On the Efficiency of Stable Matchings in Large Markets

SangMok Lee[*] Leeat Yariv[†‡]

August 22, 2018

Abstract. Stability is often the goal for matching clearinghouses, such as those matching residents to hospitals, students to schools, etc. We study the wedge between stability and utilitarian efficiency in large one-to-one matching markets. We distinguish between stable matchings' average efficiency (or, efficiency per-person), which is maximal asymptotically for a rich preference class, and their aggregate efficiency, which is not. The speed at which average efficiency of stable matchings converges to its optimum depends on the underlying preferences. Furthermore, for severely imbalanced markets governed by idiosyncratic preferences, or when preferences are sub-modular, stable outcomes may be average inefficient asymptotically. Our results can guide market designers who care about efficiency as to when standard stable mechanisms are desirable and when new mechanisms, or the availability of transfers, might be useful.

Keywords: Matching, Stability, Efficiency, Market Design.

[*]Department of Economics, Washington University in St. Louis, http://www.sangmoklee.com

[†]Department of Economics, Princeton University, http://www.princeton.edu/yariv

[‡]We thank Federico Echenique, Matthew Elliott, Aytek Erdil, Mallesh Pai, Rakesh Vohra, and Dan Walton for very useful comments and suggestions. Euncheol Shin provided us with superb research assistance. Financial support from the National Science Foundation (SES 0963583) and the Gordon and Betty Moore Foundation (through grant 1158) is gratefully acknowledged.

1. Introduction

1.1. Overview. The design of most matching markets has focused predominantly on mechanisms in which only ordinal preferences are specified: the National Resident Matching Program (NRMP), clearinghouses for matching schools and students in New York City and Boston, and many others utilize algorithms that implement a stable matching corresponding to reported rank preferences. The use of ordinal preferences contrasts with other economic settings, such as auctions, voting, etc., in which utilities are specified for market participants and serve as the primitive for the design of mechanisms maximizing efficiency. For certain matching contexts, such as those pertaining to labor markets, school choice, or real estate, to mention a few, it would appear equally reasonable to assume cardinal assessments.[1] A designer may then face a trade-off between utilitarian efficiency and stability. In this paper, we investigate this trade-off and identify environments in which it may be particularly important.

Our analysis considers large markets, as most applications of market design to matching markets entail tens or hundreds of thousands of participants. We consider two notions of utilitarian efficiency: *average efficiency*, corresponding to the average expected utility each market participant receives, and *aggregate efficiency*, corresponding to the sum of expected utilities in the market. We show that stability generates average efficient outcomes asymptotically for a wide class of preferences. However, vanishingly small efficiency losses on the individual level sum up to substantial aggregate losses—aggregate efficiency generated by stable matchings falls short of what is feasible for all classes of preferences we consider. Furthermore, for particular types of preferences, and for dramatic enough imbalances in participant volumes on each market side, even average efficiency is sub-optimal when stable outcomes are implemented. We also characterize the speeds at which stable matchings' average efficiency converges and illustrate its strong dependence on the structure of participants' preferences.[2] While we present

[1] In fact, there is a volume of work that studies matching scenarios in which agents' preferences are cardinal. E.g., in the context of the marriage market, Becker 1973, 1974 and Hitch, Hortacsu, and Ariely, 2010; in the context of decentralized matching, Lauermann, 2013 and Niederle and Yariv, 2009; in the context of assignment problems, Budish and Cantillon, 2012 and Che and Tercieux, 2018; etc.

[2] The NRMP is a leading example of a large matching market. In the 2018 installment of the NRMP, looking at matched US seniors, 48.5% of applicants were matched with their first-ranked hospital and 79.5% were matched with one of their four top-ranked hospitals, with similar figures appearing for independent applicants (see Table 15 in the NRMP's Results and Data report from the 2018 Main Residency Match). Furthermore, there is a negative association between the number of applicants in a specialty, reflecting the size of the relevant sub-market, and the average rank of matched programs (where lower ranks correspond to

the results for one-to-one matching environments, their messages extend directly to many-to-one matching markets with responsive preferences. Taken together, our results offer guidance for market designers concerned with efficiency as to when different classes of mechanisms than those often considered might be called for.

Utilitarian efficiency and stability are not completely disjoint notions. Indeed, when transfers are available, assuming match utilities are quasi-linear in monetary rewards, stability is tantamount to utilitarian efficiency (see Roth and Sotomayor, 1992). Even so, most centralized matching mechanisms in place, due to their ordinal nature, do not allow for transfers between participants. In fact, in some cases, such as organ donations or child adoption, transfers are viewed not only as "repugnant," they are banned by law (see Roth, 2007). Certainly, (ordinal) stable matchings absent transfers are appealing in many respects—it is simple to identify one of them once preferences are reported, and they are all Pareto efficient. Furthermore, some work suggests that clearinghouses that implement such stable matchings tend to be relatively persistent (see Roth, 2002; Roth and Xing, 1994). Nevertheless, the NRMP, for instance, has been subject to complaints from residents regarding the underlying mechanism's ordinal nature. These complaints culminated in an official lawsuit filed by a group of resident physicians on May of 2002.[3] From this perspective, our paper identifies environments in which restricting attention to ordinal mechanisms that ban transfers might have a substantial impact on resulting efficiency levels.

In general, stable matchings identified only by ordinal preferences, ignoring cardinal utilities (and excluding transfers), need not be utilitarian efficient. Indeed, consider a market with two firms $\{f_1, f_2\}$ and two workers $\{w_1, w_2\}$, in which any match between a firm f_i and a worker w_j generates an identical payoff to both (say, as a consequence of splitting the

more preferred programs).

[3]Details of the case can be found at http://www.gpo.gov/fdsys/pkg/USCOURTS-dcd-1_02-cv-00873 The lawsuit alleged that several major medical associations such as the NRMP and the American Council for Graduate Medical Education, as well as numerous prominent hospitals and universities, violated the Sherman antitrust act by limiting competition in the "recruitment, hiring, employment, and compensation of resident physicians" and by imposing "a scheme of restraints which have the purpose and effect of fixing, artificially depressing, standardizing, and stabilizing resident physician compensation and other terms of employment." The lawsuit highlighted the restricted ability of the NRMP to account for marginal (cardinal) preferences of participants over matches (see Crall, 2004). It inspired a flurry of work studying the potential effects the NRMP imposes on wage patterns, as well as on possible modifications to the NRMP that could potentially alleviate the issues (see Bulow and Levin, 2006, Crawford, 2008, and follow-up literature).

resulting revenue), and all participants prefer to be matched to anyone in the market over being unmatched. Payoffs are given as follows:

	w_1	w_2
f_1	**5**	4
f_2	3	**1**

where the entry corresponding to f_i and w_j is each agent's payoff for that pair if matched. In this case, the unique stable matching matches f_i with w_i, $i = 1, 2$ and generates utilitarian welfare of $2 \times (5 + 1) = 12$. However, the alternative matching, between f_i and w_j, $i \neq j$, $i = 1, 2$, generates a greater utilitarian welfare of $2 \times (3 + 4) = 14$, and would be the unique stable outcome were transfers available, assuming quasi-linear utilities in money. In this paper, we analyze the wedge between stability and efficiency in large markets. Certainly, if we just replicate the 2×2 market above, we can easily generate an arbitrarily large market in which stable matchings lead to a significantly lower utilitarian welfare than the first best, and transfers could prove useful. To obtain results on the likelihood of such cases, we introduce randomness to match utilities. As we illustrate, the conclusions are nuanced: for a rather broad class of preferences and their hybrids, substantial per-person utility losses are increasingly unlikely as market size grows. Nonetheless, when considering the total utilitarian welfare, stable matchings yield substantial inefficiencies. Furthermore, the speed at which expected per-person utility of stable matchings converges to its maximum varies for different preferences.

In more detail, we first generalize the example above to settings in which firms and workers split their match surpluses using a fixed sharing rule, which we term aligned preferences. Such settings are common in many applications (see, e.g., Sorensen, 2007, and the literature that followed). In Section 3.1, we show that the (generically unique) stable matching is asymptotically average efficient. Furthermore, when match utilities are drawn uniformly, the convergence speed is of the order of $\frac{\log n}{n}$. The proof relies on the following idea. In such aligned markets, there is always a firm and worker who are each other's most preferred. Any stable matching would therefore have them matched. Their utility is the maximal match surplus generated in the market corresponding to n^2 possible pairings. Naturally, the expected surplus from this match approaches the maximal possible as n grows large. Once that pair is matched, we are left with a market of $n-1$ firms and workers. Again, there must be a firm and worker who are each other's most preferred within this restricted market. The expected surplus they generate is very high as well. Continuing recursively, we show that a sufficiently large fraction of firms and workers receives a fairly large share of the maximal surplus they could hope for.

Similar arguments hold for assortative markets, where at least one side of the market is unified in its rankings. For instance, medical residents may share preferences over hospitals through their published rankings and hospitals may agree on what makes a medical resident desirable.[4] Such markets also generate average efficiency converging to the optimum when markets are large, with a convergence speed of $\frac{\log n}{n}$.

Turning to aggregate efficiency, when preferences are fully aligned, the benchmark maximum efficiency achievable through any matching is a solution of the *optimal assignment problem* in statistics (see, e.g., Walkup, 1979 and work that followed). In such settings, the aggregate efficiency loss relative to the maximally feasible efficiency is of the order of $\log n$. This is also the order of the average efficiency loss if the support of individual match utilities grows linearly with market size.

There are many settings in which individuals have idiosyncratic preferences over partners. For instance, employees may have idiosyncratic preferences over locations of their employers, while employers may have idiosyncratic preferences over the particular profiles of potential employees. Propositions 1 and 2 in Section 3 provide general results for settings in which the utilities of each participant pair depend on common and idiosyncratic shocks in an arbitrary manner. We show that average efficiency in *all* stable matchings is asymptotically maximal in these settings. The proof introduces some new techniques inspired by Pittel (1989, 1992). Nonetheless, idiosyncrasies reduce convergence speed. When utilities are fully idiosyncratic and uniform, the difference between average efficiency generated by stable matchings and the maximally possible average efficiency is of the order of $\frac{1}{\log n}$. Furthermore, when considering aggregate efficiency, or if the support of match utilities grows linearly with market size, stable outcomes no longer fare as well. We generalize the classical optimal assignment problem in Section 3.4 to show that, when preferences are idiosyncratic, the maximal aggregate efficiency

[4] Agarwal (2015) reports that conversations with residency program and medical school administrators indicate that, indeed, programs broadly agree on what makes a resident appealing.

is $n - O(\sqrt{n})$ while aggregate efficiency of stable matchings is of the order of $\frac{n}{\log n}$.

In Section 4, we consider arbitrary hybrids of assortative and idiosyncratic preferences. This case also corresponds to asymptotic average efficiency, but with convergence speed of the order of $n^{-1/4}$, substantially higher than that corresponding to fully idiosyncratic preferences. Aggregate efficiency of stable matchings is still substantially lower than maximally feasible, and the distance is in-between that corresponding to aligned and idiosyncratic preferences.

There are certain market features that may induce stable matchings whose average efficiency is bounded away from the maximal even in large markets. One such feature is market imbalances. Indeed, many real-world markets contain unequal volumes of participants on both sides of the market. From a theoretical perspective, recent work suggests that an imbalance in the market gives a disadvantage to the market's abundant side (see Ashlagi, Kanoria, and Leshno, 2017). In our setting, a bounded difference between the volumes on either side does not change our conclusions that stable matchings achieve asymptotically the maximal average efficiency. These results also continue to hold for unbounded volume differences when preferences are aligned or assortative. However, in Section 5.1 we show that whenever preferences are idiosyncratic and the difference between the volumes on the two sides of the market increases at least linearly in the size of the market, stable matchings may not yield the maximal average efficiency asymptotically. To get a sense of the competitive forces driving this result, consider a market with one firm and n workers, where match utilities for both the firm and the workers are independently drawn. The generically unique stable matching would match the firm to its favorite worker that views her as acceptable, not accounting for that worker's match utilities at all. Maximal efficiency, on the other hand, is achieved by matching the firm and worker that generate the greatest joint match surplus.

Another market feature that may lead to sub-optimal stable matchings, even in terms of average efficiency asymptotically, has to do with preferences. While the classes of preferences we focus on in the paper (namely, hybrids of aligned, assortative, and idiosyncratic components) are some of the most prominent in the literature, our results do not hold globally. For instance, Becker (1974) has already pointed out the impact of preference modularity on the efficiency of stable matchings. In Section 5.2 we show that with sub-modular assortative preferences, stable matchings may entail a substantial amount of average inefficiency regardless of market size.[5]

To summarize, our results can provide guidance to market designers who care about efficiency, or contemplate the introduction of some form of cardinal mechanisms, potentially including transfers between participants, as to when standard stable mechanisms are desirable. If a designer is concerned with expected utilities per participant, and markets are fairly balanced, implementing an (ordinal) stable matching is justified on efficiency grounds for sufficiently large markets. In particular, the availability of transfers will not affect welfare significantly. However, if markets are either limited in size, or severely imbalanced and entailing a prominent idiosyncratic component in participants' preferences, or if the designer worries about aggregate efficiency, commonly-used ordinal stable mechanisms may not be ideal.

1.2. Literature Review. There are several strands of literature related to this paper. Efficiency of stable matchings has been a topic of recent study. Boudreau and Knoblauch (2013) provide an upper bound on the sum of partner ranks in stable matchings when preferences exhibit particular forms of correlation. Consistent with our results, these upper bounds increase at a speed slower than the size of the market.[6]

Several papers have considered the utilitarian welfare loss stability may entail in matching markets. Anshelevich, Das, and Naamad (2013) consider finite markets and particular constellations of utilities. They provide bounds on the utilitarian welfare achieved through stability relative to that achieved by the welfare-maximizing matching. Compte and Jehiel (2008) consider a modified notion of stability taking into account a default matching and suggest a mechanism that produces an "optimal" such matching that is asymptotically efficient when fully idiosyncratic preferences are drawn from the uniform distribution, in line with our Proposition 1. Durlauf and Seshadri (2003) consider markets with assortative preferences in which agents may form coalitions, of any size, whose output depends on individuals' ability

[5]For these assortative preferences, each individual is characterized by an ability, and a pair's utility coincides with their "output," which increases in both of their abilities. Sub-modularity then means that the marginal increase in output with respect to a match's ability is decreasing in one's own ability.

[6]In a related paper, Knoblauch (2007) illustrates bounds on expected ranks for participants when one side of the market has fully idiosyncratic preferences and the other has arbitrary preferences. Liu and Pycia (2016) consider ordinally efficient mechanisms and illustrate that uniform randomizations over deterministic efficient mechanisms in which no small group of agents can substantially change the allocation of others are asymptotically ordinally efficient, thereby showing that ordinal efficiency and ex-post Pareto efficiency become equivalent in large markets, and that many standard mechanisms are asymptotically ordinally efficient.

profile. Their results imply that the average efficiency of assortative matchings depends on the presence of positive cross-partial derivatives between the abilities of the partners in the output of a marriage, in line with our results in Section 5.2.[7]

Che and Tercieux (2018) study assignment problems in which individual agents have utilities that are composed of a valuation common to all agents and idiosyncratic individual shocks, analogous to our hybrid model of assortative and idiosyncratic preferences, studied in Section 4. They show that Pareto efficient allocations are asymptotically average efficient. However, in the case of assignment problems they study, stable allocations are not necessarily Pareto efficient, so they are not necessarily average efficient. To glean intuition on the different mechanisms at play in this paper and ours, suppose that n agents' utilities from n objects are composed of a valuation common to all agents and an idiosyncratic shock. For simplicity, suppose the common component takes values of 0 or 1, equally likely and independently determined for each object. With sufficiently large n, every agent finds some high common-value objects appealing, accounting for idiosyncratic shocks. Thus, n agents compete for $n/2$ high-value objects, with different agents preferring different high-value objects, depending on their idiosyncratic shocks. The market effectively segments: some agents are assigned a high-value object ("high-tier agents"). Those who do not ("low-tier agents") are assigned a low-value object. The competition for high-value objects is similar to that in an imbalanced market (Ashlagi, Kanoria, and Leshno, 2017). Competition implies that high-value objects are assigned according to objects' priorities without accounting for agents' utilities, which leads to asymptotic average inefficiency. This kind of inefficiency is ruled out in our paper (other than in Section 5.1). By assuming a continuous distribution with full support for common values, we ensure that every agent has sufficiently many close alternatives to any potential partner in terms of common values. We also consider other kinds of preferences, including arbitrary hybrids of aligned and idiosyncratic preferences on *both* sides.

Our results focus on large markets, which have received some attention in the literature,

[7] Dizdar and Moldovanu (2016) study a matching market of fixed size in which agents are characterized by privately known, multi-dimensional attributes that jointly determine the "match surplus" of each potential partnership. They assume utilities are quasi-linear, and monetary transfers among agents are feasible. Their main result shows that the only robust rules compatible with efficient matching are those that divide realized surplus in fixed proportions, independently of the attributes of the pair's members. Several papers highlight the possible impact of incomplete information on the efficiency of commonly used assignment and matching mechanisms (Abdulkadiro ğlu, Che, and Yasuda, 2011; Bordon and Fu, 2015; Fernandez and Yariv, 2018).

mostly due to the observation that many real-world matching markets involve many participants (e.g., the NRMP that involves several tens of thousands of participants each year, schooling systems in large cities, etc.). The literature thus far has mostly focused on incentive compatibility constraints imposed by stable matching mechanisms when markets are large; see, for instance, Immorlica and Mahdian (2005), Kojima and Pathak (2009), and Lee (2017). While most of this literature focuses on balanced markets, Ashlagi, Kanoria, and Leshno (2017) have noted that imbalances in participant volumes across market sides may alleviate incentive compatibility issues, particularly in large markets. We use some of their results when analyzing the efficiency in imbalanced markets in Section 5.1.

Our paper also relates to the "price of anarchy" notion introduced in Computer Science (see Roughgarden and Tardos, 2007). In general, the price of anarchy is defined as the ratio between the utilitarian welfare of the (worst) Nash equilibrium outcome of a game and the maximum utilitarian welfare possible in that game. In our context, a natural substitute to Nash equilibrium is a stable matching. In that respect, our results characterize an analogous "price of stability" in many one-to-one matching environments. In particular, the asymptotic price of stability is 1 for a wide array of balanced markets.[8]

There is a large body of literature studying efficiency of mechanisms in other realms, such as auctions (see Chapter 3 in Milgrom, 2004) or voting (see Krishna and Morgan, 2015). The current paper provides an analogous study in the matching context.

Methodologically, our results borrow techniques introduced by Knuth (1976), Walkup (1979), Pittel (1989, 1992), and Lee (2017).

2. The Model

Consider a market of n firms $F = \{f_1, ..., f_n\}$ and n workers $W = \{w_1, ..., w_n\}$ who are to be matched with one another. At the outset, two $n \times n$ matrices $(u^f_{ij})_{i,j}$ and $(u^w_{ij})_{i,j}$ are randomly determined according to a non-atomic probability distribution $\mathcal{G}$ over $[0,1]^{2n^2}$. When firm f_i and worker w_j match, they receive match utilities u^f_{ij} and u^w_{ij}, respectively. We assume that any agent remaining unmatched receives a match utility of 0, so that all agents prefer, at

[8]Without restricting preferences in any way, and taking a worst-case point of view, Echenique and Galichon (2015) show that the price of anarchy can be arbitrarily low (i.e., for any value, one can always find a market in which stability in the non-transferable utility model produces per-person utility lower than the maximal by at least that value).

least weakly, to be matched with any agent over remaining unmatched (and this preference is strict almost always).[9] If we further assume that utilities are quasi-linear in money,[10] there is a natural benchmark of utilitarian welfare that corresponds to any stable matching with transfers. Throughout, we often refer to u_{ij}^f and u_{ij}^w as utilities.

We consider market matchings $\mu : F \cup W \to F \cup W$ such that for any $f_i \in F$, $\mu(f_i) \in W$, for any $w_j \in W$, $\mu(w_j) \in F$, and if $\mu(f_i) = w_j$, then $\mu(w_j) = f_i$. We will often abuse notation and denote $\mu(i) = j$ and $\mu(j) = i$ if $\mu(f_i) = w_j$. Denote by M the set of all market matchings. For any realized match utilities u_{ij}^f and u_{ij}^w, *a stable matching* $\mu \in M$ satisfies the following condition: For any firm and worker pair (f_i, w_j), either either $u_{i\mu(i)}^f \geq u_{ij}^f$ or $u_{\mu(j)j}^w \geq u_{ij}^w$. In other words, at least one of the members of the pair (f_i, w_j) prefer their allocated match under μ over their pair member.[11] Whenever there exist a firm and a worker that prefer matching with one another over their allocated match partners, the matching under consideration is unstable and that pair is referred to as a *blocking pair*.

In most applications, centralized clearinghouses are designed to implement stable matchings. Our focus in this paper is therefore in assessing the relative utilitarian welfare of stable matchings to the maximal utilitarian welfare achievable through any matching.

The expected maximal utilitarian welfare achievable across all matchings, which we call the *maximal aggregate efficiency*, is denoted by E_n :

$$E_n \equiv \mathbb{E}_{\mathcal{G}} \max_{\mu \in M} \sum_{i=1}^{n} \left(u_{i\mu(i)}^f + u_{i\mu(i)}^w \right).$$

Since stable matchings are not necessarily unique, and utilities of firms and workers are not necessarily symmetric, we denote the worst-case utilitarian welfare of stable matchings for firms and workers as follows:

$$S_n^f \equiv \mathbb{E}_{\mathcal{G}} \min_{\{\mu \in M | \mu \text{ is stable}\}} \sum_{i=1}^{n} u_{i\mu(i)}^f \quad \text{and} \quad S_n^w \equiv \mathbb{E}_{\mathcal{G}} \min_{\{\mu \in M | \mu \text{ is stable}\}} \sum_{i=1}^{n} u_{i\mu(i)}^w.$$

[9] We relax the assumptions that utilities are drawn from distributions with bounded support and that all agents are acceptable in the Online Appendix.

[10] That is, whenever firm f_i matches with worker w_j and transfers an amount m to the worker, the respective utilities for the firm and worker are given by: $U_{f_i}(w_j; m) = u_{ij}^f - m$ and $U_{w_j}(f_i; m) = u_{ij}^w + m$.

[11] In general, stability also entails an individual rationality constraint: no agent prefers remaining unmatched over her prescribed match. Given our assumptions on utilities, this constraint is automatically satisfied.

The expected minimal utilitarian welfare achievable by a stable matching, which we call *the aggregate efficiency* of (worst-case) stable matchings, is denoted by S_n and defined as:

$$S_n \equiv \mathbb{E}_{\mathcal{G}} \min_{\{\mu \in M | \mu \text{ is stable}\}} \sum_{i=1}^{n} \left(u^f_{i\mu(i)} + u^w_{i\mu(i)} \right) \geq S^f_n + S^w_n.$$

We call $\frac{E_n}{2n}$ and $\frac{S_n}{2n}$ *the maximal average efficiency* and *the average efficiency* (of worst-case stable matchings), respectively. Our goal is to characterize when $\frac{E_n}{2n}$ and $\frac{S_n}{2n}$, as well as E_n and S_n, become close when markets are very large, in which case we say that stable matchings are *asymptotically efficient* under the respective criterion. Notice that the average efficiency of stable matchings is always bounded by the maximal value of the support of match utilities, $\frac{S_n}{2n} \leq 1$ for all n. In particular, whenever $\frac{S_n}{2n} \to 1$, stable matchings achieve maximal average efficiency asymptotically.

A few notes on our underlying model. First, while we phrase our results with the labeling of "firms" and "workers", they pertain to pretty much any two-sided one-to-one matching environment in which a centralized clearinghouse could be utilized. Furthermore, the results extend directly to many-to-one settings, such as school choice, labor markets in which each firm seeks multiple workers, etc., as long as preferences are responsive. Formally, suppose that each firm f_i has a quota q_i and that whenever $\mu(f_i) = \{w_{j_1}, ..., w_{j_{q_i}}\}$, the firm's match utility is given by $v^f(u^{f,1}_{ij_1}, ..., u^{f,q_i}_{ij_{q_i}})$, where $u^{f,k}_{ij}$ now stands for the randomly determined value of the match of f_i with worker w_j in the k'th position, and v^f is continuous and increasing in each of its arguments. If the market has n firms and $\sum_{i=1}^{n} q_i$ workers, our results regarding asymptotic average and aggregate efficiency go through directly by considering a one-to-one market in which each firm f_i is duplicated q_i times and the duplicates inherit match utilities $u^{f,1}_{i\cdot}, ..., u^{f,q_i}_{i\cdot}$. Nonetheless, there are some details that our analysis does not handle, for example pre-match investment in perceived quality in the form of test preparation, geographic relocation, and the like (see, e.g., Cole, Mailath, and Postlewaite, 2001 and Avery and Pathak, 2017). We hope this paper opens the door to further studies incorporating these elements.

Second, for most of the paper we will consider the case of a balanced market (n agents on each side). All of our results go through when the gap between the volume of firms and workers is bounded (say, there are n firms and $n + k$ workers, where k is fixed). When the difference in volumes is increasing (say, there are n firms and $n + k(n)$ workers, where $k(n)$ is increasing in n), some subtleties arise that we discuss in Subsection 5.1.

Last, we assume throughout the paper that utilities are drawn from bounded supports[12] and are such that all participants view all partners as acceptable. In the Online Appendix, we relax these restrictions and show that our results do not hinge on them. We illustrate that as long as the support of match utilities increases sufficiently slowly with the size of the market, our insights regarding asymptotic efficiency continue to hold. If the supports increase linearly or faster with the size of the market, however, our results regarding average efficiency break down. In fact, the average efficiency corresponding to the case in which supports increase linearly with market size is effectively the aggregate efficiency we analyze in the paper (when supports are bounded). Having a random fraction of agents unacceptable makes no difference to our results. This point relates to the market thickness inherent in our environment—we effectively show that thinning the market by deeming a certain fixed fraction, even a large fraction, of participants unacceptable does not alter results qualitatively. Naturally, the speed of convergence may slow down and, even when considering average efficiency, achieving close to maximal efficiency may require larger markets.

3. General Aligned Markets with Idiosyncratic Shocks

We first focus on markets that entail aligned or common impacts on utilities (say, the revenue a worker and firm can generate together) as well as idiosyncratic ones (say, ones corresponding to the geographical location of an employer, or the precise educational background of a potential employee). We consider general markets allowing for both components. We illustrate that stable matchings in such markets yield the maximal average efficiency asymptotically, but not the maximal aggregate efficiency. We also characterize the speed at which the average efficiency of stable matchings converges to its maximal level. In particular, we show that convergence speed is substantially faster as the idiosyncratic component of preferences vanishes.

[12]This is an analogous assumption to that made in the literatures inspecting auctions and elections with large numbers of participants—the support of valuations (in auctions), or utilities from different policies or candidates (in elections), are commonly assumed independent of the number of participants.

3.1. Average Efficiency. Formally, we consider utility realizations such that each pair (f_i, w_j) receives a utility that is a combination of the pair's common surplus c_{ij} and independent utility "shocks" z_{ij}^f and z_{ij}^w. That is,

$$\begin{aligned} u_{ij}^f &\equiv \phi(c_{ij}, z_{ij}^f) \quad \text{and} \\ u_{ij}^w &\equiv \omega(c_{ij}, z_{ij}^w). \end{aligned}$$

We assume the functions $\phi(.,.)$ and $\omega(.,.)$ from $[0,1]^2$ to $[0,1]$ are continuous. We further assume that they are both either strictly increasing in the common component c_{ij} but independent of the idiosyncratic components, z_{ij}^f and z_{ij}^w; or strictly increasing in the idiosyncratic components but independent of the common component; or strictly increasing in both arguments. In this way, we capture markets characterized by fully aligned preferences, fully idiosyncratic preferences, and non-trivial hybrids of aligned and idiosyncratic preferences. Each of c_{ij}, z_{ij}^f, and z_{ij}^w is drawn independently from distributions that have positive density functions over $[0,1]$.

Our main result here shows that, asymptotically, market participants achieve, on average, their maximal conceivable match utility, regardless of which stable matchings are selected.

Proposition 1 [Efficiency of Stable Matchings]**.** *Stable matchings are asymptotically efficient:* $\lim_{n\to\infty} \frac{S_n}{2n} = 1$. *Furthermore,*

$$\lim_{n\to\infty} \frac{S_n^f}{n} = \lim_{n\to\infty} \frac{S_n^w}{n} = 1.$$

Proposition 1 illustrates that *stable matchings achieve the maximal average efficiency, even for utilities that are arbitrary combinations of common and idiosyncratic components that are realized from arbitrary continuous distributions.* An indirect consequence of the proposition is that the most efficient matching, stable or not, asymptotically achieves the maximal conceivable utility per participant.

Some prior work (most notably, Abdulkadiroglu, Pathak, and Roth, 2009 in the context of school choice) has suggested that the top trading cycle mechanism can considerably improve upon deferred-acceptance algorithms, even in large markets, when considering matched partners' ranks. These results are consistent with ours. For illustration, consider the case in which agents' preferences are fully idiosyncratic: each firm's utility from each worker, and each worker's utility from each firm are i.i.d draws from the uniform distribution over $[0,1]$. Our results suggest that for any $\varepsilon > 0$, in large enough markets, most agents will be matched with an agent achieving at least $1-\varepsilon$ in utility terms. So, while agents will not necessarily be matched with their very top candidates, they will match with fairly close substitutes.[13] As we show in the Online Appendix, this observation is not an artifact of our assumption that utilities are drawn from bounded supports, though when the support of utilities expands rapidly enough, asymptotic average efficiency fails, as we will see in Section 3.4.

We soon describe some intuition for this result. Before we do so, we discuss the speeds with which the limits in Proposition 1 are achieved.

3.2. Speed of Convergence. We now turn to the speed of convergence pertaining to the average efficiency of stable matchings. As it turns out, the structure of preferences is crucial. In order to provide a characterization of the convergence speed, we restrict attention to linear functions ϕ and ω. Namely, for each firm f_i and worker w_j,

$$\begin{aligned}\phi(c_{ij}, z_{ij}^f) &= (1-\alpha)c_{ij} + \alpha z_{ij}^f \quad \text{and} \\ \omega(c_{ij}, z_{ij}^w) &= (1-\alpha)c_{ij} + \alpha z_{ij}^w,\end{aligned}$$

where $\alpha \in [0,1]$. We further assume that c_{ij}, z_{ij}^f, and z_{ij}^w are all uniformly distributed over $[0,1]$.

Proposition 2. *1. If $\alpha = 0$, then for any $n \geq 3$,*

$$\frac{1}{2}\frac{\log n}{n} \leq 1 - \frac{S_n}{2n} \leq \frac{\log n}{n}.$$

2. If $\alpha > 0$,

$$\limsup_{n\to\infty}\left(1 - \frac{S_n^f}{n}\right)\log n = \limsup_{n\to\infty}\left(1 - \frac{S_n^w}{n}\right)\log n \leq 2.$$

[13] In fact, simulations of such markets suggest that even for a market with 1000 participants on each side, the top trading cycle mechanism improves the ranking of matched partners under deferred acceptance for over 50% of market participants, but improves average utilities of participants by about 1% within the $[0,1]$ range.

The proposition suggests that the speed of convergence is substantially faster when preferences are aligned. Indeed, when $\alpha = 0$, preferences depend solely on the common component c_{ij}. The proposition implies that the speed of convergence of $\frac{S_n}{2n}$ in this case is of the order of $\frac{\log n}{n}$. In contrast, whenever $\alpha > 0$, there is positive weight on the idiosyncratic component. In this case, there might be multiple stable matchings and we consider $\frac{S_n^f}{n}$ and $\frac{S_n^w}{n}$ separately. Their speed of convergence is at least of the order of $\frac{1}{\log n}$. In fact, when $\alpha = 1$, we show in the Online Appendix that

$$\lim_{n\to\infty}\left(1-\frac{S_n^f}{n}\right)\log n = \lim_{n\to\infty}\left(1-\frac{S_n^w}{n}\right)\log n = 1.$$

Figure 1 illustrates numerical results for the average efficiency for different levels of α. For each market size n, we run 100 simulations, each corresponding to one realization of preferences in the market. For each simulation, we compute the lowest per-participant utility by a stable matching. The solid black line, the long dashed line, and the short dashed line depict, respectively, the mean, the 95'th percentile, and the 5'th percentile of the simulated distributions of these averages across the 100 simulations.[14] The solid red line depicts the mean maximal per-participant utility feasible across the realized markets. As one might expect, greater values of α are associated with lower speeds of convergence. There are two features to note in the figure. First, even for high levels of α, the fraction of the maximal per-participant utility that is achieved through stability is substantial. For markets with around 1000 participants on each side—much smaller than many of the markets in the applications we discuss—that fraction is about 88% even when $\alpha = 1$, and much higher for lower α. Second, that fraction does not depend linearly on α. For instance, for markets with around 1000 participants on each side, that fraction is about 98% for $\alpha = 1/3$, about 95% for $\alpha = 2/3$, and, as mentioned, about 88% for $\alpha = 1$.

Ultimately, Propositions 1 and 2 combined illustrate that when preferences have either or both aligned and idiosyncratic components, any selection of stable matchings will lead to approximately average efficient matchings for sufficiently large markets. Nonetheless, the

[14] In principle, when preferences are perfectly aligned, $\alpha = 0$, we can use the formula developed in the section that follows to calculate the average efficiency. The figure gives a sense of the spread of the distribution (and the mean tracks closely that generated by the formal expression of the average efficiency).

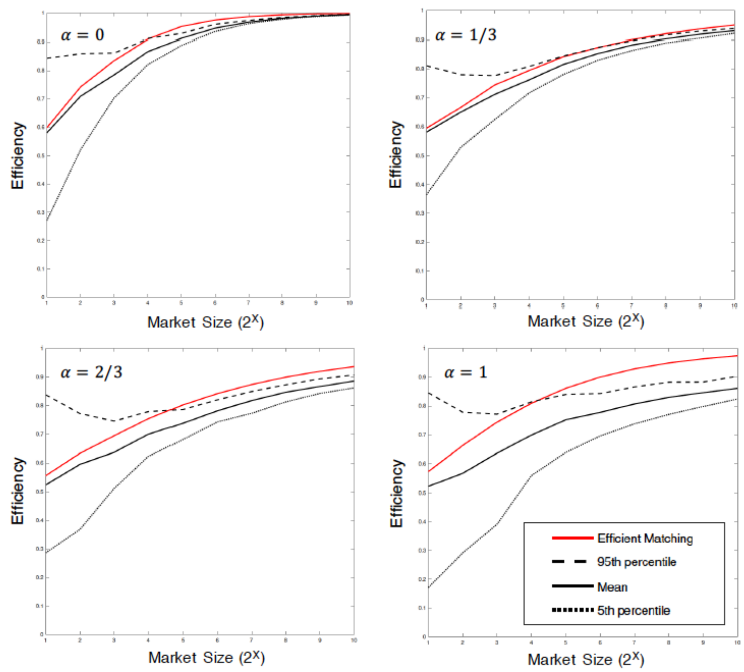

Figure 1: Convergence Speeds for Linear Hybrid Models of Aligned and Idiosyncratic Preferences (α is the Weight on the Idiosyncratic Component)

speed of convergence depends heavily on the structure of preferences: as the idiosyncratic component becomes more prominent, markets need to be larger to near the maximal average efficiency.

3.3. Intuition Underlying Propositions 1 and 2. In this section, we provide a heuristic construction underlying the proofs of Propositions 1 and 2. As it turns out, the case in which preferences are fully aligned requires qualitatively different techniques than the case involving idiosyncratic components. We thus describe them separately.

Fully Aligned Preferences. We start with the fully aligned case, where members of each matched pair receive utilities proportional to one another (e.g., a firm and a worker may split the revenues their interaction generates). As mentioned in the Introduction, such settings are common in many applications (see, e.g., Sorensen, 2007, and the literature that followed).

Formally, we assume here that the utility both firm f_i and worker w_j receive if they are matched is given by $u_{ij} \equiv u_{ij}^f = u_{ij}^w$. We assume u_{ij} are independently drawn across all pairs (i, j) from a continuous distribution over $[0, 1]$. It follows that, generically, utility realizations

$(u_{ij})_{i,j}$ entail a unique stable matching. Indeed, consider utility realizations $(u_{ij})_{i,j}$ such that no two entries coincide, and take the firm and worker pair (f_i, w_j) that achieve the maximal match utilities, $\{(i,j)\} = \arg\max_{(i',j')} u_{i'j'}$. They must be matched in any stable matching since they both strictly prefer one another over any other market participant. Consider then the restricted market absent (f_i, w_j) and the induced match utilities for the remaining participants. Again, we can find the pair that achieves the maximal match utility within that restricted market. As before, that pair must be matched in any stable matching. Continuing recursively, we construct the unique stable matching.

The proof of Proposition 1 in this case proceeds as follows. We first consider the uniform distribution. When determining match utilities, the greatest realized entry, corresponding to the first matched pair in the construction of the generically unique stable matching mentioned above, is the extremal order statistic of n^2 entries. Since each entry is uniform, the expected value of the maximal entry is given by $\frac{n^2}{n^2+1}$. In the next step of our construction, we seek the expected maximal value within the restricted market (derived by extracting the firm and worker pair that generates the highest match utility). That value is the extremal order statistic of $(n-1)^2$ uniform random numbers that are lower than the entry chosen before, and can be shown to have expected value of $\frac{n^2}{n^2+1}\frac{(n-1)^2}{(n-1)^2+1}$. Continuing recursively,

$$\frac{S_n}{2} = \frac{n^2}{n^2+1} + \frac{n^2}{n^2+1}\frac{(n-1)^2}{(n-1)^2+1} + \frac{n^2}{n^2+1}\frac{(n-1)^2}{(n-1)^2+1}\frac{(n-2)^2}{(n-2)^2+1} + ...$$

While corresponding summands become smaller and smaller as we proceed with the recursive process above, there are enough summands that are close enough to 1 so that $\lim_{n\to\infty} \frac{S_n}{2n} = 1$, which is what the proof illustrates.

We then show that our result regarding asymptotic average efficiency does not depend on the uniform distribution of utilities.[15] However, for Proposition 2, in order to show that the average efficiency of stable matchings converges to 1 at a speed of the order of $\frac{\log n}{n}$, we use the precise formulation of S_n above.

[15]This effectively relies on the speedy convergence of extremal order statistics for the distributions we consider. As mentioned, in the Online Appendix, we show that stable matchings are asymptotically average efficient even when considering a class of utility distributions that are not bounded and allowing for agents to view certain partners as unacceptable.

Fully Idiosyncratic Preferences. The other extreme case of Propositions 1 and 2 has to do with the polar case in which preferences are determined in a fully independent manner. That is, all values $(u^f_{ij})_{i,j}$ and $(u^w_{ij})_{i,j}$ are independently and identically distributed according to a continuous distribution over $[0, 1]$.

In this setting, stable matchings are not generically unique and the formal proof of Proposition 1 for this case utilizes different techniques than those employed to prove the proposition for the fully aligned case. It relies on results by Pittel (1989). To see the method of proof, suppose utilities are drawn from the uniform distribution. Notice that from the ex-ante symmetry of the market, each firm f_i (respectively, each worker w_j) has equal likelihood to be ranked at any position in any worker's (respectively, firm's) preference list. Therefore, each one of $n!$ matches of n firms and n workers has the same probability P_n of being stable. Knuth (1976) proved that

$$P_n = \underbrace{\int_0^1 \cdots \int_0^1}_{2n} \Pi_{1\leq i\neq j\leq n} \left(1 - (1 - u^f_{ii})(1 - u^w_{jj})\right) d\mathbf{u}^f_{ii} d\mathbf{u}^w_{jj},$$

where $d\mathbf{u}^f_{ii} = du^f_{11} du^f_{22} \cdots du^f_{nn}$ and $d\mathbf{u}^w_{jj} = du^w_{11} du^w_{22} \ldots du^w_{nn}$.

The intuition behind this formula is simple. The formula essentially evaluates the probability that the matching μ, with $\mu(i) = i$ for all i, is stable. For any realized market, in order for μ to be stable, utilities $(u^f_{ij}, u^w_{ij})_{1\leq i\neq j\leq n}$ must satisfy that either $u^f_{ij} \leq u^f_{ii}$ or $u^w_{ij} \leq u^w_{jj}$ for all $i \neq j$. The integrand corresponds to the probability that these restrictions hold.

Take any $\varepsilon > 0$. Let $P_{\varepsilon,n}$ be the probability that μ is stable and the sum of firms' utilities is less than or equal to $(1 - \varepsilon)n$. That is,

$$P_{\varepsilon,n} = \int_{\substack{\mathbf{0}\leq \mathbf{u}^f_{ii}, \mathbf{u}^w_{jj} \leq \mathbf{1} \\ \sum_{i=1}^n u^f_{ii} \leq (1-\varepsilon)n}} \Pi_{1\leq i\neq j\leq n} \left(1 - (1 - u^f_{ii})(1 - u^w_{jj})\right) d(\mathbf{u}^f_{ii}, \mathbf{u}^w_{jj}). \tag{1}$$

From symmetry, the probability that any matching is stable and the sum of firms' utilities is at most $(1 - \varepsilon)n$ coincides with $P_{\varepsilon,n}$. Since there are $n!$ possible matchings, it suffices to show that $n!P_{\varepsilon,n}$ converges to 0 as n increases. Our proof then uses the techniques developed in Pittel (1989) to illustrate this convergence.[16] When utilities are distributed uniformly, we

[16]The proof appearing in the Online Appendix for this extreme case circumvents the formulas described

further show in the Appendix that the convergence speed of $1-\frac{S_n^f}{n}$ is of the order of $\frac{1}{\log n}$.

We note that Proposition 1 for the hybrid model is not a direct generalization of the arguments used for the two polar cases above. In order to get a sense of the difficulty introduced by combining aligned preferences with idiosyncratic shocks, consider Equation 1 above. Roughly speaking, alignment introduces a positive correlation between match utilities (in fact, the relevant match utilities in Equation (1) are *positively associated*, see Esary, Proschan, and Walkup, 1967). This positive correlation affects both the integrand as well as the conditioning region over which the integral (or expectation) is taken. Much of the proof appearing in the Appendix handles these correlations.

3.4. Aggregate Efficiency. Up to now, we considered average efficiency, where utility is averaged across market participants. The average efficiency notion is particularly useful when the designer is concerned with expected outcomes of a clearinghouse's participants, or when contemplating individual incentives to shift from one institution to another (e.g., allowing for transfers or implementing an efficient rather than stable matching). However, market designers may also be concerned with aggregate efficiency. In this section, we study the wedge in terms of aggregate efficiency between optimal matchings, those maximizing aggregate efficiency, and stable matchings. Our results suggest a substantial welfare loss induced by stability, one that is more pronounced when preferences are idiosyncratic.

Formally, recall that we denoted by E_n the maximal aggregate efficiency across all matchings. Our goal in this Section is to characterize the aggregate efficiency loss $L_n \equiv E_n - S_n$. In order to provide precise bounds on this difference, we focus on two polar cases in our setting: fully aligned and fully idiosyncratic preferences, where utilities are drawn from the uniform distributions—the environments discussed in Subsections 3.3 and 3.3.

We denote the aggregate efficiency loss associated with fully aligned preferences (with uniformly distributed utilities over $[0,1]$) with n participants on each side by L_n^A and the aggregate efficiency loss associated with fully idiosyncratic preferences (with uniformly distributed utilities over $[0,1]$) with n participants on each side by L_n^I. The following proposition provides bounds on L_n^A and L_n^I.

---

here and utilizes more directly results from Pittel (1989).

Proposition 3 [Aggregate Efficiency Loss]**.**

1. *For any* $n \geq 3$,
$$\log n - 6 \leq L_n^A \leq 2\log n.$$

2. *The relative loss of aggregate efficiency satisfies the following:*
$$1 \leq \liminf_{n \longrightarrow \infty} \frac{(\log n)^2 L_n^I}{n L_n^A} \leq 2.$$

The proposition illustrates the substantial efficiency loss imposed by stability relative to any "optimal" matching, despite this loss having a vanishing effect on individual participants' expected payoffs. The proposition also suggests that the structure of preferences impacts significantly the speed at which this efficiency loss grows with market size, with idiosyncratic preferences exhibiting a greater loss asymptotically. Namely, the ratio between the aggregate efficiency loss with idiosyncratic preferences relative to the loss with aligned preferences is asymptotically of the order of $n/(\log n)^2$, which increases with market size.

The proof of Proposition 3 relies on two sets of results. First, notice that Proposition 2 provides bounds on the speeds at which the average efficiency of stable matchings grows for the environments we focus on here. We therefore need bounds on the speed with which the maximal average efficiency grows. As it turns out, finding the maximal aggregate efficiency is a variation of the *optimal assignment problem* in statistics. The literature on optimal assignment problems is still in flux and results are known only for particular distributions, mainly the uniform and exponential distributions. When preferences are fully aligned, we can interpret a result of Walkup (1979), which implies directly that when utilities are drawn from the uniform distribution, $2n - 6 \leq E_n \leq 2n$.[17]

Consider now markets with fully idiosyncratic preferences. That is, for each (i, j), match utilities are given by u_{ij}^f and u_{ij}^w that are distributed uniformly on $[0, 1]$. We define $\tilde{u}_{ij} \equiv \frac{u_{ij}^f + u_{ij}^w}{2}$ and consider the maximal aggregate efficiency achieved by the optimal matching corresponding to a fully aligned market with preferences specified by $(\tilde{u}_{ij})_{i,j}$. Walkup (1979)'s result

[17] Follow-up work has improved upon this bound (see, for instance, Coppersmith and Sorkin, 1999, whose work suggests that $E_n \geq 2n - 3.88$). We use Walkup's bound since it is sufficient for our conceptual message and as we use his method of proof to identify E_n when preferences are fully idiosyncratic.

cannot be used directly, however, since now each $\tilde{u}_{ij}$ is distributed according to the symmetric triangular distribution over $[0,1]$. In the Appendix, we modify the proof in Walkup (1979) and illustrate that, in this environment, $E_n \geq 2n - 3\sqrt{n}$. In fact, in the Online Appendix, we also show that $\lim_{n\to\infty} \frac{2n-E_n}{\sqrt{n}} \geq \sqrt{\frac{\pi}{2}}$ so that indeed the difference between E_n and $2n$ is of the order of $\sqrt{n}$.

4. Assortative Markets

Another class of matching markets that plays an important role in many applications allows for assortative preferences (see Becker, 1973). In such markets, one or both sides of the market agree on the ranking of the other side. For instance, medical residents may evaluate hospitals, at least to some extent, according to their publicly available rankings and hospitals may agree on the attributes that make a resident appealing (see Agarwal, 2015); similarly, potential adoptive parents may evaluate children up for adoption similarly (see Baccara et al., 2014); and so on. In this section, we illustrate that such markets, in which preferences are a combination of a common ranking across firms or workers and arbitrary idiosyncratic shocks, still entail asymptotically average efficient stable matchings.

We assume that each agent has her own intrinsic value, which we denote by $(c_i^f)_{i=1}^n$ for firms and $(c_j^w)_{j=1}^n$ for workers. When firm f_i matches with worker w_j, the firm's utility is determined by the worker's intrinsic value c_j^w and the worker's value assessed individually by the firm, the idiosyncratic component z_{ij}^f. Similarly, worker w_j's utility of matching with firm f_i is a combination of the firm's intrinsic value c_i^f and the worker's idiosyncratic assessment of the firm z_{ij}^w. That is,

$$\begin{aligned} u_{ij}^f &\equiv \Phi(c_j^w, z_{ij}^f) \quad \text{and} \\ u_{ij}^w &\equiv \Omega(c_i^f, z_{ij}^w). \end{aligned}$$

The functions $\Phi(.,.)$ and $\Omega(.,.)$ from $[0,1]^2$ to $[0,1]$ are continuous and strictly increasing in both arguments. We assume that c_i^f, c_j^w, z_{ij}^f, and z_{ij}^w are all drawn independently from distributions that have positive density functions over $[0,1]$.

Let E_n^f and E_n^w be the maximal aggregate efficiency for n firms and workers, respectively, achievable by any market matching:

$$E_n^f \equiv \mathbb{E}\left(\max_{\mu\in M}\sum_{i=1}^n u_{i\mu(i)}^f\right) \quad \text{and} \quad E_n^w \equiv \mathbb{E}\left(\max_{\mu\in M}\sum_{i=1}^n u_{i\mu(i)}^w\right) .$$

In the following proposition, we show that all stable matchings deliver approximately maximal average efficiency as market size increases.

Proposition 4 [Average Efficiency of Stable Matchings]**.** *Stable matchings in assortative markets with idiosyncratic shocks achieve maximal average efficiency asymptotically:*

$$\lim_{n\to\infty}\frac{E_n^f - S_n^f}{n} = \lim_{n\to\infty}\frac{E_n^w - S_n^w}{n} = 0.$$

The proof of Proposition 4 is a direct consequence of Lee (2017). Proposition 1 in the online appendix of Lee (2017) indicates that:

$$\lim_{n\to\infty}\mathbb{E}\left[\frac{S_n^f}{\sum_{i=1}^n \Phi(c_j^w, 1)}\right] = 1,$$

which, in turn, implies the claim of the Proposition pertaining to firms. A symmetric argument holds for the average efficiency experienced by workers.[18]

Lee (2017) suggests that in settings such as these, for any stable matching mechanism, asymptotically, there is an "almost"-equilibrium that implements a stable matching corresponding to the underlying preferences. Formally, Lee (2017) implies that for any stable matching mechanism and any $\varepsilon, \delta, \theta > 0$, there exists N such that with probability of at least $1-\delta$, a market of size $n > N$ has an ε-Nash equilibrium in which a fraction of at least $1-\theta$ of agents reveal their true preferences. Together with our results, this suggests the following.

Corollary 1 [Stable Matching Mechanisms]**.** *When preferences are hybrids of assortative and idiosyncratic components, stable matching mechanisms are asymptotically average efficient and incentive compatible.*

[18]The results of Che and Tercieux (2018) suggest average efficiency that is bounded below the maximal feasible when utility distributions are atomic. Combined with our results, we conjecture that the speed of convergence for continuous distributions that approach atomic ones becomes infinitely slow.

In order to identify the speed of convergence, we restrict attention to linear functions Φ and Ω. For each pair (f_i, w_j), we assume that:

$$\begin{aligned}\Phi(c_j^w, z_{ij}^f) &= (1-\beta)c_j^w + \beta z_{ij}^f, \quad \text{and} \\ \Omega(c_i^f, z_{ij}^w) &= (1-\beta)c_i^f + \beta z_{ij}^w,\end{aligned}$$

where $\beta \in [0,1]$, and c_i^f, c_j^w, z_{ij}^f and z_{ij}^w are independently drawn from the uniform distribution over $[0,1]$ for all i, j.

Notice that any matching generates the same expected average efficiency corresponding to the assortative component of preferences, evaluated at $\beta = 0$, given by $1/2$. Therefore, the maximum conceivable average efficiency is

$$\lim_{n\to\infty} E_n^f = \lim_{n\to\infty} E_n^w = (1-\beta)\frac{1}{2} + \beta.$$

As mentioned, when $\beta = 0$, the average efficiency of the (generically unique) stable matching is $1/2$ for all n. When $\beta = 1$, our characterization in Proposition 2 provides the speed of convergence. The following proposition characterizes the speed at which average efficiency converges to the maximum conceivable when $\beta \in (0,1)$.

Proposition 5. *For any* $\beta \in (0,1)$,

$$(1-\beta)\frac{1}{2} + \beta - \frac{S_n^f}{n} = (1-\beta)\frac{1}{2} + \beta - \frac{S_n^w}{n} = O(n^{-1/4}).$$

Proposition 5 suggests an important difference between markets entailing preferences that have an aligned component, relative to markets characterized by preferences with a dominant assortative component. The convergence speed in the latter is substantially faster.

Figure 2 corresponds to average efficiency levels in 100 simulated markets. As before, we depict a worst-case scenario, where we consider firms' per-participant utilities when the worker-optimal stable matching is implemented in each market. We also depict the maximal feasible per-participant utility, which depends on β. In each panel of the figure, corresponding to a different level of β, we also mark with a horizontal line the bound on the maximal conceivable average efficiency, labeled "efficiency limit." When markets contain about 1000

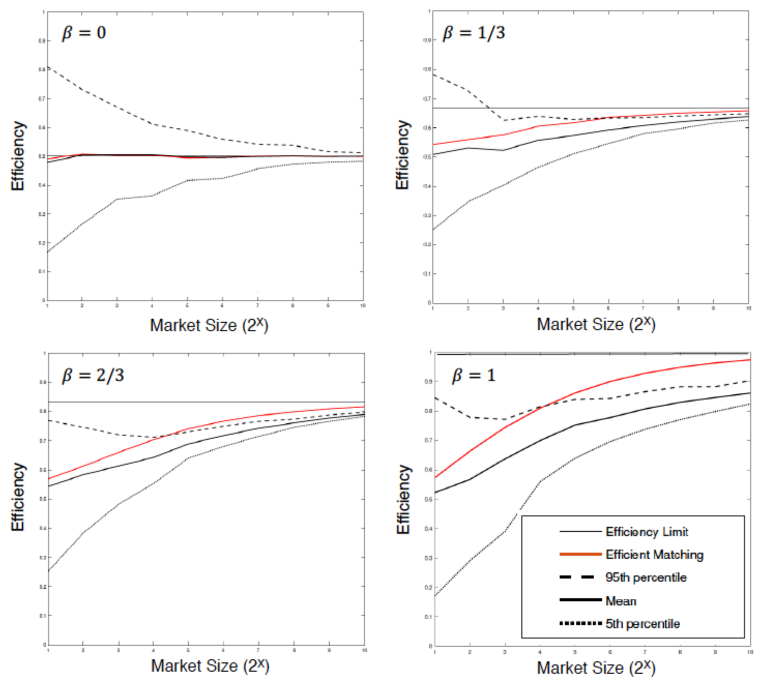

Figure 2: Convergence Speeds for Linear Hybrid Models of Assortative and Idiosyncratic Preferences (β is the Weight on the Idiosyncratic Component)

individuals on each side, the fraction of the maximal achievable per-participant utility that stable matchings yield is quite high, especially for $\beta \leq 2/3$. As for the case of aligned and idiosyncratic hybrid preferences, the dependence of that fraction on β is non-linear, with more pronounced efficiency losses appearing only for fairly high levels of β, even for these fairly small markets.[19]

[19]The case considered here lends itself to particular match utilities that involve interaction terms. Consider a simple case with interactions, where for each firm i with attributes c_i^f and worker j with attributes c_j^w:

$$\begin{aligned} u_{ij}^f &= \beta c_i^f c_j^w + \varepsilon_{ij}^f \\ u_{ij}^w &= \beta c_i^f c_j^w + \varepsilon_{ij}^w, \end{aligned}$$

where c_i^f, c_j^w, and $\varepsilon_{ij}^f, \varepsilon_{ij}^w$ are random variables and we assume that individual attributes c_i^f and c_j^w have supports bounded above 0, say of the form $[a, 1]$, where $a > 0$. Then, when supports are bounded, it suffices to consider the efficiency of stable matchings in a market with modified utilities that normalize each agent's original utility by her own attribute. Under further assumptions of distributions, these utilities are of the form we discuss here. We note that Menzel (2015) considers utility forms as such, where the set of attributes is finite and error terms follow extreme distributions, in the style of Dagsvik (2000). With those assumptions, Menzel (2015) and Peski (2017) illustrate asymptotic average efficiency in the marriage market and the roommate problem, respectively.

In fully assortative markets, all matchings of everyone in the market entail the same efficiency level. Thus, both average and aggregate efficiency levels are maximized under such markets' stable matchings. In the Online Appendix, we also analyze the asymmetric case in which workers all share the same evaluation of firms with utilities determined uniformly, while firms have independent evaluations of workers. In that case, the speed at which the average efficiency of stable matchings converges to the maximum is of the order of $\frac{\log n}{n}$ and aggregate efficiency patterns mimic those of fully aligned preferences. In particular, asymptotically, aggregate efficiency of stable matchings is bounded away from what is feasibly maximal.

5. Markets with Stable Matchings that are not Average Efficient Asymptotically

It is not very hard to find a large market where stable matchings do not achieve maximal average efficiency asymptotically. As mentioned in the Introduction, a replication of small markets in which stable matchings are inefficient generates (larger) markets that entail stable matchings that are inefficient.[20] However, in the setting studied up to now, such markets are asymptotically unlikely. In this section, we study two environments in which stable matchings are average inefficient even when markets are large.

5.1. Severely Imbalanced Markets. Throughout the paper, we assumed that markets are roughly balanced: our presentation pertained to coinciding volumes of firms and workers and, as mentioned at the outset, would carry through for bounded imbalances, e.g., if there were n firms (workers) and $n+k$ workers (firms), where k is fixed.[21] Since in many real-world

[20] One natural way to think of replicating an $m \times m$ market (such as the 2×2 market we first discussed in the introduction) characterized by utilities (u_{ij}^f, u_{ij}^w) is by considering a market of size $km \times km$, with match utilities $(\tilde{u}_{ij}^f, \tilde{u}_{ij}^w)$, where

$$\tilde{u}_{i'j'}^x = \begin{cases} u_{i' \bmod k,\ j' \bmod k}^x & i' \operatorname{div} k = j' \operatorname{div} k \\ 0 & \text{otherwise} \end{cases},$$

so that, for any $l = 0, 1, \ldots, k-1$, firms $f_{lm+1}, ..., f_{(l+1)m}$ and workers $w_{lm+1}, ..., w_{(l+1)m}$ have the same preferences over one another as in the original market, and generically prefer matching with agents in this "sub-market" over matching with anyone else in the market.

[21] In fact, the claims go through for any bounded difference in volumes – i.e., markets with n firms and $n + k(n)$ workers, where $k(n) \leq K$ for all n. The proofs for markets that entail idiosyncratic preferences need to be more carefully modified and are available from the authors. We note that the results are consistent with Ashlagi, Kanoria, and Leshno (2017). For instance, for markets with n firms, $n+1$ workers, and fully idiosyncratic preferences, with high probability, their results suggest that firms' average rank of employed workers in any stable matching is no more than $3 \log n$, whereas the workers' average rank of their employing

matching markets one side has more participants than the other, in this section, we study the robustness of our main result to the assumption that this imbalance is not too severe. This is particularly interesting in view of recent results by Ashlagi, Kanoria, and Leshno (2017) that illustrate the sensitivity of the structure of stable matchings to the relative sizes of both sides of the matching market.

We now consider markets with n firms and $n + k(n)$ workers, where $k(n)$ is increasing in n, and examine the asymptotic average efficiency for matched workers when the firm-optimal stable matching, the worker-pessimal stable matching, is implemented. The main result of this section is that average efficiency for workers may be bounded away from the maximal feasible, even asymptotically, whenever markets are severely imbalanced *and* preferences exhibit substantial idiosyncratic components.

We will assume that utilities from matching with anyone are positive almost always, whereas remaining unmatched generates zero utility. Under these assumptions, all participants of the scarce side of the market are generically matched in any stable matching. Furthermore, the Rural Hospital Theorem (see Roth and Sotomayor, 1992) assures that the set of unmatched individuals does not depend on the implemented stable matching. Since no matching can increase the number of matched individuals, a natural analogue for our average efficiency notion considers the per-person expected utility, *conditional on being matched.* As before, since there might be multiple stable matchings, we will inspect the worst-case scenario. We will continue using the term "average efficiency" for this notion. We focus on cases in which the relative volumes of participants on both sides of the market are comparable, so that $\frac{k(n)}{n}$ is bounded.[22]

Notice that the addition of workers can only improve firms' average efficiency when focusing on the extremal stable matchings (see Roth and Sotomayor, 1992). Therefore, in any balanced setting in which maximal average efficiency is achieved asymptotically, the introduction of more workers will maintain the asymptotic average efficiency of stable matchings for firms.

When markets are fully aligned or fully assortative, the proofs of Propositions 1 and 4 carry through for arbitrary increasing functions $k(n)$ and maximal average efficiency of

firms is at least $n/3 \log n$. With normalization by n, both of these bounds converge to 0.

[22]Whenever $\frac{k(n)}{n}$ explodes, the relevant efficiency statements would pertain to an insignificant fraction of firms that end up being matched.

stable matchings is achieved asymptotically.[23] We now focus on markets with idiosyncratic preferences, where we normalize the utility from remaining unmatched to be zero. Recall that S_n^w denotes workers' aggregate efficiency in the worker-pessimal stable matching.

The following proposition illustrates the impacts of market imbalances. If one side of the market is proportionally larger and preferences are fully idiosyncratic, average inefficiency may arise even when markets are large.

Proposition 6 [Imbalanced Markets with Fully Idiosyncratic Preferences]**.** *Suppose* $k(n) \geq \lambda n$ *for some* $\lambda > 0$*, and all utilities* $(u_{ij}^f)_{i,j}$ *and* $(u_{ij}^w)_{i,j}$ *are independently drawn from the uniform distribution over* $[0,1]$*. Then,*

$$\lim_{n\to\infty} \frac{S_n^w}{n} \leq \begin{cases} 1 - \frac{1}{-3\log\lambda} & \textit{for} \quad 0 < \lambda \leq 1/2 \\ 1 - \frac{1}{3\log 2} & \textit{for} \quad 1/2 < \lambda. \end{cases}$$

Notice that this indeed suggests average inefficiency of stable matchings in large markets. For each realization of a market, characterized by realized utilities $(u_{ij}^f)_{i,j}$ and $(u_{ij}^w)_{i,j}$, consider the induced fully aligned market with utilities $(\tilde{u}_{ij} \equiv \frac{u_{ij}^f + u_{ij}^w}{2})_{i,j}$. That is, in the induced market, each matched firm and worker receive their average match utilities in the original market. The average efficiency of stable matchings pertaining to aligned markets then carry through for the induced market. Since these matchings produce the same per-participant utilities in the original market, maximal average efficiency can be achieved asymptotically. The wedge identified in Proposition 6 then implies a substantial average efficiency loss due to stability, even in large markets.

To gain some intuition as to why severe imbalances can lead to average inefficiencies when preferences are idiosyncratic, consider an extreme sequence of markets comprised of one firm and n workers. In such markets, the stable matching matches the firm with its favorite worker. Therefore, the firm's expected utility in the stable matching is the maximum of n samples from the uniform distribution over $[0,1]$, which is $\frac{n}{n+1}$ and indeed converges to 1. However, since workers' utilities are drawn independently from the firm's, the worker matched under the stable matching has an expected utility of $\frac{1}{2}$. Nonetheless, for the matching that maximizes

[23] For markets with underlying preferences that are hybrids of assortative and idiosyncratic, slightly more involved arguments are required that follow directly from results in Lee (2017).

efficiency, we should look for j that maximizes $u^f_{1j} + u^w_{1j}$, which is distributed according to the symmetric triangular distribution over $[0, 2]$. Therefore, the maximal feasible aggregate, and thereby average, efficiency corresponds to the maximum of n samples from the triangular distribution, which converges to 2 as n grows large. Roughly speaking, the crux of this example is that, in stable matchings, agents in the scarce side of the market do not take into account utilities achieved by the other side of the market. In particular, a matching that implies even a minuscule loss for the firm, but a substantial increase in the utility of the matched worker will not be implemented.

5.2. Sub-modular Match Utilities. Another important class of markets in which average inefficiency arises asymptotically pertains to assortative preferences in which match utilities are sub-modular in partners' intrinsic values. For finite markets, Becker (1974) illustrated that sub-modularity in assortative markets leads to the negatively assortative unique efficient matching and the positively assortative unique stable matching.

Formally, consider a sequence of markets in which firms' intrinsic values are given by $(c^f_i = i/n)^n_{i=1}$ and workers' intrinsic values are given by $(c^w_j = j/n)^n_{j=1}$. Match utilities are determined according to an "output function" ϕ:

$$u^f_{ij} = u^w_{ij} = \phi(c^f_i, c^w_j)$$

such that

$$\frac{\partial \phi(c^f_i, c^w_j)}{\partial c^f_i} > 0, \quad \frac{\partial \phi(c^f_i, c^w_j)}{\partial c^w_j} > 0.$$

The positively assortative matching partners each f_i with w_i, and it is the unique stable matching in these markets. The negatively assortative matching partners each f_i with w_{n+1-i}.

The cross-partial derivatives of the output function ϕ are crucial in determining whether the positively assortative matching is an (average or aggregate) efficient matching or not. Indeed, when ϕ is linear, all matchings generate the same efficiency and both the positively and negatively assortative matchings are efficient. When output is super-modular, $\frac{\partial^2 \phi(c^f_i, c^w_j)}{\partial c^f_i \partial c^w_j} > 0$, the positively assortative matching (i.e., the stable matching) is an efficient matching, while when output is sub-modular, $\frac{\partial^2 \phi(c^f_i, c^w_j)}{\partial c^f_i \partial c^w_j} < 0$, the positively assortative matching is not an efficient matching, which is negatively assortative.

In order to illustrate how these features may carry through to large markets, we consider a particular class of output functions:

$$u_{ij}^f = u_{ij}^w = \phi(c_i^f, c_j^w) \equiv (c_i^f + c_j^w)^\alpha,$$

where $\alpha \in (0,1)$, so that output is sub-modular.

The maximal aggregate efficiency, from the efficient (negatively assortative) matching, is

$$E_n = 2 \cdot \sum_{i=1}^{n} \left(\frac{i + \mu(i)}{n} \right)^\alpha = 2 \cdot \sum_{i=1}^{n} \left(\frac{n+1}{n} \right)^\alpha = 2n \left(\frac{n+1}{n} \right)^\alpha .$$

Therefore,

$$\lim_{n \to \infty} \frac{E_n}{2n} = \lim_{n \to \infty} \left(\frac{n+1}{n} \right)^\alpha = 1.$$

On the other hand, aggregate efficiency from the stable (positively assortative) matching is

$$S_n = 2 \cdot \sum_{i=1}^{n} \left(\frac{i + \mu(i)}{n} \right)^\alpha = 2 \cdot \sum_{i=1}^{n} \left(\frac{2i}{n} \right)^\alpha .$$

Note that

$$\frac{1}{n} \sum_{i=1}^{n} \left(\frac{2i}{n} \right)^\alpha \geq \int_0^1 (2x)^\alpha dx = \frac{2^\alpha}{\alpha + 1} \geq \frac{1}{n} \sum_{i=0}^{n-1} \left(\frac{2i}{n} \right)^\alpha = \frac{1}{n} \sum_{i=1}^{n} \left(\frac{2i}{n} \right)^\alpha - \frac{2^\alpha}{n}.$$

Thus,

$$\lim_{n \to \infty} \frac{S_n}{2n} = \lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} \left(\frac{2i}{n} \right)^\alpha = \frac{2^\alpha}{\alpha + 1}.$$

In this example, $\frac{\partial^2 \phi(c_i^f, c_j^w)}{\partial c_i^f \partial c_j^w} = -\alpha(1-\alpha)(c_i^f + c_j^w)^{\alpha-2}$. Sub-modularity then vanishes as α approaches 0 or 1. When α is close to 0, output is very insensitive to individual values, and any matching, in particular the stable one, generates average efficiency close to the optimum. When α is close to 1, the output function is "almost linear" in individual values and, again, the stable matching is asymptotically nearly average efficient. Nonetheless, for intermediate values of α, preferences are strictly sub-modular and asymptotic average efficiency of the stable matching is bounded strictly below 1, which is achieved by the most efficient matching.

6. Conclusions

This paper illustrates that for a large class of preferences, stable matchings achieve maximal average efficiency asymptotically, but their aggregate efficiencies fall far short of the maximal feasible efficiency even when markets are large. These conclusions are particularly relevant in view of the observation that many markets entail fixed wages (see Hall and Kreuger, 2012), or are subject to legal or "moral" constraints that ban transfers (see Roth, 2007). In such settings, a market designer who implements stable outcomes using an ordinal mechanism sacrifices little in terms of average efficiency. However, a designer who aims at maximizing aggregate efficiency would potentially require mechanisms that take into account cardinal utilities, and possibly allow transfers among participants.

Our results also illustrate the speeds of convergence of the average efficiency of stable matchings to the optimum. Idiosyncratic preferences yield a substantially lower speed of convergence than those exhibited in markets with aligned or assortative preferences. This suggests that even market designers concerned with average efficiency should consider market size with special caution.

Markets with idiosyncratic preferences are also fragile to imbalances in the volumes of participants on either side. When those imbalances are severe—when the volume of one side constitutes a fixed fraction of the volume of the other—stable matchings are no longer average efficient in general, even in large markets. Again, mechanisms that take cardinal utilities into account, possibly allowing transfers, could prove beneficial.

While our results simply assess the efficiency features of stable matchings in a variety of markets, they open the door for many interesting questions regarding incentive compatibility of efficient mechanisms. When preferences combine assortative and idiosyncratic components, stable matchings are not only asymptotically average efficient, they are also asymptotically incentive compatible. Our results then serve as a rather positive defense of commonly used mechanisms such as the Gale-Shapley (1962) deferred acceptance algorithm when average efficiency is the objective—in such settings, these mechanisms are asymptotically incentive compatible and average efficient (our Corollary 1). Nonetheless, for a designer concerned with aggregate efficiency, it would be important to analyze the most efficient incentive compatible mechanisms. Furthermore, for other types of preferences, even the question of incentive compatibility of stable mechanisms in large markets is still open.

Our analysis pertains to one-to-one matching markets such as those matching doctors and residency positions, rabbis and congregations, children up for adoption and potential adoptive parents, etc. It also extends directly to many-to-one matching markets, such as school choice, if preferences are responsive. Additional details could be relevant for such an analysis, for instance parents' investments in student qualifications (see Cole, Mailath, and Postlewaite, 2001) and the interaction between matching processes and the real estate market (see Avery and Pathak, 2017). Naturally, incorporating these details would make welfare assessments more subtle in this context.

7. Appendix – Proofs

7.1. Proof of Propositions 1 and 2 for Fully Aligned Preferences. We start by proving Propositions 1 and 2 for the case of fully aligned preferences, as in Section 3.3. That is, the match utilities of each firm and worker pair depend only on a random common value. Formally, we assume the functions $\phi(.,.)$ and $\omega(.,.)$ are both strictly increasing in c_{ij} but independent of the idiosyncratic components.

We begin with the derivation of a formula for S_n suggested in the text when utilities are distributed uniformly and illustrate both asymptotic average efficiency and the speed of convergence for that case. We then generalize our asymptotic average efficiency result to arbitrary continuous distributions.

As illustrated in the text, realized utilities $(u_{ij})_{i,j}$ generically induce a unique stable matching. Denote by $u_{[k;n]}$ the k-th highest match utility of pairs matched within that unique stable matching. Therefore,

$$\frac{S_n}{2} = \mathbb{E}\left(\sum_{k=1}^{n} u_{[k;n]}\right) = \sum_{k=1}^{n} \mathbb{E}(u_{[k;n]}).$$

We use induction to show that for $k = 1, \ldots, n$,

$$\mathbb{E}(u_{[k;n]}) = \frac{n^2}{n^2+1}\frac{(n-1)^2}{(n-1)^2+1}\cdots\frac{(n-k+1)^2}{(n-k+1)^2+1}.$$

For $k = 1$, $u_{[1;n]}$ is the maximal utility achievable from all firm-worker pairs. Thus, $u_{[1;n]}$ is the highest entry from n^2 samples from the uniform distribution over $[0,1]$ and so:

$$\mathbb{E}(u_{[1;n]}) = \frac{n^2}{n^2+1}$$

Suppose the claim is shown for $k-1$. From the construction of the stable matching, $u_{[k;n]}$ is the maximal utility among all firm-worker pairs, after all firms and workers receiving the $k-1$ highest utilities within the stable matching have been removed from the market. Thus, $u_{[k;n]}$ is the highest entry from $(n-k+1)^2$ samples from the uniform distribution over $[0,1]$ restricted so that each sample has a value lower than or equal to $u_{[k-1;n]}$. Therefore,

$$\mathbb{E}(u_{[k;n]}|u_{[k-1;n]}) = u_{[k-1;n]}\frac{(n-k+1)^2}{(n-k+1)^2+1}.$$

By the law of iterated expectations,

$$\begin{aligned}\mathbb{E}(u_{[k;n]}) &= \mathbb{E}\left(\mathbb{E}(u_{[k;n]}|u_{[k-1;n]})\right) = \mathbb{E}\left(u_{[k-1;n]}\right)\frac{(n-k+1)^2}{(n-k+1)^2+1} \\ &= \frac{n^2}{n^2+1}\cdots\frac{(n-k+2)^2}{(n-k+2)^2+1}\frac{(n-k+1)^2}{(n-k+1)^2+1},\end{aligned}$$

where the last equality is from the induction hypothesis. The formula for S_n follows.

We now turn to the proof of Propositions 1 and 2 for this setting. First, denote by $C_n = 1 - \frac{S_n}{2n}$.

Notice that $C_1 = 1/2$. For any $n > 1$,

$$\frac{S_n}{2} = \frac{n^2}{n^2+1} + \frac{n^2}{n^2+1}\frac{S_{n-1}}{2},$$

which implies that

$$\begin{aligned}C_n &= 1 - \frac{S_n}{2n} = 1 - \frac{n}{n^2+1} - \frac{n}{n^2+1}\frac{S_{n-1}}{2} \\ &= \frac{1}{n^2+1} + \frac{n(n-1)}{n^2+1}\left(1 - \frac{S_{n-1}}{2(n-1)}\right) = \frac{1}{n^2+1} + \frac{n(n-1)}{n^2+1}C_{n-1}.\end{aligned}$$

Then, we can find that

$$C_3 = \frac{1}{10} + \frac{6}{10}C_2 = \frac{1}{10} + \frac{6}{10}\left(\frac{1}{5} + \frac{2}{5}\frac{1}{2}\right) = \frac{17}{50} = 0.34,$$

which is between $\frac{1}{2}\frac{\log 3}{3} = 0.18...$ and $\frac{\log 3}{3} = 0.37\ldots$.

For any $n > 3$, suppose

$$\frac{1}{2}\frac{\log(n-1)}{n-1} \le C_{n-1} \le \frac{\log(n-1)}{n-1}.$$

Then,

$$\frac{1}{n^2+1} + \frac{n}{2(n^2+1)}\log(n-1) \le C_n = \frac{1}{n^2+1} + \frac{n(n-1)}{n^2+1}C_{n-1} \le \frac{1}{n^2+1} + \frac{n}{n^2+1}\log(n-1).$$

Notice that

$$\frac{1}{n^2+1}+\frac{n}{n^2+1}\log(n-1)<\frac{1}{n^2}+\frac{1}{n}\log(n-1)<\frac{\log n}{n},$$

where the second inequality, which is equivalent to $\log\left(1-\frac{1}{n}\right)<-\frac{1}{n}$, holds from $1-\frac{1}{n}<e^{-1/n}$.

Further, we have

$$\frac{1}{n^2+1}+\frac{n}{2(n^2+1)}\log(n-1)=\frac{2+n\log(n-1)}{2(n^2+1)}>\frac{n\log n+(1/2)}{2(n^2+1)}>\frac{\log n}{2n},$$

where the first inequality, which is equivalent to $n\log\left(\frac{n}{n-1}\right)<\frac{3}{2}$, holds for $n=4$ since $4\log(4/3)=1.15\ldots$ and for $n>4$ since

$$\frac{d\left(n\log\left(\frac{n}{n-1}\right)\right)}{dn}=\log\left(\frac{n}{n-1}\right)-\frac{1}{n-1}<0.$$

Note that $\frac{n}{n-1}=1+\frac{1}{n-1}<e^{\frac{1}{n-1}}$ for $n\geq 3$.

The proof of Propositions 1 and 2 for this case then follows.

We now show that stable matchings are asymptotically average efficient even when the match utilities of firm and worker pairs are independently and identically drawn from an arbitrary continuous distribution G over $[0,1]$. Indeed, we show that for any $\varepsilon>0$,

$$\lim_{n\to\infty}\frac{S_n}{2n}\geq 1-\varepsilon.$$

We construct a uniform distribution G' such that G first order stochastically dominates G'. The support of G is $[0,1]$, so there exists γ such that $1-\varepsilon<\gamma<1$ and $G(\gamma)<1$.

Let

$$G'(x)=\frac{1-G(\gamma)}{\gamma}x+G(\gamma),$$

so that G' is the uniform distribution over $[0,\gamma]$.

Let $S_n^{'}$ be the expected aggregate efficiency derived from the stable matching when utilities are drawn from $G'(x)$. The proof above can be replicated to show

$$\lim_{n\to\infty}\frac{S_n'}{2n}=\gamma.$$

Note that G first order stochastically dominates G'. Thus, every order statistic corresponding to samples from G first order stochastically dominates the corresponding order statistic of the same number of samples from G'. For every realized utilities $(u_{ij})_{i,j}$, S_n is a sum of specific n order statistics. Thus, the aggregate efficiency from the stable matching under G first order stochastically dominates that under G'. It follows that

$$\lim_{n\to\infty} \frac{S_n}{2n} \geq \lim_{n\to\infty} \frac{S'_n}{2n} = \gamma > 1 - \varepsilon.$$

Since ε is arbitrary, $\lim_{n\to\infty} \frac{S_n}{2n} = 1$, as desired. ■

7.2. Proof of Proposition 1 for Aligned Preferences with Idiosyncratic Shocks. We now provide the proof of Proposition 1 for the case in which firms' and workers' match utilities, $\phi(.,.)$ and $\omega(.,.)$, are strictly increasing in the idiosyncratic components, z_{ij}^f and z_{ij}^w and either both are independent of the common component c_{ij} or both are strictly increasing in c_{ij}. It is without loss of generality to consider c_{ij}, z_{ij}^f, and z_{ij}^w that are all uniformly distributed over $[0, 1]$. Indeed, an appropriate change of variables will generate an equivalent setting in which the underlying distributions are uniform.[24]

The model is potentially a mixture of aligned preferences captured by the variables $c = (c_{ij})_{i,j}$ and idiosyncratic preferences captured by the variables $z^f = (z_{ij}^f)_{i,j}$ and $z^w = (z_{ij}^w)_{i,j}$. Accordingly, our proof is comprised of two parts.

For each realization (c, z^f, z^w), let

$$\bar{F}(\varepsilon; c, z^f, z^w) \equiv \{f_i | c_{i\mu^w(i)} \leq 1 - \varepsilon\}.$$

Whenever $\phi(.,.)$ and $\omega(.,.)$ are strictly increasing in the common component, we first show that for any $\varepsilon > 0$,

$$\lim_{n\to\infty} \mathbb{E}\left[\frac{|\bar{F}(\varepsilon; c, z^f, z^w)|}{n}\right] = 0. \tag{2}$$

[24]See the online appendix of Lee (2017) for details.

In the second part of the proof, we show that for any $\varepsilon > 0$,

$$P\left(\frac{\sum_{i=1}^{n} z^{f}_{i\mu^{w}(i)}}{n} \leq 1-\varepsilon\right) \rightarrow 0 \quad \text{as} \quad n \rightarrow \infty. \tag{3}$$

Proposition 1, for cases excluding fully aligned preferences, is immediate from (2) and (3). For any ε, there exist ε' such that if $\phi(c, z^f) < \phi(1,1) - \varepsilon$ then either $c < 1-\varepsilon'$ or $z^f < 1-\varepsilon'$. Therefore,

$$\frac{1}{n}\left|\{f_i | u_{i\mu^w(i)} \leq \phi(1,1)-\varepsilon\}\right| \leq \frac{1}{n}\left|\{f_i | c_{i\mu^w(i)} \leq 1-\varepsilon'\}\right| + \frac{1}{n}\left|\{f_i | z_{i\mu^w(i)} \leq 1-\varepsilon'\}\right|.$$

The right hand side converges to zero in probability by (2) and (3).

Proof of Equation (2)**.** Assume that $\phi(.,.)$ and $\omega(.,.)$ are strictly increasing in the common component c_{ij}, in which case Equation (2) is relevant for our proof.

A *graph* G is a pair (V, E), where V is a set called *nodes* and E is a set of unordered pairs (i,j) or (j,i) of $i, j \in V$ called *edges*. The nodes i and j are called the *endpoints* of (i,j). We say that a graph $G = (V, E)$ is *bipartite* if its node set V can be partitioned into two disjoint subsets V_1 and V_2 such that each of its edges has one endpoint in V_1 and the other in V_2.

A *biclique* of a bipartite graph $G = (V_1 \cup V_2, E)$ is a set of nodes $U_1 \cup U_2$ such that $U_1 \subset V_1$, $U_2 \subset V_2$, and for all $i \in U_1$ and $j \in U_2$, $(i,j) \in E$. In other words, a biclique is a complete bipartite subgraph of G. We say that a biclique is *balanced* if $|U_1| = |U_2|$, and refer to a balanced biclique with the maximal number of nodes as *a maximal balanced biclique.*

Given a partitioned set $V_1 \cup V_2$, we consider a random bipartite graph $G(V_1 \cup V_2, p)$. A bipartite graph $G = (V_1 \cup V_2, E)$ is constructed so that each pair of nodes, one in V_1 and the other in V_2, is included in E independently with probability p. We use the following proposition in the proof.

Proposition 7 [Dawande et al., 2001]**.** *Consider a random bipartite graph $G(V_1 \cup V_2, p)$, where $0 < p < 1$ is a constant, $|V_1| = |V_2| = n$, and $\delta(n) = \log n / \log \frac{1}{p}$. If a maximal balanced biclique of this graph has size $B \times B$, then*

$$Pr\left(\delta(n) \leq B \leq 2\delta(n)\right) \rightarrow 1, \quad \text{as} \quad n \rightarrow \infty.$$

By continuity of $\phi(.,.)$ and $\omega(.,.)$, there exists $\varepsilon' > 0$ such that

$$c_{ij}, z^f_{ij}, z^w_{ij} > 1 - \varepsilon' \implies \phi(c_{ij}, z^f_{ij}) > \phi(1-\varepsilon, 1) \text{ and } \omega(c_{ij}, z^w_{ij}) > \omega(1-\varepsilon, 1).$$

For each realization (c, z^f, z^w), we draw a bipartite graph such that $F \cup W$ is the set of nodes (where F and W constitute the two parts of the graph), and each pair of f_i and w_j is connected by an edge if and only if at least one of c_{ij}, z^f_{ij}, or z^w_{ij} is lower than or equal to $1 - \varepsilon'$.

Let

$$\bar{W}(\varepsilon; c, z^f, z^w) \equiv \{w_j | \mu^w(j) \in \bar{F}(\varepsilon; c, z^f, z^w)\}.$$

Then $\bar{F} \cup \bar{W}$ is a balanced biclique. If a pair (f_i, w_j) from $\bar{F} \cup \bar{W}$ is not connected by an edge, then the pair can achieve utilities $\phi(c_{ij}, z^f_{ij}) > \phi(1-\varepsilon, 1)$ and $\omega(c_{ij}, z^w_{ij}) > \omega(1-\epsilon, 1)$ because $c_{ij}, z^f_{ij}, z^w_{ij} > 1 - \varepsilon'$. The two utilities are higher than their utilities under μ^w. This contradicts μ^w being stable.

Proposition 7 then implies Equation (2).

Proof of Equation (3)**.** Let $\mu \equiv \{(i,i) | i = 1, \ldots, n\}$. By symmetry, each one of the $n!$ matchings has the same probability of being both stable and entailing a sum of firms' idiosyncratic components that is lower than or equal to $(1-\varepsilon)n$. Therefore,

$$P\left(\frac{\sum_{i=1}^n z^f_{i\mu^w(i)}}{n} \leq 1 - \varepsilon\right) \leq n! P\left(\mu \text{ is stable and } \frac{\sum_{i=1}^n z^f_{i\mu(i)}}{n} \leq 1 - \varepsilon\right).$$

For each realization (c, z^f, z^w), we consider the following profile of utilities.

$$\begin{aligned} &\tilde{u}^f_{ij} = \phi(c_{ij}, z^f_{ij}) \quad \text{and} \quad \tilde{u}^w_{ij} = \omega(c_{ij}, z^w_{ij}) \quad \text{if } i \neq j, \quad \text{and} \\ &\tilde{u}^f_{ij} = \phi(1, z^f_{ij}) \quad \text{and} \quad \tilde{u}^w_{ij} = \omega(c_{ij}, z^w_{ij}) \quad \text{if } i = j. \end{aligned}$$

Remark 1. *If preferences are fully idiosyncratic, so that ϕ and ω are independent of c_{ij}, then $\tilde{u}_{ij} = u_{ij}$ for all $f_i \in F$ and $w_j \in W$.*

If a pair of f_i and w_j with $i \neq j$ is a blocking pair of μ under match utilities $(\tilde{u}^f_{ij}, \tilde{u}^w_{ij})$, then the pair also blocks μ under the actual realized utilities. Thus,

$$P\left(\mu \text{ is stable and } \frac{\sum_{i=1}^n z^f_{i\mu(i)}}{n} \leq 1-\varepsilon\right) \leq P_{\varepsilon,n}$$

where $P_{\varepsilon,n}$ is the probability that μ is stable with respect to the utilities $(\tilde{u}^f_{ij}, \tilde{u}^w_{ij})_{i,j}$ and $\frac{\sum_{i=1}^n z^f_{i\mu(i)}}{n} \leq 1-\varepsilon$. We prove that $n!P_{\varepsilon,n}$ converges to zero as n increases.

Preparation Steps for Proof of (3). We denote by Γ^f the marginal distribution of $\tilde{u}^f_{ij}$ for pairs of (f_i, w_j) that are not matched under μ (i.e., $i \neq j$), and by Γ^w the marginal distribution of $\tilde{u}^w_{ij}$ for any pair of (f_i, w_j). We define, for all i, j, $\hat{u}^f_{ij} = \Gamma^f(\tilde{u}^f_{ij})$ and $\hat{u}^w_{ij} = \Gamma^w(\tilde{u}^w_{ij})$.

Remark 2. *The marginal distributions of $\hat{u}^f_{ij}$ for firm and worker pairs with $i \neq j$, and $\hat{u}^w_{ij}$ for all pairs are uniform over $[0,1]$. Whereas, for pairs with $i = j$, the marginal distribution of $\hat{u}^f_{ij}$ first order stochastically dominates the uniform distribution over $[0,1]$.*

For each given realization $(\hat{u}^f_{ii})^n_{i=1}$ and $(\hat{u}^w_{jj})^n_{j=1}$, the probability that μ is stable with respect to $(\tilde{u}^f_{ij}, \tilde{u}^w_{ij})_{i,j}$ is the same as

$$\Pi_{1\leq i\neq j\leq n}\left(1 - P[\hat{u}^f_{ij} > \hat{u}^f_{ii} \text{ and } \hat{u}^w_{ij} > \hat{u}^w_{jj}]\right).$$

Note that c_{ij}, z^f_{ij}, and z^w_{ij} are independently and identically distributed, so they are positively associated (See Theorem 2.1 in Esary, Proschan, and Walkup, 1967). Indeed, since $\Gamma^f(\phi(.,.))$ and $\Gamma^w(\omega(.,.))$ are non-decreasing functions of c_{ij}, z^f_{ij}, and z^w_{ij}, the covariance of the corresponding values, $\hat{u}^f_{ij}$ and $\hat{u}^w_{ij}$ for $i \neq j$ is non-negative. Thus, we have

$$P[\hat{u}^f_{ij} > \hat{u}^f_{ii} \text{ and } \hat{u}^w_{ij} > \hat{u}^w_{jj}] \geq P[\hat{u}^f_{ij} > \hat{u}^f_{ii}]P[\hat{u}^w_{ij} > \hat{u}^w_{jj}] = (1-\hat{u}^f_{ii})(1-\hat{u}^w_{jj}).$$

Last, take any $0 < \gamma < 1/2$ such that

$$\Gamma^f(\phi(1, z^f_{ii})) \leq (1-\gamma) + \gamma z^f_{ii}.$$

Then,

$$1 - \hat{u}^f_{ii} \geq \gamma(1 - z^f_{ii}).$$

Therefore, for each realization of $c_{ij}, z^f_{ij}, z^w_{ij}$ for pairs with $i = j$, the probability that μ is stable is bounded above by

$$\Pi_{1\le i\neq j\le n}\left(1-\gamma(1-z^f_{ii})(1-\hat{u}^w_{jj})\right).$$

We therefore obtain that

$$P_{\varepsilon,n} \le \int\int_{\sum_{i=1}^n z^f_{ii}\le(1-\varepsilon)n} \Pi_{1\le i\neq j\le n}\left(1-\gamma(1-z^f_{ii})(1-\hat{u}^w_{jj})\right) d\mathbf{z^f_{ii}} d\mathbf{\hat{u}^w_{jj}}.$$

Now, let $x_i = 1 - z^f_{ii}$ and $y_j = 1 - \hat{u}^w_{jj}$. Then,

$$P_{\varepsilon,n} \le \int_{\substack{\mathbf{0}\le\mathbf{x},\mathbf{y}\le 1\\ \varepsilon n\le\sum_{i=1}^n x_i}} \Pi_{1\le i\neq j\le n}\left(1-\gamma x_i y_j\right) d(\mathbf{x},\mathbf{y}).$$

Proof of Convergence.

$$\begin{aligned} P_{\varepsilon,n} &\le \int_{\substack{\mathbf{0}\le\mathbf{x},\mathbf{y}\le 1\\ \varepsilon n\le\sum_{i=1}^n x_i}} \Pi_{1\le i\neq j\le n}(1-\gamma x_i y_j) d(\mathbf{x},\mathbf{y}) \\ &= \int_{\substack{\mathbf{0}\le\mathbf{x},\mathbf{y}\le 1\\ \varepsilon n\le\sum_{i=1}^n x_i}} \Pi_{1\le j\le n}\underbrace{\left(\int_0^1 \Pi_{i\neq j}(1-\gamma x_i y_j) dy_j\right)}_{(*)} d\mathbf{x}. \end{aligned}$$

Let $t = n^{-7/8}$ and $\delta = e^t(1-\gamma t)$. As $0 < \gamma < 1/2$, for any $0 \le z \le t$, we have $1-\gamma z \le \delta e^{-z}$. Thus, when $0 \le y_j \le t$, we have $0 \le x_i y_j \le t$, so $1 - \gamma x_i y_j \le \delta \exp(-x_i y_j)$. In addition, $1 + \gamma z \le e^{\gamma z}$ for any z, so $1 - \gamma x_i y_j \le exp(-\gamma x_i y_j)$.

Therefore,

$$\begin{aligned} (*) &= \int_0^t \Pi_{i\neq j}(1-\gamma x_i y_j) dy_j + \int_t^1 \Pi_{i\neq j}(1-\gamma x_i y_j) dy_j \\ &= \int_0^t \Pi_{i\neq j}\delta\exp(-x_i y_j) dy_j + \int_t^1 \Pi_{i\neq j}\exp(-\gamma x_i y_j) dy_j \\ &= \delta\int_0^t \exp\left(-y_j\sum_{i\neq j} x_i\right) dy_j + \int_t^1 \exp\left(-\gamma y_j\sum_{i\neq j} x_i\right) dy_j. \end{aligned}$$

Let

$$s = \sum_{i=1}^{n} x_i \quad \text{and} \quad s_j = \sum_{i \neq j} x_i.$$

Then,

$$\begin{aligned} (*) &= \delta \int_0^t \exp\left(-y_j s_j\right) dy_j + \int_t^1 \exp\left(-\gamma y_j s_j\right) dy_j \\ &= \delta \frac{1 - e^{-t s_j}}{s_j} + \frac{e^{-\gamma t s_j} - e^{-\gamma s_j}}{\gamma s_j} \leq \frac{1}{s_j}\left(\delta + \frac{1}{\gamma}\exp(-\gamma t s_j)\right). \end{aligned}$$

We claim that

$$\delta + \frac{1}{\gamma}\exp(-\gamma t s_j) < \exp(n^{-6/7}) \quad \text{for sufficiently large } n. \tag{4}$$

As $s > \varepsilon n$, we have $s_j > \varepsilon n - 1$. Thus, (4) follows from

$$\exp(n^{-7/8}) + \frac{1}{\gamma}\exp(-\gamma n^{-7/8}(\varepsilon n - 1)) < \exp(n^{-6/7}) \quad \text{for any sufficiently large } n.$$

Now, we have

$$\begin{aligned} P_{\varepsilon,n} &\leq \int_{\substack{\mathbf{0} \leq \mathbf{x} \leq 1 \\ \varepsilon n \leq \sum_{i=1}^{n} x_i}} \Pi_{1 \leq j \leq n} (*) d\mathbf{x} \\ &\leq \int_{\substack{\mathbf{0} \leq \mathbf{x} \leq 1 \\ \varepsilon n \leq \sum_{i=1}^{n} x_i}} \Pi_{1 \leq j \leq n} \left(\frac{1}{s_j}\exp(n^{-6/7})\right) d\mathbf{x} = \exp(n^{1/7}) \int_{\substack{\mathbf{0} \leq \mathbf{x} \leq 1 \\ \varepsilon n \leq \sum_{i=1}^{n} x_i}} \Pi_{1 \leq j \leq n} \left(\frac{1}{s_j}\right) d\mathbf{x}. \end{aligned}$$

Note that $\left(\log \frac{1}{s_j}\right)' = -\frac{1}{s_j}$. Thus,

$$\sum_{j=1}^{n} \log \frac{1}{s_j} = \sum_{j=1}^{n} \left(\log \frac{1}{s} + \log \frac{s}{s_j}\right) = n \log \frac{1}{s} + \sum_{j=1}^{n} \log \frac{s}{s - x_j}.$$

In the last term,

$$\frac{s}{s - x_j} \leq \frac{s}{s - 1} \leq \frac{\varepsilon n}{\varepsilon n - 1}, \quad \text{for any sufficiently large } n.$$

Thus,

$$\sum_{j=1}^{n} \log \frac{1}{s_j} \leq n \ln \frac{1}{s} + n \log \frac{\varepsilon n}{\varepsilon n - 1}.$$

Moreover

$$n \log \frac{\varepsilon n}{\varepsilon n - 1} = \log \left(\left(1 + \frac{1}{\varepsilon n - 1} \right)^n \right) \to \frac{1}{\varepsilon} \quad \text{as} \quad n \to \infty,$$

which implies that for any $c > \frac{1}{\varepsilon}$,

$$\sum_{j=1}^{n} \log \frac{1}{s_j} \leq n \log \frac{1}{s} + c, \quad \text{for any sufficiently large } n.$$

Therefore,

$$n! P_{\varepsilon,n} \leq n! \exp(n^{1/7}) \int_{\varepsilon n \leq s} \exp \left(n \log \frac{1}{s} + c \right) f_n(s) ds,$$

where $f_n(s)$ is the probability distribution function of s.

We show the convergence of the right hand side of the above inequality by using the following Lemma.[25]

Lemma 1 [Pittel, 1989]**.** *Let $x_1, \ldots, x_{n-1}$ be i.i.d samples from the Uniform distribution over $[0,1]$. Denote by $x_{(k)}$ the k'th highest of these samples. We define a random variable*

$$r_n \equiv \max_{0 \leq i \leq n-1} \{ x_{(i)} - x_{(i+1)} \},$$

where $x_{(0)} \equiv 1$ and $x_{(n)} \equiv 0$.

Then,

$$f_n(s) = \frac{s^{n-1}}{(n-1)!} Pr(r_n \leq s^{-1}),$$

and

$$Pr(r_n \leq x) \leq \exp \left(-n e^{-x(n + n^{9/14})} \right) + O \left(e^{-\frac{n^{2/7}}{2}} \right).$$

[25] The Lemma follows from Lemma 1 combined with the first two equations on the top of page 548 in Pittel (1989).

By applying Lemma 1, we get

$$\begin{aligned} n!P_{\varepsilon,n} &\leq n!\exp(n^{1/7})\int_{\varepsilon n\leq s}\exp\left(n\log\frac{1}{s}+c\right)\frac{s^{n-1}}{(n-1)!}Pr(r_n\leq s^{-1})ds \\ &\leq e^c n\exp(n^{1/7})Pr\left(r_n\leq\frac{1}{\varepsilon n}\right)\int_{\varepsilon n}^{n}\frac{1}{s}\,ds \\ &= e^c n\exp(n^{1/7})Pr\left(r_n\leq\frac{1}{\varepsilon n}\right)(-\log\varepsilon). \end{aligned}$$

Now,

$$\begin{aligned} n\exp(n^{1/7})Pr\left(r_n\leq\frac{1}{\varepsilon n}\right) &\leq n\exp(n^{1/7})\left(\exp\left(-ne^{-\frac{1+n^{-\frac{5}{14}}}{\varepsilon}}\right)+O(e^{-\frac{n^{2/7}}{2}})\right) \\ &= \exp\left(\log n+n^{1/7}-ne^{-\frac{1+n^{-\frac{5}{14}}}{\varepsilon}}\right)+O\left(n\exp(n^{1/7}-\frac{1}{2}n^{2/7})\right). \end{aligned}$$

Both of the last two terms converge to 0. This completes the proof. ■

7.3. Proof of Proposition 2. We focus here on the case in which match utilities depend non-trivially on the idiosyncratic components, i.e. when $\alpha > 0$. The case where utilities are fully aligned, $\alpha = 0$, was shown in Section 7.1 in this Appendix.

The model is a mixture of aligned preferences and idiosyncratic preferences. As such, our proof is comprised of two parts.

Fix $\varepsilon_n = \frac{2}{\log n}$.

For each market realization (c, z^f, z^w), let

$$\bar{F}(\varepsilon_n; c, z^f, z^w) \equiv \{f_i | c_{i\mu^w(i)} \leq 1 - \varepsilon_n\}.$$

Whenever $\alpha < 1$, so that utilities depend non-trivially on the common component, we first show that

$$\eta_n \equiv \mathbb{E}\left[\frac{|\bar{F}(\varepsilon_n; c, z^f, z^w)|}{n}\right] = o(n^{-1/2}) \quad \text{as} \quad n \to \infty. \tag{5}$$

In the second part of the proof, we show that

$$\zeta_n \equiv P\left(\frac{\sum_{i=1}^n z^f_{i\mu^w(i)}}{n} \leq 1 - \varepsilon_n\right) = o(e^{-n^{1/4}}) \quad \text{as} \quad n \to \infty. \tag{6}$$

Proposition 2 is immediate from (5) and (6). Note that

$$\begin{aligned}
1-\frac{S_n^f}{n} &= \mathbb{E}\left[\frac{\sum_{i=1}^n (1-u_{i\mu^w(i)})}{n}\right] \\
&= (1-\alpha)\mathbb{E}\left[\frac{\sum_{i=1}^n (1-c_{i\mu^w(i)})}{n}\right] + \alpha\mathbb{E}\left[\frac{\sum_{i=1}^n (1-z^f_{i\mu^w(i)})}{n}\right] \\
&\leq (1-\alpha)(\eta_n + \varepsilon_n(1-\eta_n)) + \alpha(\zeta_n + \varepsilon_n(1-\zeta_n)) \\
&\leq \varepsilon_n + (1-\alpha)\eta_n + \alpha\zeta_n .
\end{aligned}$$

Therefore,

$$\limsup_{n\to\infty}\left(1-\frac{S_n^f}{n}\right)\log n \leq \limsup_{n\to\infty} \varepsilon_n \log n = 2.$$

Proof of Equation (5)**.** The proof of Equation (5) is useful when the common component enters match utilities non-trivially, i.e. when $\alpha < 1$. In this case, we use a result in Dawande et al., (2001, page 396). Consider a random bipartite graph $G(V_1 \cup V_2, p)$, where $0 < p < 1$ is a constant, $|V_1| = |V_2| = n$, and $\delta_n = \log n / \log \frac{1}{p}$. Let Z_b be the number of bicliques of size $b \times b$. The result shows that

$$Pr(Z_b \geq 1) \leq \frac{1}{(b!)^2}, \quad \text{for every } n > 1.$$

If a maximal balanced biclique of this graph has size $B \times B$, then

$$Pr\left(B \geq \delta_n\right) \leq \frac{1}{(\lceil \delta_n \rceil !)^2}, \quad \text{for every } n > 1. \tag{7}$$

For each realization (c, z^f, z^w), we draw a bipartite graph such that $F \cup W$ is the set of nodes (where F and W constitute the two parts of the graph), and each pair of f_i and w_j is connected by an edge if and only if at least one of c_{ij}, z^f_{ij}, or z^w_{ij} is lower than or equal to $1-(1-\alpha)\varepsilon_n$.

Let

$$\bar{W}(\varepsilon_n; c, z^f, z^w) \equiv \{w_j | \mu^w(j) \in \bar{F}(\varepsilon_n; c, z^f, z^w)\}.$$

Then $\bar{F} \cup \bar{W}$ is a balanced biclique. If a pair (f_i, w_j) from $\bar{F} \cup \bar{W}$ is not connected by an edge, then $c_{ij}, z_{ij}^f, z_{ij}^w > 1-(1-\alpha)\varepsilon_n$ implies

$$\begin{aligned} \phi(c_{ij}, z_{ij}^f) &> 1-(1-\alpha)\varepsilon_n = \phi(1-\varepsilon_n, 1), \quad \text{and} \\ \omega(c_{ij}, z_{ij}^w) &> 1-(1-\alpha)\varepsilon_n = \omega(1-\varepsilon_n, 1). \end{aligned}$$

The pair can achieve utilities higher than their utilities under μ^w. This contradicts μ^w being stable.

The approximate upper bound of the sizes of bicliques (7) implies

$$Pr\left(|\bar{F}(\varepsilon_n; c, z^f, z^w)| \geq \delta_n\right) \leq \frac{1}{(\lceil \delta_n \rceil !)^2}, \quad \text{for every } n > 1,$$

where

$$\delta_n = \frac{\log n}{-\log p_n} \quad \text{and} \quad p_n = 1-(1-\alpha)^3\varepsilon_n^3.$$

Thus,

$$\mathbb{E}\left[\frac{|\bar{F}(\varepsilon_n; c, z^f, z^w)|}{n}\right] \leq \frac{\delta_n}{n} + \frac{1}{(\lceil \delta_n \rceil !)^2} \quad \text{for every } n > 1. \tag{8}$$

To prove Equation (5), we need to show $\frac{\delta_n}{n} = o(n^{-1/2})$ and $\frac{1}{(\lceil \delta_n \rceil !)^2} = o(n^{-1/2})$.

First, as $-\log p_n \geq 1 - p_n$,

$$\frac{\delta_n}{n} = \frac{\log n}{(-\log p_n)n} \leq \frac{\log n}{(1-p_n)n} = \frac{\log n}{(1-\alpha)^3\varepsilon_n^3 n} = \frac{(\log n)^4}{8(1-\alpha)^3 n} = o(n^{-1/2}).$$

Second, from Stirling's formula,[26]

$$\frac{1}{(\lceil \delta_n \rceil !)^2} \leq \frac{1}{\left(\sqrt{2\pi\delta_n}\,(\delta_n/e)^{\delta_n}\right)^2} = \frac{e^{\delta_n}}{2\pi\delta_n^{\delta_n+1}} \quad \text{for every } n > 1.$$

For every sufficiently large n, since $p_n \to 1$,

$$\delta_n = \frac{\log n}{-\log p_n} > \log n,$$

[26]For every $n \geq 1$,

$$n! = \sqrt{2\pi n}\left(\frac{n}{e}\right)^n e^{r_n} \quad \text{with} \quad \frac{1}{12n+1} \leq r_n \leq \frac{1}{12n}.$$

and, since $\delta_n \to \infty$,

$$\log\left(\frac{e^{\delta_n}}{2\pi\delta_n^{\delta_n+1}}\right) = -\log(2\pi) - \log\delta_n - \delta_n(\log\delta_n - 1) \le -\delta_n.$$

Therefore,

$$\frac{1}{(\lceil\delta_n\rceil!)^2} \le \frac{1}{e^{\delta_n}} \le \frac{1}{e^{\log n}} = \frac{1}{n} \quad \text{for any sufficiently large } n.$$

Proof of Equation (6)**.** Let $\mu \equiv \{(i,i)|i = 1,\ldots,n\}$. By symmetry, each one of the $n!$ matchings has the same probability of being both stable and entailing a sum of firms' idiosyncratic components that is lower than or equal to $(1-\varepsilon_n)n$. Therefore,

$$P\left(\frac{\sum_{i=1}^n z^f_{i\mu^w(i)}}{n} \le 1-\varepsilon_n\right) \le n!P\left(\mu \text{ is stable and } \frac{\sum_{i=1}^n z^f_{i\mu(i)}}{n} \le 1-\varepsilon_n\right).$$

For each realization (c, z^f, z^w), we consider the following profile of utilities:

$$\tilde{u}^f_{ij} = (1-\alpha)c_{ij} + \alpha z^f_{ij} \quad \text{and} \quad \tilde{u}^w_{ij} = (1-\alpha)c_{ij} + \alpha z^w_{ij} \quad \text{if } i \ne j, \quad \text{and}$$
$$\tilde{u}^f_{ij} = (1-\alpha) + \alpha z^f_{ij} \quad \text{and} \quad \tilde{u}^w_{ij} = (1-\alpha)c_{ij} + \alpha z^w_{ij} \quad \text{if } i = j.$$

Note that $\tilde{u}^f_{ij} > u^f_{ij}$, generically for all $i = j$.

If a pair (f_i, w_j) with $i \ne j$ is a blocking pair of μ under match utilities $(\tilde{u}^f_{ij}, \tilde{u}^w_{ij})$, then this pair also blocks μ under the actual realized utilities. Thus,

$$P\left(\mu \text{ is stable and } \frac{\sum_{i=1}^n z^f_{i\mu(i)}}{n} \le 1-\varepsilon_n\right) \le P_{\varepsilon_n,n}$$

where $P_{\varepsilon_n,n}$ is the probability that μ is stable with respect to the utilities $(\tilde{u}^f_{ij}, \tilde{u}^w_{ij})_{i,j}$ and $\frac{\sum_{i=1}^n z^f_{i\mu(i)}}{n} \le 1-\varepsilon_n$. We now show that $n!P_{\varepsilon_n,n}$ converges to zero as n increases and identify the speed at which it converges, which yields Equation (6).

Preparation Steps for Proof of Equation (6)**.** We denote by Γ^f the marginal distribution of $\tilde{u}^f_{ij}$ for pairs of (f_i, w_j) who are not matched under μ (i.e., $i \ne j$), and by Γ^w the marginal distribution of $\tilde{u}^w_{ij}$ for any pair of (f_i, w_j).

We define $\hat{u}^f_{ij} \equiv \Gamma^f(\tilde{u}^f_{ij})$ and $\hat{u}^w_{ij} \equiv \Gamma^w(\tilde{u}^w_{ij})$. For each given realization $(\hat{u}^f_{ii})_{i=1}^n$ and $(\hat{u}^w_{jj})_{j=1}^n$, the probability that μ is stable with respect to $(\tilde{u}^f_{ij}, \tilde{u}^w_{ij})_{i,j}$ is

$$\Pi_{1\leq i\neq j\leq n}\left(1 - P[\hat{u}^f_{ij} > \hat{u}^f_{ii} \text{ and } \hat{u}^w_{ij} > \hat{u}^w_{jj}]\right).$$

Note that c_{ij}, z^f_{ij}, and z^w_{ij} are independently and identically distributed, so they are positively associated. Then, for $i \neq j$, the covariance of $\hat{u}^f_{ij}$ and $\hat{u}^w_{ij}$ is non-negative because both $\Gamma^f(\phi(.,.))$ and $\Gamma^w(\omega(.,.))$ are non-decreasing functions of c_{ij}, z^f_{ij}, and z^w_{ij}.

Thus, we have

$$P[\hat{u}^f_{ij} > \hat{u}^f_{ii} \text{ and } \hat{u}^w_{ij} > \hat{u}^w_{jj}] \geq P[\hat{u}^f_{ij} > \hat{u}^f_{ii}]P[\hat{u}^w_{ij} > \hat{u}^w_{jj}] = (1-\hat{u}^f_{ii})(1-\hat{u}^w_{jj}).$$

Since

$$1 - \hat{u}^f_{ii} = 1 - \phi(1, z^f_{ii}) = \alpha(1 - z^f_{ii}),$$

for each realization of $c_{ij}, z^f_{ij}, z^w_{ij}$ for pairs with $i = j$, the probability that μ is stable is bounded above by

$$\Pi_{1\leq i\neq j\leq n}\left(1 - \alpha(1-z^f_{ii})(1-\hat{u}^w_{jj})\right).$$

We therefore obtain that

$$P_{\varepsilon_n,n} \leq \int\int_{\sum_{i=1}^n z^f_{ii}\leq(1-\varepsilon_n)n} \Pi_{1\leq i\neq j\leq n}\left(1-\alpha(1-z^f_{ii})(1-\hat{u}^w_{jj})\right) d\mathbf{z^f_{ii}}d\mathbf{\hat{u}^w_{jj}}.$$

Now, let $x_i = 1 - z^f_{ii}$ and $y_j = 1 - \hat{u}^w_{jj}$. Then,

$$P_{\varepsilon_n,n} \leq \int_{\substack{\mathbf{0}\leq\mathbf{x},\mathbf{y}\leq 1\\ \varepsilon_n n\leq\sum_{i=1}^n x_i}} \Pi_{1\leq i\neq j\leq n}\left(1-\alpha x_i y_j\right) d(\mathbf{x},\mathbf{y}).$$

Proof of Convergence.

$$\begin{aligned} P_{\varepsilon_n,n} &\leq \int_{\substack{\mathbf{0}\leq\mathbf{x},\mathbf{y}\leq 1\\ \varepsilon_n n\leq\sum_{i=1}^n x_i}} \Pi_{1\leq i\neq j\leq n}(1-\alpha x_i y_j)d(\mathbf{x},\mathbf{y}) \\ &= \int_{\substack{\mathbf{0}\leq\mathbf{x},\mathbf{y}\leq 1\\ \varepsilon_n n\leq\sum_{i=1}^n x_i}} \Pi_{1\leq j\leq n}\underbrace{\left(\int_0^1 \Pi_{i\neq j}(1-\alpha x_i y_j)dy_j\right)}_{(*)} d\mathbf{x}. \end{aligned}$$

We can assume, without loss of generality, that $0 < \alpha < 1/2$ as $(*)$ above decreases in α. Let $t = n^{-7/8}$ and $\delta = e^t(1-\alpha t)$. Note that $1 - \alpha z \le \delta e^{-z}$ if $0 \le z \le t$. Thus, when $0 \le y_j \le t$, we have $0 \le x_i y_j \le t$, so $1 - \alpha x_i y_j \le \delta \exp(-x_i y_j)$. In addition, $1 + \alpha z \le e^{\alpha z}$ for any z, so $1 - \alpha x_i y_j \le \exp(-\alpha x_i y_j)$.

Therefore,

$$\begin{aligned}(*) &= \int_0^t \Pi_{i\neq j}(1-\alpha x_i y_j)dy_j + \int_t^1 \Pi_{i\neq j}(1-\alpha x_i y_j)dy_j \\ &= \int_0^t \Pi_{i\neq j}\delta \exp(-x_i y_j)dy_j + \int_t^1 \Pi_{i\neq j}\exp(-\alpha x_i y_j)dy_j \\ &= \delta \int_0^t \exp\left(-y_j \sum_{i\neq j} x_i\right) dy_j + \int_t^1 \exp\left(-\alpha y_j \sum_{i\neq j} x_i\right) dy_j.\end{aligned}$$

Let

$$s = \sum_{i=1}^n x_i \quad \text{and} \quad s_j = \sum_{i\neq j} x_i.$$

Then,

$$\begin{aligned}(*) &= \delta \int_0^t \exp\left(-y_j s_j\right) dy_j + \int_t^1 \exp\left(-\alpha y_j s_j\right) dy_j \\ &= \delta \frac{1-e^{-t s_j}}{s_j} + \frac{e^{-\alpha t s_j} - e^{-\alpha s_j}}{\alpha s_j} \le \frac{1}{s_j}\left(\delta + \frac{1}{\alpha}\exp(-\alpha t s_j)\right).\end{aligned}$$

We claim that

$$\delta + \frac{1}{\alpha}\exp(-\alpha t s_j) < \exp(n^{-6/7}) \tag{9}$$

for sufficiently large n.

As $s > \varepsilon_n n$, we have $s_j > \varepsilon_n n - 1$. Thus, (9) follows from

$$\exp(n^{-\frac{7}{8}}) + \frac{1}{\alpha}\exp(-\alpha n^{-7/8}(\varepsilon_n n - 1)) < \exp(n^{-6/7}) \quad \text{with sufficiently large } n.$$

Now, we have

$$\begin{aligned} P_{\varepsilon_n,n} &\leq \int_{\substack{\mathbf{0}\leq\mathbf{x}\leq 1\\ \varepsilon_n n\leq\sum_{i=1}^n x_i}} \Pi_{1\leq j\leq n}(*)d\mathbf{x} \\ &\leq \int_{\substack{\mathbf{0}\leq\mathbf{x}\leq 1\\ \varepsilon_n n\leq\sum_{i=1}^n x_i}} \Pi_{1\leq j\leq n}\left(\frac{1}{s_j}\exp(n^{-6/7})\right)d\mathbf{x} = \exp(n^{1/7})\int_{\substack{\mathbf{0}\leq\mathbf{x}\leq 1\\ \varepsilon_n n\leq\sum_{i=1}^n x_i}} \Pi_{1\leq j\leq n}\left(\frac{1}{s_j}\right)d\mathbf{x}. \end{aligned}$$

Note that $\left(\log\frac{1}{s_j}\right)' = -\frac{1}{s_j}$. Thus,

$$\sum_{j=1}^n \log\frac{1}{s_j} = \sum_{j=1}^n\left(\log\frac{1}{s}+\log\frac{s}{s_j}\right) = n\log\frac{1}{s}+\sum_{j=1}^n\log\frac{s}{s-x_j}.$$

In the last term,

$$\frac{s}{s-x_j}\leq\frac{s}{s-1}\leq\frac{\varepsilon_n n}{\varepsilon_n n-1}\quad\text{for any sufficiently large } n.$$

Thus,

$$\sum_{j=1}^n\log\frac{1}{s_j}\leq n\ln\frac{1}{s}+n\log\frac{\varepsilon_n n}{\varepsilon_n n-1}.$$

Moreover

$$\varepsilon_n n\log\frac{\varepsilon_n n}{\varepsilon_n n-1} = \log\left(\left(1+\frac{1}{\varepsilon_n n-1}\right)^{\varepsilon_n n}\right)\to 1,\quad\text{as}\quad n\to\infty,$$

which implies that

$$\sum_{j=1}^n\log\frac{1}{s_j}\leq n\log\frac{1}{s}+\frac{2}{\varepsilon_n} = n\log\frac{1}{s}+\log n.$$

Therefore,

$$n!P_{\varepsilon_n,n}\leq n!\exp(n^{1/7})\int_{\varepsilon_n n\leq s}\exp\left(n\log\frac{1}{s}+\log n\right)f_n(s)ds,$$

where $f_n(s)$ is the probability distribution function of s.

From Lemma 1,

$$\begin{aligned} n!P_{\varepsilon_n,n} &\leq n!\exp(n^{1/7})\int_{\varepsilon_n n\leq s}\exp\left(n\log\frac{1}{s}+\log n\right)\frac{s^{n-1}}{(n-1)!}\,Pr(r_n\leq s^{-1})ds \\ &\leq n^2\exp(n^{1/7})Pr\left(r_n\leq\frac{1}{\varepsilon_n n}\right)\int_{\varepsilon_n n}^{n}\frac{1}{s}\,ds \\ &= n^2\exp(n^{1/7})Pr\left(r_n\leq\frac{1}{\varepsilon_n n}\right)(-\log\varepsilon_n) \\ &\leq n^2\exp(n^{1/7})Pr\left(r_n\leq\frac{1}{\varepsilon_n n}\right)(\log(\log n)). \end{aligned}$$

Now,

$$\begin{aligned} & n^2\exp(n^{1/7})Pr\left(r_n\leq\frac{1}{\varepsilon_n n}\right)(\log(\log n)) \\ \leq\; & \exp\left(2\log n+n^{\frac{1}{7}}-ne^{-\frac{1+n^{-\frac{5}{14}}}{\varepsilon_n}}\right)\log(\log n)+o\left(e^{-n^{1/4}}\right) \\ \leq\; & \exp\left(2\log n+n^{\frac{1}{7}}-n^{\frac{1}{2}-\frac{n^{-\frac{5}{14}}}{2}}\right)\log(\log n)+o\left(e^{-n^{1/4}}\right)=o\left(\exp(-n^{1/4})\right). \end{aligned}$$

■

7.4. Aggregate Efficiency – Proof of Proposition 3. 1. As mentioned in the text, Walkup (1979) implies that $E_n \geq 2n-6$. Since, by definition, $E_n \leq 2n$, combining these observations with the bounds on S_n provided by Proposition 2, the claim follows.

2. We provide a bound on E_n for this environment. For each pair of firm f_i and worker w_j, suppose u_{ij}^f and u_{ij}^w are distributed uniformly on $[0,1]$. We define $\tilde{u}_{ij}\equiv\frac{u_{ij}^f+u_{ij}^w}{2}$, which has a triangular distribution on $[0,1]$. We show that

$$E_n = 2\cdot\max_{\mu\in M}\sum_{i=1}^{n}\tilde{u}_{i\mu(i)}\geq 2n-3\sqrt{n}\quad\text{for every}\quad n\geq 2.$$

We consider two random variables $\tilde{v}_{ij}^f$ and $\tilde{v}_{ij}^w$ with cumulative distribution functions

$$H(x)=\begin{cases} 0 & \text{for}\quad 0\leq x<1-1/\sqrt{2} \\ \sqrt{1-2(1-x)^2} & \text{for}\quad 1-1/\sqrt{2}\leq x\leq 1. \end{cases}$$

Notice that

$$P\left(\max\{\tilde{v}^f_{ij}, \tilde{v}^w_{ij}\} \leq x\right) = \begin{cases} 0 & \text{for} \quad 0 \leq x < 1 - 1/\sqrt{2} \\ 1 - 2(1-x)^2 & \text{for} \quad 1 - 1/\sqrt{2} \leq x \leq 1 \end{cases}, \text{and}$$

$$P\left(\tilde{u}_{ij} \leq x\right) = \begin{cases} 2x^2 & \text{for} \quad 0 \leq x < 1/2 \\ 1 - 2(2-x)^2 & \text{for} \quad 1/2 \leq x \leq 1 \end{cases}.$$

Therefore,

$$P\left(\tilde{u}_{ij} \leq x\right) \leq P(\max\{\tilde{v}^f_{ij}, \tilde{v}^w_{ij}\} \leq x) \quad \text{for} \quad 0 \leq x \leq 1.$$

That is, $\tilde{u}_{ij}$ first order stochastically dominates $\max\{\tilde{v}^f_{ij}, \tilde{v}^w_{ij}\}$.

We denote by $\tilde{v}^f_{i(k)}$ the k'th highest value of $(\tilde{v}^f_{ij})_{j=1}^n$. As $H(\cdot)$ is a concave function on the support of the distribution, Jensen's inequality implies that, for any $k = 1, \ldots, n$,

$$H\left(\mathbb{E}[\tilde{v}^f_{i(k)}]\right) \geq \mathbb{E}\left[H(\tilde{v}^f_{i(k)})\right].$$

In addition, $H(\tilde{v}^f_{i(k)})$ is equal to the k-th highest value of $\{H(\tilde{v}^f_{ij})\}_{j=1}^n$, and $H(\tilde{v}^f_{ij})$ is distributed uniformly on $[0,1]$. Thus,

$$H\left(\mathbb{E}[\tilde{v}^f_{i(k)}]\right) \geq \mathbb{E}\left[H(\tilde{v}^f_{i(k)})\right] = \frac{n+1-k}{n+1}.$$

Therefore,

$$\mathbb{E}\left[\tilde{v}^f_{i(k)}\right] \geq H^{-1}\left(\frac{n+1-k}{n+1}\right).$$

Identical calculations hold for $\{\tilde{v}^w_{ij}\}_{i=1}^n$ and the corresponding value $\tilde{v}^w_{j(k)}$.

Consider now a random directed bipartite graph with F and W serving as our two classes of nodes, denoted by G. Each firm f_i has arcs to two workers with the highest realized values of $\tilde{v}^f_{ij}$. Similarly, each worker w_j has arcs to two firms generating the highest realized values of $\tilde{v}^w_{ij}$.

Let $\mathcal{B}$ denote the set of all directed bipartite graphs containing at least one perfect matching. Let α_G denote the maximum aggregate efficiency achievable by matchings in G. We have

$$E_n \geq \mathbb{E}\left[\alpha_G | G \in \mathcal{B}\right] \cdot P\left(G \in \mathcal{B}\right).$$

Each pair of firm f_i and worker w_j matched in the efficient matching in G has utility $\tilde{u}_{ij}$ which is no less than either $\tilde{v}^f_{i(2)}$ or $\tilde{v}^w_{j(2)}$. Both have expected values no less than $H^{-1}\left(\frac{n+1-k}{n+1}\right)$, which is equal to $1-\frac{\sqrt{2n}}{n+1}$.

Walkup (1979) illustrates that

$$P(G \in \mathcal{B}) \geq 1 - \frac{1}{5n}.$$

Therefore, we have

$$\begin{aligned} E_n &\geq 2n \cdot \left(1 - \frac{\sqrt{2n}}{n+1}\right) \cdot \left(1 - \frac{1}{5n}\right) \\ &\geq 2n - 3\sqrt{n}. \end{aligned}$$

In the Online Appendix we show that for idiosyncratic preferences, a slightly tighter characterization of the speed of convergence holds. Namely,

$$\lim_{n\to\infty}\left(1-\frac{S^f_n}{n}\right)\log n = \lim_{n\to\infty}\left(1-\frac{S^w_n}{n}\right)\log n = 1.$$

This, together with the last inequality, completes the claim. ■

7.5. Convergence Speeds with Assortative Preferences – Proof of Proposition 5. We provide a proof for the case of deterministic common values and assume that

$$(c^f_1, c^f_2, \ldots, c^f_n) = (c^w_1, c^w_2, \ldots, c^w_n) = \left(\frac{n-1}{n}, \frac{n-2}{n}, \ldots, \frac{1}{n}, 0\right).$$

The assumption of deterministic common values is without loss of generality since the distribution of deterministic common values and the empirical distribution of common values from the uniform distribution converge to one another at an exponential rate (see Fact 4 in the Online Appendix of Lee, 2017).

Let $\varepsilon_n = 6(1-\beta)n^{-1/4}$ and define

$$B_F(\varepsilon_n; z^f, z^w) \equiv \{f_i \in F \mid u^f_{i\mu_W(i)} \leq (1-\beta)c^f_i + \beta - (3/5)\varepsilon_n\}.$$

We will use some results in the Online Appendix of Lee (2017) and show that

$$P\left(\frac{|B_F(\varepsilon_n; z^f, z^w)|}{n} > \theta_n\right) \leq \delta_n, \tag{10}$$

with some sequences $\theta_n = O(n^{-1/4})$ and $\delta_n = o(e^{-n^{1/2}})$. Thus,

$$(1-\beta)\frac{1}{2} + \beta - \frac{S_n^f}{n} \leq \delta_n + (1-\delta_n)(\theta_n + (1-\theta_n)(3/5)\varepsilon_n) = O(n^{-1/4}).$$

A similar argument holds for workers.

Proof of (10)**.** Let F_n denote the set of firms in markets with n participants on each side. We partition the set F_n into $K_n = \lceil n^{1/4} \rceil$ "tiers".[27] For each tier $k = 1, ..., K_n$, the firms in tier k are given by $F_{k;n}$, where

$$\begin{aligned} F_{k;n} &\equiv \left\{f_i \in F_n \mid (k-1)n^{3/4} < i \leq kn^{3/4}\right\} \\ &= \left\{f_i \in F_n \mid 1 - kn^{-1/4} < c_i^f < 1 - (k-1)n^{-1/4}\right\} \quad (\text{since } c_i^f = 1 - \frac{i}{n}). \end{aligned}$$

For any $k = 1, \ldots, K_n$, define

$$B_{F_{k;n}}(\varepsilon_n; z^f, z^w) \equiv \{f \in F_{k,n} \mid u^f_{i\mu_W(i)} \leq (1-\beta)(1-kn^{-1/4}) + \beta - (2/5)\varepsilon_n\}.^{28}$$

If $f_i \in B_F(\varepsilon_n; z^f, z^w)$ and $f_i \in F_{k,n}$, then

$$\begin{aligned} u^f_{i\mu_W(i)} &\leq (1-\beta)c_i^f + \beta - (3/5)\varepsilon_n \\ &< (1-\beta)(1-(k-1)n^{-1/4}) + \beta - (3/5)\varepsilon_n \\ &< (1-\beta)(1-kn^{-1/4}) + \beta - (2/5)\varepsilon_n. \end{aligned}$$

Therefore,

$$B_F(\varepsilon_n; z^f, z^w) \subset \bigcup_{k=1}^{K_n} B_{F_{k;n}}(\varepsilon_n; z^f, z^w).$$

[27] For $x \in \mathbb{R}$, $\lceil x \rceil$ is the smallest integer that is not smaller than x.

[28] There is a typo in the definition of $B_{F_{K;n}}$ on page 13 of the Online Appendix of Lee (2017): $n^{-1/2}$ should be replaced with $n^{-1/4}$.

Note that

$$\sum_{k=K_n-2}^{K_n} |F_{k;n}| \leq 3n^{3/4}.$$

With arguments similar to those in the Online Appendix of Lee (2017), we can show that

$$P\left(\sum_{k=1}^{K_n-3} |B_{F_{k;n}}| > (K_n - 3)\phi_n\right) \leq 1 - (1 - \psi_n)^{K_n - 3},$$

with some sequences $\phi_n = O((\log n)n^{1/2})$ and $\psi_n = o(e^{-\phi_n})$. By taking into account $K_n \geq n^{1/4}$, we obtain

$$P\left(\sum_{k=1}^{K_n} |B_{F_{k;n}}| > \frac{1}{n}(n^{1/4} - 3)\phi_n + 3n^{-1/4}\right) \leq 1 - (1 - \psi_n)^{n^{1/4} - 3}.$$

We can then show that

$$\begin{aligned}\theta_n &\equiv \frac{1}{n}(n^{1/4} - 3)\phi_n + 3n^{-1/4} = O(n^{-1/4}), \quad \text{and} \\ \delta_n &\equiv 1 - (1 - \psi_n)^{n^{1/4} - 3} = o(e^{-n^{1/2}}).\end{aligned}$$

■

7.6. Severely Imbalanced Markets – Proof of Proposition 6. For any matching μ, let $R_j^w(\mu)$ denote the rank of worker w_j's partner: $R_j^w(\mu) = 1$ if worker w_j is matched with the most preferred firm, $R_j^w(\mu) = 2$ if worker w_j is matched with the second most preferred firm, etc.

We use Theorem 1 in Ashlagi, Kanoria, and Leshno (2017), which implies that for $0 < \lambda \leq 1/2$,

$$P_n \equiv P\left(\sum_{i=1}^{n} R^w_{\mu^f(i)}(\mu^f) \geq \frac{n}{-3 \log \lambda}\right) \to 1 \quad \text{as } n \to \infty.$$

Thus, for $0 < \lambda \leq 1/2$,

$$
\begin{aligned}
\frac{S_n^w}{n} &= \frac{\mathbb{E}\left[\sum_{i=1}^n u^w_{i\mu^f(i)}\right]}{n} \\
&= \frac{\mathbb{E}_{\succ}\left[\mathbb{E}_{u|\succ}\left[\sum_{i=1}^n u^w_{i\mu^f(i)} | \sum_{i=1}^n R^w_{\mu^f(i)}(\mu^f)\right]\right]}{n} \\
&= \frac{\mathbb{E}_{\succ}\left[\sum_{i=1}^n 1 - \frac{R^w_{\mu^f(i)}(\mu^f)}{n+1}\right]}{n} = 1 - \frac{\mathbb{E}\left[\sum_{i=1}^n R^w_{\mu^f(i)}(\mu^f)\right]}{n+1} \\
&\leq 1 - \frac{nP_n}{-3(n+1)\log\lambda} \to 1 - \frac{1}{-3\log\lambda} \quad \text{as} \quad n \to \infty.
\end{aligned}
$$

A market with $\lambda > 1/2$ is a market with more workers than those available in a market with $\lambda = 1/2$. Crawford (1991) shows that every worker becomes weakly worse off in μ^f as more workers enter the market, which concludes our proof. ■

References

[1] Abdulkadiroglu, Atila, Yeon-Koo Che, and Yosuke Yasuda. 2011. "Resolving Conflicting Preferences in School Choice: The "Boston Mechanism" Reconsidered," *The American Economic Review*, 101(1), 399-410.

[2] Abdulkadiroglu, Atila, Parag Pathak, and Alvin E. Roth. 2009. "Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match," *The American Economic Review*, 99(5), 1954-1978.

[3] Agarwal, Nikhil. 2015. "An Empirical Model of the Medical Match," *The American Economic Review*, 105(7), 1939-1978.

[4] Avery, Christopher and Parag Pathak. 2017. "The Distributional Consequences of Public School Choice," mimeo.

[5] Anshelevich, Elliot, Sanmay Das, and Yonatan Naamad. 2013. "Anarchy, Stability, and Utopia: Creating Better Matchings," *Autonomous Agents and Multi-Agent Systems*, 26(1), 120-140.

[6] Ashlagi, Itai, Yashodhan Kanoria, and Jacob D. Leshno. 2017. "Unbalanced Random Matching Markets: The Stark Effect of Competition," *Journal of Political Economy*, 125(1), 69-98.

[7] Baccara, Mariagiovanna, Allan Collard-Wexler, Leonardo Felli, and Leeat Yariv. 2014. "Child-adoption Matching: Preferences for Gender and Race," *American Economic Journal: Applied Economics*, 6(3), 133-158.

[8] Becker, Gary S. 1973. "A Theory of Marriage: Part 1," *The Journal of Political Economy*, 81(4), 813-846.

[9] Becker, Gary S. 1974. "A Theory of Marriage: Part II," *The Journal of Political Economy*, 82(2), S11-S26.

[10] Bordon, Paola and Chao Fu. 2015. "College-Major Choice to College-Then-Major Choice," *Review of Economic Studies,* 82(4), 1247-1288.

[11] Boudreau, James W. and Vicki Knoblauch. 2013. "Preferences and the price of stability in matching markets," *Theory and Decision*, 74(4), 565-589.

[12] Budish, Eric and Estelle Cantillon. 2012. "The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard," *The American Economic Review*, 102(5), 2237-2271.

[13] Bulow, Jeremy and Jonathan Levin. 2006. "Matching and Price Competition," *The American Economic Review*, 96(3), 652-668.

[14] Che, Yeon-Koo and Olivier Tercieux. 2018. "Efficiency and Stability in Large Matching Markets," *Journal of Political Economy*, forthcoming.

[15] Cole, Harold, George Mailath, and Andrew Postlewaite. 2001. "Efficient Non-contractible Investments in Finite Economies," *Advances in Theoretical Economics*, 1(1), Article 2.

[16] Compte, Olivier and Philippe Jehiel. 2008. "Voluntary Participation and Re-Assignment in Two-Sided Matching," mimeo.

[17] Coppersmith, Don and Gregory B. Sorkin. 1999. "Constructive Bounds and Exact Expectations for the Random Assignment Problem," Random Structures and Algorithms, 15(2), 113-144.

[18] Crall, Heather S. 2004. "Unreasonable Restraints: Antitrust Law and the National Residency Matching Program," *Washington University Law Review*, 82(1), 245-274.

[19] Crawford, Vincent P. 1991. "Comparative statics in matching markets," *Journal of Economic Theory*, 54(1), 389-400.

[20] Crawford, Vincent P. 2008. "The Flexible-Salary Match: A proposal to increase the salary flexibility of the National Resident Matching Program," *The Journal of Economic Behavior and Organizations*, 66, 149-160.

[21] Dagsvik, John K. 2000. "Aggregation in Matching Markets," *International Economic Review*, 41(1), 27-58.

[22] Dawande, Milind, Pinar Keskinocak, Jayashankar M. Swaminathan, and Sridhar Tayur. 2001. "On bipartite and multipartite clique problems," *Journal of Algorithms*, 41(2), 388-403.

[23] Dizdar, Deniz, and Benny Moldovanu. 2016. "On the Importance of Uniform Sharing Rules for Efficient Matching," *Journal of Economic Theory*, 165, 106-123.

[24] Durlauf, Steven N., and Ananth Seshadri. 2003. "Is assortative matching efficient?," *Economic Theory*, 21, 475-493.

[25] Esary, James D., Frank Proschan, and David W. Walkup. 1967. "Association of random variables, with applications," *The Annals of Mathematical Statistics*, 38(5), 1466-1474.

[26] Fernandez, Marcelo A. and Leeat Yariv. 2018. "Centralized Matching with Incomplete Information," mimeo.

[27] Goemans, Michel X. and Muralidharan S. Kodialam. 1993. "A Lower Bound on the Expected Cost of an Optimal Assignment," *Mathematics of Operations Research*, 18(2), 267-274.

[28] Hall, Robert E. and Alan B. Krueger. 2012. "Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-the-Job Search," *The American Economic Journal: Macroeconomics*, 4(4), 56-67.

[29] Hitch, Günter J., Ali Hortaçsu, and Dan Ariely. 2010. "Matching and Sorting in Online Dating," *The American Economic Review*, 100(1), 130-163.

[30] Immorlica, Nicole and Mohammad Mahdian. 2005. "Marriage, Honesty, and Stability," *SODA Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 53-62.

[31] Kojima, Fuhito, and Parag A. Pathak. 2009. "Incentives and Stability in Large Two-Sided Matching Markets," *The American Economic Review*, 99(3), 608-627.

[32] Knoblauch, Vicki. 2007. "Marriage Matching: A Conjecture of Donald Knuth,"mimeo.

[33] Knuth, Donald E. 1976. *Mariages stables et leurs relations avec d'autres problmes combinatoires: introduction l'analyse mathamatique des algorithmes*,Presses de l'Universite de Montreal.

[34] Krishna, Vijay and John Morgan. 2015. "Majority Rule and Utilitarian Welfare," *American Economic Journal: Microeconomics*, 7(4), 339-375.

[35] Lauermann, Stephan. 2013. "Dynamic Matching and Bargaining Games: A General Approach," *The American Economic Review*, 103(2), 663-689.

[36] Lazarus, Andrew J. 1993. "Certain Expected Values in the Random Assignment Problem," *Operations Research Letters*, 14, 207-214.

[37] Lee, SangMok. 2017. "Incentive Compatibility of Large Centralized Matching Markets," *Review of Economic Studies,* 84(1), 444-463.

[38] Liu, Qingmin, and Marek Pycia. 2016. "Ordinal Efficiency, Fairness, and Incentives in Large Markets," mimeo.

[39] Menzel, Konrad. 2015. "Large Matching Markets as Two-sided Demand Systems," *Econometrica*, 83(3), 897-941.

[40] Milgrom, Paul. 2004. Putting Auction Theory to Work, Churchill Lectures in Economics, Cambridge University Press.

[41] Neiderle, Muriel and Leeat Yariv. 2009. "Decentralized Matching with Aligned Preferences," mimeo.

[42] Peski, Marcin. 2017. "Large Rommate Problem with Non-transferable Random Utility," *Journal of Economic Theory*, 168, 432-471.

[43] Pittel, Boris. 1989. The average number of stable matchings," *SIAM Journal on Discrete Mathematics*, 2(4), 530-549.

[44] Pittel, Boris. 1992. "On likely solutions of a stable marriage problem," *The Annals of Applied Probability*, 2(2), 358-401.

[45] Roth, Alvin E. 2002. "The economist as engineer: Game theory, experimentation, and computation as tools for design economics," *Econometrica*, 70(4), 1341-1378.

[46] Roth, Alvin E. 2007. "Repugnance as a Constraint on Markets," *Journal of Economic Perspectives*, 21(3), 37-58.

[47] Roth, Alvin E., and Marilda A. Oliveira Sotomayor. 1992. *Two-sided Matching: A Study in Game-theoretic Modeling and Analysis*, Cambridge University Press.

[48] Roth, Alvin E., and Xiaolin Xing. 1994. "Jumping the gun: imperfections and institutions related to the timing of market transactions," *The American Economic Review*, 84(4), 992-1044.

[49] Roughgarden, Tim and Eva Tardos (2007), "Introduction to the inefficiency of equilibria," In N. Nisan, T. Roughgarden, E. Tardos & V. V. Vazirani (Eds.), *Algorithmic Game Theory*, 443-459, Cambridge University Press.

[50] Sorensen, Morten. 2007. "How Smart is Smart Money? A Two-sided Matching Model of Venture Capital," *Journal of Finance*, 62(6), 2725-2762.

[51] Walkup, David W. 1979. "On the Expected Value of a Random Assignment Problem," *SIAM Journal on Computing*, 8(3), 440-442.