Online Appendix for "Strategic Decentralized Matching: The Effects of Information Frictions"

Andrew Ferdowsian* Muriel Niederle † Leeat Yariv‡

This version: December 10, 2023

Abstract

In this Online Appendix, we provide further simulations on markets' duration of operation, considering non-strategic stabilization dynamics that vary in agents' sophistication. We also show ways by which our richness assumption can be relaxed while still ensuring implementation of stable outcomes. Last, we analyze special cases of our environment in which either firms or workers are informed of the realized market at the outset.

A The Effect of Updating and Strategic Sophistication on Stabilization Duration

In Section 4.5 of the main text, we present simulations to illustrate the difference in convergence times between strategic and non-strategic stabilization dynamics. Our key finding is that strategic interactions lead markets to their final matching more rapidly than non-strategic dynamics. There are two related reasons for this result. First, in the naïve, non-strategic dynamics, agents do not incorporate any information accrued over time. Second, non-strategic agents do not optimize which blocking pairs they pursue, even in the short run. In this section, we consider ways by which naïveté can be relaxed. Nonetheless, without full-fledged strategic agents, the duration required to achieve stability remains relatively long.

*Department of Economics, Princeton University.
†Department of Economics, Stanford University and NBER.
‡Department of Economics, Princeton University, CEPR, and NBER.

To inspect whether the assessed increased duration to convergence of non-strategic dynamics is due to the limited incorporation of information, we consider the dynamics from Roth and Vande Vate (1990). The following algorithm describes the exact structure placed on these dynamics:

1. An agent is chosen uniformly at random from all agents.
2. The set of all blocking pairs involving the chosen agent is enumerated. If the chosen agent is not part of any blocking pair, the algorithm returns to the previous step.
3. One such blocking pair is selected uniformly at random and implemented: The agents forming the blocking pair are matched, and their previous partners are unmatched.
4. The period counter is incremented by one.
5. The algorithm terminates if the current matching is stable, so that no further rematching is possible; otherwise, the algorithm returns to step 1.

This algorithm can be thought of as the complete information analogue of our original simulations. Not only does an agent know the state of the market, but agents are also aware of all other agents' preference profiles. Therefore, agents can correctly identify all their possible blocking partners, eliminating the possibility of rejected proposals. In that respect, agents' choices are "better responses" than the benchmark considered in the paper. These dynamics are a special case of those dynamics considered by Roth and Vande Vate (1990), assuming a uniform selection of blocking pairs.[1] Figure 1 contains the results of these dynamics run over 15,000 simulations. As a reminder, the strategic dynamics converge in ≈ 7.5 periods on average when the market size is 100. Additionally, the function mapping the convergence time under strategic dynamics to market size is concave. Under the better-response dynamics, the number of periods required for convergence appears to be convex. Importantly, these non-strategic dynamics also take considerably longer to converge than the strategic dynamics we investigate.

The naïveté assumed in Roth and Vande Vate (1990) is in many ways extreme: it entails both limited foresight and a limited ability to select a desirable blocking pair myopically.

[1] The label of "better responses" is in line with the terminology in Ackermann et al. (2011).

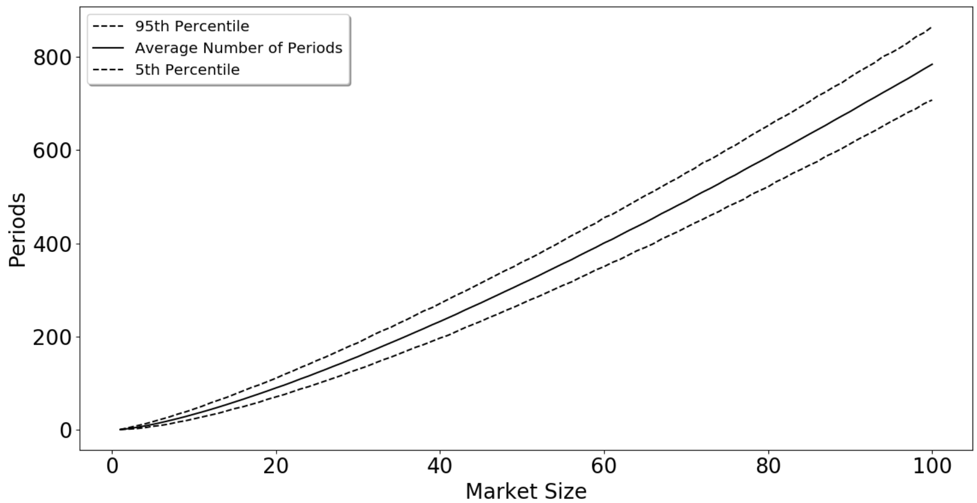

Figure 1: Myopic better-response dynamics

We relax the latter. Following Ackermann et al. (2011), we consider a "best-response" dynamics, whereby agents select their *best* blocking partners. Specifically, the algorithm utilized is identical to that corresponding to the "better-response" dynamics, except for the following. Not only do agents correctly identify their set of possible blocking pairs, but they also select the blocking partner that maximizes their current (myopic) match payoff. Figure 2 displays the results from 15,000 simulations of these dynamics. Myopic best responses yield far shorter stabilization durations: a comparison with Figure 1 suggests that for markets with 100 agents on each side, convergence takes roughly one sixth of the expected duration with full naïveté. Nonetheless, convergence durations are still substantially longer than those expected when agents are fully strategic. Thus, both statistical and strategic sophistication play an important role in reducing stabilization duration.

B Robustness of the Richness Assumption

Our richness assumption is, in many ways, strong—it requires that *all* possible aligned preferences have positive probability in the economy.[2] In the Appendix of the main text, we

[2]While admittedly not an innocuous assumption, it resembles the assumption commonly made in the large-markets literature, where preferences are often drawn uniformly at random, see Immorlica and Mahdian

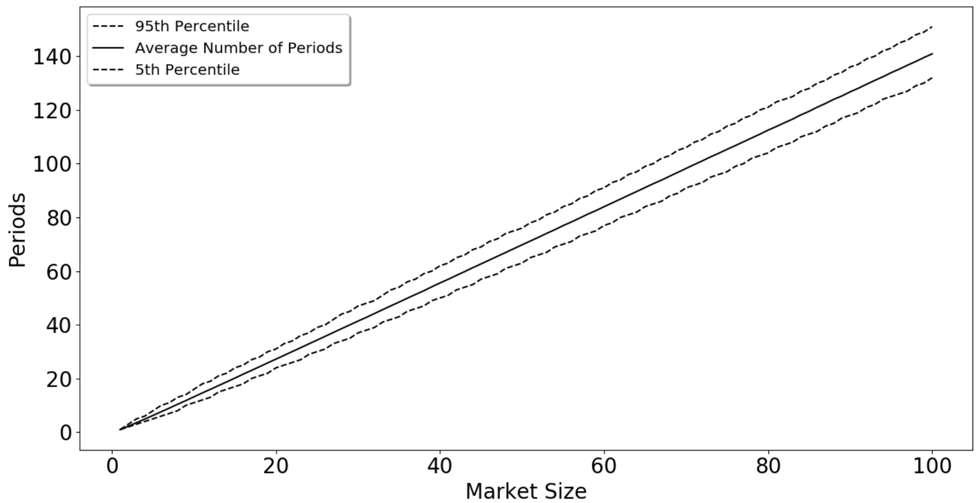

Figure 2: Myopic best-response dynamics

present a general characterization of economies that allow for the implementation of stable outcomes. These can generally support far fewer markets than our rich economies. Here, we consider a more minimal robustness criterion. We ask whether eliminating markets that represent one (arbitrary) preference profile in a rich economy alters our conclusions.

Formally, let $\mathcal{E}$ be a rich economy. We call two markets **equivalent** if every agent in the economy has the same ordinal preference rankings in both markets. We can then partition the economy $\mathcal{E}$'s markets into equivalence classes. The equivalence class of a market M is denoted $\mathcal{M}$.

We define $\mathcal{E} \setminus \mathcal{M}$ as the economy containing all markets in $\mathcal{E}$ other than those in $\mathcal{M}$. The prior over these markets is a renormalization of the prior in the economy $\mathcal{E}$ over the corresponding markets. That is, we eliminate an arbitrary equivalence class of markets from the economy, while maintaining the relative prior weights on all other markets. We show that if $\mathcal{E}$ is rich, then there exists an equilibrium in decentralized minimal DA strategies in $\mathcal{E} \setminus \mathcal{M}$ and implementation of stable matchings, market by market, is possible.[3]

(2005) and Kojima and Pathak (2009).

[3]As a reminder, a reduced DA strategy is minimal for a firm i if it "skips" any workers that i knows cannot be stable match partners in equilibrium. A reduced DA strategy is minimal for a worker j if any firm i ranked by j below all of j's stable match partners is listed as unacceptable.

Proposition B (Robustness). *Let $\mathcal{E}$ be rich and* $\min\{F, W\} > 1$. *Then, for any market M, for sufficiently high δ, there exists an equilibrium in decentralized minimal DA strategies in $\mathcal{E} \setminus \mathcal{M}$. In particular, the resulting outcome is stable in every market supported in $\mathcal{E} \setminus \mathcal{M}$.*

To glean some intuition, assume that $F > 2$. By the top-top match property, in the first period, some firm i is matched to its most preferred worker j and exits the market. As it turns out, after such an exit, the market proceeds as if our economy was rich. Indeed, consider firm i and worker j's preferences. Their preference rankings over other participants do not affect their behavior, as long as they remain a top-top match. Furthermore, when $F, W > 2$, there are always multiple preference rankings for firm i and worker j that are consistent with them being a top-top match. Thus, richness implies that there exists a market identical—in terms of preference rankings—to the removed markets, except for the preferences of those involved in the original top-top match: agents i and j's ranking of others can be arbitrary.[4] Therefore, after the first round, the market restricted to agents other than i and j is also rich. The arguments in the main text can then be used to achieve the result.

Proof of Proposition B

Consider a rich economy $\mathcal{E}$ and an arbitrary market M supported in $\mathcal{E}$. Choose a corresponding ordinal potential Φ such that every element of Φ is unique.[5] We show that decentralized minimal DA strategies are mutual best responses in $\mathcal{E} \setminus \mathcal{M}$.

Assume, first, that $F = 2$. In the economy $\mathcal{E}$, suppose all agents other than worker j use decentralized minimal DA strategies. In any period, if worker j receives an offer from a single firm, j accepts the offer even if the proposing firm is not his top match. Indeed, alignment implies that a top-top match has occurred. Therefore, if worker j does not receive an offer from its most-preferred firm, he can conclude that his favorite firm is part of that period's top-top match. Then, holding or rejecting the other firm's offer either delays the match or runs the risk of the firm approaching a different worker. In particular, worker j best responds by using a decentralized minimal DA strategy.

[4]For $F = W = 2$, the argument is somewhat different. In that case, the restriction that i and j are one another's top-top match pins down their ordinal preferences.

[5]The absence of preference indifferences implies that this is possible.

Now consider the economy corresponding to $\mathcal{E} \setminus \mathcal{M}$. In the original economy $\mathcal{E}$, suppose worker j exits or is matched in the last active period under decentralized minimal DA strategies when some market in $\mathcal{M}$ is realized.[6] Reordering j's preferences over the two firms (in all markets in $\mathcal{M}$) does not impact outcomes. Indeed, such a reordering would have an impact only when the worker receives offers from both firms. When agents other than worker j use decentralized minimal DA strategies, worker j can receive offers from both firms only in period 1. The worker would then optimally accept the offer from his favorite firm, thereby rejecting the other. Since under decentralized minimal DA, there will be at least one worker available in period 2, the firm worker j rejects would necessarily make an additional offer then, contradicting the fact that worker j was the last to exit. Thus, there exists another market, not in $\mathcal{M}$, that generates the same outcomes and induces the same incentives for all others workers and all firms. Furthermore, worker j follows his decentralized minimal DA strategy when immediately accepting his top available offer. It follows that for any rich economy $\mathcal{E}$ and market M, there exists an equilibrium in decentralized minimal DA strategies in the economy $\mathcal{E} \setminus \mathcal{M}$.

Now suppose $F > 2$ and consider the economy $\mathcal{E}$ and market M with corresponding ordinal potential Φ. Suppose all players are following decentralized minimal DA strategies. Alignment implies that the unique stable matching corresponds to a sequence of top-top matches that can be ordered by their corresponding ordinal potential values. Relabel agents so that firm i and worker i form the i-th top-top match, according to this ordering. We use $\Gamma(\Phi, fw, f'w')$ to denote a new ordinal potential defined as follows:

$$\Gamma(\Phi, fw, f'w')_{ij} = \begin{cases} \Phi_{ij} & \text{if } ij \neq fw \\ \Phi_{f'w'} + \epsilon & \text{if } ij = fw \ \& \ \Phi_{fw} < \Phi_{f'w'} \\ \Phi_{f'w'} - \epsilon & \text{if } ij = fw \ \& \ \Phi_{fw} > \Phi_{f'w'}, \end{cases}$$

where ϵ is a constant that is less than the minimum difference between any two of Φ's ordinal values. The newly-defined potential $\Gamma(\Phi, fw, f'w')$ coincides with the original ordinal potential Φ for all firm-worker pairs other than (f, w). For the pair (f, w), $\Gamma(\Phi, fw, f'w')_{fw}$'s

[6]There could be multiple such workers. We focus on an arbitrary one. Since decentralized minimal DA strategies are used, any market within $\mathcal{M}$ yields the same timing over matches.

value is chosen to reverse the ordering of Φ_{fw} and $\Phi_{f'w'}$. We suppress the arguments of Γ when they are clear from context. We show that the appropriate choice of pairs (f, w) and (f', w') yields a new ordinal potential that generates incentives for all agents that are identical to those in markets governed by Φ, but corresponds to a different equivalence class of markets than $\mathcal{M}$. There are two cases:

1. *Under decentralized minimal DA strategies, in market M, f_2 and w_2 match in the first period:*

   In this case, we show that there exists an ordinal potential identical to Φ, except at (f_1, w_2). This new ordinal potential maintains the mutual best response property of the decentralized minimal DA strategy profile under Φ. Richness then yields the desired result. We do so by considering markets corresponding to potentials that differ from Φ only in the value of Φ_{12}. The distinction between markets corresponding to Φ and those corresponding to the constructed potential are then unobservable by agents other than f_1 and w_2.

   Consider $\Gamma(12, 32)$. Only f_1 and w_2 face altered preference rankings relative to those induced by Φ. We now show that Φ and $\Gamma(12, 32)$ generate identical public histories. Observe that the relabelling of agents by top-top matches implies that $\Phi_{11} > \Phi_{22}$. Furthermore, $\Phi_{22} > \max\{\Phi_{12}, \Phi_{32}\}$. Otherwise, f_2 and w_2 would not have matched immediately under Φ. These two facts combined imply that $\Phi_{11} > \max\{\Phi_{12}, \Phi_{32}\}$. Therefore, f_1 still prefers w_1 and w_2 prefers f_2 under $\Gamma(12, 32)$. Thus, the public history and the set of final matches remain unchanged. This immediately implies that the decentralized minimal DA strategies under $\mathcal{E}$ are mutual best responses for all agents that are not f_1 or w_2 in $\mathcal{E} \setminus \mathcal{M}$.

   Of course, the difference between Φ and $\Gamma(12, 32)$ is observable by f_1 and w_2 themselves. By construction, under $\Gamma(12, 32)$ both still find it optimal to follow their decentralized minimal DA strategies under $\mathcal{E}$, as it ensures they end up with their top choice in the first period. We now show that even when the market's governing potential is $\Gamma(12, 32)$, neither f_1 nor w_2 wish to deviate from their decentralized minimal DA strategy under $\mathcal{E}$. To do so, take any ordinal potential where f_1 and w_2 have the same preference rankings as under Φ and have the same top-top partners as under

Φ. Since this new potential is supported in the economy $\mathcal{E} \setminus \mathcal{M}$, it is a best response for f_1 and w_2 to follow the decentralized minimal DA strategy they use in $\mathcal{E}$ in the first period. After that first period, if either has not yet been matched, then Φ could not have been the ordinal potential of the realized market. Therefore, the removal of $\mathcal{M}$ is no longer relevant. In particular, the strategy profile under $\mathcal{E}$ is a mutual best responses.

2. *Under decentralized minimal DA strategies, in market M, w_2 and f_2 match in the second period:*

   We use a series of arguments similar to those above. We begin by finding ordinal potentials that maintain the public history and therefore incentives, but correspond to markets outside of $\mathcal{M}$. Then, since the decentralized minimal DA strategy profile is an equilibrium under $\mathcal{E}$, it must also be an equilibrium under $\mathcal{E} \setminus \mathcal{M}$.

   We break this case down into two further subcases. First, suppose that f_2 does not make an offer to w_2 in the first period in $\mathcal{E}$ when any market in $\mathcal{M}$ is realized. Then, we can again utilize $\Gamma(12,32)$. Since w_2 never receives an offer from f_1 in either Φ or $\Gamma(12,32)$, after the first period, the public history and incentives are equivalent in both markets. An argument identical to the first case implies that the decentralized minimal DA strategy profile under $\mathcal{E}$ is an equilibrium.

   Second, suppose f_2 makes an offer to w_2 in the first period in $\mathcal{E}$ when any market in $\mathcal{M}$ is realized. To start, note that $\Phi_{22} > \Phi_{33}$ implies $\Phi_{22} > \Phi_{32}$; otherwise, by transitivity, $\Phi_{32} > \Phi_{33}$, which implies that w_2 and f_3 are matched under decentralized minimal DA strategies. A similar argument also implies $\Phi_{22} > \Phi_{23}$: otherwise, f_2 and w_3 would match. Together, these inequalities imply that f_2 prefers w_2 to w_3 and w_2 prefers f_2 to f_3.

   For f_2 and w_2 not to match in the first period, w_2 must prefer f_1 to f_2, implying that $\Phi_{12} > \Phi_{22}$. The fact that f_2 makes an offer to w_2 in the first period also implies that f_2 prefers w_2 to w_1, or that $\Phi_{21} < \Phi_{22}$. If $\Phi_{31} < \Phi_{21}$, consider $\Gamma(21,31)$ as an ordinal potential. Under $\Gamma(21,31)$, w_1 still prefers f_1 above all other firms, while the entry corresponding to f_2 and w_1 is lower than under Φ. Therefore, f_2 still makes an offer to w_2. If $\Phi_{31} > \Phi_{21}$, $\Phi_{22} > \Phi_{21}$, and $\Phi_{22} > \Phi_{3j}$ for any $j \neq 1$. We consider $\Gamma(\Gamma(\Phi, 31, 22), 21, 31)$. This new ordinal potential maintains a match history identical to that of Φ; namely, f_3's preference ranking remains unchanged, w_2 is still f_2's favorite worker, and f_1 is still w_1's favorite. As such, all agents except w_1 cannot observe the difference between markets corresponding to Φ and to $\Gamma(21, 31)$ or $\Gamma(\Gamma(\Phi, 31, 22), 21, 31)$ depending on the case. Therefore, the decentralized minimal DA strategy profile under $\mathcal{E}$ entails mutual best responses for all agents except w_1 in $\mathcal{E} \setminus \mathcal{M}$. In addition, under the new potentials—$\Gamma(21, 31)$ or $\Gamma(\Gamma(\Phi, 31, 22), 21, 31)$—worker w_1 immediately matches with his top choice. Thus, when all other agents use their decentralized minimal DA strategies under $\mathcal{E}$, w_1 best responds with his decentralized minimal DA strategy under $\mathcal{E}$.

We have shown that the original strategy profile under $\mathcal{E}$ is an equilibrium under $\mathcal{E} \setminus \mathcal{M}$. Furthermore, for a given agent i, every stable match partner under $\mathcal{E}$ is also a stable match partner under $\mathcal{E} \setminus \mathcal{M}$. Then, since the original strategy profile was minimal under $\mathcal{E}$, it must also be minimal under $\mathcal{E} \setminus \mathcal{M}$.

This completes the proof: for any removed equivalence class of markets $\mathcal{M}$, the economy $\mathcal{E} \setminus \mathcal{M}$ still admits an equilibrium in decentralized minimal DA strategies. ■

C One-Sided Incomplete Information

As discussed in the text, in some environments, one side of the market may have superior information about agents' realized preferences. Here, we analyze the polar cases in which either firms or workers are perfectly informed of the realized market.

C.1 Informed Firms

We start with the setting in which firms are informed of the realized market. We show that equilibrium implementation of stable outcomes is possible for any degree of discounting. Specifically,

Proposition C.1 (Informed Firms)**.** *When all firms have complete information, the minimal DA strategy profile constitutes a Bayesian Nash equilibrium of the decentralized market game* *in strategies that survive iterated elimination of weakly dominated strategies. This profile implements the unique stable matching in each supporting market.*

Proof of Proposition C.1 When firms have complete information, they know the unique stable matching of any realized market at the outset of the decentralized game.

The minimal DA strategy profile takes the following form. Since all firms know the unique stable matching, they immediately apply to their partner in that matching. In the first period, firms who are unmatched in the unique stable matching exit immediately. Workers, recognizing this, always accept the top-ranked acceptable firm who makes them an offer. Workers who receive no offers exit if $W > F$ and remain in the market otherwise. In any subsequent period, remaining firms whose previous offers have been rejected redetermine the unique stable matching corresponding to the remaining participants and apply or exit accordingly.

The minimal DA strategy profile yields the stable matching in each realized market. We now show that it constitutes a Bayesian Nash equilibrium. To see why firms cannot profitably deviate, consider a firm f that would prefer to match with a worker w rather than her stable partner. Because f prefers w, stability requires that w prefers his stable partner f' over f. However, in the minimal DA strategy profile, w's stable partner f' makes w and offer in the first period. In particular, f cannot benefit by making an offer to w: she will either lose or delay her intended equilibrium assignment.

Similar arguments show that firms have no incentives to deviate in later rounds. Indeed, if the market does not conclude in the first period, a deviation must have occurred. Given the monitoring available to firms, the prescribed behavior is optimal with the belief that all remaining firms withheld offers in prior periods.

Workers are also best responding. Under the candidate equilibrium, a worker who refuses the top-ranked acceptable firm may end up unmatched.[7]

Furthermore, minimal DA strategies survive iterated elimination of weakly dominated

[7] Formally, there are two cases to consider. When a worker receives no offer or a unique offer, in equilibrium, a deviation can either delay or harm the worker's outcome. If a worker receives more than one offer, he can infer that a deviation has occurred. However, the prescribed behavior is consistent with a belief that all other firms behave according to equilibrium and leave the market in that period. Accepting the offer from the highest-ranked acceptable firm is therefore the best response.

strategies. By the top-top match property, there exists a worker w_1 and a firm f_1 that form a top-top match. Because f_1 knows it is in a top-top match, it knows that its partner w_1 must accept an offer from her as long as he uses an undominated strategy. Since an offer to w_1 will be accepted, f_1 must make him an offer in the first period in any strategy surviving iterated elimination of weakly dominated strategies. Knowing this, the remaining firms can consider the submarket without w_1 and f_1. The top-top match property implies that there exists a worker w_2 and a firm f_2 that constitute a top-top match in this submarket. Firm f_2 making an offer to worker w_2 in the first period survives iterated elimination of weakly dominated strategies. Indeed, if w_2 prefers f_2 to f_1, then f_2 must be w_2 most preferred acceptable firm. If w_2 does not use a dominated strategy, he must accept the offer immediately. Alternatively, if w_2 prefers f_1 to f_2, he may wait for a period. However, exiting the market—either alone or with an inferior firm—would be dominated for w_2 and f_2's offer would then be accepted in the second period. We can continue recursively to encompass all firms and workers. ■

Uniqueness for high δ For sufficiently high δ, *any* pure Bayesian Nash equilibrium that survives iterated elimination of weakly dominated strategies implements the stable matching. To see why, suppose an unstable matching is implemented in a pure equilibrium. Consider any firm in a top-top match. She knows that if she makes an offer to her top-ranked worker, that worker will accept immediately, yielding her maximal utility. Therefore, she must make the offer immediately.

By the second period, firms in top-top matches have exited, and any firm in a top-top match of the current submarket knows that her stable match partner will also immediately accept an offer. Consider the first firm that does not match with her stable match partner. If she repeatedly makes offers to her stable match partner, all more preferable firms will have matched within at most F periods, and so she will match by period $F+1$ at the latest. Since she is the first firm not to match with her stable match partner, her current match is less preferable to her stable match partner. For sufficiently high δ, deviating to repeatedly making offers to her stable match partner, which ensures that they are matched within $F+1$ periods, must be profitable.

C.2 Informed Workers

We now turn to the other extreme case, where workers are fully informed of the realized market. A result similar to Proposition C.1 holds. However, in order to ensure equilibrium implementation of the stable matching, we now need the additional restriction that δ is sufficiently high. Otherwise, firms have an incentive to skip plausible stable match partners that are sufficiently unlikely in order to speed up their eventual matching.

Proposition C.2 (Informed Workers)**.** *When all workers have complete information and δ is sufficiently high, the minimal DA strategy profile constitutes a Bayesian Nash equilibrium of the decentralized market game in strategies that survive iterated elimination of weakly dominated strategies. This profile implements the unique stable matching in each supporting market.*

Proof of Proposition C.2

The minimal DA strategy profile takes the following form. Upon observing their match utilities, firms construct their list of plausible stable match partners. Each period, every firm applies to her highest-ranked plausible stable match partner still present in the market, as long as her list of plausible stable match partners is non-empty. Based on what each firm observes, she narrows her list of plausible stable match partners to be consistent with her information. If there are no remaining plausible stable match partners, and the number of remaining firms exceeds that of remaining worker, the firm exits; otherwise, she makes an offer to her highest-ranked worker present in the market. Workers only accept offers from their stable match partners or better. If that partner has exited, they immediately accept their highest-ranked offer. If they have no offers, they exit if the number of remaining workers exceeds the number of remaining firms; they stay in the market otherwise.

The minimal DA strategy profile yields the stable matching in each realized market. We now show that it constitutes a Bayesian Nash equilibrium. Firms clearly cannot benefit from deviating. An offer to a worker ranked higher than the highest-ranked plausible stable match partner is never accepted, and can only delay the final match. Skipping the highest-ranked plausible stable match partner can lead to a loss of a match with the highest-ranked plausible stable partner, which is never profitable for sufficiently high δ.

Similarly, workers cannot benefit from deviating. In the candidate equilibrium, a firm never makes an offer to a worker who prefers the firm to his stable match partner: the firm would necessarily prefer that worker to her resulting equilibrium partner, contradicting stability of the outcome. Accepting an offer from a firm below any worker's stable match partner is clearly suboptimal. Deviating by rejecting or holding an offer cannot be beneficial either. Indeed, firms only ever narrow their list of plausible stable match partners through the updating process. Alignment implies that, while workers can manipulate the beliefs of firms through the strategic holding or rejection of offers, workers cannot trigger preferable offers by doing so. Furthermore, such deviations cannot speed up the receipt of preferred offers for similar reasons.

It follows immediately that strategies for both firms and workers survive iterated elimination of weakly dominated strategies. Therefore, the claim holds. ■

Multiplicity of Equilibrium Outcomes When workers have complete information, multiple different matchings may be implemented in equilibrium. The example offered toward the end of Section 4.3 of the main text offers an illustration. Specifically, consider an economy in which one market is far more likely than others, which all occur with equal (small) probability. Assume further that payoffs are such that, in any market, all agents strictly prefer to match with any agent for sure rather than take a fair lottery between their most favored partner and staying unmatched. Suppose workers are fully informed of the realized market, while firms only know the underlying distribution. The following strategy profile is than a Bayesian Nash equilibrium. Firms make an offer to their stable match partner in the most likely market. Workers accept their best offer that is superior to their stable match partner in that market, and exit otherwise. Similarly, any firm that is not accepted at $t = 1$ exits immediately. When the probability of the most likely market is sufficiently high, and the economy is large enough, this profile constitutes an equilibrium surviving iterated elimination of weakly dominated strategies.

References

Ackermann, H., P. W. Goldberg, V. S. Mirrokni, H. Röglin, and B. Vöcking (2011): "Uncoordinated two-sided matching markets," *SIAM Journal on Computing*, 40, 92–106.

Immorlica, N. and M. Mahdian (2005): "Marriage, honesty, and stability," in *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, Society for Industrial and Applied Mathematics, 53–62.

Kojima, F. and P. A. Pathak (2009): "Incentives and stability in large two-sided matching markets," *American Economic Review*, 99, 608–27.

Roth, A. E. and J. H. Vande Vate (1990): "Random paths to stability in two-sided matching," *Econometrica*, 1475–1480.