

Supplementary Material for “Optimal Dynamic Matching”

Mariagiovanna Baccara

SangMok Lee

Leeat Yariv

June 9, 2018

Abstract

We provide bounds on match utilities that assure individual rationality is satisfied in the matching protocols discussed in the paper. We also detail extensions to the model described in the body of the paper for alternative matching protocols (a simple fixed-window centralized protocol and a discretionary setting with random priority), asymmetric environments, richer type sets, and different arrival processes. Last, we present additional proofs associated with the paper’s results.

1 Individual Rationality

Throughout the paper we assumed that agents leave the market only after being matched. Agents do not leave the market unmatched, regardless of their expected utility from staying in the market. Certainly, B -squares or β -rounds sometimes stay in the market simply for lack of available agents who will match with them. Therefore, the expected continuation payoff for some agents can sometimes be less than zero even when $U_x(y) > 0$ for all x, y . We now provide a bound on the value of remaining unmatched, or the value of an outside option agents have, that assures all the matching protocols we discuss are individually rational. Notice that since, from Corollary 3, the discretionary threshold is higher than the optimal threshold, it suffices to find such a bound for the discretionary process.

In the discretionary process under FIFO, A -squares or α -rounds always have the possibility of matching with β -rounds or B -squares instantaneously when they decide to wait in the market. Therefore, when deciding to wait they expect an even greater utility and individual rationality always holds for them, as long as the outside option is not greater than zero. Consider now β -rounds (analogously, B -squares). A β -round who is k -th in line can always declare all squares as acceptable in each period. By construction, $k \leq \bar{k}^{fifo}$. The time between

arrivals of B -squares is distributed geometrically with probability $1 - p$. Therefore, with such a strategy, the expected time for the β -round to match with a B -square is at most $k/(1 - p)$, yielding a match utility of $U_\beta(B)$. The wait time till matching with a B -square could be even shorter if β -rounds who precede the β -round in question are willing to match only with A -squares. Furthermore, the β -round could end up matching with an A -square before k B -squares arrive at the market. It follows that such a strategy guarantees an expected utility of at least

$$U_\beta(B) - \frac{kc}{1 - p} \geq U_\beta(B) - \frac{\bar{k}^{fifo}c}{1 - p} \geq U_\beta(B) - \frac{p}{1 - p} (U_\alpha(A) - U_\alpha(B)) \equiv U^{\min}.$$

If β -rounds follow a different strategy in equilibrium, their expected utility must be at least as high. Therefore, as long as the value of remaining unmatched is lower than U^{\min} , individual rationality holds under both the optimal and the discretionary processes (analogous calculations follow for B -squares and, under full symmetry of utilities, the bound corresponding to them is also U^{\min}). In addition, we show later in this Online Appendix that both the LIFO protocol and the optimal fixed-window mechanism entail less waiting than under the FIFO protocol. Thus, the construction above guarantees that this bound on remaining unmatched assures individual rationality for the LIFO protocol and the fixed-window protocol as well.

2 Alternative Matching Protocols

In this section we study two alternative matching protocols: in the first one, a centralized clearinghouse matches individuals at fixed time intervals. In the second one, agents form matches in a discretionary setting in which, at every period, they are ranked according to a uniformly random priority rule. We compare the welfare performance of these protocols to that generated by the optimal mechanism as well as the discretionary protocols (FIFO and LIFO) analyzed in the paper.

2.1 Matching with Fixed Windows

We consider the class of mechanisms that are identified by a fixed-window size—every fixed number of periods, the efficient matching for the participants who have arrived at the market in that time window is formed and the market is cleared. Larger windows then allow for thicker markets and potentially greater match surplus. However, larger windows also correspond to

longer expected waiting times for participants. We analyze the optimal fixed window. While the welfare it generates is still lower than that produced by the optimal mechanism, it may be substantially greater than that produced by the discretionary process under FIFO and LIFO. The proofs of these results are presented in Section 6 below.

The arrival of squares and rounds is the same as that discussed so far. When fixed matching windows are used, a window of some size n governs the process. Namely, every n periods, the most efficient matching pertaining to the n squares and n rounds who arrived within that window is implemented.

Suppose that k_A and k_α are the number of A -squares and α -rounds who arrived during a window of n time periods, respectively. Given our assumptions on match utilities, efficient matchings correspond to a unique distribution of pair-types (number of (A, α) pairs, (A, β) pairs, etc.) and generate the maximal number of (A, α) and (B, β) pairs. The total matching surplus generated by a matching as such is

$$S(k_A, k_\alpha) \equiv \begin{cases} k_\alpha U_{A\alpha} + (k_A - k_\alpha) U_{A\beta} + (n - k_A) U_{B\beta} & \text{if } k_A \geq k_\alpha, \\ k_A U_{A\alpha} + (k_\alpha - k_A) U_{B\alpha} + (n - k_\alpha) U_{B\beta} & \text{otherwise.} \end{cases}$$

Consider now the expected waiting costs when the window size is n . The first square and round to arrive wait for $n - 1$ periods, the second square and round to arrive wait for $n - 2$ periods, etc. Thus, the total waiting cost is

$$2c[(n - 1) + (n - 2) + \dots + 0] = cn(n - 1).$$

Therefore, the net expected welfare for each square-round pair generated by a window size n is

$$W_n \equiv \frac{1}{n} \sum_{0 \leq k_\alpha, k_A \leq n} \binom{n}{k_A} \binom{n}{k_\alpha} p_A^{k_A} (1 - p_A)^{n - k_A} p_\alpha^{k_\alpha} (1 - p_\alpha)^{n - k_\alpha} S(k_A, k_\alpha) - c(n - 1).$$

Notice that for any window size n , the matching surplus per pair $S(k_A, k_\alpha)/n$ is at most $U_{A\alpha}$, while the expected waiting costs per pair are $c(n - 1)$. Denote by n^{\max} the largest window size n such that $U_{A\alpha} \geq c(n - 1)$. Every window of size $n > n^{\max}$ would then generate a lower welfare than that generated by a window of size 1, corresponding to instantaneously matching individuals. In particular, an optimal window size exists within the finite set $\{1, 2, \dots, n^{\max}\}$.

2.1.1 Optimal Window Size

A characterization of the precise optimal window size is difficult to achieve, so we identify bounds on the optimal window size. Later, we will find bounds on the welfare generated by

using fixed windows and compare them to the welfare generated by the optimal mechanism as well as the discretionary process.

Consider the ex-ante marginal benefit produced by increasing the window size from n to $n+1$ that is incurred by the first n square-round pairs who arrive at the market. Suppose that an efficient matching corresponding to the first n pairs generates at least some mismatches (i.e., (A, β) or (B, α) pairs). The $(n+1)$ -th pair could be beneficial for the first n pairs by correcting a mismatch. If there is no mismatch among the first n pairs, expanding the window size only leads to additional waiting costs for the first n pairs.

Denote the probability that any efficient matching with the first n squares and rounds has a mismatch (i.e., there is an unequal number of A -squares and α -rounds) by

$$\Pr(|k_A - k_\alpha| \geq 1; n).$$

Since we assumed an identical distribution of types of squares and rounds, an efficient matching of n pairs has mismatches of type (A, β) or (B, α) with equal probability. A mismatch of type (A, β) is corrected by a new $(n+1)$ -th pair of type (B, α) , which occurs with probability $p(1-p)$, and the total benefit for the pair of originally mismatched square and round is

$$U_A(\alpha) + U_\beta(B) - U_A(\beta) - U_\beta(A).$$

A similar derivation follows for the benefit of “correcting” a mismatched pair of type (B, α) . Conditional on any efficient matching entailing a mismatch, the expected benefit to the first n square-round pairs is then

$$\begin{aligned} & \frac{p(1-p)}{2} (U_A(\alpha) + U_\beta(B) - U_A(\beta) - U_\beta(A)) \\ + & \frac{p(1-p)}{2} (U_\alpha(A) + U_B(\beta) - U_\alpha(B) - U_B(\beta)), \end{aligned}$$

which is equal to

$$\frac{p(1-p)U}{2}.$$

Therefore, the ex-ante marginal welfare obtained from expanding the window size from n to $n+1$ for each of the n first square-round pairs is:

$$\Delta_+ W_n \equiv \frac{p(1-p)U}{2n} \cdot \Pr(k_A \neq k_\alpha; n) - 2c.$$

Then, a necessary condition for the optimal window size n^o is

$$\Delta_+ W_{n^o} \leq 0 \leq \Delta_+ W_{n^o-1}. \quad (1)$$

As mentioned above, it is difficult to obtain a closed-form solution for the optimal window size. The following proposition utilizes inequality (1) to establish bounds on n^o .¹

Proposition B1 (Optimal Window Size)

1. An upper bound for the optimal window size is given by

$$n^o \leq \frac{p(1-p)U}{4c}.$$

2. There exists c^* such that, for any $c < c^*$, a lower bound for the optimal window size is

$$n^o > (2p(1-p))^{\frac{1}{3}} \left(\frac{U_{A\alpha} - U_{B\beta}}{c} \right)^{\frac{2}{3}} \geq (2p(1-p))^{\frac{1}{3}} (U/c)^{\frac{2}{3}}.^2$$

Recall that U captures the extent to which preferences exhibit super-modularity, the welfare advantage of an assortative matching relative to an anti-assortative one. Intuitively, the bounds on the optimal window size increase with U and decrease with the cost of waiting.

2.1.2 Welfare Bounds for Fixed Windows

If the waiting cost c is small so that the optimal window size is large, we expect to have an approximate fraction p of A -squares and α -rounds, in which case per-pair surplus is close to S_∞ . The bounds on the optimal window size allow us to provide an approximation of how far the match surplus is from S_∞ and how costly the wait is. The following proposition illustrates the resulting bounds on the welfare generated by the optimal window size.

¹The difficulty stems from the fact that inequality (1) depends crucially on

$$Pr(k_A \neq k_\alpha; n) = 1 - \sum_{l=0}^n \binom{n}{l} p^l (1-p)^{n-l}^2,$$

and state-of-the-art combinatorics has little to say about this function's behavior with changes in n .

Proposition B2 (Fixed-Window Welfare)

1. A lower bound on the optimal fixed-window welfare is given by

$$W^{fix}(c) \geq S_\infty - \min_{n \in \{[m], \lceil m \rceil\}} \left(\sqrt{\frac{p(1-p)}{2n}} (U_{A\alpha} - U_{B\beta}) + c(n-1) \right),$$

where $m \equiv \frac{1}{2} (p(1-p))^{1/3} \left(\frac{U_{A\alpha} - U_{B\beta}}{c} \right)^{2/3}$.³

2. There exists c^* such that, for any $c < c^*$, an upper bound on the optimal fixed-window welfare is given by

$$W^{fix}(c) \leq S_\infty - 2^{-2/3} (p(1-p)c)^{1/3} (U_{A\alpha} - U_{B\beta})^{2/3} + c$$

Notice that the optimal window size is asymptotically efficient, as we have

$$\lim_{c \rightarrow 0} W^{fix}(c) \geq \lim_{c \rightarrow 0} S_\infty - (3/2)(p(1-p)c)^{1/3} (U_{A\alpha} - U_{B\beta})^{2/3} + c = S_\infty.$$

However, the convergence occurs at a lower speed compared to the optimal mechanism and the discretionary process under LIFO. From Corollary 1, the proof of Corollary 5 (see equations (7) and (10) in Section 6 below), and Proposition B2:

Corollary B1 (Relative Performance of Fixed Window Mechanisms) *We have*

$$\liminf_{c \rightarrow 0} \frac{S_\infty - W^{fix}(c)}{S_\infty - W^{opt}(c)} c^{1/6} \geq \frac{(U_{A\alpha} - U_{B\beta})^{2/3}}{2^{7/6} (p(1-p))^{1/6} U^{1/2}}, \text{ and}$$

$$\liminf_{c \rightarrow 0} \frac{S_\infty - W^{fix}(c)}{S_\infty - W^{lifo}(c)} c^{1/6} \geq \frac{3}{2} \left(\frac{U_{A\alpha} - U_{B\beta}}{p(1-p)} \right)^{1/6}.$$

The corollary suggests that the optimal fixed-window mechanism provides a substantial improvement over the discretionary process under FIFO but is inferior to the discretionary process under LIFO.

³As $c \rightarrow 0$, the lower bound becomes arbitrarily close to

$$S_\infty - \sqrt{\frac{p(1-p)}{2m}} (U_{A\alpha} - U_{B\beta}) + c(m-1) = S_\infty - (3/2)(p(1-p)c)^{1/3} (U_{A\alpha} - U_{B\beta})^{2/3} + c.$$

2.2 Uniformly Random Priority Protocol

We now consider a discretionary setting governed by a random priority rule. In particular, agents play a game similar to the one described in Section 4 of the paper, but, after a new pair of agents arrive at the market, agents are ranked according to a uniform random protocol. After priorities realize, agents submit demands sequentially according to their ranks, and the market clears.⁴

We characterize stationary* strategy profiles that satisfy certain conditions necessary for equilibrium. To make our analysis tractable, we make the following two exceptions to the uniformly random protocol:

Assumption B1

1. Upon arrival of a congruent pair (either (A, α) or (B, β)), both new agents are ranked at the top of the queues of their respective types, and
2. Upon arrival of an incongruent pair (either (A, β) or (B, α)), both new agents are ranked at the bottom of the queues of their respective types.

The first part of Assumption B1 implies that, if an (A, α) pair arrives, the new agents match with one another and leave. Therefore, the queue of agents waiting remains the same as in the previous period. Note also that an arrival of a (B, β) pair does not affect the existing A -squares' (and α -rounds') decisions, because an A -square always has a β -round available to match with. Therefore, the first part of Assumption B1 implies that the arrival of a congruent pair leaves the preexisting A -squares' and α -rounds' positions unaffected. The second part of Assumption B1 guarantees that once an A -square decides to wait upon arrival, she will continue to wait in the following period upon an arrival of an (A, β) pair. In fact, when an A -square arrives with a β -round and finds too many A -squares present in the market, it is the new A -square who may decide not to wait and demand a β -round, while all existing A -squares continue to wait and demand α -rounds.

Suppose that all A -squares play a stationary strategy ψ_A with a threshold \bar{k}_A . At each period t , after a new pair arrives and priority ranks are realized, let n_α be the number of

⁴The assumption that priorities are realized before demands are submitted allows us to characterize each player's decision as a (MDP), where an agent's per-period payoff function is deterministic.

α -rounds available. Then, an A -square i who is ranked as q_i -th demands an α -round if and only if $q_i \leq n_\alpha + \bar{k}_A$.

We use an *absorbing Markov chain* to compute each A -square's, say player i 's, expected total payoff. The Markov chain is defined by the value of k_A^τ , where the event time τ starts at 0 and increases for each arrival of (A, β) or (B, α) . The state space is $\{1, 2, \dots, \bar{k}_A, \alpha\}$ with \bar{k}_A transient states and one absorbing state α . The probability transition matrix is

$$P = \begin{bmatrix} Q & R \\ \mathbf{0} & 1 \end{bmatrix}, \quad \text{where } Q = \begin{bmatrix} 0 & \frac{1}{2} & 0 & \cdots & 0 \\ \frac{1}{4} & 0 & \frac{1}{2} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & \frac{1}{2} \left(1 - \frac{1}{\bar{k}_A - 1}\right) & 0 & \frac{1}{2} \\ 0 & \cdots & 0 & \frac{1}{2} \left(1 - \frac{1}{\bar{k}_A}\right) & \frac{1}{2} \end{bmatrix} \quad \text{and } R = \begin{bmatrix} \frac{1}{2} \\ \frac{1}{4} \\ \vdots \\ \frac{1}{2(\bar{k}_A - 1)} \\ \frac{1}{2\bar{k}_A} \end{bmatrix}.$$

Let $N \equiv (I_{\bar{k}_A} - Q)^{-1}$ and $T \equiv N \cdot 1$. The absorbing Markov chain with an initial state $k \in \{1, 2, \dots, \bar{k}_A\}$ is absorbed by state α in T_k expected number of event-time periods. Note that T_k is increasing in k as player i with a larger initial state k expects to wait longer. In particular, $T_{\bar{k}_A}$ is increasing in \bar{k}_A .

An A -square decides to wait upon arrival as long as

$$U_A(\alpha) - \frac{T_{\bar{k}_A} c}{2p(1-p)} \geq U_A(\beta), \quad (2)$$

or equivalently,

$$T_{\bar{k}_A} \leq \frac{2p(1-p)(U_A(\alpha) - U_A(\beta))}{c}.$$

Let \bar{k}^{uni} be the maximal integer satisfying (2), or $\bar{k}^{uni} = 0$ if the inequality does not hold even for $\bar{k}_A = 1$. Therefore, the queue for A -squares can increase up to \bar{k}^{uni} . In any period, when the number of A -squares is strictly larger than \bar{k}^{uni} , the A -square i who is ranked last (i.e., $q_i > n_\alpha + \bar{k}^{uni}$) benefits from demanding and matching with a β -round. On the other hand, the queue for A -squares grows to at least $\bar{k}^{uni} - 1$ whenever the number of A -squares arriving is sufficiently larger than the number of α -rounds. In fact, if all A -squares use a stationary strategy ψ_A with threshold $\bar{k}_A < \bar{k}^{uni}$, a new A -square i , who finds no available α -round but \bar{k}_A existing A -squares, has a strict incentive to deviate perpetually by using the threshold $\bar{k}_A + 1$. Agent i 's expected continuation payoff from the deviation is then even higher than what it would be were all other A -squares used the threshold $\bar{k}_A + 1 \leq \bar{k}^{uni}$.




Figure 1: Thresholds of Optimal, LIFO, Random Priority, and FIFO Matching

While in the absence of a closed-form representation for \bar{k}^{uni} the comparisons between \bar{k}^{uni} , the optimal threshold \bar{k}^{opt} , and the equilibrium thresholds under FIFO and LIFO (\bar{k}^{fifo} and \bar{k}^{lifo}) are difficult to analyze, we run simulations for the same parameter values corresponding to Figures 2 and 3 in the main text: $U_A(\alpha) = U_\alpha(A) = 3$, $U_A(\beta) = U_\alpha(B) = U_B(\alpha) = U_\beta(A) = 1$, $U_B(\beta) = U_\beta(B) = 0$, and $p = 0.3$. Figures 1 and 2 here suggest that the random priority generates outcomes, in terms of both behavior and welfare, between those generated by FIFO and LIFO. In particular, random priority generates less waiting than the FIFO protocol, and more waiting than the LIFO protocol (Figure 1). The welfare gap relative to the optimal protocol is also between that generated by LIFO and FIFO (Figure 2).

3 Asymmetric Markets

In the body of the paper, we assumed a symmetric environment in terms of waiting costs and type distributions. As we discuss below, for our characterization of the optimal mechanism, asymmetries in waiting costs across market sides are not crucial. However, for the discretionary setting, the additional symmetry in utility that we imposed (namely, the assumption that $U_A(\alpha) - U_A(\beta) = U_\alpha(A) - U_\alpha(B)$), as well as the assumption of identical waiting costs for squares and rounds, could be important. In this section, we consider a market with asymmetric type distributions, utilities, and waiting costs. We characterize the optimal mechanism in such an environment and analyze a simpler one-threshold mechanism that approximates the optimal

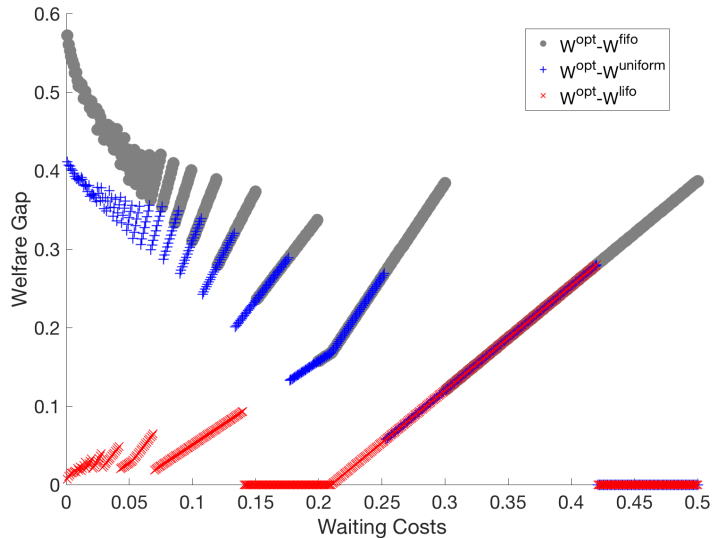


Figure 2: Welfare Gaps of LIFO, Random Priority, and FIFO relative to the Optimal Mechanism

mechanism with small waiting costs. Finally, we illustrate that the comparison between centralized and discretionary processes described in the paper carries through in this more general environment.

In what follows we allow for different type distributions for squares and rounds. Specifically, we assume the probability that a square is an A -square is p_A , while the probability that a round is an α -round is p_α . Without loss of generality, we assume that $p_A \geq p_\alpha$. Furthermore, we allow for waiting costs to differ across market sides: we denote by c_S the per-period cost experienced by squares and by c_R that experienced by rounds. We place no restrictions on match utilities other than that they are assortative and super-modular.

3.1 Optimal Dynamic Mechanism

As seen in the body of the paper, when $p_A = p_\alpha$, asymmetries in utilities play no role in the characterization of the optimal mechanism, whose welfare depends on joint match surpluses of the form $U_{xy} = U_x(y) + U_y(x)$, for $x = A, B$ and $y = \alpha, \beta$. Similarly, the optimal mechanism accounts for waiting costs incurred by *pairs*, $c_S + c_R$. An optimal mechanism can then be derived from an optimal mechanism corresponding to an environment in which waiting costs for squares and rounds coincide and are equal to $c \equiv \frac{c_S + c_R}{2}$. Our focus here is, therefore, on the impact of asymmetries in type distributions on our results, the case in which $p_A > p_\alpha$.

As in the symmetric market, (A, α) and (B, β) pairs are matched immediately when available, and we focus on dynamic mechanisms that are identified by a pair of thresholds $(\bar{k}_A, \bar{k}_\alpha)$. These thresholds do not necessarily coincide when type distributions differ for squares and rounds. Intuitively, since A -squares are more prevalent than α -rounds, it is more valuable for the mechanism designer to hold on to (B, α) pairs in the hopes of A -squares appearing in the market than it is to hold on to (A, β) pairs. We now replicate the analysis of the paper for arbitrary type distributions characterized by p_A, p_α , where $p_A \geq p_\alpha$. As in the paper, given a pair of thresholds $(\bar{k}_A, \bar{k}_\alpha)$, we find the resulting net expected time-average welfare at the steady state. We look for the pair $(\bar{k}_A^{opt}, \bar{k}_\alpha^{opt})$ that maximizes this objective.

Recall that $s_{A\alpha}^t$ denotes the value of the (signed) length of the A - α queue at the beginning of time t . $\mathbf{x}^t \in \{0, 1\}^{\bar{k}_A + k_\alpha + 1}$ is the timed vector such that x_i^t takes the value of 1 if the state is $s_{A\alpha}^t$ and 0 otherwise. Then,

$$\mathbf{x}^{t+1} = \mathbf{T}_{\bar{k}_A, \bar{k}_\alpha} \mathbf{x}^t,$$

where

$$\mathbf{T}_{\bar{k}_A, \bar{k}_\alpha} = \begin{pmatrix} 1 - (1 - p_A)p_\alpha & p_A(1 - p_\alpha) & \dots & 0 & 0 \\ (1 - p_A)p_\alpha & p_A p_\alpha + (1 - p_A)(1 - p_\alpha) & \dots & 0 & 0 \\ 0 & (1 - p_A)p_\alpha & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & p_A(1 - p_\alpha) & 0 \\ 0 & 0 & \dots & p_A p_\alpha + (1 - p_A)(1 - p_\alpha) & p_A(1 - p_\alpha) \\ 0 & 0 & \dots & (1 - p_A)p_\alpha & 1 - p_A(1 - p_\alpha) \end{pmatrix}.$$

Since the above Markov chain is ergodic, the corresponding matching process reaches a unique steady state with a distribution $\boldsymbol{\pi} \equiv (\pi_{\bar{k}_A}, \pi_{\bar{k}_A - 1}, \dots, \pi_{-\bar{k}_\alpha})$ that we now identify.

Denote

$$\begin{aligned} \eta &\equiv p_A(1 - p_\alpha) + (1 - p_A)p_\alpha, \quad \text{and} \\ \phi &\equiv \frac{(1 - p_A)p_\alpha}{p_A(1 - p_\alpha)} \quad (< 1). \end{aligned}$$

The Markov transition matrix is then

$$\mathbf{T}_{\bar{k}_A, \bar{k}_\alpha} = \begin{pmatrix} 1 - \eta + \frac{\eta}{\phi + 1} & \frac{\eta}{\phi + 1} & \dots & 0 & 0 \\ \frac{\eta\phi}{\phi + 1} & 1 - \eta & \dots & 0 & 0 \\ 0 & \frac{\eta\phi}{\phi + 1} & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & \frac{\eta}{\phi + 1} & 0 \\ 0 & 0 & \dots & 1 - \eta & \frac{\eta}{\phi + 1} \\ 0 & 0 & \dots & \frac{\eta\phi}{\phi + 1} & 1 - \eta + \frac{\eta\phi}{\phi + 1} \end{pmatrix}.$$

In particular,

$$\begin{aligned}
\pi_{\bar{k}_A} &= \left(1 - \eta + \frac{\eta}{\phi + 1}\right) \pi_{\bar{k}_A} + \frac{\eta}{\phi + 1} \pi_{\bar{k}_A - 1} \implies \pi_{\bar{k}_A - 1} = \phi \pi_{\bar{k}_A}, \\
\pi_{\bar{k}_A - 1} &= \frac{\eta \phi}{\phi + 1} \pi_{\bar{k}_A} + (1 - \eta) \pi_{\bar{k}_A - 1} + \frac{\eta}{\phi + 1} \pi_{\bar{k}_A - 2} \implies \pi_{\bar{k}_A - 2} = \phi \pi_{\bar{k}_A - 1} = \phi^2 \pi_{\bar{k}_A}, \\
&\vdots \\
\pi_{-\bar{k}_A \alpha} &= \frac{\eta \phi}{\phi + 1} \pi_{-\bar{k}_\alpha + 1} + \left(1 - \eta + \frac{\eta \phi}{\phi + 1}\right) \pi_{\bar{k}_A} + \frac{\eta}{\phi + 1} \pi_{\bar{k}_A - 1} \implies \pi_{-\bar{k}_\alpha} = \phi^{\bar{k}_A + \bar{k}_\alpha} \pi_{\bar{k}_A}.
\end{aligned}$$

Since $\sum_{k=0}^{\bar{k}_A + \bar{k}_\alpha} \phi^k \pi_{\bar{k}_A} = 1$, it follows that

$$\pi_{\bar{k}_A} = \frac{1 - \phi}{1 - \phi^{\bar{k}_A + \bar{k}_\alpha + 1}}.$$

Therefore,

$$\pi_{\bar{k}_A - k} = \frac{(1 - \phi) \phi^k}{1 - \phi^{\bar{k}_A + \bar{k}_\alpha + 1}} \quad \text{for every } k = 0, 1, \dots, \bar{k}_A + \bar{k}_\alpha.$$

Then the expected time-average match surplus at the steady state is:

$$\begin{aligned}
S(\bar{k}_A, \bar{k}_\alpha) &= p_A p_\alpha U_{A\alpha} + (1 - p_A)(1 - p_\alpha) U_{B\beta} \\
&+ \mathbf{1}\{k_A > 0\} \left(\sum_{k=1}^{\bar{k}_A} \phi^{\bar{k}_A - k} \pi_{\bar{k}_A} (1 - p_A) p_\alpha (U_{A\alpha} + U_{B\beta}) \right) \\
&+ \mathbf{1}\{k_\alpha > 0\} \left(\sum_{k=1}^{\bar{k}_\alpha} \phi^{\bar{k}_A + k} \pi_{\bar{k}_A} p_A (1 - p_\alpha) (U_{A\alpha} + U_{B\beta}) \right) \\
&+ \pi_{\bar{k}_A} p_A (1 - p_\alpha) U_{A\beta} + \phi^{\bar{k}_A + \bar{k}_\alpha} \pi_{\bar{k}_A} (1 - p_A) p_\alpha U_{B\alpha}.
\end{aligned}$$

The expected time-average waiting costs at the steady state are:

$$C(\bar{k}_A, \bar{k}_\alpha) = 2c \pi_{\bar{k}_A} \left(\sum_{k=0}^{\bar{k}_A} k \phi^{\bar{k}_A - k} + \sum_{k=0}^{\bar{k}_\alpha} k \phi^{\bar{k}_A + k} \right).$$

The optimal dynamic mechanism is identified by the pair of thresholds $(\bar{k}_A^{opt}, \bar{k}_\alpha^{opt})$ that maximize the expected time-average welfare, $S(\bar{k}_A, \bar{k}_\alpha) - C(\bar{k}_A, \bar{k}_\alpha)$.

3.2 One-Threshold Mechanisms

When A -squares are strictly more likely to arrive than α -rounds ($p_A > p_\alpha$), there is a relatively small chance that many α -rounds arrive at the market and are not matched with A -squares.

In other words, the (signed) length of the A - α queue is unlikely to reach very negative values. Therefore, we can consider a simpler mechanism, which only limits the length of the queue of A -squares. It turns out that the most efficient one-threshold mechanism, in spite of being less efficient than the optimal mechanism, is asymptotically efficient as waiting costs, c_S and c_R , vanish.

We find the expected total welfare for one period of time of the two-threshold dynamic mechanism $(\bar{k}_A, \bar{k}_\alpha)$ as \bar{k}_α becomes infinitely large.⁵

In the limit,

$$\pi_{\bar{k}_A - k} = (1 - \phi)\phi^k \quad \text{for every } k = 0, 1, 2, \dots$$

By applying this limit steady-state distribution to the formulations of $S(\bar{k}_A, \bar{k}_\alpha)$ and $C(\bar{k}_A, \bar{k}_\alpha)$ above, we obtain the corresponding limit match surplus and costs $S(\bar{k}_A, \infty)$ and $C(\bar{k}_A, \infty)$, respectively.

$$\begin{aligned} S(\bar{k}_A, \infty) &= p_A p_\alpha U_{A\alpha} + (1 - p_A)(1 - p_\alpha)U_{B\beta} \\ &+ \mathbf{1}\{k_A > 0\} \left(\sum_{k=0}^{\bar{k}_A - 1} (1 - \phi)\phi^k (1 - p_A)p_\alpha (U_{A\alpha} + U_{B\beta}) \right) \\ &+ \left(\sum_{k=1}^{\infty} (1 - \phi)\phi^{\bar{k}_A + k} p_A (1 - p_\alpha) (U_{A\alpha} + U_{B\beta}) \right) \\ &+ (1 - \phi)p_A (1 - p_\alpha)U_{A\beta}. \end{aligned}$$

and

$$C(\bar{k}_A, \infty) = 2c(1 - \phi) \left(\sum_{k=0}^{\bar{k}_A} k\phi^{\bar{k}_A - k} + \sum_{k=0}^{\infty} k\phi^{\bar{k}_A + k} \right).$$

We can simplify the above expressions to achieve, for every $\bar{k}_A = 0, 1, 2, \dots$,

$$\begin{aligned} S(\bar{k}_A, \infty) &= p_\alpha U_{A\alpha} + (1 - p_A)U_{B\beta} + (p_A - p_\alpha)U_{A\beta} = S_\infty, \quad \text{and} \\ C(\bar{k}_A, \infty) &= 2c(1 - \phi)\phi^{\bar{k}_A} \left(\frac{\phi}{(1 - \phi)^2} + \sum_{k=0}^{\bar{k}_A} k\phi^{-k} \right). \end{aligned}$$

The expected time-average welfare is $W(\bar{k}_A) \equiv S(\bar{k}_A, \infty) - C(\bar{k}_A, \infty)$.

⁵Technically, a one-threshold mechanism defines a Markov chain with a countable state space $\{\dots, -1, 0, 1, \dots, \bar{k}_A\}$. However, when transitions toward state \bar{k}_A occur with probability strictly higher than that of transitions away from state \bar{k}_A (i.e., $p_A(1 - p_\alpha) > p_\alpha(1 - p_A)$), the steady-state probabilities for the truncated Markov chain defined by a two-threshold mechanism $(\bar{k}_A, \bar{k}_\alpha)$ approach the steady-state probabilities for the untruncated Markov chain as \bar{k}_α increases.

We inspect the marginal time-average welfare with respect to the length of the queue of A -squares $\Delta_+W(\bar{k}_A) \equiv W(\bar{k}_A + 1) - W(\bar{k}_A)$ and find the most efficient one-threshold \bar{k}_A^{**} from

$$\Delta_+W(\bar{k}_A^{**}) \leq 0 \leq \Delta_+W(\bar{k}_A^{**} - 1). \quad (3)$$

Now, to derive a closed-form solution for \bar{k}_A^{**} , notice that the expected total surplus $S(\bar{k}_A, \infty)$ is a constant function of \bar{k}_A . Therefore,

$$\begin{aligned} \Delta_+W(\bar{k}_A) &= C(\bar{k}_A, \infty) - C(\bar{k}_A + 1, \infty) \\ &= 2c(1 - \phi) \left(\phi^{\bar{k}_A} - \phi^{\bar{k}_A+1} \right) \left(\frac{\phi}{(1 - \phi)^2} + \sum_{k=0}^{\bar{k}_A} k\phi^{-k} \right) \\ &\quad + 2c(1 - \phi)\phi^{\bar{k}_A+1} \left(\sum_{k=0}^{\bar{k}_A} k\phi^{-k} - \sum_{k=0}^{\bar{k}_A+1} k\phi^{-k} \right) = 2c(2\phi^{\bar{k}_A+1} - 1). \end{aligned}$$

The most efficient one-threshold mechanism is identified from (3) as:

$$\bar{k}_A^{**} = \left\lfloor -\frac{\log 2}{\log \phi} \right\rfloor = \left\lfloor -\frac{\log 2}{\log(1 - p_A) + \log p_\alpha - \log p_A - \log(1 - p_\alpha)} \right\rfloor.$$

The most efficient one-threshold mechanism \bar{k}_A^{**} does not depend on c . In fact, every fixed one-threshold mechanism is asymptotically efficient with vanishingly small waiting costs. Intuitively, in the one-threshold mechanism, an incongruent pair leaves the market only when the state is $k_{A\alpha} = \bar{k}_A$, which always occurs with probability $1 - \phi$ at the steady state. Therefore, all one-threshold mechanisms result in the same expected fraction of incongruent pairs matched in the steady state. In fact, the expected total time-average match surplus is S_∞ regardless of the threshold \bar{k}_A . For any fixed threshold \bar{k}_A , as waiting costs approach zero, the expected total time-average waiting costs approach zero and efficiency is achieved.

As a corollary, it follows that the optimal (two-threshold) mechanism is approximately efficient as waiting costs vanish.

3.3 Discretionary Matching

We focus on *regular* environments in which

$$p_\alpha(U_A(\alpha) - U_A(\beta)) \neq kc_S \quad \text{and} \quad p_A(U_\alpha(A) - U_\alpha(B)) \neq kc_R$$

for every $k \in \mathbb{Z}_+$.

The decisions of an A -square (and, analogously, an α -round) remain as described in the body of the paper. Namely, when an A -square arrives at the market and an α -round is available, an (A, α) pair is formed immediately. If an α -round is not available, the arriving A -square decides to wait in the queue based on the number of A -squares already in the queue. Since an α -round is not available, this implies that the A -square arrived with a β -round. As all β -rounds in the market are willing to match with A -squares, the newly arrived A -square will wait if and only if the gain $U_A(\alpha) - U_A(\beta)$ exceeds the expected waiting costs till matching with an α -round.

In analogy with Lemma 1 in the paper, the (signed) length of the A - α queue at the beginning of a period, $k_{A\alpha} \equiv k_A - k_\alpha$, will then be bounded as follows

$$-\bar{k}_\alpha^{ffo} \leq k_{A\alpha} \leq \bar{k}_A^{ffo},$$

where

$$\begin{aligned} \bar{k}_A^{ffo} &\equiv \max \left\{ k \in \mathbb{Z}_+ \mid \frac{k_{CS}}{p_\alpha} < U_A(\alpha) - U_A(\beta) \right\}, \quad \text{and} \\ \bar{k}_\alpha^{ffo} &\equiv \max \left\{ k \in \mathbb{Z}_+ \mid \frac{k_{CR}}{p_A} < U_\alpha(A) - U_\alpha(B) \right\}. \end{aligned}$$

A β -round (similarly, a B -square) may decide to wait to match with an A -square if the queue of A -squares is long and expected to hit the threshold \bar{k}_A^{ffo} within a sufficiently short time. In contrast with the symmetric case, B -squares and β -rounds may now wait simultaneously in equilibrium. Intuitively, consider an environment in which both types of rounds are nearly indifferent between matching with A -squares or B -squares and therefore match with whomever is available immediately. In such an environment, a B -square, who is first in line, may decide to wait in the market, even when arriving with a β -round, in the hopes of a (B, α) pair arriving in the next period. In other words, in general asymmetric markets, Lemma 2 of the paper does not hold.

A full characterization of the equilibrium requires the analysis of a rather complex random process of the 3-dimensional vector $(k_{A\alpha}, k_B, k_\beta)$. As such, in this Online Appendix we study a one-dimensional Markov process of $k_{A\alpha}$ only, with a transition matrix as described in Section 3.1 in order to achieve bounds on the equilibrium welfare in the asymmetric discretionary process.

In equilibrium, as well as under the one-dimensional protocol discussed above, the expected time-average surplus is bounded above by $S_\infty = p_\alpha U_{A\alpha} + (1 - p_\alpha)U_{B\beta} + (p_A - p_\alpha)U_{A\beta}$. In

the one-dimensional process with thresholds $\bar{k}_\alpha^{fif\o}$ and $\bar{k}_A^{fif\o}$, at each state $k_{A\alpha}$, either $k_{A\alpha}$ A -squares (and at least as many β -rounds) or $|k_{A\alpha}|$ α -rounds (and at least as many B -squares) incur waiting costs. Since in equilibrium there might be additional waiting costs incurred through the simultaneous waiting of B -squares or β -rounds, it follows that the resulting per period welfare $W^{fif\o}(c_S, c_R)$ can be bounded as follows:

$$W^{fif\o}(c_S, c_R) \leq S_\infty - (c_S + c_R)\pi_{\bar{k}_A^{fif\o}} \left(\sum_{k=0}^{\bar{k}_A^{fif\o}} k\phi_{\bar{k}_A^{fif\o}-k} + \sum_{k=0}^{\bar{k}_\alpha^{fif\o}} k\phi_{\bar{k}_A^{fif\o}+k} \right).$$

After some algebraic manipulation, we can show that

$$\begin{aligned} \lim_{(c_S, c_R) \rightarrow (0,0)} W^{fif\o}(c_S, c_R) &\leq S_\infty - \lim_{(c_S, c_R) \rightarrow (0,0)} (c_S + c_R)\bar{k}_A^{fif\o} \leq S_\infty - \lim_{c_S \rightarrow 0} c_S \bar{k}_A^{fif\o} \\ &\leq S_\infty - p_\alpha(U_A(\alpha) - U_A(\beta)). \end{aligned}$$

This echoes Corollary 3 in the main text of the paper. The bound on the welfare wedge between the discretionary protocol and the optimal mechanism exhibits similar comparative statics to those described in the paper, increasing in p_α and in $U_A(\alpha) - U_A(\beta)$.

4 Richer Type Sets

The paper considers a simplified setting in which types are binary—there are only two types of squares and two types of rounds. This assumption allows us to illustrate the forces acting in centralized and discretionary dynamic matching markets in a simple and transparent way. A natural extension of our setting is to an environment in which each square can take one of multiple types S_1, \dots, S_l and each round can take one of multiple types R_1, \dots, R_m with distributions p and q , respectively. Denote by $U_{S_i}(R_j)$ the utility a square of type S_i gets from matching with a round of type R_j and by $U_{R_j}(S_i)$ the utility a round of type R_j gets from matching with a square of type S_i . Assume that preferences are assortative, so that:

$$\begin{aligned} U_{S_i}(R_1) &> U_{S_i}(R_2) > \dots > U_{S_i}(R_m) \quad \text{for all } i; \text{ and} \\ U_{R_j}(S_1) &> U_{R_j}(S_2) > \dots > U_{R_j}(S_l) \quad \text{for all } j. \end{aligned}$$

Assume also that preferences are super-modular, so that:

$$\begin{aligned} U_{S_i}(R_v) - U_{S_i}(R_w) &> U_{S_j}(R_v) - U_{S_j}(R_w) \quad \text{and} \\ U_{R_v}(S_i) - U_{R_v}(S_j) &> U_{R_w}(S_i) - U_{R_w}(S_j) \quad \text{for all } i < j \text{ and } v < w. \end{aligned}$$

These assumptions mirror the assumptions asserted by Becker (1974) for assortative super-modular markets and correspond to a set of environments that contain the one studied in the body of the paper as a special case.

The full characterization of the optimal mechanism in this more general setting requires some new tools and is left for future research. However, our analysis of the binary-type environment does provide a set of necessary restrictions on the optimal mechanism in this setting. Certainly, since utilities are super-modular, whenever an S_1 -square and an R_1 -round, or an S_l -square and an R_m -round, are available in the market, they are to be matched immediately. Next, consider the potential match of a pair (S_y, R_z) , with $1 \leq y \leq l, 1 \leq z \leq m$, available on the market. A decision of *not* matching the pair can be justified by either one of the following scenarios:

1. planning to match S_y to $R_{z'}$ with $z < z'$ and R_z to $S_{y'}$ with $y' < y$, or
2. planning to match S_y to $R_{z'}$ with $z > z'$ and R_z to $S_{y'}$ with $y' > y$.

In any other scenario, both S_y and R_z match with better partners than each other, or both match with worse partners than each other. Since utilities are super-modular, this would lower overall welfare and entail additional waiting costs incurred by S_y and R_z . Therefore, the pair (S_y, R_z) is matched immediately when neither scenario 1 nor scenario 2 hold.

To obtain necessary conditions for the optimal mechanism in this setting, we analyze two induced binary-type markets, such that (S_y, R_z) constitutes an incongruent pair in both. As for the first binary-type market, if $1 \leq y < l$ and $1 < z \leq m$, consider the following partition of the types of squares and rounds: $\mathbf{A} = \{S_1, \dots, S_y\}$, $\mathbf{B} = \{S_{y+1}, \dots, S_l\}$, $\boldsymbol{\alpha} = \{R_1, \dots, R_{z-1}\}$, and $\boldsymbol{\beta} = \{R_z, \dots, R_m\}$. In particular, an S_y -square in the original market is classified as an \mathbf{A} -square in this induced binary market and an R_z -round in the original market is classified as a $\boldsymbol{\beta}$ -round in the induced binary market. Accordingly, denote

$$p_{\mathbf{A}} \equiv \sum_{i=1}^y p_i \quad \text{and} \quad p_{\boldsymbol{\alpha}} \equiv \sum_{j=1}^{z-1} q_j$$

so that squares are of either type \mathbf{A} or \mathbf{B} (with probability $p_{\mathbf{A}}$ or $1 - p_{\mathbf{A}}$) and rounds are of either type $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ (with probability $p_{\boldsymbol{\alpha}}$ or $1 - p_{\boldsymbol{\alpha}}$). Finally, we choose utilities in the first binary-type market that are consistent with the partition of squares and rounds we defined

above and *maximize* the extent of super-modularity, i.e., the efficiency gain from matching agents assortatively. Specifically, define

$$(\bar{x}, \bar{w}) \in \arg \max_{\substack{\{(x, w) \mid x = y + 1, \dots, l, \\ w = 1, \dots, z - 1\}}} U_{S_y R_w} + U_{S_x R_z} - U_{S_y R_z} - U_{S_x R_w}.$$

In words, \bar{x} and \bar{w} are chosen in $\{y + 1, \dots, l\}$ and in $\{1, \dots, z - 1\}$, respectively, to maximize supermodularity. However, for any $x > y$ and $w \in \{1, \dots, z - 1\}$, we have $U_{S_y R_1} - U_{S_x R_1} \geq U_{S_y R_w} - U_{S_x R_w}$ because of supermodularity, which guarantees that $\bar{w} = 1$. Therefore, we have

$$\bar{x} \in \arg \max_{x \in \{y+1, \dots, l\}} U_{S_x R_z} - U_{S_x R_1},$$

the solution of which, since $U_{S_x R_z} - U_{S_x R_1} \leq 0$, is $\bar{x} = l$. Then, we choose:

$$\begin{aligned} U_A(\alpha) &\equiv U_{S_y}(R_1), & U_A(\beta) &\equiv U_{S_y}(R_z), \\ U_B(\alpha) &\equiv U_{S_l}(R_1), & U_B(\beta) &\equiv U_{S_l}(R_z), \\ U_\alpha(A) &\equiv U_{R_1}(S_y), & U_\alpha(B) &\equiv U_{R_1}(S_l), \\ U_\beta(A) &\equiv U_{R_z}(S_y), & U_\beta(B) &\equiv U_{R_z}(S_l). \end{aligned}$$

As for the second binary-type market, if $1 < y \leq l$ and $1 \leq z < m$, we consider the binary-type partition, $\mathbf{A} = \{S_1, \dots, S_{y-1}\}$, $\mathbf{B} = \{S_y, \dots, S_l\}$, $\alpha = \{R_1, \dots, R_z\}$, and $\beta = \{R_{z+1}, \dots, R_m\}$. We define type distributions, lengths of queues, and utilities to maximize the extent of super-modularity similarly to the first binary-type market, which in this case are defined as:

$$\begin{aligned} U_A(\alpha) &\equiv U_{S_1}(R_z), & U_A(\beta) &\equiv U_{S_1}(R_m), \\ U_B(\alpha) &\equiv U_{S_y}(R_z), & U_B(\beta) &\equiv U_{S_y}(R_m), \\ U_\alpha(A) &\equiv U_{R_z}(S_1), & U_\alpha(B) &\equiv U_{R_z}(S_y), \\ U_\beta(A) &\equiv U_{R_m}(S_1), & U_\beta(B) &\equiv U_{R_m}(S_y). \end{aligned}$$

Notice that any state in our original multiple-type market can be mapped into a state in each of the two induced binary markets, for which our analysis identifies the optimal mechanism, see Section 3.1 in this Online Appendix. We are now ready to state necessary conditions for the optimal mechanism in the richer type-sets case as follows:

Proposition B3 *In any scenario in which both the match (A, β) and the match (B, α) are created in the first and second binary market, respectively, then the match (S_y, R_z) is created immediately by any optimal mechanism in the original market.*

Following a similar logic, we can identify another set of necessary conditions of the optimal mechanism. Specifically, in the binary-type market with partitions $\mathbf{A} = \{S_1, \dots, S_y\}$, $\mathbf{B} = \{S_{y+1}, \dots, S_l\}$, $\boldsymbol{\alpha} = \{R_1, \dots, R_{z-1}\}$, and $\boldsymbol{\beta} = \{R_z, \dots, R_m\}$, we define the utilities in the first binary-type market as before, but now \bar{x} and \bar{w} are chosen to *minimize* supermodularity as follows:

$$(\bar{x}, \bar{w}) \in \arg \min_{\substack{\{(x, w) \mid x = y + 1, \dots, l, \\ w = 1, \dots, z - 1\}}} U_{S_y R_w} + U_{S_x R_z} - U_{S_y R_z} - U_{S_x R_w}.$$

For arguments similar to the ones above, it is easy to see that $\bar{x} = y + 1$ and $\bar{w} = z - 1$. In addition, we can define the second binary-type market as $\mathbf{A} = \{S_1, \dots, S_{y-1}\}$, $\mathbf{B} = \{S_y, \dots, S_l\}$, $\boldsymbol{\alpha} = \{R_1, \dots, R_z\}$, and $\boldsymbol{\beta} = \{R_{z+1}, \dots, R_m\}$, with corresponding supermodularity-minimizing utilities (where now $\bar{x} = y - 1$ and $\bar{w} = z + 1$). Then, the utility specification of the two binary-type markets assures the following necessary conditions for the optimal mechanism:

Proposition B4 *In any scenario in which the matches (A, β) and (B, α) are not formed in any of the binary-type markets, then the match (S_y, R_z) is not created by any optimal mechanism in the original market.*

5 Independent Arrivals

We have assumed that at each period a square and a round arrive at the market. As mentioned in Section 2 of the paper, allowing for pairs to arrive at random times (say, through a Poisson arrival process) would not change any of the analysis. The only modification would be the effective length of a period—instead of a unit of time, it would correspond to the expected time until a new pair arrives.

Allowing for independent arrivals of squares and rounds requires some changes in our analysis, however. To see this, consider a symmetric environment in which a square arrives with probability q each period and, similarly, a round arrives with probability q each period. As we assumed throughout the paper, a newly arrived square is of type A or B with probability p or $1 - p$, respectively. Likewise, a newly arrived round is of type α or β with probability p or $1 - p$, respectively. In this setting, there is a probability that a long queue of squares (or, similarly, rounds) would form with no round (or square) available, of whichever type.

For any fixed value of the outside option, there would be a sufficiently high threshold above which it would be optimal to retire a square and allow her to benefit from the outside option instead of experiencing a prohibitively long wait. In addition to our symmetry assumptions on utilities, suppose that A -squares and α -rounds face equivalent outside options and that, similarly, B -squares and β -rounds face equivalent outside options. It can be shown that the optimal mechanism is identified by a threshold \bar{k}^* such that whenever the number of A -squares exceeds \bar{k}^* , A -squares are matched with β -rounds if those are available. Similarly, whenever the number of α -rounds exceeds \bar{k}^* , α -rounds are matched with B -squares if those are available. There is also a second threshold $\bar{k}^H \geq \bar{k}^*$ such that any excess of A -squares beyond \bar{k}^H and any excess of α -rounds beyond \bar{k}^H is retired from the market immediately (in the absence of β -rounds or B -squares to match them with). Last, the mechanism specifies when to retire B -squares or β -rounds, depending potentially on the number of available α -rounds or A -squares, respectively.

6 Additional Proofs

6.1 Proofs Regarding Conditions 1 and 2

We present the proof of Lemma A1 from the Appendix of the main text, followed by Theorem B1, which illustrates that the restriction to SD-mechanisms satisfying Conditions 1 and 2 we imposed in the main text is without loss of generality.

Proof of Lemma A1: We prove that for any *deterministic* mechanism μ , there exists another *deterministic* mechanism μ' such that Conditions 1 and 2 hold and $v(\mu') \geq v(\mu)$. The result extends to random mechanisms, which are essentially convex combinations of deterministic mechanisms.

(1) Take any (deterministic) mechanism μ that may hold some (A, α) pairs after some histories. Consider another mechanism μ' that creates the same set of matches as μ at every history, except that (i) when μ holds an (A, α) pair, say agents (i, j) , μ' matches the pair as soon as they are available, (ii) μ' does not create any match that μ creates involving either i or j , and (iii) if μ matches i to a round $r (\neq j)$ in period $t' > t$, and matches j to a square $s (\neq i)$ in period $t'' > t$, then μ' forms the match (s, r) in period $\max\{t', t''\}$. It is clear from the construction that $v(\mu') \geq v(\mu)$ since μ' involves strictly lower waiting costs and a weakly

higher match surplus than μ does for every finite time horizon. The argument that we can further improve μ' by matching (B, β) pairs immediately follows analogously.

(2) Take a (deterministic) mechanism μ , and assume, without loss of generality from the previous step, that (i) μ matches a newly arriving pair of (A, α) or (B, β) immediately. Assume additionally that, (ii) μ matches A -squares (or β -rounds) who are held in the market to α -rounds (respectively, B -squares) on a first-in-first-out (FIFO) basis, and to β -rounds (respectively, A -squares) on a last-in-first-out (LIFO) basis. That is, when μ matches an existing A -square to an α -round upon an arrival of (B, α) , it selects the A -square who arrived first among all existing A -squares. However, when μ matches an existing A -square to a β -round, it selects the A -square who arrived last. This assumption does not affect $v(\mu)$ since agents of the same type are interchangeable from a welfare perspective.

Suppose that μ holds $m = \lfloor \frac{U}{2c} \rfloor$ or more (A, β) pairs at some histories. We construct another mechanism μ' that creates the same set of matches as μ at every history, except that whenever μ holds m or more (A, β) pairs, μ' holds only $\{(A_i, \beta_i)\}_{i=1, \dots, m}$ (where the index i marks the order of arrival at the market, with lower indices denoting more recent arrivals) and matches $\{(A_i, \beta_i)\}_{i > m}$ immediately. Naturally, at later periods, when μ creates some matches with agents who have already left under μ' , the mechanism μ' does not create those matches. Also, for any (A, β) pair, say (i, j) , that μ' creates but μ does not (namely, $i > m$), if μ matches i to a round $r (\neq j)$, and j to a square $s (\neq i)$ at some later periods, μ' matches (r, s) as soon as they become available.

We claim that $v(\mu') \geq v(\mu)$. To see this, take any (A, β) pair, say (i, j) , held by the mechanism μ , but not μ' , at some period t . Assumptions (i) and (ii) above and the fact that only one pair arrives at each period guarantee that either of the following options occur: (a) μ matches the pair (i, j) at some period after t , or (b) μ matches i to an α -round $r (\neq j)$, and j to a B -square $s (\neq i)$ at the same period $t' \geq t + m$ upon the arrival of (r, s) .

In both cases (a) and (b), μ generates a lower average welfare than μ' for every finite time horizon. Case (a) is clear. For case (b), note that the additional waiting costs generated by μ are strictly higher than $(\frac{U}{2c})(2c)$ and therefore exceed the highest surplus gain U that can be generated by holding i and j and matching them with others: i.e., $U_{ir} + U_{sj} - U_{ij} - U_{rs} \leq U$.

We omit an analogous proof showing that we can further improve μ' by not holding more than $\frac{U}{2c}$ number of (B, α) pairs on the market. ■

As described in the main text, Lemma A1 allows us to simplify our problem using the following Markov decision problem with agents arriving in incongruent pairs, a finite set of states, and a finite set of actions:

$$(MDP, s^0) \equiv \{T, S, s^0, K, (r(s, k), p(\cdot|k))_{s \in S, k \in A_s}\},$$

where s^0 denotes a particular initial state. Each component is defined as follows:

1. $T \equiv \{0, 1, 2, \dots\}$ is *the set of decision event times*. As described in the body of the paper, event times correspond to times at which incongruent pairs (A, β) or (B, α) arrive. Since the probability of an incongruent pair arriving at any period is $2p(1-p)$, the expected time between event times is $\frac{1}{2p(1-p)}$.
2. $S \equiv \{z \in \mathbb{Z} : -(U/2c) - 1 \leq z \leq (U/2c) + 1\}$ is *the set of possible states (or stocks)*. Each state $s_{A\alpha} \equiv s_A - s_\alpha \in S$ represents the (signed) number of incongruent pairs of type (A, β) or (B, α) in the market. Since we restrict our attention to mechanisms that do not hold more than $U/2c$ squares (and rounds), a state, which takes a new arriving pair into account, has to belong to the set $\{-\lfloor U/2c \rfloor - 1, \dots, \lfloor U/2c \rfloor + 1\}$.
3. $s^0 = 0$ is *the initial state*. Initially, no agent waits.
4. $K \equiv \{z \in \mathbb{Z} : -U/2c \leq z \leq U/2c\}$ is *the set of available actions*. Each $k \in K$ represents the (signed) number of incongruent pairs held in the market from one period to the next.
5. $r(s, k)$ is *the reward function*: for every $s \in S, k \in K$,

$$r(s, k) = \begin{cases} (s - k)U_{A\beta} - \frac{kc}{2p(1-p)} & \text{if } s \geq k \geq 0 \\ (|s| - |k|)U_{B\alpha} - \frac{|k|c}{2p(1-p)} & \text{if } s \leq k \leq 0 \\ -\infty & \text{otherwise.} \end{cases}$$

The expected waiting cost incurred to any agent who waits for one event time is $\frac{c}{2p(1-p)}$. The reward function returns $-\infty$ if an action is infeasible. For all feasible actions, the values of the reward function are in the interval $\left[-\frac{U}{4p(1-p)}, (\frac{U}{2c} + 1)U_{A\alpha}\right]$.

6. $p(s, k)$ is *the transition probability*, the probability the system is in state $s \in S$ at any time $\tau + 1$, after the action k has been chosen at time τ . In particular,

$$p(s, k) = \begin{cases} 1/2 & \text{for } s = k - 1, k + 1, \\ 0 & \text{otherwise.} \end{cases}$$

(MDP, s^0) is stationary in the sense that the reward function $r(s, k)$ and the transition probability function $p(s, k)$ do not depend on time, or event times, explicitly. A *policy* of (MDP, s^0) is any rule, deterministic or randomized, governing the choice of actions. Such a rule may, in principle, be history-dependent. The value of a policy μ is then,

$$v(\mu) \equiv \liminf_{T \rightarrow \infty} \frac{1}{T} E_{\mu} \left[\sum_{\tau=1}^T r(s^{\tau}, k^{\tau}) \right].$$

A stationary and deterministic policy, which we call a *SD-policy*, of (MDP, s^0) applies the same deterministic decision rule $\mu^{SD} : S \rightarrow K$ regardless of the history. The value of μ^{SD} is then

$$v(\mu^{SD}) = \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\sum_{\tau=1}^T r(s^{\tau}, \mu^{SD}(s^{\tau})) \right].$$

The limit exists, as guaranteed, for example, by Proposition 8.1.1(b) in Puterman (2005).

We now show that restricting attention to mechanisms satisfying Conditions 1 and 2 in the main text is without loss of generality.

Theorem B1 *There exists an optimal SD-mechanism that satisfies Conditions 1 and 2.*

Proof of Theorem B1: In what follows, we prove that there exists an optimal SD-policy of (MDP, s^0) . By Lemma A1, an optimal SD-policy of (MDP, s^0) defines an optimal SD-mechanism of our matching problem. The following result from Ross (2014) will be useful.

Theorem Ross 1 (V.2.1 in Ross, 2014) *If there exists a bounded function $h(s)$, $s \in S$, and a constant g such that*

$$g + h(s) = \max_{k \in K} \left[r(s, k) + \sum_{s' \in S} p(s', k) h(s') \mid s \in S \right] \text{ for all } s \in S, \quad (4)$$

1. *there exists an SD-policy μ^* such that*

$$g = v(\mu^*) = \sup_{\mu} v(\mu);$$

2. *μ^* is any SD-policy that, for each $s \in S$, prescribes an action k that maximizes the RHS of (4).*

We extend the notion of (MDP, s^0) by allowing the initial state to be an arbitrary $s \in S$, and we denote the Markov decision problem with an arbitrary initial state by (MDP) . It is straightforward to extend the definition of a policy and a SD -policy to (MDP) . Let E_μ represent an expectation conditional on policy μ being implemented. For any $0 < \delta < 1$, initial state $s \in S$, and a policy μ of (MDP) , define

$$v_\delta(\mu; s) \equiv E_\mu \left[\sum_{\tau=0}^{\infty} r(s^\tau, k^\tau) \delta^\tau \mid s \right].$$

For each $s \in S$, $v_\delta(\mu; s)$ represents the expected total discounted return earned when the policy μ of (MDP) is employed. Since the reward function is bounded by the interval $\left[-\frac{U}{4p(1-p)}, \frac{U}{2c}U_{A\alpha}\right]$, and $0 < \delta < 1$, the expectation is well-defined for all policies that implement feasible actions. For any $\delta \in (0, 1)$ and initial state $s \in S$, let

$$v_\delta(s) \equiv \sup_{\mu} v_\delta(\mu; s).$$

The following second theorem from Ross (2014) describes sufficient conditions for applying Theorem Ross 1, and therefore guaranteeing that there exists an optimal SD -policy of (MDP) , which is also an optimal SD -policy of (MDP, s^0) and defines an optimal SD -mechanism for our dynamic matching problem.

Theorem Ross 2 (V.2.2 in Ross, 2014) *If there exists $M < \infty$ such that*

$$|v_\delta(s) - v_\delta(0)| < M \text{ for all } \delta \in (0, 1) \text{ and } s \in S,$$

then there exists a bounded function $h(s)$ and a constant g satisfying (4).

By Theorems Ross 1 and Ross 2, the following claim is sufficient to complete the proof of Theorem B1.

Claim: *There exists $M < \infty$, such that, for any $s \in S$, $\delta \in (0, 1)$, and policy μ of (MDP) , there exists another policy μ' of (MDP) with*

$$|v_\delta(\mu; s) - v_\delta(\mu'; 0)| < M.$$

Proof of Claim: Take any initial state $s \in S$, $\delta \in (0, 1)$, and policy μ of (MDP) . Let $s > 0$ (a similar proof, which we omit, applies for the case of $s < 0$). We assume, as above, that

μ matches agents on a FIFO basis. Let $m \equiv \lfloor (U/2c) + 1 \rfloor$ denote the maximum number of incongruent pairs that μ would hold on the market, after a new pair's arrival. During the first $3m$ event-time periods, $3m$ incongruent pairs arrive at the market. Let n be the number of (A, β) pairs arriving during the first $3m$ event-time periods, so $3m - n$ is the number of (B, α) pairs arriving. Suppose that $n < 2m$, so at least m number of (B, α) pairs arrive. As the policy μ matches agents on a FIFO basis and never holds (A, α) or (B, β) pairs, all agents, both A -squares and β -rounds, who were initially in the market would be matched by μ within the first $3m$ event-time periods. Next, suppose $n \geq 2m$. As μ holds at most m incongruent pairs at any time, it would hold at most m number of (A, β) pairs at the end of event time $3m$. Because of the FIFO protocol, all (A, β) pairs held by μ at the end of event-time $3m$ must have arrived after the initial event-time period.

We construct another policy μ' of (MDP) that differs from μ only when the initial state is 0. If the initial state is 0, in each of the first $3m$ event-time periods, μ' holds all agents who arrived after the initial event-time period and would have been held by μ if the initial state were s and the same types of agents arrived. In these $3m$ event-time periods, the policy μ' with the initial state 0 matches all other agents arbitrarily and immediately.

From our discussion above, after event time $3m$, the policy μ' with initial state 0 creates the same matches as μ were the initial state s . Therefore, the rewards generated by μ' with initial state 0 differ from those generated by μ with initial state s only for the first $3m$ event-time periods. Thus,

$$|v_\delta(\mu; s) - v_\delta(\mu'; 0)| \leq 3m \left(\frac{U}{2c} U_{A\alpha} + \frac{U}{4p(1-p)} \right),$$

where the inequality is guaranteed by the fact that the reward function is bounded in $\left[-\frac{U}{4p(1-p)}, \left(\frac{U}{2c} + 1 \right) U_{A\alpha} \right]$. Note that the right hand side is independent of s and δ . Therefore, the claim holds for $M \equiv 3 \left(\frac{U}{2c} + 1 \right) \left(\frac{U}{2c} U_{A\alpha} + \frac{U}{4p(1-p)} \right)$. \blacksquare

6.2 Proofs Regarding the LIFO Protocol

Proof of Lemma 3: For an A -square, say player i , let $\theta_i = (\mathbf{s}, q_i)$ denote her augmented state, where q_i now denotes her rank under the LIFO protocol. We define a *threshold strategy* as a SD-strategy ψ_A such that, with some $\bar{k}_A \in \mathbb{Z}_+$,

$$\psi_A(\theta_i) = \begin{cases} \alpha & \text{if } q_i \leq \bar{k}_A \\ \beta & \text{if } q_i > \bar{k}_A + 1 \end{cases} . \quad (5)$$

Similarly, we define a threshold strategy for α -rounds with the threshold denoted by \bar{k}_α .

Suppose that all A -squares play a threshold strategy ψ_A with threshold $\bar{k}_A \in \mathbb{Z}_+$. We use an *absorbing Markov chain* to compute the expected total payoff for an A -square, say player i , whose augmented state is $\theta_i^t = (\mathbf{s}^t, q_i^t)$ in some period t .

The state space of the absorbing Markov chain is $\{1, 2, 3, \dots, \bar{k}_A, \alpha, \beta\}$ where integer *transient states* denote player i 's ranking q_i^τ (1 if there are no A -squares who arrived after i that are waiting), and each of the two *absorbing states* α and β denote the type of player i 's match partner. The event time τ starts from 0 and increases for each arrival of an incongruent pair.⁶ In expectation, an increment of τ takes $\frac{1}{2p(1-p)}$ periods. The matrix of transition probabilities p_{ij} from state i to j is

$$P = \begin{bmatrix} Q & R \\ \mathbf{0} & I \end{bmatrix}, \text{ where } Q = \begin{bmatrix} 0 & 1/2 & \cdots & 0 \\ 1/2 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & 1/2 \\ 0 & \cdots & 1/2 & 0 \end{bmatrix}, R = \begin{bmatrix} 1/2 & 0 \\ 0 & 0 \\ \vdots & \vdots \\ 0 & 1/2 \end{bmatrix}, \text{ and } I = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

Q represents the transitions between transient states. For any $1 < k < \bar{k}_A$, the state changes upon an arrival of either an (A, β) or a (B, α) pair, each of which occurs with conditional probability $1/2$. R_{11} represents the probability of a transition from $q_i = 1$ to an absorbing state α caused by an arrival of a (B, α) pair. $R_{\bar{k}_A 2}$ represents the transition from $q_i = \bar{k}_A$ to an absorbing state β caused by an arrival of an (A, β) pair.

Let $N \equiv (I_{\bar{k}_A} - Q)^{-1}$, $T \equiv N \cdot \mathbf{1}$, and $B \equiv NR$. The absorbing Markov chain with initial state $k \in \{1, 2, \dots, \bar{k}_A\}$ is absorbed in T_k expected number of steps. It is absorbed by state α (or β) with probability $B_{k\alpha}$ (or $B_{k\beta}$, respectively). It is easy to verify that N is a symmetric matrix with $N_{ij} = \frac{2^j(k-i+1)}{k+1}$ for all $i \geq j$, $T_k = T_{\bar{k}_A+1-k} = \sum_{i=1}^k (\bar{k}_A - 2(i-1))$ for all $k \leq \bar{k}_A/2$, and $B_{k2} = 1 - B_{k1} = k/(\bar{k}_A + 1)$ for $k = 1, \dots, \bar{k}_A$.⁷ The expected total payoff

⁶An arrival of (A, α) or (B, β) does not change player i 's position in line. In particular, if an (A, α) pair arrives, the new players match with each other immediately under LIFO.

⁷For example, if $k = 5$,

$$N = \frac{1}{6} \begin{bmatrix} 10 & 8 & 6 & 4 & 2 \\ 8 & 16 & 12 & 8 & 4 \\ 6 & 12 & 18 & 12 & 6 \\ 4 & 8 & 12 & 16 & 8 \\ 2 & 4 & 6 & 8 & 10 \end{bmatrix}, \quad T = \begin{bmatrix} 5 \\ 5+3 \\ 5+3+1 \\ 5+3 \\ 5 \end{bmatrix}, \quad \text{and} \quad B = \frac{1}{6} \begin{bmatrix} 5 & 1 \\ 4 & 2 \\ 3 & 3 \\ 2 & 4 \\ 1 & 5 \end{bmatrix}.$$

for player i in period t with initial condition $q_i^t = k$ is

$$B_{k1}U_A(\alpha) + B_{k2}U_A(\beta) - T_k \frac{c}{2p(1-p)}. \quad (6)$$

This payoff is strictly decreasing in k , implying that a β -round with rank \bar{k}_A has the highest incentive to deviate from ψ_A by demanding β , among all A -squares who are supposed to demand α according to ψ_A . The total expected payoff for player i with $q_i^t = \bar{k}_A$ is

$$\frac{1}{\bar{k}_A + 1}U_A(\alpha) + \frac{\bar{k}_A}{\bar{k}_A + 1}U_A(\beta) - \frac{\bar{k}_A c}{2p(1-p)}.$$

This payoff is strictly decreasing in \bar{k}_A . Thus, there exists a maximum threshold, which we denote by \bar{k}^{lif0} , such that player i 's payoff exceeds $U_A(\beta)$.⁸ After some algebraic steps, one can verify that

$$\bar{k}^{lif0} \equiv \left\lfloor \sqrt{\frac{2p(1-p)(U_A(\alpha) - U_A(\beta))}{c} + \frac{1}{4} - \frac{1}{2}} \right\rfloor. \quad (7)$$

Next, we show that if $\Psi = (\psi_A, \psi_B, \psi_\alpha, \psi_\beta)$ is a stationary* equilibrium in which ψ_A (and ψ_α) is a threshold strategy with a threshold \bar{k}_A (respectively, \bar{k}_α), then, $\bar{k}_A = \bar{k}_\alpha = \bar{k}^{lif0}$.

(i) Suppose, toward a contradiction, that $\bar{k}_A > \bar{k}^{lif0}$. Take any A -square i whose augmented state in some period t satisfies $q_i^t = \bar{k}_A$. Her expected total payoff from period t by playing ψ_A is strictly lower than $U_A(\beta)$. Therefore, player i has an incentive to deviate and demand a β -round.

(ii) Suppose, toward a contradiction, that $\bar{k}_A < \bar{k}^{lif0}$. Take any A -square i whose augmented state in some period t satisfies $q_i^t = \bar{k}$. We will show that player i has an incentive to deviate and use the threshold $\bar{k}_A + 1$ instead of \bar{k}_A perpetually till matching. To show that this deviation is profitable, consider the following absorbing Markov chain. The state space is $\{1, 2, \dots, \bar{k}_A, \bar{k}_A + 1, \alpha\}$, where each integer transient state denotes player i 's ranking, and the absorbing state α represents the only possible match partner for i , a match with an α -round. Indeed, the queue for A -squares never exceeds the threshold $\bar{k}_A + 1$, because all other A -squares use the threshold \bar{k}_A . Therefore, player i will never match with a β -round.

The event time τ increases for each arrival of an incongruent pair. The transition proba-

⁸The last payoff is never equal to $U_A(\beta)$ because of the regularity assumption.

bility matrix is

$$P = \begin{bmatrix} Q & R \\ \mathbf{0} & 1 \end{bmatrix}, \quad \text{where} \quad Q = \begin{bmatrix} 0 & 1/2 & 0 & \cdots & 0 & 0 \\ 1/2 & 0 & 1/2 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1/2 & 1/2 \end{bmatrix} \quad \text{and} \quad R = \begin{bmatrix} 1/2 \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

To understand $Q_{(\bar{k}_A+1)(\bar{k}_A+1)} = 1/2$, suppose that player i 's augmented state in some period τ satisfies $q_i^\tau = \bar{k}_A + 1$. If an (A, β) pair arrives in the following period, one of other A -squares who play ψ_A with threshold \bar{k}_A demands a β -round and leaves the market. Player i 's rank (i.e., the state in the absorbing Markov chain) will remain at $\bar{k}_A + 1$. Let $N \equiv (I_{\bar{k}_A+1} - Q)^{-1}$ and $T \equiv N \cdot 1$. The absorbing Markov chain with initial state $k \in \{1, 2, \dots, \bar{k}_A + 1\}$ is absorbed by state α in T_k expected number of steps. It is easy to verify that N is a symmetric matrix with $N_{ij} = 2j$ for all $i \geq j$, and $T_k = 2 \sum_{i=1}^k (\bar{k}_A + 2 - i)$.⁹ Therefore, when player i 's augmented state in period τ is $q_i^\tau = \bar{k}_A + 1$, she would deviate from ψ_A by increasing the threshold to $\bar{k}_A + 1$ permanently because

$$U_A(\alpha) - \frac{(\bar{k}_A + 1)(\bar{k}_A + 2)c}{2p(1-p)} \geq U_A(\alpha) - \frac{(\bar{k}^{lifo})(\bar{k}^{lifo} + 1)c}{2p(1-p)} > U_A(\beta).$$

Thus, a stationary* strategy ψ_A with threshold $\bar{k}_A < \bar{k}^{lifo}$ cannot be a stationary* equilibrium strategy. ■

Proof of Lemma 4: First, consider the decisions of A -squares. Suppose that A -squares play a stationary* (threshold) strategy ψ_A with threshold \bar{k}^{lifo} . We prove that, for each A -square, say player i , $\psi_i = \psi_A$ is an optimal policy of the (MDP) (without restrictions on her initial state), defined by other A -squares' strategy ψ_A . It follows that ψ_A is each A -square's best-response.

Given any initial augmented state $\theta_i = (\mathbf{s}, q_i)$ and other A -squares' strategy ψ_A , define the value of policy $\psi_i (= \psi_A)$ as

$$v_i(\theta_i; \psi_i, \psi_A) \equiv E_{\psi_i} \left[\sum_{t=0}^{\infty} u_i(\psi_i(\theta_i^t), \theta_i^t) : \theta_i^0 = \theta_i \right].$$

⁹For example, if $\bar{k}_A + 1 = 3$,

$$N = \begin{bmatrix} 2 & 2 & 2 \\ 2 & 4 & 4 \\ 2 & 4 & 6 \end{bmatrix}, \quad T = \begin{bmatrix} 6 \\ 10 \\ 12 \end{bmatrix}.$$

From equation (6), we obtain that

$$v_i(\theta_i; \psi_i, \psi_A) = \begin{cases} U_A(\alpha) & \text{if } q_i \leq s_\alpha \\ \left(1 - \frac{k}{\bar{k}^{lifo} + 1}\right) U_A(\alpha) + \frac{k}{\bar{k}^{lifo} + 1} U_A(\beta) - T_k \frac{c}{2p(1-p)} & \text{if } k \equiv q_i - s_\alpha \in \{1, \dots, \bar{k}^{lifo}\} \\ U_A(\beta) & \text{if } q_i - s_\alpha > \bar{k}^{lifo}. \end{cases}$$

It is easy to verify that $v_i(\theta_i; \psi_i, \psi_A)$ solves the optimality equation

$$v(\theta_i) = \max_{d_i \in \{\alpha, \beta\}} \left[u_i(d_i, \theta_i) + \sum_{\theta'_i \in \Theta_i} p(\theta'_i : \theta_i, d_i) V(\theta'_i) \right] \quad \text{for all } \theta_i \in \Theta_i.$$

Then, by Theorem Puterman 2 appearing in the Appendix of the main text of the paper, ψ_i is an optimal SD-policy of the Markov decision problem, defined by other A -squares' stationary* strategy ψ_A .

Let us now turn to the β -rounds' decisions. Suppose that A -squares (and α -rounds) play the stationary* strategy with threshold \bar{k}^{lifo} . Then, only if an (A, β) pair arrives, there may exist an A -square (in fact, exactly one A -square) who may demand a β -round, and she matches with the last arriving β -round. Thus, if a β -round remains unmatched after the first period at the market, he won't match with an A -square ever again. Thus, every β -round has an incentive to leave immediately. Therefore, we have that

$$\psi_\beta(\mathbf{s}, q_i) = \begin{cases} \alpha & \text{if } s_{A\alpha} \geq \bar{k}^{lifo} + 1 \text{ and } q_i = 1 \\ \beta & \text{otherwise,} \end{cases} \quad (8)$$

is a best-response for all β -rounds with any initial augmented state. ■

Proof of Corollary 5:

1. Ignoring integer constraints, we have

$$\begin{aligned} \lim_{c \rightarrow 0} \frac{\bar{k}^{opt}}{\bar{k}^{lifo}} &= \lim_{c \rightarrow 0} \frac{\sqrt{\frac{p(1-p)U}{2c}}}{\sqrt{\frac{2p(1-p)(U_A(\alpha) - U_A(\beta))}{c} + \frac{1}{4}} - \frac{1}{2}} \\ &= \frac{\sqrt{p(1-p)U}}{\sqrt{4p(1-p)(U_A(\alpha) - U_A(\beta))}} \\ &= \frac{1}{2} \sqrt{\frac{U}{U_A(\alpha) - U_A(\beta)}} < 1, \end{aligned}$$

where the last inequality is due to $U < 4(U_A(\alpha) - U_A(\beta))$ under our symmetry assumption. Furthermore, we have $\bar{k}^{lifo} < \bar{k}^{fifo}$ if and only if

$$\sqrt{\frac{2p(1-p)(U_A(\alpha) - U_A(\beta))}{c} + \frac{1}{4}} - \frac{1}{2} < \frac{p(U_A(\alpha) - U_A(\beta))}{c}. \quad (9)$$

Let $x \equiv \frac{2p(1-p)(U_A(\alpha) - U_A(\beta))}{c}$ and $z \equiv \frac{p(U_A(\alpha) - U_A(\beta))}{c}$, so that inequality (9) is equivalent to $\sqrt{x + \frac{1}{4}} < z + \frac{1}{2}$, or $x < z^2 + z$. Thus, (9) is satisfied if and only if

$$\frac{2p(1-p)(U_A(\alpha) - U_A(\beta))}{c} < \frac{p^2(U_A(\alpha) - U_A(\beta))^2}{c^2} + \frac{p(U_A(\alpha) - U_A(\beta))}{c},$$

or equivalently

$$1 - 2p < \frac{p(U_A(\alpha) - U_A(\beta))}{c}.$$

Therefore, if $p \geq \frac{1}{2}$, then $1 - 2p \leq 0$, and $\bar{k}^{lifo} < \bar{k}^{fifo}$ for any $c > 0$. If $p < \frac{1}{2}$, then $1 - 2p > 0$, and $\bar{k}^{lifo} < \bar{k}^{fifo}$ for any $c < \frac{p(U_A(\alpha) - U_A(\beta))}{1 - 2p}$.

2. As $\bar{k}^{lifo} = \sqrt{\frac{2p(1-p)(U_A(\alpha) - U_A(\beta))}{c} + \frac{1}{4}} - \frac{1}{2}$ and

$$W^{lifo}(c) = S_\infty - \frac{p(1-p)U}{2\bar{k}^{lifo} + 1} - \frac{2\bar{k}^{lifo}(\bar{k}^{lifo} + 1)}{2\bar{k}^{lifo} + 1}c, \quad (10)$$

it is easy to verify that $\lim_{c \rightarrow 0} W^{lifo}(c) = S_\infty$. ■

6.3 Proofs Regarding the Fixed-Window Protocol

Proof of Proposition B1:

1. Notice that

$$\Delta_+ W_n = \frac{p(1-p) \cdot U}{2n} \cdot Pr(k_A \neq k_\alpha; n) - 2c \leq \frac{p(1-p)U}{2n} - 2c.$$

As the upper bound of $\Delta_+ W_n$ decreases in n , let x be the unique solution of

$$\frac{p(1-p)U}{2x} - 2c = 0.$$

It follows that

$$n^o \leq x = \frac{p(1-p)U}{4c}.$$

2. We find a lower bound of $\Delta_+ W_n$ that will give us a lower bound on the optimal window size. Fix a window size n , and define a multinomial random variable X_t for each time $t = 1, 2, \dots, n$ such that

$$X_t = \begin{cases} 1 & \text{if } (A, \beta) \text{ arrives,} \\ -1 & \text{if } (B, \alpha) \text{ arrives,} \\ 0 & \text{otherwise.} \end{cases}$$

Thus, for all t , X_t takes the values of 1, -1 , or 0 with probabilities $p(1-p)$, $p(1-p)$, or $1-2p(1-p)$, respectively. Notice that

$$k_A - k_\alpha = \sum_{t=1}^n X_t.$$

Let $\Phi(\cdot)$ be the cumulative distribution function of the standard normal distribution. We use the following Berry-Esseen Theorem on the speed of convergence in the Central Limit Theorem (see Feller, 1972 and Tyurin, 2010).

Theorem (Berry-Esseen) *Let Y_1, Y_2, \dots, Y_n be i.i.d. random variables with mean 0 and variance σ^2 , and $Z_n = \frac{\sum_{i=1}^n Y_i}{n}$, and \widehat{G}_n be the cumulative distribution of $\frac{\sqrt{n}Z_n}{\sigma}$. Then,*

$$\sup_y \left| \widehat{G}_n(y) - \Phi(y) \right| \leq \frac{E[|Y_1|^3]}{2\sigma^3\sqrt{n}}.$$

Now, note that we have

$$\text{Var}[X_1] = E[|X_1|^3] = 2p(1-p).$$

Let \widehat{F}_n be the empirical cumulative distribution function of $\frac{k_A - k_\alpha}{\sqrt{2np(1-p)}}$. By the Berry-Esseen Theorem,

$$\begin{aligned} \text{Pr}(k_A = k_\alpha; n) &= \widehat{F}_n\left(1/\sqrt{2np(1-p)}\right) - \widehat{F}_n(0) \\ &\leq \Phi(1/\sqrt{2np(1-p)}) - \Phi(0) + \frac{1}{\sqrt{2np(1-p)}} \\ &\leq \frac{1}{2\sqrt{\pi np(1-p)}} + \frac{1}{\sqrt{2np(1-p)}}. \end{aligned}$$

Thus,

$$\begin{aligned} \Delta_+ W_n &\geq \frac{p(1-p)U}{2n} \left(1 - \left(\frac{1}{2\sqrt{\pi}} + \frac{1}{\sqrt{2}} \right) \frac{1}{\sqrt{np(1-p)}} \right) - 2c \\ &= \frac{p(1-p)U}{2n} \left(1 - \frac{\sqrt{2\pi} + 1}{2\sqrt{\pi np(1-p)}} \right) - 2c. \end{aligned} \tag{11}$$

We now show that there exists $c_1 > 0$ such that for all $c < c_1$, the optimal fixed window size n° is greater than 2 by illustrating that $\Delta_+ W_1 > 0$ and $\Delta_+ W_2 > 0$. Indeed,

$$\Delta_+ W_1 = \frac{p(1-p)U}{2} \text{Pr}(k_A \neq k_\alpha; 1) - 2c = p^2(1-p)^2U - 2c > 0,$$

and

$$\begin{aligned}
\Delta_+ W_2 &= \frac{p(1-p)U}{2} Pr(k_A \neq k_\alpha; 2) - 2c \\
&= \frac{p(1-p)U}{4} \left(1 - (p^2 + (1-p)^2)^2 - 2p^2(1-p)^2 \right) - 2c \\
&\geq \frac{p(1-p)U}{4} (1 - p^2 - (1-p)^2) - \frac{p^3(1-p)^3U}{2} - 2c \\
&= \frac{p^2(1-p)^2U}{2} - \frac{p^3(1-p)^3U}{2} - 2c \geq \frac{3p^2(1-p)^2U}{8} - 2c > 0
\end{aligned}$$

whenever $c < c_1 \equiv \frac{3p^2(1-p)^2U}{16}$.

We now consider $n \geq 3$. From (11), $\Delta_+ W_n \geq h(r_n)$, where $r_n = n^{-1/2}$ and

$$h(r) \equiv \frac{p(1-p)Ur^2}{2} \left(1 - \frac{(\sqrt{2\pi} + 1)r}{2\sqrt{\pi p(1-p)}} \right) - 2c.$$

We use the following observations (i)-(v), assuming that $c < c_2 \equiv \frac{4\pi p^2(1-p)^2U}{27(\sqrt{2\pi} + 1)^2} < c_1$:

(i) $\Delta_+ W_3 \geq h(r_3) = \frac{p(1-p)U}{6} \left(1 - \frac{\sqrt{2\pi} + 1}{2\sqrt{3\pi p(1-p)}} \right) - 2c$. To show that $h(r_3) > 0$, it is sufficient to prove that

$$c < \frac{4\pi p^2(1-p)^2U}{27(\sqrt{2\pi} + 1)^2} < \frac{p(1-p)U}{12} \left(1 - \frac{\sqrt{2\pi} + 1}{2\sqrt{3\pi p(1-p)}} \right).$$

Indeed, the right inequality holds whenever

$$\frac{\pi}{27(\sqrt{2\pi} + 1)^2} < \frac{1}{12} - \frac{\sqrt{2\pi} + 1}{24\sqrt{3\pi}},$$

which holds if and only if

$$\frac{12\pi}{27(\sqrt{2\pi} + 1)^2} + \frac{\sqrt{2\pi} + 1}{2\sqrt{3\pi}} = 0.4437422... < 1.$$

(ii) $h'(r) = 0$ at $r = 0$ and $\frac{4\sqrt{\pi p(1-p)}}{3(\sqrt{2\pi} + 1)}$.

(iii) $h(0) < 0$ and $h\left(\frac{4\sqrt{\pi p(1-p)}}{3(\sqrt{2\pi} + 1)}\right) = \frac{8\pi p^2(1-p)^2U}{27(\sqrt{2\pi} + 1)^2} - 2c > 0$.

(iv) Let $c_3 \equiv \left(\frac{p(1-p)}{4(U_{A\alpha} - U_{B\beta})}\right)^2 \left(\frac{U}{3}\right)^3 \leq c_2$. For every $c < c_3$, let $r^o \equiv \left(\frac{1}{2p(1-p)}\right)^{1/6} \left(\frac{c}{U_{A\alpha} - U_{B\beta}}\right)^{1/3}$.

Then,

$$\begin{aligned}
h(r^o) &= \frac{p(1-p)U(r^o)^2}{2} - \frac{p(1-p)U(\sqrt{2\pi} + 1)(r^o)^3}{2\sqrt{\pi p(1-p)}} - 2c \\
&> \frac{p(1-p)U(r^o)^2}{2} - 3c \\
&= \left(\frac{p(1-p)c}{4(U_{A\alpha} - U_{B\beta})}\right)^{2/3} U - 3c > 0.
\end{aligned}$$

(v) $h(r)$ is a cubic function and the leading coefficient is negative. Thus, for any $c < c_3$ and $r \in [r^o, 1/\sqrt{3}]$, we have $h(r) > 0$. It follows that, if window size n satisfies $r^o \leq 1/\sqrt{n} \leq 1/\sqrt{3}$, then $h(1/\sqrt{n}) > 0$. On the other hand, $0 \geq \Delta_+ W_{n^o} \geq h(1/\sqrt{n^o})$. Therefore,

$$n^o > (r^o)^{-2} = (2p(1-p))^{1/3} \left(\frac{U_{A\alpha} - U_{B\beta}}{c} \right)^{2/3}.$$

■

Proof of Proposition B2: We denote by S_n the ex-ante per-pair surplus when the window size is n . We find bounds on S_n starting from the following inequalities:

$$\begin{aligned} S_n &\geq E \left[\min \left\{ \frac{k_A}{n}, \frac{k_\alpha}{n} \right\} \right] \cdot U_{A\alpha} + \left(1 - E \left[\min \left\{ \frac{k_A}{n}, \frac{k_\alpha}{n} \right\} \right] \right) \cdot U_{B\beta} \quad \text{and} \\ S_n &\leq E \left[\max \left\{ \frac{k_A}{n}, \frac{k_\alpha}{n} \right\} \right] \cdot U_{A\alpha} + \left(1 - E \left[\max \left\{ \frac{k_A}{n}, \frac{k_\alpha}{n} \right\} \right] \right) \cdot U_{B\beta}. \end{aligned}$$

Notice that

$$\begin{aligned} E \left[\min \left\{ \frac{k_A}{n}, \frac{k_\alpha}{n} \right\} \right] &= \frac{E[|k_A + k_\alpha|] - E[|k_A - k_\alpha|]}{2n} = p - \frac{1}{2} \sqrt{\frac{(E[|k_A - k_\alpha|])^2}{n^2}} \\ &\geq p - \frac{1}{2} \sqrt{\frac{E[(k_A - k_\alpha)^2]}{n^2}} \\ &= p - \frac{1}{2} \sqrt{\frac{E[k_A^2] + E[k_\alpha^2] - 2E[k_A k_\alpha]}{n^2}} \\ &= p - \frac{1}{2} \sqrt{2 \left(p^2 + \frac{p(1-p)}{n} \right) - 2p^2} \\ &= p - \sqrt{\frac{p(1-p)}{2n}}, \end{aligned}$$

where the first inequality follows from Jensen's inequality, and the fourth equality from $E[k_A^2] = \text{Var}(k_A) + (E[k_A])^2$. Similarly, we obtain

$$E \left[\max \left\{ \frac{k_A}{n}, \frac{k_\alpha}{n} \right\} \right] \leq p + \sqrt{\frac{p(1-p)}{2n}}.$$

Therefore, for every window size n ,

$$S_\infty - \sqrt{\frac{p(1-p)}{2n}} (U_{A\alpha} - U_{B\beta}) \leq S_n \leq S_\infty + \sqrt{\frac{p(1-p)}{2n}} (U_{A\alpha} - U_{B\beta}).$$

It follows that

$$S_\infty - \sqrt{\frac{p(1-p)}{2n}} (U_{A\alpha} - U_{B\beta}) - c(n-1) \leq W_n \leq S_\infty + \sqrt{\frac{p(1-p)}{2n}} (U_{A\alpha} - U_{B\beta}) - c(n-1).$$

We now turn to show the two parts of the claim.

1. We have

$$W^{fix}(c) \geq W_n \geq S_\infty - \sqrt{\frac{p(1-p)}{2n}}(U_{A\alpha} - U_{B\beta}) - c(n-1), \quad \text{for all } n \in \mathbb{Z}_+.$$

Note that the above lower bound is strictly concave in n . With ignoring the integer constrain on n , the lower bound is maximized at $m > 0$ such that

$$-\frac{1}{2m^{3/2}} \sqrt{\frac{p(1-p)}{2}}(U_{A\alpha} - U_{B\beta}) + c = 0.$$

That is, $m = (1/2)(p(1-p))^{1/3} \left(\frac{U_{A\alpha} - U_{B\beta}}{c}\right)^{2/3}$. With the integer constraint on window size, the value of the lower bound is maximized at either $\lfloor m \rfloor$ or $\lceil m \rceil$.

2. From the above inequalities, we have

$$W^{fix}(c) \leq S_\infty + \sqrt{\frac{p(1-p)}{2n^o}}(U_{A\alpha} - U_{B\beta}) - c(n^o - 1).$$

From the proof of Proposition B1, when $c < c_3$, we have $n^o > (2p(1-p))^{1/3} \left(\frac{U_{A\alpha} - U_{B\beta}}{c}\right)^{2/3}$. Therefore,

$$\begin{aligned} W^{fix}(c) &< S_\infty + \sqrt{\frac{p(1-p)}{2}}(U_{A\alpha} - U_{B\beta})r^o - c(r^o)^{-2} + c \\ &= S_\infty - 2^{-2/3}(p(1-p)c)^{1/3}(U_{A\alpha} - U_{B\beta})^{2/3} + c. \end{aligned}$$

■