commitment, and the other firms are playing σ, w rejects the offer made by f.

(ii) $\hat{o}_f^{\widetilde{t}} = w$, $u_f(w) > u_f(\mu(f))$, and $\mu^{\widetilde{t}-1}(w) = w$:: Since μ is stable and agents have been playing σ until stage $t-1$, we have that $0 = u_w(\mu(w)) > u_w(w)$. Then, w rejects the offer made by f.

(iii) $\hat{o}_f^{\widetilde{t}} = f$, meaning that f deviates by making no offer:: In period $\widetilde{t}+1$, since the rest of the agents are playing σ, the best outcome for f is to obtain $\mu(f)$, thus

$$\sum_{t=1}^{\widetilde{t}-1} \delta_f^{t-1} u_f(\mu(f)) + \sum_{t=\widetilde{t}+1}^{\infty} \delta_f^{t-1} u_f(\mu(f)) < \sum_{t=1}^{\infty} \delta_f^{t-1} u_f(\mu(f)).$$

Therefore, the firm does not benefit by not making offers in period $\widetilde{t}$.

Then, from cases (i)–(iii), firm f does not benefit by deviating in period $\widetilde{t}$ and, therefore, follows the strategy σ making offer $o_f^{\widetilde{t}} = \mu(f)$.

Now, consider the strategies of the workers. Firms do not deviate, meaning $o_f = \mu(f)$ for all $f \in F$. For each $w \in W$, in every period t, we have $\mu(w) \in O_w^t$. Since μ is stable, the offer that provides the highest utility to each w is $\mu(w)$. Therefore, the workers' strategies are optimal, as without commitment, they accept the offer that maximizes their utility.

Therefore, σ is a stationary equilibrium in constant strategies, whose outcome is μ. □

The strategy σ considered in Theorem 3.5 assumes that firms behave in a strategically constant manner. We now turn to the fact that any stable matching can also be sustained as the outcome of a stationary equilibrium in responsive strategies, provided that the discount factor is suitably chosen. In what follows, we lay out the specific requirements that the discount factor must satisfy.

Let $f \in F$ and μ a stable matching. Define c^f as follows:

$$c^f = \left(\frac{u_f(\mu(f))}{u_f(\mu_F(f))} \right).$$

The following theorem guarantees that if the discount factor of each firm f is bounded by c^f, then the stable matching μ results from a different stationary strategy, where firms actions are strategically responsive.

Let μ be a stable matching, and consider the following strategy profile σ at each period t in the dynamic game:

(i) For $t = 1$, each firm f makes an offer a $\mu(f)$, i.e. $o_f^1 = \mu(f)$. For $t > 1$, if each firm f observe that $\mu^{t-1} = \mu$, then $o_f^t = \mu(f)$. Otherwise (if each firm f observe $\mu^{t-1} \neq \mu$), then $o_f^t = \mu_F(f)$.
(ii) Each worker w accepts the offer that provides the highest utility.

Note that condition (i) for $t > 1$ states that, if firm f observes that in the previous period, the resulting matching is not μ, the offer it makes is $\mu_F(f)$, i.e. its partner under the firm-optimal stable matching.

Theorem 3.6. *Given a dynamic matching game where workers hold commitment, let μ be a stable matching. If $\delta_f \in (0, c^f)$ for each $f \in F$, then σ is a stationary equilibrium in responsive strategies supporting the stable matching μ.*

Proof. Let μ be a stable matching. We will prove that the strategy σ is a stationary equilibrium in responsive strategies that supports μ when $\delta_f \in (0, c^f)$.

Assume that until period $\widetilde{t} - 1 > 2$, all agents play the strategy profile σ. Also, assume that there is a firm that deviates at stage $\widetilde{t}$, i.e., $\hat{o}_f^{\widetilde{t}} \neq \mu(f)$. Now, we have three cases to consider:

- **(i) $\hat{o}_f^{\widetilde{t}} = w$, $u_f(w) > u_f(\mu(f))$, and $\mu^{\widetilde{t}-1}(w) \neq w$::** Since workers hold commitment, and given that other firms in this period are playing σ and the outcome matching from the previous period is μ, each firm $\hat{f} \neq f$ makes an offer to $\mu(\hat{f})$, i.e., $o_{\hat{f}}^{\widetilde{t}} = \mu(\hat{f})$. Consequently, w rejects the offer $\hat{o}_f^{\widetilde{t}}$.
- **(ii) $\hat{o}_f^{\widetilde{t}} = w$, $u_f(w) > u_f(\mu(f))$, and $\mu^{\widetilde{t}-1}(w) = w$::** Since μ is stable and agents have been playing σ until stage $\widetilde{t} - 1$, we have that $0 = u_w(\mu(w)) > u_w(f)$. Then, w rejects the offer $\hat{o}_f^{\widetilde{t}}$.
- **(iii) $\hat{o}_f^{\widetilde{t}} = f$, meaning that f deviates by making no offer::** In period $\widetilde{t}+1$, since the rest of the agents are playing σ and the resulting matching in period $\widetilde{t}$ differs from μ, each firm $\hat{f} \neq f$ makes an offer to $\mu_F(\hat{f})$, i.e., $o_{\hat{f}}^{\widetilde{t}+1} = \mu_F(\hat{f})$. Firm f would benefit from deviating by not making an offer, and the best strategy it can follow is $\mu_F(f)$ if

$$\sum_{t=1}^{\infty} \delta_f^{t-1} u_f(\mu(f)) < \sum_{t=1}^{\widetilde{t}-1} \delta_f^{t-1} u_f(\mu(f)) + \sum_{t=\widetilde{t}+1}^{\infty} \delta_f^{t-1} u_f(\mu_F(f)). \tag{2}$$

By the resolution of the Geometric Series, we have[6]

$$u_f(\mu(f))\frac{1}{1-\delta_f} < u_f(\mu(f))\left(\frac{1-\delta_f^{\widetilde{t}-1}}{1-\delta_f}\right) + u_f(\mu_F(f))\frac{\delta_f^{\widetilde{t}}}{1-\delta_f}.$$

By operating the previous inequality, we obtain

$$u_f(\mu(f)) < u_f(\mu(f))\left(1-\delta_f^{\widetilde{t}-1}\right) + u_f(\mu_F(f))\delta_f^{\widetilde{t}},$$

and thus

$$u_f(\mu(f)) - u_f(\mu(f)) + u_f(\mu(f))\delta_f^{\widetilde{t}-1} < u_f(\mu_F(f))\delta_f^{\widetilde{t}}.$$

Then, operating again, we obtain

$$\frac{u_f(\mu(f))}{u_f(\mu_F(f))} < \delta_f.$$

That is, firm f will satisfy (2) if their discount factor meets the condition $\delta_f > \frac{u_f(\mu(f))}{u_f(\mu_F(f))}$. Since by hypothesis we have $\delta_f \in (0, c^f)$, where $c^f = \left(\frac{u_f(\mu(f))}{u_f(\mu_F(f))}\right)$, firm f does not benefit from deviating from σ.

Then, from cases (i)–(iii), firm f does not benefit by deviating in period $\widetilde{t}$ and, therefore, follows the strategy σ making offer $o_f^{\widetilde{t}} = \mu(f)$.

Now, consider the strategies of the workers. Firms do not deviate, meaning $o_f^t = \mu(f)$ for each $f \in F$ and each period t. Hence, $\mu(w) \in O_w^t$ for each $w \in W$. Since μ is stable, the offer that provides the highest utility to each w is $\mu(w)$. Therefore, the workers' strategies are optimal, as in the case that workers have no commitment, they accept the offer that maximizes their utility. Therefore, σ is a stationary equilibrium in responsive strategies, whose outcome is μ. □

4. **Concluding remarks.** This paper extends the static matching model by [9] to a dynamic setting where firms and workers engage repeatedly, examining how different commitment structures and levels of patience among agents impact long-term stability in decentralized markets. We model a non-cooperative dynamic game where firms periodically offer positions, and workers choose to accept or reject these offers, without forming beliefs about others' actions. The study considers three commitment scenarios—both sides, firms-only, and workers-only—and introduces stationary equilibrium as the solution concept. We find that stable matchings can be supported as stationary equilibria under specific commitment settings, providing insights into how varying levels of patience and commitment affect equilibrium and stability outcomes over time.

If we consider a potential extension to a many-to-one model, where firms can hire multiple workers, the game becomes asymmetric regarding which agents make the offers. Due to this asymmetry, our results do not extend directly, even assuming responsive utilities for firms. For example, in cases where firms hold commitment, it is not possible to determine the exact number of stages required for re-stabilization; only a lower bound can be provided [5]. Therefore, considering these market characteristics, an indirect extension may be feasible for many-to-one markets where firms exhibit responsive or even substitutable utility functions.

[6] For the resolution of Geometric Series see [19].

Finally, our assumption that agents do not form beliefs about others' actions plays a key role in simplifying the analysis and ensuring equilibrium existence. Nevertheless, incorporating belief formation—through learning or reputation—could enrich the model's strategic structure and shed light on how anticipatory behavior affects dynamic stability. Exploring this direction, possibly through equilibrium refinements such as those proposed by [20], constitutes a promising avenue for future research.

REFERENCES

[1] J. Alcalde, D. Pérez-Castrillo, A. Romero-Medina, et al., *Hiring Procedures to Implement Stable Allocations*, Instituto Valenciano de Investigaciones Económicas, (1997).

[2] J. Alcalde and A. Romero-Medina, Simple mechanisms to implement the core of college admissions problems, *Games and Economic Behavior*, **31** (2000), 294-302.

[3] A. Amieva, A. G. Bonifacio and P. Neme, Mechanisms for a dynamic many-to-many school choice problem, arXiv preprint, `arXiv:2411.07851`, (2024).

[4] Y. Blum, A. E. Roth and U. G. Rothblum, Vacancy chains and equilibration in senior-level labor markets, *Journal of Economic Theory*, **76** (1997), 362-411.

[5] A. G. Bonifacio, N. Guiñazu, N. Juarez, P. Neme and J. Oviedo, Counting steps for re-stabilization in a labor matching market, arXiv preprint, `arXiv:2405.07084`, (2024).

[6] E. Diamantoudi, E. Miyagawa and L. Xue, Decentralized matching: The role of commitment, *Games and Economic Behavior*, **92** (2015), 1-17.

[7] L. Doval, Dynamically stable matching, *Theoretical Economics*, **17** (2022), 687-724.

[8] L. Doval and V. Skreta, Mechanism design with limited commitment: Markov environments, *European Economic Review*, **167** (2024), 104777.

[9] D. Gale and L. S. Shapley, College admissions and the stability of marriage, *The American Mathematical Monthly*, **69** (1962), 9-15.

[10] G. Haeringer and M. Wooders, Decentralized job matching, *International Journal of Game Theory*, **40** (2011), 1-28.

[11] H. Konishi and M. Sapozhnikov, Decentralized matching markets with endogenous salaries, *Games and Economic Behavior*, **64** (2008), 193-218.

[12] D. G. McVitie and L. B. Wilson, Stable marriage assignment for unequal sets, *BIT Numerical Mathematics*, **10** (1970), 295-309.

[13] A. Nicolò, P. Salmaso and R. Saulle, Dynamic one-sided matching, *Available at SSRN 4352130*, (2023).

[14] M. Niederle and L. Yariv, Matching through decentralized markets, *Discussion Paper, Stanford University*, (2007).

[15] J. S. Pereyra, A dynamic school choice model, *Games and Economic Behavior*, **80** (2013), 100-114.

[16] A. Roth, The evolution of the labor market for medical interns and residents: A case study in game theory, *Journal of Political Economy*, **92** (1984), 991-1016.

[17] A. E. Roth and M. A. O. Sotomayor, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambidge University Press, Cambridge, 1990.

[18] R. Saulle, A. Nicolò and P. Salmaso, Rationalizable conjectures in dynamic matching, Available at SSRN: `https://ssrn.com/abstract=5070813`, (2024).

[19] J. Stewart, *Calculus: Early Transcendentals*, Cengage Learning, 2012.

[20] E. Van Damme, *Stability and Perfection of Nash Equilibria*, vol. 339, Springer, 1991.

Received May 2025; revised July 2025; early access July 2025.