

Dynamic Matching

Mariagiovanna Baccara Leeat Yariv

July 15, 2021

1 Introduction

Many matching environments are inherently dynamic—participants arrive at the market over time, or interact dynamically before forming matches. For example, in the child-adoption process, children become available progressively and often wait for a match while being cared for by social services, incurring wait costs in terms of well-being and long-term outcomes. Similarly, potential families pay attorney and agency fees while waiting to be matched to a child. In kidney donation, the U.S. Department of Health and Human Services reports that a new patient is added to the kidney transplant list every 14 minutes and about 3000 patients are added each month. Since health conditions of potential recipients can potentially deteriorate as time passes, the timing of matches is crucial for minimizing lost lives: in 2014, 4761 patients died while waiting for a kidney transplant, and another 3668 people became too sick to receive one. In the realm of public housing, families are often placed on waitlists before obtaining housing units, which become available stochastically. And so on and so forth.

To set the stage for our analysis, consider a simple setting in which squares \mathcal{S} and rounds \mathcal{R} arrive over time. For simplicity, suppose any square and round are compatible, or agreeable, to one another with some probability p .¹ Consider then a bi-partite graph connecting k squares with k rounds, where a link appears only when the corresponding pair is compatible, see Figure 1. This is a random graph, with each link occurring with probability p .

A bi-partite graph as such induces a *perfect matching* if there exists a matching $\mu : \mathcal{S} \rightarrow \mathcal{R}$ that is injective and surjective, where $\mu(s) = r$ implies that s and r are linked in the graph (but not necessarily the converse: only some links get implemented when constructing the induced matching). The following is a well-known result.

Proposition 1 (Erdős and Rényi, 1964) *As long as the graph is connected enough, namely as long as p approaches 0 slower than $\frac{\log k}{k}$, there is a perfect matching with probability approaching 1 as k grows large.*

¹In principle, we could allow for heterogeneity among squares and rounds and have the compatibility probability depend on agents' types. Indeed, in the organ-donation context, some blood types are more common than others. The probability p we consider here could be thought of as the minimal compatibility probability across types.

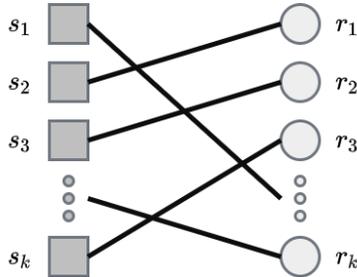


Figure 1: Compatibility-based Random Preferences

Thus, even if compatibility rates are incredibly low, a large enough population of market participants would allow us to nearly guarantee a perfect match for everyone. Similar results hold when agents are heterogeneous and some matches generate greater surplus than others: in a very large population, it is almost always possible to create the maximal number of efficient matches, and minimize the loss due to less desirable ones.

In reality, when a market populates over time or when interactions between agents occur dynamically, waiting for a match is costly. Thus, in general, we cannot rely on the desirable asymptotic features of such markets. At the heart of the work on optimal design of clearinghouses in such settings is then the trade-off between market thickness, which allows for high-quality matches, and costly waiting.

In this chapter, we illustrate how this trade-off affects the optimal design of matching mechanisms in different settings, from one-sided matching or allocation environments in which one side of the market has preferences over items constituting the other market side, to two-sided markets, in which both sides of the market have preferences over matched partners.

2 Dynamic One-Sided Allocations

We start with a description of one-sided allocation problems in which scarce items—public housing, daycare spots, organs from deceased donors, etc.—arrive over time and are allocated to waiting agents. Importantly, there is heterogeneity in agents’ valuations of items. In the following analysis, we consider two potential impediments to the socially optimal allocation: first, the social planner may be unable to impose item assignments on unwilling agents, who would rather wait for more preferable options; second, agents’ preferences over items may not be transparent to the social planner.

2.1 Priority Protocols in Discretionary Settings

In this section we compare alternative priority protocols that a planner can implement when agents maintain discretion over the acceptance of an item offered

to them. We study this problem in a setting in which the number of agents allowed to wait in line is fixed, and a new participant can join the line as soon as an agent waiting is assigned an item and leaves the market.

Consider a waitlist that, for simplicity, consists of only two agents, who are ranked according to their order of arrival: $\rho = 1, 2$. At every period, one item becomes available and is offered to the agents according to an independently-determined priority order. Specifically, the item is offered to the agents according to their order of arrival ρ with probability $r \in [\frac{1}{2}, 1]$, and according to the reverse order with probability $1 - r$. Hence, $r = 1$ represents the first-in-first-out (FIFO) protocol, while $r = \frac{1}{2}$ represents an equal-weight lottery. If the first agent the item is offered to declines it, the item is offered to the other agent. If the second agent also passes, the item goes to waste. If the $\rho = 1$ agent accepts the item, she leaves the market and the other agent gets her slot, namely is labeled with $\rho = 1$, while a new agent joins the waitlist with a $\rho = 2$ label. If the second, $\rho = 2$ agent is assigned the item, he leaves the market and a new agent replaces him with $\rho = 2$. Each period, agents incur a waiting cost $c > 0$. Assume that the reservation utility is sufficiently low so that agents entering the waitlist never leave without being assigned an item.

We start with the *private value case* in which each agent's valuation for each item is independent, equals 1 with probability $p \in (0, 1)$, and 0 otherwise. Any agent i always accepts an item that she values at 1. Let $q(i)$ denote agent i 's probability of acceptance of an item she values at 0. For $i = 1, 2$, agent i 's continuation value, the difference between her expected value and her expected cost, is:

$$V(i) = \frac{A(i)}{B(i)} - \frac{ic}{B(i)},$$

where

$$\begin{aligned} A(1) &= p[r + (1 - r)(1 - p)(1 - q(2))], \\ A(2) &= p[2 - p - (1 - p)(rq(1) + (1 - r)q(2))], \\ B(1) &= [p + (1 - p)q(1)][r + (1 - r)(1 - q(2))(1 - p)], \\ B(2) &= [p + (1 - p)q(1)][r + (1 - r)(1 - p)(1 - q(2))] \\ &\quad + [p + (1 - p)q(2)][1 - r + r(1 - p)(1 - q(1))]. \end{aligned}$$

For example, $A(1)$ represents the expected value of the item picked by agent 1. The item is valued at 1 for agent 1 with probability p , in which case agent 1 accepts it whenever offered. The agent is offered the item first with probability r . The item is offered to agent 1 after being rejected by agent 2 with probability $(1 - r)(1 - p)(1 - q(2))$.

The probabilities of misallocation and waste are

$$\begin{aligned} \mu &= p(1 - p)[rq(1) + (1 - r)q(2)] \text{ and} \\ v &= (1 - p)^2(1 - q(1))(1 - q(2)), \end{aligned}$$

respectively. Unsurprisingly, the probability of waste v is independent of the queuing protocol r . It is easy to check that there are three equilibria comprised

of pure-strategy Markov strategies, $\{q^j(1), q^j(2)\}_{j=1}^3$, where $q^1(1) = q^1(2) = 0$, $q^2(1) = 0$, $q^2(2) = 1$, and $q^3(1) = q^3(2) = 1$.

In all three equilibria, an increase in r leads to an increase in $V(1)$. We now focus on the effect of an increase in r on $V(2)$. Intuitively, there are two effects at work. An increase in r decreases the probability with which agent 2 is offered the item first, and therefore decreases agent 2's continuation value. However, an increase in r also increases the continuation value of agent 1, and therefore makes agent 1 more selective. This latter effect benefits agent 2, since every item rejected by agent 1 is offered to agent 2 next, potentially generating a strictly positive value for him. As it turns out, the two effects cancel one another and, in all three equilibria, any increase in r leaves $V(2)$ unchanged. We can conclude that an increase in r weakly increases the equilibrium values of both agents, and leaves the probability of waste unchanged. Finally, the probability of misallocation is negatively correlated with agents' equilibrium values, and it is weakly decreasing in r .

Next, consider the *common value case* in which the values of the two agents are perfectly correlated. We can express the continuation value as in the private value case, where now

$$\begin{aligned} A(1) &= rp, \\ A(2) &= p, \\ B(1) &= r[p + (1-p)q(1)] + (1-r)(1-p)q(1)(1-q(2)), \\ B(2) &= 1 - (1-p)(1-q(1))(1-q(2)). \end{aligned}$$

For example, as before, $A(1)$ represents the expected value of the item picked by agent 1. Since values are correlated, the only case in which agent 1 is offered an item of value 1 is when she is first in line. The expected value is then rp .

Since the two agents value all items in the same way, there is never misallocation. The probability of waste is

$$v = (1-p)(1-q(1))(1-q(2)),$$

which is again independent of r . As in the private value case, it is easy to check that there are three pure-strategy Markov equilibria $\{\tilde{q}^j(1), \tilde{q}^j(2)\}$, where $\tilde{q}^1(1) = \tilde{q}^1(2) = 0$, $\tilde{q}^2(1) = 0$, $\tilde{q}^2(2) = 1$, and $\tilde{q}^3(1) = \tilde{q}^3(2) = 1$. An increase in r always increases $V(1)$ and does not affect $V(2)$. The reasoning behind the effect on $V(2)$ is slightly different from that pertaining to the private value case. Specifically, in the common value case, the only scenario in which an agent can make a positive payoff is by being offered the item first. An increase in r lowers the probability with which agent 2 is approached first. However, it also increases the rate at which agent 1 accepts an item and leaves the market, freeing her spot on the waitlist for agent 2. Again, the two effects balance one another, leaving $V(2)$ unaffected by changes in r .

To summarize, this example illustrates two settings in which the FIFO queuing protocol dominates a lottery from both participants' and the social planner's perspectives.

2.2 Buffer-Queues Mechanism with Private Preferences

Consider now a setting in which a large set of agents is present at time $t = 0$, and at each period, a new item arrives at the market. The items can be of two sorts: an A -item with probability p_A , and a B -item with probability $p_B \equiv 1 - p_A$. Agents can also be of two types. Agents of type α prefer A -items, while agents of type β prefer B -items. Each agent is of type α with probability p_α and of type β with probability $p_\beta \equiv 1 - p_\alpha$. For simplicity, assume that the system is *balanced*; that is, $p_A = p_\alpha \equiv p$. Each agent gets utility $v > 0$ from being assigned her preferred item, and 0 from being assigned a different item. As before, the wait costs agents incur before being matched are linear, and the per-period cost is $c > 0$. For example, this setup fits the allocation process of public housing units to families that vary in their preferences over housing units' attributes: location, floor, etc.

Without observing agents' types, the social planner needs to select a mechanism μ to allocate each arriving item to an agent, and we allow the social planner to impose assignments on agents if needed. Since the system is assumed to be overloaded, an item is assigned at every period—i.e., there is no waste. Therefore, the total wait cost is constant across allocations, and the social planner's goal is to minimize the welfare loss due to items' misallocation. Given an assignment μ , if ξ_t is an indicator equal to 1 if the arriving item at t is misallocated, the long-run misallocation rate is $\xi = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \xi_t$. Therefore, the welfare loss from misallocation is

$$WFL = v\xi.$$

As a convenient benchmark, consider a sequential assignment mechanism that assigns the arriving item to an arbitrary agent, without agents having discretion on whether to accept or decline any item. It is easy to see that such a mechanism generates a misallocation rate of $\xi^{SA} = 2p(1 - p)$. Can we do better than this mechanism by inducing agents to report their preferences?

It is possible to induce agents to reveal their type by allowing them discretion over whether to accept an item offered immediately and a different item at some future time. Such a mechanism yields an endogenous separation of the agents into two sets: ones that have not yet been approached, and ones that have been approached in the past and have decided to wait for their preferred item. Those agents who are waiting after having been approached form the **buffer queue** (or BQ). Agents in the buffer queue are divided into ones that are waiting for an A -item, in the A -buffer queue (or the $A - BQ$), and ones that are waiting for a B -item, in the B -buffer queue (or the $B - BQ$). Since items' arrival is stochastic, agents in the buffer queue face uncertainty regarding the time at which they will receive their desired item. As buffered agents accumulate, the mechanism needs to take their presence into consideration as new offers are made to subsequent agents. For $x = A, B$, a **buffer-queue policy** for x -items governs the maximal number of agents that are allowed in the $x - BQ$ and how each new x -item is allocated to the agents waiting in the $x - BQ$.

Definition 1 (Buffer-Queue Policy) For $x = A, B$, a $\langle \bar{k}^x, \varphi^x \rangle$ buffer-queue policy for x consists of a threshold \bar{k}^x of the number of agents in the x -BQ and, for any length $k \leq \bar{k}^x$ of the x -BQ, a probability $\varphi_k^x(i)$ with which an agent in position $1 \leq i \leq k$ of the x -BQ is assigned the item x . Therefore, $\varphi_k^x(i) \geq 0$ for any $1 \leq i \leq k \leq \bar{k}^x$, and $\sum_{i=1}^k \varphi_k^x(i) = 1$ for any $1 \leq k \leq \bar{k}^x$.

Definition 2 (Buffer-Queue Mechanism) A buffer-queue mechanism $\mathcal{M} = \langle \bar{k}^A, \varphi^A, \bar{k}^B, \varphi^B \rangle$ specifies a buffer-queue policy for each item and, for $x = A, B$, if an x -item arrives, it implements the following steps:

- (1) If the x -BQ is not empty, the x -item is assigned to an agent in the x -BQ according to φ^x .
- (2) If the x -BQ is empty, the mechanism sequentially approaches new agents until the x -item is assigned. For each approached agent:
 - (a) Let $y \neq x$. If the y -BQ currently has $k - 1 < \bar{k}^y$, the mechanism offers the agent the choice of (i) taking the x -item immediately, or (ii) declining the x -item and joining the y -BQ in the k -th position to receive a y -item in the future according to φ^y . If the agent chooses (i), the period ends, and if the agent chooses (ii), the mechanism approaches another new agent.
 - (b) If the y -BQ currently has \bar{k}^y agents, then the new agent is assigned the x -item and the period ends.

A buffer-queue mechanism is incentive compatible if, whenever a buffer queue is not full, any agent chooses to join that buffer queue rather than accept a less desirable item; that is, agents are *truthful*. For example, consider an unapproached agent of type α who is offered a B -item. The agent has the choice of either taking that item or revealing that she is mismatched and being placed in the A -BQ in position k . In this case, suppose w_k^A is the expected number of periods the agent will have to wait until receiving an A -item. Given w_k^A , the agent prefers to join the A -BQ if $v - cw_k^A \geq 0$, or $w_k^A \leq \frac{v}{c} \equiv \bar{w}$. If the agent believes that other agents are truthful, the expected wait w_k^A depends only on $\langle \bar{k}^A, \varphi^A \rangle$. Therefore, \mathcal{M} is referred to as **incentive compatible** if $w_k^x \leq \bar{w}$ for all $k \leq \bar{k}^x$ and $x = A, B$; put differently, if it is an equilibrium for all agents to be truthful.

The dynamics of a buffer-queue mechanism when agents are truthful are captured by an ergodic Markov chain with state space $\{-\bar{k}^B, \dots, \bar{k}^A\}$, where the state $k \geq 0$ indicates k agents waiting in the A -BQ, and $k \leq 0$ indicates $|k|$ agents waiting in the B -BQ. Note that at most one of the queues can be non-empty at any given time.

Buffer queues can be used to study the performance of a single waiting list governed by a FIFO priority order. All agents wait in an ordered line and know their position. Each arriving item is offered to the first agent in the queue. If that agent declines, the item is offered to the second agent in line, and so on. All agents who decline an item keep their positions in the queue. The following result allows us to quantify the equilibrium welfare loss associated with this protocol.

Proposition 2 (Equilibrium of a FIFO Waitlist) *The single FIFO waitlist has a unique equilibrium in which the outcome is identical to a buffer-queue mechanism $\mathcal{M} = \langle \bar{k}^A, \varphi^A, \bar{k}^B, \varphi^B \rangle$ when agents are truthful, φ^A and φ^B follow a FIFO order, and $\bar{k}^A = \bar{k}^B = \lfloor p\bar{w} \rfloor$. The welfare loss is given by:*

$$WFL^{FIFO} = \frac{2vp(1-p)}{(1-p)\lfloor p\bar{w} \rfloor + p\lfloor (1-p)\bar{w} \rfloor + 1}.$$

The intuition behind Proposition 2 is simple. Consider an agent of type α in position k of a single waitlist. If the agent is offered an A -item, she will accept. If she is offered a B -item, it must be the case that $k-1$ agents before her declined the same item, since they are waiting for A . Therefore, the agent is expected to wait k/p periods before being offered an A -item. This implies that the agent prefers to wait if and only if $v - kc/p \geq 0$, or $k \leq p\bar{w}$. The Markov-chain structure of the buffer queues allows an easy computation of the welfare loss. As the wait cost approaches zero, the welfare loss of the single FIFO waitlist in the balanced case vanishes.

Finally, we can characterize the welfare-maximizing incentive-compatible buffer-queue mechanism. To do so, we first define a **Load Independent Expected Wait (LIEW)** policy for item $x = A, B$. This is a buffer-queue policy $\langle \bar{k}^x, \varphi^x \rangle$ in which, when agents are truthful, the expected wait w for agents in the queue does not depend on their position in the queue, or on the number of other agents in the queue. A mechanism $\mathcal{M} = \langle \bar{k}^A, \varphi^A, \bar{k}^B, \varphi^B \rangle$ is a LIEW mechanism if $\langle \bar{k}^x, \varphi^x \rangle$ is LIEW for $x = A, B$. We have

Proposition 3 (Optimal Buffer Queue) *Let $\mathcal{M}^* = \langle \bar{k}^A, \varphi^A, \bar{k}^B, \varphi^B \rangle$ be a LIEW mechanism such that $\bar{k}^A = \bar{k}^B = \bar{k}^{LIEW}$, where*

$$\bar{k}^{LIEW} = \lfloor 2p\bar{w} \rfloor - 1.$$

Then \mathcal{M}^ is incentive compatible and achieves weakly higher welfare when agents are truthful than any other incentive-compatible buffer-queue mechanism.*

To understand the ideas underlying Proposition 3, observe that, since misallocations decrease in \bar{k}^A and \bar{k}^B , welfare maximization requires the maximization of \bar{k}^A and \bar{k}^B subject to the incentive-compatibility constraint. Now, we can establish that any buffer-queue policy with $\bar{k} > \bar{k}^{LIEW}$ is not incentive compatible. To see this, notice that for a policy to be incentive compatible, it must be that for any $k \leq \bar{k}$, any agent that joins the buffer queue at position k must expect a wait w_k such that $w_k \leq \bar{w}$, so that $E[w] \leq \bar{w}$. In the balanced case, the average number of people in the buffer queue is $L = \frac{K+1}{2}$. Little's law implies that $E[w] = \frac{L}{p} = \frac{\bar{k}+1}{2p} \leq \bar{w}$, or $\bar{k} \leq [2p\bar{w}] - 1$. Therefore, if $\bar{k} > \bar{k}^{LIEW}$, there must be at least one position k' such that $w_{k'} > \bar{w}$, thereby violating incentive compatibility.

It is possible to show that \mathcal{M}^* reduces the welfare loss to almost half that of the FIFO mechanism. Intuitively, any buffer-queue mechanism compensates agents that reject a mismatch by promising them a better match in the future, albeit at the cost of additional wait time. In a FIFO mechanism, agents who join the buffer queue when it is relatively short obtain larger net payoffs than those joining the buffer queue later on, for whom the incentive-compatibility constraint binds. A LIEW mechanism induces the same expected wait time for all agents. It is then able to accommodate more agents in the buffer queue while maintaining incentive compatibility.

3 Dynamic Two-Sided Matching

3.1 Dynamic Matching with Fixed Participants

Many real markets have a fixed set of participants interacting in a dynamic fashion. For instance, every year, new economics graduate students enter the job market for academic positions. The set of candidates and the available positions are, by and large, set at the start of each year and interactions between market participants occur dynamically: universities often invite candidates to interviews sequentially and the generation of offers often spans several months. Many other markets share those features: the market for new law clerks, freshly-minted rabbis, etc. Even the medical match, while famously governed by a centralized clearinghouse, is preceded by interactions—namely, interviews—that occur over time and are by and large decentralized. How do such decentralized markets work? What matches can they achieve? Will they be stable?

One natural way to answer these questions is to describe precisely a two-sided **decentralized matching market game** in which participants interact over time. The main ingredients of such a game are naturally the underlying preference distribution of participants, and the information available to them.

Consider the following simple setup. A market corresponds to a triplet $(\mathcal{S}, \mathcal{R}, U)$, where $\mathcal{S} = \{1, \dots, S\}$ is a finite set of squares—say, hiring firms—and $\mathcal{R} = \{1, \dots, R\}$ is a finite set of rounds—say, potential employees. Match utilities

can then be described as follows:

$$U = \left\{ \begin{array}{cc} \left\{ u_{ij}^s \right\}, & \left\{ u_{ij}^r \right\} \\ \text{square } i\text{'s utility from matching with } j & \text{round } j\text{'s utility from matching with } i \end{array} \right\}.$$

For simplicity, we can assume that remaining unmatched generates a utility of 0 for any participant and that all match utilities are strictly positive.

Certainly, if a market has multiple stable matchings, any dynamic interaction would conceivably suffer from coordination problems: even if participants aimed at establishing a stable outcome, they would need to agree on which one. Consider then the simple case in which utilities are such that there is a unique stable matching, which we denote by μ_M .

One way to model market interactions is via a dynamic version of the Deferred Acceptance (DA) algorithm. At every period $t = 1, 2, \dots$ there are two stages. First, squares simultaneously decide whether and to whom to make an offer, where an unmatched square can have at most one offer out. Then, each round j who has received an offer from square i can accept, reject, or hold the offer. If such an offer is accepted at period t , square i is matched to round j irreversibly. Square i then receives a payoff of $\delta^t u_{ij}^r$, and round j receives a payoff of $\delta^t u_{ij}^s$, where $\delta \leq 1$ is the market discount factor. Unmatched agents receive 0 throughout the game. Importantly, in contrast with the way a direct-revelation version of DA operates, squares need not make offers in order of their preference lists and rounds need not hold only offers that are their favorites. In particular, squares can approach rounds multiple times, much like in labor-market applications, where some individuals may receive repeat offers from particular employers.

As for market monitoring, assume that squares and rounds observe receipt, rejection, and deferral only of offers they are involved in. However, whenever an offer is accepted, the whole market is informed of the union. Similarly, whenever there is market exit, all participants are informed. We make these assumptions for their realism. While individuals or firms are privy to details of offers they engage with, they are unlikely to know the ins and outs of all offers made in the market. Nonetheless, theoretically, one could consider various alternatives to this monitoring structure.

We now consider the Nash equilibria of such a market game. As it turns out, we can always implement the unique stable outcome μ_M through equilibrium:

Proposition 4 (Equilibrium with Fixed Participants) *There exists a Nash equilibrium in strategies that are not weakly dominated that generates the unique stable matching.*

The intuition of this proposition is straightforward. Indeed, consider the following profile. At $t = 1$, each square i makes an offer to round $\mu_M(i)$. Each round j accepts the highest-ranked square that is at least as good as $\mu_M(j)$,

breaking any ties in favor of $\mu_M(j)$. The round leaves immediately if she receives no offers and all squares are matched or exited. Otherwise, off the equilibrium path, squares and workers revert to strategies that emulate the DA algorithm.

Nonetheless, there can be other (unstable) equilibrium outcomes, as the following example illustrates.

Example 1 (Multiplicity) Suppose $\mathcal{S} = \{s_1, s_2, s_3\}$, $\mathcal{R} = \{r_1, r_2, r_3\}$, and that utilities induce the following ordinal preferences:

$$\begin{array}{ll} s_1 : & r_2 \succ r_1 \succ r_3 & r_1 : & s_1 \succ s_3 \succ s_2 \\ s_2 : & r_1 \succ r_2 \succ r_3 & r_2 : & s_2 \succ s_1 \succ s_3 \\ s_3 : & r_1 \succ r_2 \succ r_3 & r_3 : & s_1 \succ s_3 \succ s_2. \end{array}$$

We continue assuming that all participants are acceptable.

These preferences induce a unique stable matching μ_M such that $\mu_M(s_i) = r_i$ for all i .

As it turns out, we can induce another matching μ , where s_1 and s_2 swap their partners and match with their favorite rounds, as long as the discount factor δ is high enough. Namely, we can implement in equilibrium $\mu(s_1) = r_2$, $\mu(s_2) = r_1$, and $\mu(s_3) = r_3$.

How can that be done in equilibrium? Consider the following profile of strategies. In period 1, s_3 makes an offer to r_3 , who accepts any offer from a square ranked at least as high as r_3 . Squares s_1 and s_2 make no offers, while rounds r_1 and r_2 accept offers only from their favorite squares. In round 2, each square s_i , with $i = 1, 2$, makes an offer to $\mu(s_i)$. Rounds r_1 and r_2 accept any offer. Upon any observable deviation, all remaining agents revert to emulating (square-proposing) DA strategies.

Why is this profile an equilibrium? Notice that s_3 and r_3 are bound to match with one another. They therefore best respond by doing so in the first period. In fact, they have a strict preference for doing so with any discount factor smaller than 1. With s_3 and r_3 out of the way, the resulting “sub-market” exhibits the following restricted preferences:

$$\begin{array}{ll} s_1 : & r_2 \succ r_1 & r_1 : & s_1 \succ s_2 \\ s_2 : & r_1 \succ r_2 & r_2 : & s_2 \succ s_1 \end{array} .$$

In particular, this sub-market entails two stable matchings: one matching s_i with r_i for $i = 1, 2$ and one matching s_i with r_{3-i} for $i = 1, 2$. What transpires from period 2 and on is essentially the profile we used to prove Proposition 1: each square makes an offer to their most-preferred stable partner. Notice that the rounds are not using weakly dominated strategies in this example. In particular, if square s_i makes an offer to r_i in period 1, for $i = 1, 2$, that offer would be accepted immediately. Nonetheless, for high enough discount factors, these squares would prefer to wait for one period to get their most preferred partner.

This example hinges on the dynamic nature of interactions. Agents are making *contingent* offers: conditional on s_3 and r_3 leaving the market, s_1 and s_2

target their most favored stable partners. For example, in the job-market context, this suggests that certain participants could be placed in a “waitlist” and approached with an offer only after other participants are matched. Clearly, what allows for this example to occur is the fact that, despite the overall market having a unique stable matching, one of the sub-markets has multiple stable matchings. Ruling out the possibility that sub-markets exhibit multiple stable matchings eliminates such examples when combined with appropriate refinements.²

3.2 Dynamic Matching with Evolving Participants

In many two-sided matching processes, such as child adoption and kidney exchanges, participants arrive over time. Likewise, many labor markets entail unemployed workers and job openings that become available at different periods. Such settings open the door for new questions regarding the operations of both decentralized and centralized interactions.

3.2.1 Dynamic Stability

When market participants arrive over time, certain matches might be created along the way—patients get transplants, parents adopt children, individuals get public housing, etc. In such settings, attempts to block a market matching are constrained by the fact that only a subset of individuals is available at any point in time. We start by modifying the standard notion of stability for such settings.

For illustration purposes, we consider a particularly simple setting. Suppose there are only two periods, $t = 1, 2$. A finite set of squares \mathcal{S} all arrive at $t = 1$. A finite set of rounds \mathcal{R} arrive in two installments. A subset of rounds $\mathcal{R}_1 \subset \mathcal{R}$ arrives at $t = 1$ and a subset $\mathcal{R}_2 \subset \mathcal{R}$ arrives at $t = 2$, where

$$\mathcal{R}_1 \cap \mathcal{R}_2 = \emptyset \quad \text{and} \quad \mathcal{R}_1 \cup \mathcal{R}_2 = \mathcal{R}.$$

Thus, only one side of the market appears in increments. Furthermore, there is certainty on future arrivals.

We assume that squares are discounted-utility maximizers. That is, for each $s \in \mathcal{S}$, there is a utility $u(s, \cdot) : \mathcal{R} \rightarrow \mathbb{R}$ and a discount factor $\delta_s \in [0, 1]$, such that s 's utility from matching with round r at time t is given by $\delta_s^t u(s, r)$. We also assume that remaining unmatched generates 0 utility: $u(s, s) = 0$. We make similar assumptions when it comes to rounds, but distinguish between rounds arriving at period 1 and rounds arriving at period 2. Namely, each round $r \in \mathcal{R}$ is associated with a utility $v(\cdot, r) : \mathcal{S} \rightarrow \mathbb{R}$, where we assume that $v(r, r) = 0$. Furthermore, each round r in \mathcal{R}_1 uses a discount factor $\delta_r \in [0, 1]$. Agents in \mathcal{R}_2 experience no discounting as they exist for only one period in the market. At any

²One can show that, with *aligned* preferences, where there is no “preference cycle” in the matrix of match surpluses, iterated elimination of weakly dominated strategies generates a unique equilibrium prediction, which is stable. As a note about refinements, due to our assumptions on the structure of monitoring, subgame perfection has little bite. The only public monitoring occurs through market exits, which limits the set of proper subgames.

time t , only available agents can match with one another. Consequently, we can define a **period- t matching** as an injective map $m_t : \mathcal{S} \cup \bigcup_{\tau=1}^t \mathcal{R}_\tau \rightarrow \mathcal{S} \cup \bigcup_{\tau=1}^t \mathcal{R}_\tau$ such that (i) for all $s \in \mathcal{S}$, $m_t(s) \in \{s\} \cup \bigcup_{\tau=1}^t \mathcal{R}_\tau$, and (ii) for all $r \in \bigcup_{\tau=1}^t \mathcal{R}_\tau$, $m_t(r) \in \mathcal{S} \cup \{r\}$. Let M_t denote the set of all period- t matchings.

Naturally, period-1 matchings imposes constraints on period-2 matchings. Namely, if an agent is matched in period 1, she cannot be rematched in period 2. Formally, a pair $(m_1, m_2) \in M_1 \times M_2$ is **feasible** if for all $s \in \mathcal{S}$, if $m_1(s) \neq s$ then $m_2(s) = m_1(s)$ (and since m_2 is injective, $m_2(m_1(s)) = s$). Let M denote the set of all feasible matchings.

Even when the set of participants is fixed, we already saw that dynamics allow for matchings to be *contingent* on prior market interactions. The same occurs when market participation evolves: period-2 matchings can depend on the matchings implemented in period 1. Our object of analysis is therefore a **contingent matching**. Such a matching specifies the selection of a period-1 matching and, for each matching in period 1, the selection of a period-2 matching. Formally, a **contingent matching** μ is a map

$$\mu : \{\emptyset\} \cup M_1 \rightarrow M_1 \cup M_2$$

such that $\mu(\emptyset) \in M_1$ and, for all $m_1 \in M_1$, $\mu(m_1) \in M_2$ and $(m_1, \mu(m_1))$ is feasible.

We are now ready to define dynamic stability of a contingent matching. It entails two conditions. First, once a matching is formed in period 1, the matching in period 2 must be stable among the remaining agents and the new entrants, for much the same reasons as those described in the introductory chapter for static markets. The period-2 matching should then be individually rational and entail no blocking pairs. Second, taking the outcomes of period 2 as given, no group of agents in period 1 can beneficially deviate from the prescribed matching.

Definition 3 (Dynamic Stability) *A contingent matching μ is dynamically stable if:*

1. *For each $m_1 \in M_1$, the resulting period-2 matching $\mu(m_1)$ is stable, entailing no blocking individuals or pairs;*
2. *There is no set $A \subseteq \mathcal{S} \cup \mathcal{R}_1$ that can implement m_1 such that all agents in A prefer $(m_1, \mu(m_1))$ to $(\mu(\emptyset), \mu(\mu(\emptyset)))$. Namely, there is no group of agents in $t = 1$ that can improve their outcomes by changing who they are matched with at $t = 1$, waiting to match at $t = 2$, or both.*

Suppose a contingent matching μ is individually rational but not dynamically stable. Then either there exists a pair of contemporary agents that prefer matching to one another over their prescribed partners under μ , or there exists a group of $t = 1$ agents who want to block by waiting to be matched.

We can also define the **core** analogously to the way it is defined for static matching markets. Namely, a contingent matching μ is in the core if there is no agent who would rather remain single than match according to μ , and there is no pair that would prefer to generate a (feasible) match at some point over the prescribed match under μ . Formally, no blocking pair means that there is no $s \in \mathcal{S}$ and $r \in \mathcal{R}_t$ such that

$$\delta_s^{1(r \in \mathcal{R}_2)} u(s, r) > U(s, m_\mu) \quad \text{and} \quad v(s, r) > V_t(r, m_\mu),$$

where U is the discounted utility for the squares and V_t is the utility for rounds arriving at time t .

What is the difference between the core and dynamically stable contingent matchings? There are two main differences. First, in the core, even if $b \in \mathcal{R}_2$, $\{a, b\}$ can form a blocking coalition at the outset. Essentially, there is no concern for the timing of arrival of different agents. Second, blocking coalitions in the core compare the payoffs they obtain by blocking with the payoffs they obtain from the matching prescribed by μ . In contrast, dynamic stability requires that a coalition that blocks by waiting in period 1 compares its payoff under μ to the payoffs generated in the *continuation matching* originated by the block.

Unfortunately, dynamically stable matchings do not always exist, as the following example illustrates.

Example 2 Suppose $\mathcal{S} = \{\text{Erdős, Kuhn, Gale}\}$, $\mathcal{R}_1 = \{\text{Renyi, Tucker}\}$, and $\mathcal{R}_2 = \{\text{Shapley, Nash}\}$.

In what follows, we denote rankings of any square s over elements of the form (r, t) , where if s ranks (r, t) over (r', t') , then $\delta_s^t u(s, r) > \delta_s^{t'} u(s, r')$. We further assume that all agents prefer to match with specified agents over remaining unmatched.³

Erdős:	(Shapley, 0)	(Shapley, 1)	(Renyi, 0)	(Renyi, 1)
Kuhn:	(Shapley, 1)	(Nash, 0)	(Tucker, 0)	(Nash, 1)
Gale:	(Tucker, 0)	(Tucker, 1)	(Shapley, 0)	(Shapley, 1)

and

Renyi:	Erdős
Tucker:	Kuhn Gale
Shapley:	Gale Erdős Kuhn
Nash:	Kuhn

Assume that for all r , δ_r is sufficiently high so that $v(s, r) > v(s', r)$ implies that $\delta_r v(s, r) > v(s', r)$.

If μ is dynamically stable, then Gale has to be matched under μ . Otherwise, he would be unmatched at $t = 2$ and could block with Shapley.

³In addition, we assume non-trivial discounting so that $(r, 0)$ is always preferable to $(r, 1)$ for any r .

Thus, Gale has to be matched, either with Tucker at $t = 1$ or with Shapley at $t = 2$.

Let $m_1 = \mu(\emptyset)$. Suppose Gale is matched with Shapley. It has to be that $m_1(\text{Kuhn}) = \text{Tucker}$ and $m_1(\text{Erdős}) = \text{Renyi}$. Indeed:

1. Kuhn cannot be unmatched because he would block with Nash at $t = 2$, and
2. Kuhn blocks a match with Nash by matching early with Tucker since, importantly, he does not like to wait for Nash.

However, the unique stable matching at $t = 2$, when only Erdős is matched with Renyi at $t = 1$, matches Kuhn with Shapley, whom Kuhn prefers to Tucker. Hence, Kuhn blocks any such contingent matching μ by waiting.

Similar reasoning rules out the case in which Gale is matched with Tucker at $t = 1$.

What drives non-existence in the example? It is discounting: Kuhn prefers Nash to Tucker, but does not want to wait for him. That is why Kuhn cannot be matched to Tucker in period 2.

Consider then a market with *trivial discounting* so that if $u(s, r) > u(s, r')$, then $\delta_s u(s, r) > u(s, r')$ for any square s and rounds r, r' .

Proposition 5 (Existence) *If preferences satisfy trivial discounting, the set of dynamically stable contingent matchings is nonempty.*

The root of the impact of discounting are cycles. In the example above, Kuhn prefers Shapley at $t = 2$, Shapley prefers Gale, Gale prefers Tucker immediately, who prefers Kuhn. So, if Kuhn were matched to Tucker and Gale were matched to Shapley, both Kuhn and Shapley would want to swap their partners. As seen, this allows for deviations such as those in which Kuhn beneficially waits for period 2 to match.

In general, a **simultaneous preference cycle** is an alternating sequence of squares and rounds, $r_1, s_1, r_2, s_2, \dots, r_N$, where $r_1 = r_N$, such that, when considering period-1 preferences,

1. Each square s_i prefers $r_{(i+1) \bmod N}$ to r_i , and both of these rounds are acceptable to s_i ;
2. Each round r_i prefers s_i to s_{i-1} , and both of these squares are acceptable to r_i .⁴

We can now summarize the impacts of such cycles on the structure and existence of dynamically stable matchings.

Proposition 6 (Preference Cycles) *If a core matching is not part of a dynamically stable contingent matching, then there is a preference cycle. If there are no preference cycles, the set of dynamically stable contingent matchings coincides with the core.*

⁴We interpret $s_0 = s_N$ so that, for r_1 , this implies that s_1 is preferable to s_N .

3.2.2 A Simple Model of Dynamic Centralized Design

The notion of dynamic stability offers a decentralized benchmark for considering markets in which agents arrive in sequence. It is also natural to consider the optimal design of clearinghouses in such settings. We now analyze a simple model in which the optimal centralized clearinghouse can be characterized. As already seen, one of the hurdles an evolving market presents is that agents may prefer partners that arrive later than they do. Waiting for them comes at a cost.

Suppose that, each period, one square and one round arrive at the market. Assume that each square is of a desirable, or high type H , with probability p and of a less desirable, or low type L , with probability $1 - p$. Similarly, suppose that each round is of a desirable type h with probability p and of a less desirable type l with probability $1 - p$. Let U_{xy} denote the surplus that any type x of square and any type y of round generate, $x = H, L, y = h, l$.⁵

For simplicity, assume super-modular preferences, so that:

$$U \equiv U_{Hh} + U_{Ll} - U_{Hl} - U_{Lh} > 0.$$

In particular, the efficient matching entails the maximal number of (H, h) and (L, l) pairs.

We assume $p \in (0, 1)$ so that, conceivably, a social planner might want market participants to wait in order to generate more efficient matchings. This assumption guarantees that thicker markets generate greater expected efficiency, or overall match surplus.

Assume that all agents suffer a cost of $c > 0$ for each period on the market without a match. In particular, any pair that is held in the market for a period generates a loss of $2c$ to the planner.

Suppose agents depart the market only upon matching. As long as remaining unmatched generates sufficiently low utilities, this restriction would be consistent with individual rationality.⁶

Consider general dynamic mechanisms, where the social planner can create matches between available agents at every period. Formally, at any time t , before a new square-round pair arrives, a queue is represented by (k_H, k_h, k_L, k_l) , where the length of queues of squares are given by k_H and k_L for H -squares and L -squares, respectively. Similarly, the length of queues of rounds are given by k_h and k_l for h -rounds and l -rounds, respectively.

We focus on stationary and deterministic mechanisms. At time t , after a new square-round pair enters the market, there is a queue $\mathbf{n}^t = (n_H^t, n_h^t, n_L^t, n_l^t)$, where indices correspond to types as before. A mechanism is characterized

⁵These surpluses can naturally be generated by individual agents' utilities from matching with agents of different types.

⁶Individual rationality can be defined once individual utilities from matches are specified. The idea that remaining unmatched yields extremely low payoffs is relevant for many of the applications such a model speaks to. For instance, parents seeking to adopt a child may suffer greatly by not finding a match. Similarly, patients seeking an organ donor may suffer dire consequences from not being matched.

by a mapping $\mu : \mathbb{Z}_+^4 \rightarrow \mathbb{Z}_+^4$ such that for every $\mathbf{n} \in \mathbb{Z}_+^4$, $\mu(\mathbf{n}) = \mathbf{m} = (m_{Hh}, m_{Hl}, m_{Lh}, m_{Ll})$ is a feasible profile of matches, with m_{xy} corresponding to the number of matches generated between any type x of square and any type y of round.⁷ Once matches are created, a new queue $\mathbf{k} = (k_A, k_B, k_\alpha, k_\beta)$ contains the remaining agents:

$$\begin{aligned} k_x &= n_x - (m_{xh} + m_{xl}) & \text{for } x \in \{H, L\}, \\ k_y &= n_y - (m_{Hy} + m_{Ly}) & \text{for } y \in \{h, l\}. \end{aligned}$$

The generated surplus by matches \mathbf{m} is:

$$S(\mathbf{m}) \equiv \sum_{(x,y) \in \{H,L\} \times \{h,l\}} m_{xy} U_{xy}.$$

Waiting costs incurred by retaining agents \mathbf{k} are:

$$C(\mathbf{n}, \mathbf{m}) \equiv c \left(\sum_{x \in \{A, B, \alpha, \beta\}} k_x \right).$$

The welfare generated at time t is then:

$$w(\mathbf{n}^t, \mathbf{m}^t) \equiv S(\mathbf{m}^t) - C(\mathbf{n}^t, \mathbf{m}^t).$$

The social planner assesses the performance of a mechanism using the *average welfare*, defined as:

$$W(\mu) \equiv \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T w(\mathbf{n}_t, \mu(\mathbf{n}_t)) \right].$$

The average welfare is well-defined in that it can be shown that the limit exists for every mechanism μ .

An **optimal mechanism** is a mechanism achieving the maximal average welfare. An optimal mechanism exists since there is only a finite number of stationary and deterministic mechanisms leading to a bounded stock of agents in each period.

Lemma 1 (Congruent Matches) *Any optimal mechanism requires (H, h) and (L, l) pairs to be matched as soon as they become available.*

The intuition for this lemma is the following. If the social planner holds on to an (H, h) pair, it is only for the hope of matching the relevant agents with a future L -square and an l -round. However, our supermodularity assumption implies that this would entail an efficiency loss. Hence, the planner may as well

⁷Naturally, the social planner cannot match more agents than are available, so we must have that $m_{xh} + m_{xl} \leq n_x$ for $x \in \{H, L\}$ and, similarly, $m_{Hy} + m_{Ly} \leq n_y$ for $y \in \{h, l\}$.

match the (H, h) pair immediately. Symmetric logic follows for any available (L, l) pair.

Thus, the optimal mechanism potentially holds on to agents only when they form incongruent pairs. Intuitively, the optimal mechanism cannot hold an exceedingly large number of agents since waiting costs would be prohibitive. In fact, the following holds:

Proposition 7 (Optimal Mechanism) *An optimal dynamic mechanism is identified by a pair of thresholds $(\bar{k}_H, \bar{k}_h) \in Z_+$ such that*

1. *whenever more than \bar{k}_H H -squares are present, $n_H - \bar{k}_H$ pairs of type (H, l) are matched immediately, and*
2. *whenever more than \bar{k}_h h -rounds are present, $n_h - \bar{k}_h$ pairs of type (L, h) are matched immediately.*

Since the environment here is fully symmetric, we can assume, without substantial loss of generality, that $\bar{k}_H = \bar{k}_h = \bar{k}$.⁸

Denote by $k_{Hh} = k_H - k_h$. The value of k_{Hh} captures both the number of agents the social planner holds on to and their type: when $k_{Hh} > 0$, there are H -squares and l -rounds waiting, while when $k_{Hh} < 0$, there are L -squares and h -rounds waiting. It follows a Markov process, where states correspond to values $-\bar{k} \leq k_{Hh} \leq \bar{k}$.

Notice that for any $-\bar{k} < k_{Hh} < \bar{k}$, the stock of agents held by the planner does not change if an (H, h) or an (L, l) pair arrives, which occurs with probability $p^2 + (1 - p)^2$. The stock changes, up or down, if an (H, l) or an (L, h) pair arrives, each occurring with a probability of $p(1 - p)$.

We can therefore characterize the Markov chain associated with k_{Hh} through the transition formula $\mathbf{x}^{t+1} = \mathbf{T}_{\bar{k}} \mathbf{x}^t$, where \mathbf{x}^t is a $(2\bar{k} + 1)$ -dimensional vector such that its k -th entry equals 1 whenever, at time t , $k_{Hh} = k$, and 0 otherwise.

$$T_{\bar{k}} = \begin{bmatrix} 1 - p(1 - p) & p(1 - p) & 0 & & 0 \\ p(1 - p) & p^2 + (1 - p)^2 & & & \\ 0 & & \ddots & & \\ & & & p^2 + (1 - p)^2 & p(1 - p) \\ 0 & & & p(1 - p) & 1 - p(1 - p) \end{bmatrix}.$$

The resulting process is ergodic, implying a unique steady-state distribution, which is the uniform distribution over states. That is, in the steady state, k_{Hh} takes each of its possible $2\bar{k} + 1$ values with probability $\frac{1}{2\bar{k} + 1}$.

We now use the characterization of the steady-state distribution to identify costs and benefits for any \bar{k} . In state j , $2|j|$ participants are present on the

⁸For almost all parameters of the environment, there is a unique optimal mechanism with such identical thresholds. However, for a negligible set of parameters, which will soon become apparent, additional asymmetric thresholds generate the same level of average welfare.

market. The expected total waiting costs are therefore:

$$C(\bar{k}) = \frac{1}{2\bar{k} + 1} \left(\sum_{j=-\bar{k}}^{\bar{k}} 2|j| \right) c = \frac{\bar{k}(\bar{k} + 1)c}{2\bar{k} + 1}.$$

To calculate the expected benefits, suppose an (H, l) pair (similarly for an (L, h) pair) arrives at the market when the state is j . If $-\bar{k} \leq j < 0$, the optimal mechanism creates one (H, h) match and one (L, l) match, generating a surplus of $U_{Hh} + U_{Ll}$. If $0 \leq j < \bar{k}$, the mechanism creates no matches, generating no match surplus. If $j = \bar{k}$, the mechanism creates an (H, l) match, generating a surplus of U_{Hl} . Any congruent pair arriving at the market is matched immediately and generates its corresponding match surplus. Thus, after algebraic manipulations, the expected per-period total match surplus is:

$$B(\bar{k}) = pU_{Hh} + (1 - p)U_{Ll} - \frac{p(1 - p)U}{2\bar{k} + 1}.$$

The optimal \bar{k} maximizes $B(\bar{k}) - C(\bar{k})$. We can therefore fully characterize the optimal mechanism in this setting.

Proposition 8 (Optimal Threshold) *The threshold*

$$\bar{k}^* = \left\lfloor \sqrt{\frac{p(1 - p)U}{2c}} \right\rfloor$$

identifies an optimal dynamic mechanism, and it is generically unique.

The optimal threshold \bar{k}^* balances market thickness and waiting costs. It decreases in c and is positive only when costs are sufficiently low, namely when $c \leq \frac{p(1-p)U}{2}$. The optimal threshold increases as $p(1 - p)$, the probability of an incongruent pair's arrival, increases and maximized at $p = 1/2$.⁹

When the optimal threshold is implemented, the resulting welfare can then be calculated as follows:

$$W^*(c) = pU_{Hh} + (1 - p)U_{Ll} - \frac{p(1 - p)U}{2\bar{k}^* + 1} - \frac{\bar{k}^*(\bar{k}^* + 1)c}{2\bar{k}^* + 1}.$$

From super-modularity, the maximal conceivable welfare, were there no waiting costs and an infinitely thick market, is given by:

$$S_\infty = pU_{Hh} + (1 - p)U_{Ll}.$$

Naturally, $W^*(c) \leq S_\infty$. In fact, this inequality is strict for any $c > 0$. We can now consider the comparative statics of $W^*(c)$ with respect to costs. First,

⁹Uniqueness of the optimal threshold breaks down only when $\sqrt{\frac{p(1-p)U}{2c}}$ is an integer, which occurs for a zero-measure set of parameters.

the optimal welfare decreases in c . Indeed, for $c_1 > c_2$, the social planner can emulate the mechanism designed for c_1 when waiting cost c_2 is in place. That would generate the same matching surplus under both costs, but a lower average waiting cost under c_2 . Second, the welfare loss is concave in c . To see this, we utilize the fact that the optimal threshold \bar{k}^* decreases in c . Therefore, as c increases, fewer individuals wait and the effect of a marginal cost increase is smaller. We therefore have the following:

Corollary 1 (Optimal Welfare) *The welfare under the optimal mechanism is given by $W^*(c) = S_\infty - \Theta(c)$, where $\Theta(c)$ is continuous, increasing, and concave in c , $\lim_{c \rightarrow 0} \Theta(c) = 0$, and $\Theta(c) = p(1-p)U$ for all $c \geq \frac{p(1-p)U}{2}$.*

As a direct consequence, for vanishingly small waiting costs, the optimal mechanism achieves approximately the maximal conceivable welfare.

3.2.3 Other Considerations

Organ donation is a natural environment in which participants—patients seeking an organ and donors willing to give an organ—arrive over time and need to be paired. Furthermore, waiting for a transplant is costly for patients. A substantial fraction of organ donation is from cadavers. The allocation then is in many ways simpler. Preferences of patients are by and large observable: their blood and tissue type, their urgency, demographics, etc. The system can then generate (Pareto) efficient allocations without much concern for incentive-compatibility constraints.¹⁰

In kidney exchange, patients arrive with a live donor—a family member, a friend, etc.—who is not necessarily a compatible match. Can there be beneficial swaps between patients and their donors that would induce agents to enter the system to begin with? The timing of such swaps is also important. Matching all compatible pairs reduces waiting costs. However, keeping some desirable donors in the pool—say, those with O blood type, who are blood-type compatible with any patient—can have future benefits, in terms of facilitating other exchanges, or finding immediate matches for future patients in particularly critical health conditions.

In the first analysis of this issue, Ünver (2010) showed that if only two-way exchanges are allowed, every optimal mechanism matches all compatible pairs immediately. Nonetheless, when multi-way exchanges are allowed, the efficient mechanism may require holding some available matches to keep relatively scarce donors available for future use.

Along similar lines, one can consider environments in which agents can at some point become critical, and drop out of the system if not matched immediately. Consider random compatibility between pairs, which can be formulated in graph-theoretic terms.¹¹ Participating (incompatible) patient-donor pairs con-

¹⁰See Section 7 of Chapter 30 for a related discussion.

¹¹Two patient-donor pairs are compatible if a swap of donors and patients yields compatibility in terms of blood type, tissue type, etc.

stitute the graph's nodes and arrive at a Poisson rate of $m \geq 1$ in an interval of time $[0, T]$. Any two such pairs are compatible with probability $p = \frac{d}{m}$, with $d > 0$. Moreover, pairs become critical at a Poisson rate normalized to 1. If they are not matched immediately, critical pairs perish. Otherwise, there is no waiting cost. Given a matching policy, the resulting expected number of pairs that perish can be thought of as the *loss* of the policy. A planner, observing the set of pairs that become critical, seeks to minimize loss. While characterizing the optimal matching algorithm of this model is computationally difficult, it is easy to see that it has to satisfy the following two conditions: (i) since there are no waiting costs, two connected pairs are matched only if at least one is critical, and (ii) if a critical pair is connected to someone else, it is always matched immediately.

Furthermore, it is possible to obtain quantitative insights on the performance of the optimal algorithm by considering the following two simple algorithms:

Definition 4 (Greedy Algorithm) *If any new pair enters the market at time t , match them randomly with any existing compatible pair, if it exists.*

Definition 5 (Patient Algorithm) *If a pair becomes critical, match them randomly with any compatible pair. Otherwise, hold on to pairs.*

Both these algorithms are obviously suboptimal since they do not use any information regarding the underlying graph. Denote by $L(G)$ and $L(P)$ the loss associated with the Greedy and Patient algorithms, respectively, over the horizon $[0, T]$.

Proposition 9 (Loss Bounds) *For $d \geq 2$, as $T, m \rightarrow \infty$, we have:*

$$L(G) \geq \frac{1}{2d+1} \text{ and } L(P) \leq \frac{1}{2}e^{-\frac{d}{2}}.$$

Proposition 9 suggests that the Patient algorithm's loss is exponentially small, while the Greedy algorithm's loss is not. That is, the option value of waiting before matching pairs is large. To see why, suppose there are z pairs in the market. If a new pair enters the market under the Greedy algorithm, or becomes critical under the Patient algorithm, the probability that no pair on the market is compatible is $(1 - \frac{d}{m})^z$. However, the number of pairs in the market depends on the algorithm under consideration. As more pairs wait under the Patient algorithm, the market is thicker, which reduces the probability of any critical pair being unmatched. To estimate the performance of the Patient algorithm, it is useful to establish a lower bound on the loss achieved by the optimal algorithm, which we denote by L^* .

Proposition 10 (Efficiency of the Patient Algorithm) *Let A be any algorithm that observes the set of critical pairs with associated loss $L(A)$. Then, for $d \geq 2$, as $T, m \rightarrow \infty$,*

$$L^* \geq \frac{e^{-\frac{d}{2}(1+L(A))}}{d+1}.$$

Substituting the Patient algorithm for A , we obtain $L^* \geq \frac{e^{-\frac{d}{2}(1+\frac{\epsilon}{2})}}{d+1}$ or, as $T, m \rightarrow \infty$,

$$L(P) - L^* \leq e^{-\frac{d}{2}} \left[\frac{1}{2} - \frac{e^{-\frac{d}{4}} e^{-\frac{d}{2}}}{d+1} \right],$$

indicating that the performance of the optimal algorithm is close to that of the Patient algorithm. This suggests that the benefit of allowing the market to thicken before matching pairs is substantial even if the implemented mechanism does not fully exploit all the information contained in the network structure.

4 Notes

The model described in Section 2.1 relies on Bloch and Cantala (2017). The model described in Section 2.2 is analyzed by Leshno (2019). While that model assumes agents are heterogenous in what items, say public-housing units, they prefer, recent work considers settings in which agents agree on the ranking of the items, but differ in their preference intensities, see Ortoleva et al. (2021). In addition, when the items are thought of as services—e.g., medical services by junior or senior physicians, legal aid from rookie or experienced lawyers, etc.—the distribution of these services can be endogenized. Namely, junior service providers become senior providers after attending to a sufficient number of tasks. In such settings, the balance between service quality and wait times needs to account for the training possibilities that affect the future distribution of available services. See Baccara et al. (2021) for details.

Our discussion of decentralized market games in Section 3.1 relies on a model offered by Ferdowsian et al. (2021), while the main example in that section relates to an example appearing in Echenique et al. (2016). Ferdowsian et al. (2021) also study decentralized market games in settings in which agents have incomplete information about others' preferences. They show that incomplete information introduces another hurdle for (complete-information) stability, even when market participants are very patient and interactions offer ample opportunities for learning. Dynamic decentralized interactions have been studied using lab experiments as well, see Echenique et al. (2021) and Agranov et al. (2021). While stability has strong drawing power with complete information, the introduction of transfers or incomplete information impedes stability.

The notion of dynamic stability described in Section 3.2.1, as well as the example in that section, were suggested by Doval (2021). She offers a definition of stability for far more general dynamic settings than those sketched here, allowing for multiple periods and uncertainty about agents' arrival over time.

Our discussion in section 3.2.2 of optimal clearinghouses in environments in which agents on two market sides arrive over time relies on Baccara et al. (2020). They also analyze a discretionary counterpart in which participants *choose* whether to wait for a more desirable partner.¹² They illustrate that

¹²Individual utilities are assumed to be such that all squares prefer h -rounds and all rounds

agents wait excessively, not internalizing the externalities on other agents who arrive after them.

The literature on dynamic matching arguably started from the consideration of organ donation, see Part III of this book. Ünver (2010) illustrated the value of multi-way swaps. The results described in Section 3.2.3 appear in Akbarpour et al. (2020).¹³

The discussion in this chapter illustrates the different ways by which waiting costs are modeled: as flow costs, discounted match values, or through the likelihood of urgently needing a match. The use of each naturally depends on the application. In decentralized interactions, costs can take either form. From a market-design perspective, however, discounting introduces several challenges. First, the market designer may need to keep track of the *arrival time* of participants, placing a heavy computational and potentially logistical burden. Second, agents who have waited for a long time exhibit low discounted match values and would then receive lower weight in the market designer’s considerations. In contrast, in many applications, seniority lends an advantage—for example, patients waiting for a long time for an organ or families queuing for public housing are prioritized.

There are many natural directions by which the models described in this chapter could be extended. While transfers are banned or limited in many applications such as organ donation, child adoption, and public housing, they are present in many others, particularly when considering dynamic labor markets. Their consideration could enrich our models substantially. Incomplete information regarding the underlying preferences in the market would also be an interesting direction to pursue further in this area.

References

- AGRANOV, M., A. DIANAT, L. SAMUELSON, AND L. YARIV (2021): “Paying to Match: Decentralized Markets with Information Frictions,” *mimeo*.
- AKBARPOUR, M., S. LI, AND S. O. GHARAN (2020): “Thickness and information in dynamic matching markets,” *Journal of Political Economy*, 128, 783–815.
- BACCARA, M., A. COLLARD-WEXLER, L. FELLI, AND L. YARIV (2014): “Child-adoption matching: Preferences for gender and race,” *American Economic Journal: Applied Economics*, 6, 133–58.

prefer H -squares.

¹³The decentralized settings considered here and in the context of organ donation differ from those of the search and matching literature presented in Chapter 22. Those models commonly assume a stationary type distribution. Stationarity might be reasonable in very large markets, but is harder to square with smaller markets in the applications this chapter speaks to. For example, data from one U.S. child adoption facilitator who links adoptive parents and birth mothers indicates a rate of about 11 new potential adoptive parents and 13 new birth mothers entering the facilitator’s operation each month, see Baccara et al. (2014).

- BACCARA, M., S. LEE, AND L. YARIV (2020): “Optimal dynamic matching,” *Theoretical Economics*, 15, 1221–1278.
- (2021): “Task Allocation and On-the-job Training,” *mimeo*.
- BLOCH, F. AND D. CANTALA (2017): “Dynamic assignment of objects to queuing agents,” *American Economic Journal: Microeconomics*, 9, 88–122.
- DOVAL, L. (2021): “Dynamically stable matching,” *mimeo*.
- ECHENIQUE, F., A. ROBINSON-CORTÉS, AND L. YARIV (2021): “An experimental study of decentralized matching,” *mimeo*.
- ECHENIQUE, F., A. J. WILSON, AND L. YARIV (2016): “Clearinghouses for two-sided matching: An experimental study,” *Quantitative Economics*, 7, 449–482.
- ERDŐS, P. AND A. RÉNYI (1964): “On random matrices,” *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 3, 455–460.
- FERDOWSIAN, A., M. NIEDERLE, AND L. YARIV (2021): “Decentralized matching with aligned preferences,” *mimeo*.
- LESHNO, J. (2019): “Dynamic matching in overloaded waiting lists,” *mimeo*.
- ORTOLEVA, P., E. SAFONOV, AND L. YARIV (2021): “Who cares more? Allocation with diverse preference intensities,” *mimeo*.
- ÜNVER, M. U. (2010): “Dynamic kidney exchange,” *The Review of Economic Studies*, 77, 372–414.