On the Efficiency of Stable Matchings in Large Markets

SangMok Lee* Leeat Yariv†‡

January 23, 2014

Abstract. Stability is often the goal for clearinghouses in matching markets, such as those matching residents to hospitals, students to schools, etc. Stable outcomes need not be utilitarian efficient, generating a tradeoff that market designers face. We study the wedge between stability and efficiency in large one-to-one matching markets. We show stable matchings are efficient asymptotically for a large class of preferences. In these environments, stability remains an appealing objective even on efficiency grounds, and monetary transfers are not necessary to induce efficient outcomes. When markets are severely imbalanced and preferences entail sufficient idiosyncrasies, stable outcomes may be inefficient even asymptotically.

Keywords: Matching, Stability, Efficiency, Market Design.

*Department of Economics, University of Pennsylvania, http://www.sas.upenn.edu/~sangmok/

†Division of the Humanities and Social Sciences, Caltech, http://www.hss.caltech.edu/~lyariv

‡We thank Federico Echenique, Aytek Erdil, and Rakesh Vohra for very useful conversations 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. Most centralized matching mechanisms do not allow for transfers between participants as an institutional constraint: the National Resident Matching Program (NRMP), clearinghouses for matching schools and students in New York City and Boston, and many others utilize algorithms that ban transfers and implement a stable matching corresponding to reported rank preferences. 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). Even without transfers, stable matchings 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 ordinal nature of the mechanism underlying the matching process. These complaints culminated in an official lawsuit filed by a group of resident physicians on May of 2002. 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 effectively highlights the restricted ability of the NRMP to account for marginal (cardinal) preferences of participants over matches (see Crall, 2004).[1] 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 highlighted issues (see Bulow and Levin, 2006, Crawford, 2008, and follow-up literature).

This raises several natural questions: would the availability of transfers affect dramatically the outcomes of stable mechanisms such as the NRMP? Would banning transfers yield stable matchings that are far from efficient? The current paper addresses these questions in the context of large markets. In particular, we characterize a wide class of environments in which

[1] Details of the case can be found at http://www.gpo.gov/fdsys/pkg/USCOURTS-dcd-1_02-cv-00873

stable matchings without transfers are asymptotically efficient. In such settings the use of stable mechanisms based on ordinal reports are justified on efficiency grounds and produce outcomes that may be rather close to those produced by stable mechanisms allowing for transfers.

Our results also relate more generally to the question of proper objective functions in the design of matching markets. The literature thus far focused predominantly on mechanisms in which only ordinal preferences are specified. In many other economic settings in which market design has been utilized, such as auctions, voting, etc., 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.[2] A designer may then face a tradeoff between efficiency and stability, the availability of transfers providing one way by which to overcome it. Our analysis suggests the types of environments in which this tradeoff is not substantial.

In general, when preferences can be represented in utility terms, stable matchings need not be 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 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**

In this case, the unique stable matching matches f_i with w_i, $i = 1, 2$ and generates a utilitarian efficiency of $2 \times (5 + 1) = 12$. However, the alternative matching, between f_i and $w_{3-i}, i = 1, 2$ generates a greater utilitarian efficiency of $2 \times (3 + 4) = 14$, and would be the unique stable outcome were transfers available.

When transfers are available, stability is tantamount to utilitarian efficiency (see Roth and

[2] 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, 2013; etc.

Sotomayor, 1992). In all that follows we will refer to stable matchings without transfers as simply stable. The goal of the current paper is then to analyze the wedge between stability and utilitarian efficiency in large markets with sufficient randomness of utilities. Certainly, if we just replicated the 2×2 market above, we could easily generate an arbitrarily large market in which stable matchings lead to a significantly lower utilitarian efficiency than the first best, and transfers could prove useful. The underlying message of the paper is that for a rather general class of markets, enough randomness assures that stable matchings are asymptotically utilitarian efficient unless markets are severely imbalanced and preferences are sufficiently idiosyncratic. In that respect, designing mechanisms that exploit the array of benefits stability entails does not come at a great utilitarian cost.

In order to glean some intuition for our results, consider the case in which there are n firms and n workers, all preferring to be matched to someone in the market over remaining unmatched, and that the utility of any firm f_i and worker w_j coincides for both and is randomly and independently determined according to a uniform distribution between 0 and 1. In other words, a market is specified by the realization of n^2 uniform variables corresponding to every firm and worker pair.

Now, for $\varepsilon > 0$, construct the following bipartite graph, where firms correspond to one type of nodes, and workers correspond to the other type of nodes. Any worker and firm are linked if the utility their matching generates is greater than $1 - \varepsilon$. For instance, Figure 1 corresponds to a 4×4 market, where match utilities are described in the left panel, and the induced bipartite graph for $\varepsilon = 0.25$ is depicted on the right with bold lines. (Bold numbers in the panel corresponds to the pairs with utilities greater than $1 - \varepsilon$.)

The probability of a link between any firm f_i and worker w_j is given by ε. Therefore, an old result from graph theory due to Erdos and Renyi (1964) suggests that, as n becomes arbitrarily large, if ε is decreasing with n at a rate slower than $\frac{\log n}{n}$, so that the induced networks are sufficiently connected, they will contain a *perfect matching*, one corresponding to a one-to-one matching connecting all firms and workers, with probability converging to 1 (in Figure 1, the squiggly connections correspond to a matching that is almost perfect).

In such a perfect matching, the maximal benefit any agent can gain by connecting with an alternative participant is at most ε and so the matching is 'almost stable.' In addition, the

	w_1	w_2	w_3	w_4
f_1	1	0.9	0.7	0.3
f_2	0.5	0.2	0.8	0.4
f_3	0.7	0.1	0.6	0.5
f_4	0.3	0.9	0.7	0.2

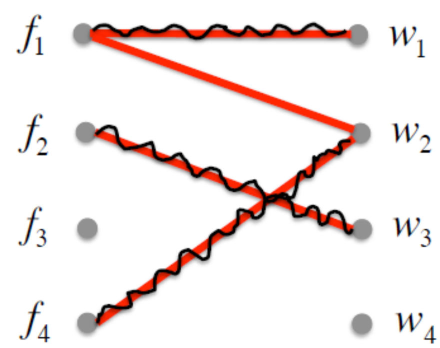

Figure 1: Heuristic Construction of Almost Efficient and Stable Matchings

per-person utility is greater than $1 - \varepsilon$, so that the matching is also 'almost efficient.'[3]

Our first result, which we view as motivating our study, makes this intuition crisper. In Proposition 1, we illustrate that asymptotically, *the fully stable matching is efficient.* Furthermore, we illustrate that the convergence speed is given by $\frac{\log n}{n}$.

While the full alignment case in which any pair of participants experiences the same utility upon matching is common in many applications, 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 portfolios of potential employees. Propositions 2 and 3 generalize our motivating case to settings in which the utilities each pair of participants experiences is a function of both a common shock, as well as individual idiosyncratic shocks, in a fairly arbitrary manner (in particular, shocks are not necessarily uniformly random). We show that *all* stable matchings are asymptotically efficient in these settings. Nonetheless, idiosyncrasies reduce the speed of convergence. When utilities are fully idiosyncratic and uniform, the difference between efficiency generated by stable matchings and the maximal feasible efficiency is of the order of $\frac{1}{\log n}$.

The last case we consider is that in which one side of the market may share preferences

[3] An alternative proof establishing this observation can be deduced from the appendix of Compte and Jehiel (2008).

over members of the other side of the market. For example, medical residents may share preferences over hospitals that are due to their published rankings and hospitals may agree on what makes a medical resident desirable.[4] Proposition 4 shows that in such settings as well, the unique stable matchings are asymptotically efficient. Furthermore, the speed of convergence coincides with that corresponding to fully aligned preferences.

In such settings, a designer using stability as an objective would be justified on efficiency grounds as well, at least when markets are large. In particular, a ban on transfers would not come at a great efficiency cost.

However, there are a few caveats to the message conveyed by our first main results on the asymptotic efficiency of stable matchings. First, there are certain market features that may induce stable matchings whose efficiency is bounded away from the optimum even when markets are large. 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 an advantage to the scarce side of the market (see Ashlagi, Kanoria, and Leshno, 2013). In our setting, a fixed difference between the volumes on either side does not change our conclusions that stable matchings are asymptotically efficient. These results also continue to hold for greater volume differences when preferences are aligned or assortative. However, in Section 7.1 we show that whenever utilities are determined in a fully random fashion 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 be asymptotically efficient. The other feature of markets that may make stable matchings be inefficient even 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, in Section 7.2 we show that in assortative markets with sub-modular preferences, stable matchings may entail a substantial amount of inefficiency regardless of the size of the market.[5]

[4] Agarwal (2013) reports that conversations with residency program and medical school administrators indicate that, indeed, programs broadly agree on what makes a resident appealing.

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

The second caveat has to do with the notion of efficiency one uses. For most of the paper, we consider an expected efficiency notion, where utilitarian efficiency is averaged across participants. This notion gives a normalized assessment of the expected outcomes participants are likely to experience in stable mechanisms. In particular, it reflects individual incentives to switch from one potential clearinghouse to another (e.g., a switch from clearinghouse allowing transfers to ones that ban them). Nonetheless, market designers may also be concerned about overall efficiency, reflecting, say, the overall costs of banning transfers in clearinghouses such as the NRMP. In Section 8 we analyze the unnormalized asymptotic efficiency of stable matchings. The baseline maximum efficiency achievable through any matching is a solution of a variation of the random assignment problem in statistics (see, e.g., Walkup, 1979 and work that followed). We show that for the classes of preferences we study, when considering unnormalized efficiency, stable matchings are substantially less efficient than the optimum for any market size. Furthermore, idiosyncratic preference components seem to generate greater wedges between the most efficient matchings and the stable ones.

To summarize, taken together our results can provide guidance to market designers who care about efficiency, or contemplate the introduction of some form of transfers between participants, as to when standard stable mechanisms are desirable. If a designer's concern regards expected outcomes per participant, and markets are fairly balanced, stable mechanisms are justified on efficiency grounds for sufficiently large markets. In particular, the availability of transfers will not affect outcomes significantly. However, if markets are not very large, or severely imbalanced and entailing a prominent idiosyncratic component of participants' preferences, or if the designer worries about overall (unnormalized) outcomes, commonly used stable mechanisms may not be ideal.

1.2. Literature Review. There are several strands of literature related to this paper. Efficiency of stable matching has been a topic of recent study. Boudreu and Knoblauch (2010) provide an upper bound on the sum of partner ranks in stable matchings when preferences exhibit particular forms of correlation. These upper bounds increase with the size of the market.[6]

[6] In a related paper, Knoblauch (2007) illustrates bounds on expected ranks for participants when one side of the market has uniformly random preferences and the other has arbitrary preferences. Liu and Pycia (2013) consider ordinally efficiency mechanism and illustrate that uniform randomizations over deterministic

Several papers have considered the utilitarian welfare loss stability may entail in matching markets. Anshelevich, Das, and Naamad (2009) consider finite markets and particular constellations of utilities and provide bounds on the utilitarian efficiency achieved through stability relative to that achieved by the efficiency-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 preferences are fully idiosyncratic and drawn from the uniform distribution (in line with our Proposition 2). Durlauf and Seshadri (2003) consider assortative markets in which agents may form coalitions, of any size, whose output depends on individuals' ability profile. Their results imply that the (utilitarian) 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 (see Section 7.2).

Che and Tercieux (2013) is possibly the most related paper to the current one. They 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 5). They show that Pareto efficient allocations are asymptotically utilitarian efficient. However, in the case of assignment problems they study, stable allocations are not necessarily Pareto efficient, so they are not necessarily utilitarian efficient. In contrast, we study general one-to-one matching markets (allowing, in particular, for aligned and/or idiosyncratic preferences on both sides, as well as imbalanced markets). In our setting of two-sided matching, stable matchings are utilitarian efficient.

Our results focus on large markets, which have received some attention in the literature, 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 (2013).

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.

While most of this literature focuses on balanced markets, recently Ashlagi, Kanoria, and Leshno (2013) have noted that imbalances in the volume of participants on both sides of the market may alleviate incentive compatibility issues, particularly when markets are large. We use some of their results when we discuss the efficiency in imbalanced markets in Section 7.1.

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 social utility of the (worst) Nash equilibrium outcome of a game and the maximum social utility possible in that game. In our context, a natural substitute to Nash equilibrium is a stable matching. In that respect, our results characterize the price of anarchy in many one-to-one matching environments. In particular, when considering normalized efficiency, the asymptotic price of anarchy is 1 for a wide array of balanced markets.[7]

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, 2012). The current paper provides an analogous study in the context of one-to-one matching.

Methodologically, our results borrow techniques introduced by Knuth (1976), Pittel (1989, 1992), and Lee (2013) and the intuition for our main insights utilizes graph theoretical results first noted by Erdos and Renyi (1964).

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 $(u_{ij}^f)_{i,j}$ and $(u_{ij}^w)_{i,j}$ 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 utilities u_{ij}^f and u_{ij}^w, respectively. We assume that any agent remaining unmatched receives a utility of 0, so that all agents prefer, at least weakly, to be matched with any agent over remaining unmatched (and this preference is strict almost always).

We consider market matchings $\mu : F \cup W \rightarrow F \cup W$ 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

[7] Without restricting preferences in any way, and taking a worst-case point of view, Echenique and Galichon (2013) 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 efficiency lower than the optimum by at least that value).

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^f_{ij} and u^w_{ij}, *stable matchings* are a subset of M that satisfies the following condition: For any firm and worker pair (f_i, w_j), either $u^f_{i\mu(i)} \geq u^f_{ij}$ or $u^w_{\mu(j)j} \geq u^w_{ij}$. 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.[8] Whenever there exist a firm and a worker that prefer being matched to 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 efficiency of stable matching to the maximal utilitarian efficiency achievable in a market matching.

The expected maximal utilitarian efficiency achievable across all market matchings is denoted by E_n :

$$E_n = \mathbb{E}_{\mathcal{G}} \max_{\mu \in M} \sum_{i=1}^{n} \left(u^f_{i\mu(i)} + u^w_{i\mu(i)} \right).$$

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

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

The minimal utilitarian efficiency achievable by implementing a stable matching for any realized match utilities is denoted by S_n and defined as:

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

Our goal is to characterize settings in which $\frac{S_n}{E_n} \to 1$ as the market becomes asymptotically large, in which case we say that stable matching are *asymptotically efficient*. Notice that the per-person expected utility when a stable matching is implemented is always bounded by the

[8] In general, stability also entails an individual rationality restriction (so that no agent prefers remaining unmatched over her prescribed match). Given our assumptions on utilities, this restriction is automatically satisfied.

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 matching are asymptotically efficient.

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. 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 there is a fixed wedge between the volume of firms and workers (say, there are n firms and $n+k$ workers, where k is fixed) with essentially identical proofs. 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 which we discuss in Subsection 7.1. For presentation simplicity, we initially focus our analysis on balanced markets.

3. A Motivating Case – Full Alignment

We start with one instance of our model that is commonly studied: the case of aligned preferences. In this case, members of each matched pair receive utilities proportional to one another (e.g., a firm and a worker may be splitting the revenues their interaction generates). We illustrate that such markets have unique stable matchings which are asymptotically efficient. We also characterize the speed at which the per-person utilitarian efficiency of stable matching converges 1.

Formally, we assume here that the utility both firm f_i and worker w_j receive if they are matched is given by $u_{ij} \equiv u^f_{ij} = u^w_{ij}$. We assume u_{ij} are independently drawn across all pairs (i,j) from a continuous distribution G over $[0,1]$. It follows that, generically, utility realizations (u_{ij}) entail a unique stable match. Indeed, consider realized utility realizations $(u_{ij})_{i,j}$ such that no two entries coincide. Consider 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 on the remaining participants. Again, we can find the pair that achieve the maximal match utility within that restricted market. As before, they must be matched in any stable matching. Continuing recursively, we construct the unique stable matching.

As it turns out, this construction lends itself to the following result.

Proposition 1 [Full Alignment – Asymptotic Efficiency and Convergence Speed]. *In fully aligned markets,*

1. *Stable matchings are asymptotically efficient:* $\lim_{n\to\infty} \frac{S_n}{2n} = 1$.
2. *For any* $n \geq 3$*, when* G *is uniform,*

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

The appendix contains the formal proof that relies on the construction of the stable matching described above. Roughly, the proof proceeds as follows. We first consider the uniform distribution. When determining the match utilities, the greatest realized entry, that corresponds 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 generate 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 the expected value $\frac{n^2}{n^2+1}\frac{(n-1)^2}{(n-1)^2+1}$. Continuing recursively, it can be shown that:

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

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 efficiency does not depend on the uniform distribution of utilities.[9] However, in order to show that the utilitarian efficiency of stable matchings converges to 1 at a speed of the order $\frac{\log n}{n}$, we use the precise formulation of S_n above.

[9] This effectively relies on the speedy convergence of extremal order statistics for distributions as the ones we consider here.

4. Idiosyncratic Preference Shocks

While the full alignment case analyzed in Section 3 is useful for many applications, there are also many applications in which individuals may have idiosyncratic preferences over their matching partners. For instance, individuals may have idiosyncratic preferences over geographical locations of their employers in the job market, or idiosyncratic preferences over partner's characteristics in the marriage market, etc. In the current section we analyze a more general model that allows for fairly arbitrary correlation between the utilities of any two matched individuals.

4.1. Fully Independent Preferences. In order to illustrate the forces underlying the asymptotic features of markets entailing idiosyncratic preference shocks, we start by describing the polar case in which preferences are determined in a fully independent manner. That is, all values $(u_{ij}^f)_{i,j}$ and $(u_{ij}^w)_{i,j}$ are independently and identically distributed according to a continuous distribution G over $[0,1]$.

Notice first that the maximal conceivable efficiency level (of 1 per person) is achievable asymptotically. Indeed, for any market realization $(u_{ij}^f, u_{ij}^w)_{i,j}$ consider the matrix of utility realizations $\tilde{u}_{ij} \equiv \frac{u_{ij}^f + u_{ij}^w}{2}$. The distribution G then induces a continuous distribution $\tilde{G}[0,1]$ from which $\tilde{u}_{ij}$ are independently drawn. We can now view $(\tilde{u}_{ij})_{i,j}$ as match utilities shared by firms and workers as in our analysis in Section 3. Proposition 1 implies that the unique stable matching of the induced fully aligned market is efficient asymptotically. Note that any matching μ generates the same levels of efficiency when each firm f_i and worker w_j both receive $\tilde{u}_{ij}$ when matching to one another, or when the firm receives u_{ij}^f and the worker receives u_{ij}^w. We therefore have the following corollary.

Corollary 1 [Full Independence – First Best Efficiency]**.** *In fully independent markets, the maximal conceivable efficiency level is achievable asymptotically. That is,*

$$\lim_{n\to\infty} \frac{E_n}{2n} = 1.$$

Notice that in this setting, stable matchings are not generically unique. As it turns out, any selection of stable matchings generates the maximal achievable efficiency asymptotically. That is,

Proposition 2 [Full Independence – Asymptotic Efficiency and Convergence Speed]. *In fully independent markets,*

1. Stable matchings are asymptotically efficient:

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

and, in particular, $\lim_{n\to\infty} \frac{S_n}{2n} = 1$.

2. When G is uniform,

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

In order to first glean some intuition for why idiosyncratic individual shocks to preferences may still be consistent with asymptotic efficiency, consider a similar construction to that described in the introduction for the case of full alignment. We look at the bipartite graph in which firms and workers constitute the two types of nodes. For simplicity, consider the case in which G is uniform over $[0, 1]$. Now, for any $\varepsilon > 0$, a link between firm f_i and worker w_j is formed if both $u_{ij}^f > 1 - \varepsilon$ and $u_{ij}^w > 1 - \varepsilon$. In particular, a link is formed with probability ε^2. Using Erdos and Renyi (1964), there is an asymptotic perfect matching in the induced graph as long as ε^2 converges to 0 sufficiently slowly so that the graph is 'connected enough'. Namely, as long as ε^2 converges to 0 slower than $\frac{\log n}{n}$, or ε converges to 0 slower than $\sqrt{\frac{\log n}{n}}$, the probability of the induced graph containing a perfect matching converges to 1. Notice that in such a matching almost all agents can gain at most ε from pairing with someone other than their allocated match partner. In other words, this is a way to construct an almost stable and almost efficient matching. Nonetheless, an 'almost stable' matching may be very far, in terms of number of blocking pairs, overlap in matched pairs, etc. from any of the market's stable matchings. As mentioned in the introduction, this distinction could be particularly important given that most centralized clearinghouses utilized in matching markets are designed to implement a fully stable matching (assuming participants report truthfully their preferences).

The formal proof of Proposition 2, which appears in the Appendix, utilizes different techniques than those employed to prove Proposition 1. It relies on results by Pittel (1989). In order to see the basis of the proof, consider a formula derived by Knuth (1976) for the probability that any matching is stable. Specifically, 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

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

where $d\mathbf{u}_{ii}^f = du_{11}^f du_{22}^f \cdots du_{nn}^f$ and $d\mathbf{u}_{jj}^w = du_{11}^w du_{22}^w \ldots du_{nn}^w$.

The intuition behind this formula is simple. The formula essentially evaluates the probability that the match μ, with $\mu(i)$ equal to i for all i, is stable. Indeed, for any realized market, in order for μ to be stable, utilities $(u_{ij}^f, u_{ij}^w)_{1\leq i\neq j\leq n}$ must satisfy that either $u_{ij}^f < u_{ii}^f$ or $u_{ij}^w < u_{jj}^w$ 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,

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

From symmetry, the probability that any other 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.[10]

When utilities are distributed uniformly, the formulas above simplify and, as we show in the Appendix, the convergence speed of $1-\frac{S_n^f}{n}$ is $\frac{1}{\log n}$. Figure 2 illustrates numerical results for the utilitarian efficiency per market participants for different classes of preferences. For each

[10] The proof appearing in the appendix circumvents the formulas described here and utilizes more directly results from Pittel (1989). An alternative proof starting from these formulas is available from the authors upon request.

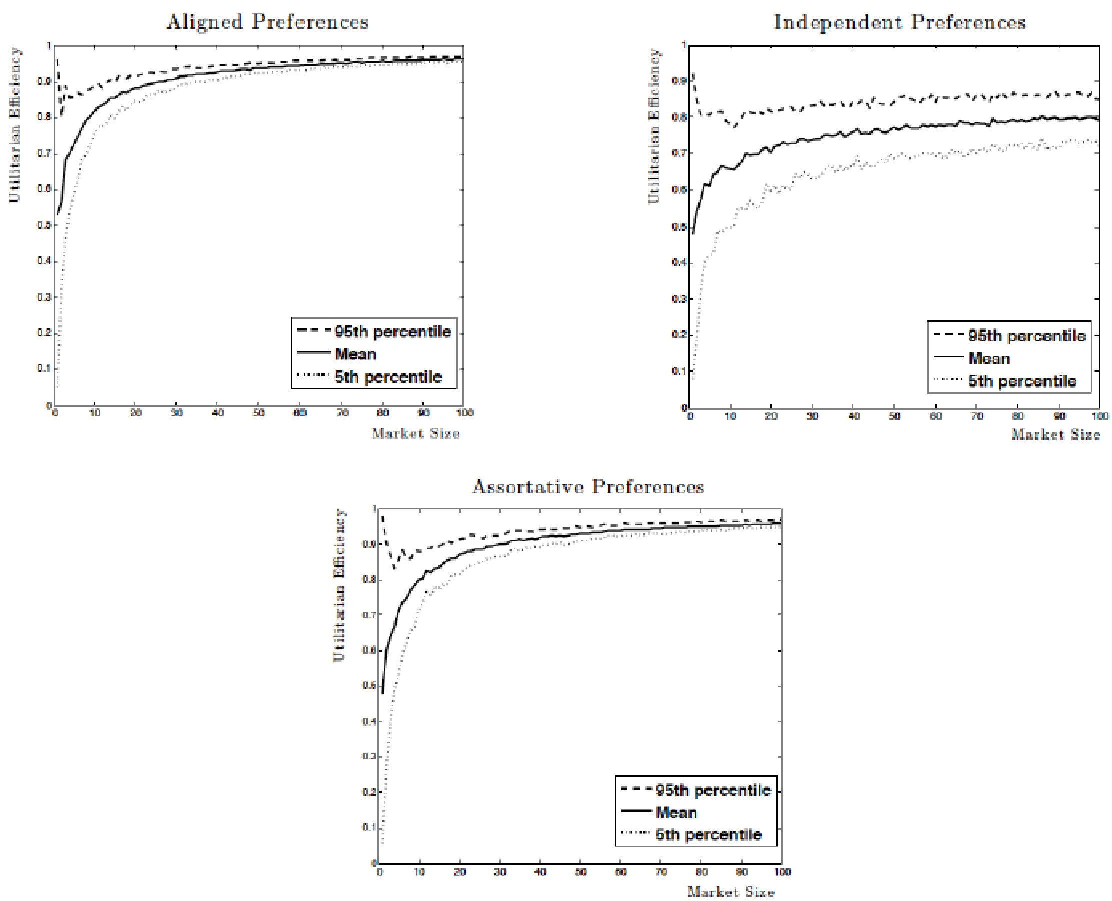

Figure 2: Comparison of Convergence Speeds Across Market Types

market size n, we run 100 simulations, each corresponding to one realization of preferences in the market. For each simulation, we compute the average per-participant utility induced by their least preferred stable match partners. The solid line, the long dashed line, and the short dashed line depict, respectively, to the the mean, the 95th percentile, and the 5th percentile of the simulated distribution of these averages across the 100 simulations.[11] The top left panel corresponds to the case of perfectly aligned preferences with uniformly drawn utilities analyzed in Section 3. The top right panel corresponds to the case analyzed in this section. Even with only 100 participants on each side, efficiency levels are quite high, for perfectly aligned preferences, around 95% of the maximal levels, while for the perfectly independent preferences, around 80% of the maximal levels. Furthermore, the spread of average efficiencies decreases with market size. Nonetheless, the speed of convergence of efficiency of stable matchings is indeed markedly slower in markets with fully independent preferences than when preferences are perfectly aligned.

4.2. General Aligned Markets with Idiosyncratic Shocks. Realistically, many markets may entail 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 courses a potential employee took in college). We now consider general markets allowing for both impacts and illustrate that the results pertaining to the asymptotic efficiency of stable matchings still hold.

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 $\mathbb{R}_+^2$ to $\mathbb{R}_+$, are continuous and strictly increasing for both arguments. Each of c_{ij}, z_{ij}^f, and z_{ij}^w is drawn independently from a non-trivial

[11] In principal, when preferences are perfectly aligned, we can use the formula developed in Section 3 to calculate the expected per-person expected utility. The figure gives a sense of the spread of the distribution (and the mean tracks closely that generated by the formal expression of expected efficiency).

distribution over $\mathbb{R}_+$, and the three distributions have positive density functions and bounded supports.

Our main result in this section shows that, asymptotically, participants of the market achieve, on average, their maximal conceivable match utility, regardless of how stable matchings are selected.

Proposition 3 [Efficiency of Stable Matchings]**.** *Let $\bar{c}$, $\bar{z}^f$, and $\bar{z}^w$ be upper bounds for the distributions of c_{ij}, z_{ij}^f, and z_{ij}^w, respectively. Then,*

$$\lim_{n\to\infty} \frac{S_n^f}{n} = \phi(\bar{c}, \bar{z}^f) \quad \textit{and} \quad \lim_{n\to\infty} \frac{S_n^w}{n} = \omega(\bar{c}, \bar{z}^w).$$

Therefore, $\lim_{n\to\infty} \frac{S_n}{n} = \phi(\bar{c}, \bar{z}^f) + \omega(\bar{c}, \bar{z}^w)$.

Proposition 3 illustrates that, in fact, *fully stable matchings are asymptotically efficient, 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 asymptotically achieves the maximal conceivable utility per participants.

We note that Proposition 3 is not a direct generalization of Propositions 1 and 2 as the distributions of c_{ij}, z_{ij}^f, and z_{ij}^w are assumed to be non-trivial. 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.

Propositions 1-3 combined illustrate that when match utilities have either or both aligned and idiosyncratic components, any selection of stable matchings will lead to matchings that are approximately utilitarian efficient for sufficiently large markets.

5. Assortative Markets

The last class of matching markets we consider is that allowing for some 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, 2013); similarly, parents may evaluate schools according to the performance of current and past students; 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 efficient stable matchings.

In order to gain intuition on the driving force behind our main result in this section, we first consider a polar case in which workers all share the same evaluation of firms with utilities determined uniformly, while firms have independent evaluations of workers. Formally, we assume for now that u_{ij}^f are independently and identically distributed according to the uniform distribution over $[0, 1]$. Unlike before, we assume that for any j and j', $u_i^w \equiv u_{ij}^w = u_{ij'}^w$ for all i, so that all workers agree on the ranking and the valuations of the firms. Again, we assume that all utilities are independently determined according to the uniform distribution over $[0, 1]$.

Generically, realized markets will entail a unique stable matching. Indeed, the unique stable matching is assortative: The most desirable firm matches with her highest ranked worker (indeed, they are each other's favorite partner in the market); the second most desirable firm then matches with her highest ranked worker of the remaining $n - 1$; and so on.

We use the same notation as before for the efficiency experienced by firms and workers under the (generically) unique stable matching. Since match utilities are determined independently, one of the workers will receive a utility corresponding to the highest entry of n uniform variables (from matching with the highest ranked firm), one will receive a utility corresponding to the next highest entry of n uniform variables, etc. Therefore,

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

For firms, the most desirable firm's expected utility is the expectation of the highest entry of n uniform random variables (corresponding to the n workers), i.e., $\frac{n}{n+1}$. The second most preferred firm's expected utility is the expectation of the highest entry of $n-1$ uniform variables, $\frac{n-1}{n}$, etc. Therefore,

$$S_n^f = \frac{n}{n+1} + \frac{n-1}{n} + ... + \frac{1}{2} = n - \sum_{k=1}^{n} \frac{1}{k+1}.$$

Notice that

$$\log(n+2) - \ln 2 = \int_2^{n+2} \frac{1}{x} dx \leq \sum_{k=1}^{n} \frac{1}{k+1} \leq \int_1^{n} \frac{1}{x} dx = \log n$$

Now, any match involves the same expected payoff for workers. It follows that stability achieves the first best in terms of utilitarian efficiency for workers. Moreover, the expected value of utilitarian welfare of firms converges to 1.[12] Formally,

Proposition 4 [Assortative Markets – Asymptotic Efficiency and Convergence Speed]**.** *For all n,*

$$\frac{S_n^w}{n} = \frac{1}{2} \quad \textit{and} \quad \frac{\log(n+2) - \log 2}{n} \leq 1 - \frac{S_n^f}{n} \leq \frac{\log n}{n}.$$

In particular, $\lim_{n\to\infty} \frac{S_n^f}{n} = 1$ and stable matchings are asymptotically efficient.

The speed of convergence in this setting is of the order of $\frac{\log n}{n}$, as in the case of fully aligned (and uniform) utilities. Indeed, the bottom panel of Figure 2 depicts the per-participant efficiency from a 100 simulations of the assortative market analyzed above for different values of n. The results mirror those corresponding to the case of perfectly aligned preferences (and convergence speeds are therefore more rapid than for the case of fully independent preferences). In particular, even for $n = 100$ participants on each side of the market, the per-participant efficiency reaches approximately 95% of the maximal achievable level.

[12] Naturally, had we assumed that firms share the same evaluations of workers (so that the market was ex-ante symmetric), the expected utilitarian efficiency of workers' would experience in the generically unique stable matching would converge to 1 as well.

We now generalize the above example to allow for idiosyncratic preference shocks, as we had done in the previous section. 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 match 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, proxied for by 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, proxied for by 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 $\mathbb{R}_+^2$ to $\mathbb{R}_+$ 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 non-atomic continuous distributions over $\mathbb{R}_+$, that all have bounded supports.

Let E_n^f and E_n^w be the expected maximal utilitarian welfare 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 utilitarian welfare as market size increases.

Proposition 5 [Efficiency of Stable Matchings]**.** *Stable matchings in assortative markets with idiosyncratic shocks are asymptotically efficient:*

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

The proof of Proposition 5 is a direct consequence of Lee (2013). Indeed, notice that it is without loss of generality to consider variables c_i^f, c_j^w, z_{ij}^f, and z_{ij}^w are all uniformly distributed over $[0, 1]$, as we can always transform the aggregating utilities Φ and Ω monotonically in a consistent manner. Proposition 1 in the online appendix of Lee (2013) implies 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, guarantees that $\lim_{n\to\infty} \frac{S_n^f}{E_n^f} = 1$. A symmetric argument holds for the utilitarian efficiency experienced by workers.

We note that Che and Tercieux (2013) provide a similar result to Proposition 5, though they allow for atomic distributions of c_j^f or c_i^w, a knife-edge case, which the proposition achieves only as a limit of continuous distributions that are concentrated at certain values.[13] Allowing for atoms in the distributions of the common component may entail stable matchings that are not efficient even when markets are large.[14]

Lee (2013) suggested 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 (2013) implies that for any stable matching mechanism and any $\varepsilon, \delta > 0$, there exists N such in any market of size $n > N$, with probability $1 - \delta$, there is an ε-equilibrium implementing a stable matching corresponding to the underlying realized utilities. Together with our results, this suggests the following corollary.

Corollary 2 [Stable Matching Mechanisms]**.** *Stable matching mechanisms are asymptotically efficiency and incentive compatible.*

6. Matching with Transfers and Price of Anarchy

Our results thus far illustrate the asymptotic efficiency of stable matchings for a large class of environments. Suppose that participants have quasi-linear utilities composed of a linear utility for money, and a potentially non-linear utility that belongs to one of the classes of utilities we have been studying (namely, allowing for components that are aligned, assortative, or

[13] Our analysis of the uniform distribution allows to characterize the speed of convergence, that is difficult to characterize more generally, in particular in the Che and Tercieux (2013) setting.

[14] For example, if c_j^f takes values of 0 or 1 with, say, equal probability for all j, and c_i^w is fixed at some constant c with probability 1 for all i, then the resulting market is equivalent to a severely imbalanced market with idiosyncratic preferences that we show to be inefficient asymptotically in Section 7.1.

idiosyncratic). Our results suggest that, in large markets, a ban on transfers does not come at a significant utility cost to individuals. This is 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).[15]

Another way to view these results is by noting that using stability as an objective in market design does not entail substantial efficiency costs when markets are large. In fact, since former results suggest that stable matching mechanisms are asymptotically incentive compatible, at least when preferences are hybrids of assortative and idiosyncratic components, our results serve as a positive defense of commonly used mechanisms such as the Gale-Shapley (1962) deferred acceptance algorithm – in such settings, they are asymptotically incentive compatible and efficient (our Corollary 2).

Our results thus far also offer insights into the 'price of anarchy' in matching markets. As mentioned in our Introduction, the price of anarchy is defined as the ratio between the social utility of the (worst) equilibrium outcome of a game and the maximum social utility possible in that game (see Roughgarden and Tardos, 2007). In our context, a natural substitute to Nash equilibrium is a stable matching, in which potential deviations are made by either pairs or individuals. Our results so far suggests that, when considering normalized efficiency, the asymptotic price of anarchy is 1 within a large class of balanced markets.

7. Markets with Asymptotically Inefficient Stable Matchings

It is not very hard to find a large market where stable matchings are asymptotically utilitarian inefficient. We have discussed in the introduction that randomness of utilities is necessary: a replication of small markets in which stable matchings are utilitarian inefficient generates (larger) markets that entail stable matchings that are inefficient.[16] In this section, we study

[15] Dizdar and Moldovanu (2013) study a two-sided matching model 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. Our results illustrates that, when markets are large, transfers can be banned altogether and outcomes arbitrarily close to efficient can be achieved.

[16] 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

two classes of environments in which even a fair bit of randomization of preferences leads to markets in which stable matchings are not asymptotically efficient.

7.1. Severely Imbalanced Markets. Throughout the paper, we have 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 whenever the imbalance were bounded, e.g., if there were n firms (workers) and $n+k$ workers (firms), where k is fixed. Since in many real-world 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 (2013) that illustrate the sensitivity of the structure of stable matchings to the relative sizes of both sides of the matching market.

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 measure of efficiency considers the per-person expected utility, conditional on being matched. As before, since there might be multiple stable matchings, we will consider the worst-case scenario.

We consider markets with n firms and $n+k(n)$ workers and examine the asymptotic efficiency for matched workers when the firm-optimal stable matching, the worker-pessimal stable matching, is implemented. 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.[17]

utilities $(\tilde{u}^f_{ij}, \tilde{u}^w_{ij})$, where

$$\tilde{u}^x_{i'j'} = \begin{cases} u^x_{i' \bmod k \; j \bmod k} & i' \operatorname{div} k = j' \operatorname{div} k \\ -1 & \text{otherwise} \end{cases},$$

so that, for any $l = 1, 2, \ldots, k$, firms $f_{l1}, \ldots, f_{lm}$ and workers $w_{l1}, \ldots, w_{lm}$ have the same preferences over one another (and over being unmatched) as in the original market, and prefer staying unmatched over matching with anyone else in the market.

[17] We believe these are the cases that are interesting to consider from an applied perspective. Furthermore, whenever $\frac{k(n)}{n}$ explodes, the relevant efficiency statements would pertain to an insignificant fraction of firms that end up being matched.

Notice that the addition of workers can only improve firms' expected utility. Therefore, in any balanced setting in which asymptotic efficiency is achieved, the introduction of more workers guarantees the asymptotic efficiency of stable matchings for firms.

The main insight of this section is that asymptotic efficiency may break whenever markets are severely imbalanced *and* preferences exhibit substantial idiosyncratic components.

When markets are perfectly aligned or perfectly assortative, the proofs of Propositions 1 and 4 carry through for arbitrary increasing functions $k(n)$ and asymptotic efficiency of stable matchings still holds.[18] We now focus on the case studied in Section 4.1 in which preferences are idiosyncratic shocks, where we normalize the utility from remaining unmatched to be zero. Recall that S_n^w denotes the expected worst-case utilitarian efficiency experienced by workers in any stable matching.

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

Proposition 6 [Imbalanced Fully Independent Markets]. *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,*

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

Notice that this indeed suggests inefficiency 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 perfectly aligned market with utilities $(u_{ij}^* \equiv \frac{u_{ij}^f + u_{ij}^w}{2})_{i,j}$. That is, in the induced market, each matched firm and worker receive their average match utilities in the original market. The efficiency results pertaining to aligned markets then carry through for the induced market, and in fact the stable matchings in the induced markets are asymptotically efficient. Since these matchings produce the same per-participant utilities in the original markets, maximal efficiency can be achieved asymptotically. The wedge identified in Proposition 6 then implies a substantial asymptotic inefficiency generated by stability.

[18]The arguments required to illustrate efficiency in purely assortative markets follow directly our proof when there are n firms and $n + k(n)$ workers. For the reverse case specified here, slightly more involved arguments are required that follow directly from results in Lee (2013).

7.2. Sub-modularity in Match Qualities. Another important case in which inefficiency arises asymptotically pertains to assortative matching markets in which match utilities are sub-modular in partners' intrinsic values. For finite markets, Becker (1974) illustrated that sub-modularity in assortative markets implies that the utilitarian efficient matching is negatively assortative, whereas the unique stable matching is positively assortative.

Formally, consider a sequence of $n \times n$ assortative markets in which firms' intrinsic values are given by $(c_i^f = \frac{i}{n})_{i=1}^n$ and workers' intrinsic values are given by $(c_j^w = \frac{j}{n})_{j=1}^n$. Match utilities are determined according to an 'output function' ϕ:

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

such that

$$\frac{\partial \phi(c_i^f, c_j^w)}{\partial c_i^f} > 0, \quad \frac{\partial \phi(c_i^f, c_j^w)}{\partial c_j^w} > 0.$$

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 efficient matching or not. Indeed, when ϕ is linear, all matchings generate the same utilitarian efficiency and both the positively and negatively assortative matchings are utilitarian efficient. When output is super-modular, $\frac{\partial^2 \phi(c_i^f, c_j^w)}{\partial c_i^f \partial c_j^w} > 0$, the positively assortative matching is an efficient matching, while when output is sub-modular, $\frac{\partial^2 \phi(c_i^f, c_j^w)}{\partial c_i^f \partial c_j^w} < 0$, the positively assortative matching (i.e., stable 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.

Utilitarian welfare from the efficient matching (i.e., the 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, utilitarian welfare from the stable matching (i.e., the positively assortative matching) is

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

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,

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

When α is close to 0, output is very insensitive to individual abilities, and any matching, in particular the stable one, generates efficiency close to the optimum. When α is close to 1, the output function is 'almost linear' in individual abilities and, again, the stable matching is asymptotically nearly efficient. Nonetheless, for intermediate values of α, asymptotic efficiency is bounded strictly below 1, which is achieved by the most efficient matching.

8. Overall Efficiency

Throughout most of the paper we consider average, or normalized, efficiency, where utility is averaged across market participants. The notion of normalized efficiency we have been studying is particularly useful when the designer is concerned with the 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 about overall efficiency. In this section, we study the wedge in terms of overall efficiency between optimal matchings (those maximizing overall efficiency) and stable matchings. In terms of overall efficiency, 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 expected maximal utilitarian efficiency across all matchings. Our goal in this Section is to characterize the wedge $E_n - S_n$. In order to provide precise bounds on this difference, we focus on two polar cases in our setting: fully aligned and fully independent (or idiosyncratic) markets, where utilities are drawn from the uniform distributions (the environments studied in Sections 3 and 4.1, respectively).

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

Proposition 7 [Efficiency Loss without Normalization]**.**

1. For any $n \geq 3$,

$$\log n - 6 \leq L_n^A \leq 2 \log n.$$

2. The relative efficiency loss satisfies the following:

$$1 \leq \liminf_{n \longrightarrow \infty} \frac{L_n^I / L_n^A}{n / (\log n)^2} \leq 2.$$

The proposition illustrates the substantial welfare 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 welfare loss grows with market size, idiosyncratic preferences exhibiting a greater loss asymptotically. Namely, the ratio between the efficiency lost in idiosyncratic markets relative to that lost in aligned markets is asymptotically of the order of $n / (\log n)^2$, which increases with market size.

The proof of Proposition 7 relies on two sets of results. First, notice that Propositions 1 and 2 provide bounds on the speeds at which the expected efficiency of stable matchings grows for the environments we focus on here. We therefore need bounds on the speed with which the efficiency of optimal matchings grows. 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, $2n - 6 \leq E_n \leq 2n$.[19]

Consider now markets with fully independent preferences, drawn from the uniform distribution. 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 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 cannot be used directly, however, since now $\tilde{u}_{ij}$ 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, $E_n \geq 2n - 3\sqrt{n}$.

9. Conclusions

The paper illustrates that for a large class of preference distributions, stable matchings are asymptotically efficient, at least when considering expected efficiency across market participants. In these environments transfers are not necessary to implement efficient outcomes. In fact, in such settings, a market designer faces no tradeoff between efficiency and stability when the market is large, and the objective of stability is viable even on efficiency grounds.

Our results also illustrate the speeds of convergence of the efficiency of stable matchings to the optimum. Idiosyncratic preferences generate a substantially lower speed of convergence than those exhibited in markets with aligned or assortative preferences. This suggests that market designers concerned with efficiency might treat stability as an objective 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 side), stable matchings are no longer efficient in general, even when markets are large, and transfers could prove beneficial.

Another caveat pertaining to our positive results regarding the asymptotic efficiency of stable matchings regards the notion of efficiency one uses. Indeed, if one focuses on overall efficiency, rather than that averaged across participants, stability and efficiency display a

[19] Follow-up work has improved upon this bound (see, for instance, Coppersmith and Sorkin, 1999, whose work suggests that $2n - 3.88 \leq E_n$). 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 independent.

pronounced wedge across all of the preference constellations we consider (hybrids of assortative, idiosyncratic, and assortative preferences). An approach aiming at maximizing overall efficiency would then potentially require mechanisms that are not necessarily stable or ones allowing transfers.

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 efficient, they are also asymptotically incentive compatible. However, for a designer concerned with overall efficiency, it would be important to analyze the most efficient incentive compatible mechanisms. Furthermore, for other types of preferences, particularly those entailing assortative components, even the question of incentive compatibility of stable mechanisms in large markets is still open.

10. Appendix

10.1. Full Alignment – Proof of Proposition 1. We start by deriving the formula for S_n suggested in the text when utilities are distributed uniformly and illustrate both asymptotic efficiency and the speed of convergence of that case. We then generalize our asymptotic 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 and worker pairs, after all firms and workers receiving the $k-1$ highest utility 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 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 Proposition 1. First, denote by $C_n = 1 - \frac{S_n}{2n}$.

For $n \geq 1$,

$$\log n = \int_1^n \frac{1}{x}dx,$$

which, using the convexity of $\frac{1}{x}$, implies that

$$\sum_{k=2}^{n}\frac{1}{k} \leq \log n \leq \sum_{k=1}^{n-1}\frac{1}{k} \qquad \text{for all} \quad n = 2,3,\ldots.$$

Thus, it suffices to show that

$$\frac{13}{25}\frac{1}{n}\sum_{k=1}^{n-1}\frac{1}{k} \leq C_n \leq \frac{1}{n}\sum_{k=2}^{n}\frac{1}{k} \qquad \text{for all} \quad n = 3,4,\ldots.$$

We prove the above two inequalities separately by induction.

Claim 8.

$$C_n \leq \frac{1}{n}\sum_{k=2}^{n}\frac{1}{k} \qquad \textit{for} \quad n = 3,4,\ldots.$$

Proof of Claim 8. For $n \geq 2$,

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

Therefore,

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

Especially,

$$C_3 = \frac{1}{10} + \frac{3\cdot 2}{10}C_2 = \frac{1}{10} + \frac{3\cdot 2}{10}\left(\frac{1}{5} + \frac{2}{5}\cdot\frac{1}{2}\right) = \frac{13}{50} < \frac{1}{3}\sum_{k=2}^{3}\frac{1}{k} = \frac{5}{18}.$$

Now, for $n > 3$, suppose that

$$C_{n-1} \leq \frac{1}{n-1}\sum_{k=2}^{n-1}\frac{1}{k}.$$

Then,

$$\begin{aligned} C_n &= \frac{1}{n^2+1} + \frac{n(n-1)}{n^2+1}C_{n-1} \leq \frac{1}{n^2+1} + \frac{n(n-1)}{n^2+1}\frac{1}{n-1}\sum_{k=2}^{n-1}\frac{1}{k} \\ &= \frac{1}{n^2+1} + \frac{n}{n^2+1}\sum_{k=2}^{n-1}\frac{1}{k} \leq \frac{1}{n^2} + \frac{1}{n}\sum_{k=2}^{n-1}\frac{1}{k} = \frac{1}{n}\sum_{k=2}^{n}\frac{1}{k}, \end{aligned}$$

and the proof of Claim 8 follows. ■

Claim 9.

$$\frac{1}{2}\frac{1}{n}\sum_{k=1}^{n-1}\frac{1}{k} \leq C_n \qquad \textit{for} \quad n = 3, 4, \ldots.$$

Proof of Claim 9. Recall that $C_3 = \frac{13}{50}$, which is larger than $\frac{1}{2}\frac{1}{3}\sum_{k=1}^{2}\frac{1}{k} = \frac{1}{4}$.

For $n > 3$, suppose that

$$C_{n-1} \geq \frac{1}{2}\frac{1}{n-1}\sum_{k=1}^{n-2}\frac{1}{k}.$$

Then,

$$C_n = \frac{1}{n^2+1} + \frac{n(n-1)}{n^2+1} C_{n-1} \geq \frac{1}{n^2+1} + \frac{n(n-1)}{n^2+1} \frac{1}{2} \frac{1}{n-1} \sum_{k=1}^{n-2} \frac{1}{k} \geq \frac{1}{2n} \sum_{k=1}^{n-1} \frac{1}{k}.$$

The last inequality follows from

$$2 \geq \frac{1}{n}\left(1+\log(n-1)\right) + \frac{n}{n-1} \geq \frac{1}{n} \sum_{k=1}^{n-1} \frac{1}{k} + \frac{n}{n-1} \quad \text{for every} \quad n > 3,$$

which implies that

$$2 + n \sum_{k=1}^{n-2} \frac{1}{k} \geq \left(n + \frac{1}{n}\right) \sum_{k=1}^{n-1} \frac{1}{k}.$$

This completes the proof of Claim 9.

The proof of Proposition 1 then follows.

We now show that asymptotic efficiency of stable matchings holds even when the match utilities of firm and worker pairs are independently and identically drawn from an arbitrary 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 is β such that $1-\varepsilon < \beta < 1$ and $G(\beta) < 1$.

Let

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

so that G' is the uniform distribution over $\left[\frac{-G(\beta)\beta}{1-G(\beta)}, \beta\right]$.

Let $S_n^{'}$ and $S_n^{''}$ be the expected utilitarian efficiencies derived from the stable matchings when utilities are drawn from $G'(x)$ and the uniform distribution over $[0,1]$, respectively. We get

$$\begin{aligned} S'_n &= \left(\beta - \frac{-G(\beta)\beta}{1-G(\beta)}\right) S''_n + \frac{-G(\beta)\beta}{1-G(\beta)} \\ &= \left(\frac{\beta}{1-G(\beta)}\right) S''_n + \frac{-G(\beta)\beta}{1-G(\beta)}. \end{aligned}$$

Therefore,

$$\lim_{n\to\infty} \frac{S'_n}{2n} = \beta.$$

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 orders of utilities $(u_{ij})_{i,j}$, S_n is a sum of specific n order statistics. Thus, the utilitarian efficiency from the stable matching under G first order stochastically dominates that under G'. It follows that

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

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

10.2. Idiosyncratic Shocks – Proof of Proposition 2. 1. We assume that all utilities are drawn from the uniform distribution on $[0,1]$. A similar technique to that used in the proof of Proposition 1 for general distributions can be used to generalize the result to arbitrary continuous distributions.

We want to show that

$$\frac{S_n^f}{n} = \frac{\mathbb{E}\left[\sum_{i=1}^n u_{i\mu_W(i)}\right]}{n} \to 1.$$

We consider a two-step procedure of generating idiosyncratic preferences. First, an ordinal preference profile, denoted by $\succ$, is drawn from the uniform distribution over the set of all possible preference profiles. That is, each firm has a preference list that is drawn from the uniform distribution over the set of permutations of n workers. Each worker's preference is similarly generated. Given a realization of $\succ$, cardinal utilities for each agent are generated as follows. n numbers are drawn from the uniform distribution on $[0,1]$. The highest number is then the match utility resulting from a match with the agent's most preferred partner, the second highest is the match utility resulting from matching with the second preferred partner, etc.

This two-step procedure then implies:

$$S_n = \mathbb{E}\left[\sum_{i=1}^{n} u_{i\mu_W(i)}\right] = \mathbb{E}_{\succ}\left[\mathbb{E}_{u|\succ}\left[\sum_{i=1}^{n} u_{i\mu_W(i)} | \succ\right]\right].$$

Let $R_i^{\mu_W}$ denote the rank number of firm i's worker-optimal stable matching partner. If the firm is matched to its most preferred worker, the rank number is 1. Also, let $u_{[k;n]}$ be k'th highest value from a sample of size n from the uniform distribution on $[0,1]$.

As the preference profile determines the rank number of the worker-optimal stable matching partners, and since the expected k'th highest value corresponding to the uniform distribution is given by $1-\frac{k}{n+1}$, we can write

$$\begin{aligned} S_n &= \mathbb{E}_{\succ}\left[\mathbb{E}_{u|\succ}\left[\sum_{i=1}^{n} u_{i\mu_W(i)} | \succ\right]\right] = \mathbb{E}_{\succ}\left[\mathbb{E}_{u|\succ}\left[\sum_{i=1}^{n} u_{[R_i^{\mu_W};n]} | \succ\right]\right] \\ &= \mathbb{E}_{\succ}\left[\sum_{i=1}^{n}\left(1-\frac{R_i^{\mu_W}}{n+1}\right)\right] = n - \frac{\mathbb{E}_{\succ}\left[\sum_{i=1}^{n} R_i^{\mu_W}\right]}{n+1}. \end{aligned}$$

We want to prove that

$$\frac{\mathbb{E}_{\succ}\left[\sum_{i=1}^{n} R_i^{\mu_W}\right]}{n(n+1)} \to 0.$$

We use Theorem 2 in Pittel (1989) showing that

$$\frac{\left(\sum_{i=1}^{n} R_i^{\mu_W}\right)\log n}{n^2} \xrightarrow{p} 1.$$

It is then immediate that $\frac{\sum_{i=1}^{n} R_i^{\mu_W}}{n^2} \xrightarrow{p} 0$. As $\frac{\sum_{i=1}^{n} R_i^{\mu_W}}{n^2}$ is bounded above by 1 with probability 1, we have $\frac{\mathbb{E}\left[\sum_{i=1}^{n} R_i^{\mu_W}\right]}{n^2} \to 0$ by Lebesgue's dominated convergence theorem. This completes the proof.

2. We now show that

$$\lim_{n\to\infty}\left(1-\frac{S_n^f}{n}\right)\log n = \lim_{n\to\infty}\frac{\left(\mathbb{E}_{\succ}\left[\sum_{i=1}^{n} R_i^{\mu_W}\right]\right)\log n}{n(n+1)} = 1.$$

We use two results from which Pittel (1989) obtains its main Theorem.

First, Equation (4.4) in Pittel (1989) assures that for any small $\rho > 0$ and $\delta \in (0, e^{\rho} - 1)$,

$$P\left(\sum_{i=1}^{n} R_i^{\mu_W} \leq \frac{n^2}{\log n}(1 + \frac{\log\log n + \rho}{\log n})\right) \geq 1 - O(n^{-\delta}).$$

Thus,

$$\frac{\mathbb{E}\left[\sum_{i=1}^{n} R_i^{\mu_W}\right]}{n^2} \leq (1 - O(n^{-\delta}))(\frac{1}{\log n})(1 + \frac{\log\log n + \rho}{\log n}) + O(n^{-\delta}),$$

which implies that

$$\lim_{n\to\infty} \frac{\mathbb{E}\left[\sum_{i=1}^{n} R_i^{\mu_W}\right]\log n}{n^2} \leq 1.$$

In addition, the result on page 545 of Pittel (1989) assures that

$$\mathbb{E}\left[\sum_{i=1}^{n} R_i^{\mu_W}\right] \geq \frac{n^2}{\log n}(1 + O(\frac{1}{\log n})),$$

which implies that

$$\lim_{n\to\infty} \frac{\mathbb{E}\left[\sum_{i=1}^{n} R_i^{\mu_W}\right]\log n}{n^2} \geq 1.$$

■

10.3. Aligned Preferences with Idiosyncratic Shocks – Proof of Proposition 3. It is without loss of generality to consider that c_{ij}, z_{ij}^f, and z_{ij}^w that are all uniformly distributed over $[0, 1]$. Indeed, an appropriate change of variables (or, equivalently, a monotone transformation of utilities), will generate an equivalent setting in which the underlying distributions are uniform.[20]

The model is a mixture of aligned preferences captured by the variables $c = (c_{ij})_{i,j}$ and independent 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 any realized market, we denote by μ^w the stable matching preferred by all workers, and least preferred by all firms (see Roth and Sotomayor, 1992). In terms of utilitarian efficiency, μ^w generates the lowest utility for firms across the set of stable matchings.

[20]See the online appendix of Lee (2013) for details.

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

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

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

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

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

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

Proposition 3 is immediate from (2) and (3). For any ε, there exist ε' such that $\phi(c, z^f) < \phi(1, 1) - \varepsilon$ implies 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)**.** By continuity of $\phi(.,.)$ and $\omega(.,.)$, there exists $\varepsilon' > 0$ such that

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

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 10 [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 $\beta(n) = \log n / \log \frac{1}{p}$. If a maximal balanced biclique of this graph has size $B \times B$, then*

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

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

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

Then $\bar{F} \cup \bar{W}$ is a 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 $\phi(c_{ij}, z^f_{ij})$ and $\omega(c_{ij}, z^w_{ij})$ 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 10 then implies that

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

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,

$$P\left(\frac{\sum_{i=1}^n z^f_{i\mu^w(i)}}{n} \leq 1 - \varepsilon\right) \leq n! P\left(\mu \text{ is stable and } \frac{\sum_{i=1}^n z^f_{i\mu(i)}}{n} \leq 1 - \varepsilon\right). \tag{4}$$

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

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

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 $(\tilde{u}^f, \tilde{u}^w)$ 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. We denote by Γ^f the marginal distribution of $\tilde{u}^f_{ij}$ for pairs of (f_i, w_j) who are not matched in 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

$$\begin{aligned}
\hat{u}^f_{ij} &= \Gamma^f(\tilde{u}^f_{ij}) = \Gamma^f\left(\phi(c_{ij}, z^f_{ij})\right) \\
\hat{u}^w_{ij} &= \Gamma^w(\tilde{u}^w_{ij}) = \Gamma^w\left(\omega(c_{ij}, z^w_{ij})\right).
\end{aligned}$$

Remark 11. *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 on $[0,1]$.*

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 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 (See Theorem 2.1 in Esary, Proschan, and Walkup, 1967). Indeed, since both $\Gamma^f(\phi(.,.))$ and $\Gamma^w(\omega(.,.))$ are non-decreasing functions of c_{ij}, z^f_{ij}, and z^w_{ij}, the covariance of the corresponding values of $\hat{u}^f_{ij}$ and $\hat{u}^w_{ij}$ is non-negative.

Thus, we have

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

Lastly, take $\beta > 0$ such that

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

Then,

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

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

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

We therefore obtain that

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

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

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

Proof of convergence.

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

Let $t = \frac{1}{\sqrt{n}}$ and $\alpha = e^t(1-\beta t)$. Note that $1 + \beta z \leq \alpha e^z$ if $-t \leq z \leq 0$. Thus, when $-t \leq y_j \leq 0$, we have $-t \leq -x_i y_j \leq 0$, so $1 - \beta x_i y_j \leq \alpha \exp(-x_i y_j)$. In addition, $1 + \beta z \leq e^{\beta z}$

for any z, so $1-\beta x_i y_j \le e^{-\beta x_i y_j}$.

Therefore,

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

Let

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

Then,

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

We claim that for any $0<r<1$,

$$\alpha+\frac{1}{\beta}\exp(-\beta t s_j) < \exp(n^{-r}) \tag{5}$$

for every sufficiently large n.

As $s>\varepsilon n$, we have $s_j>\varepsilon n-1$. Thus, (5) follows from

$$\exp(n^{-\frac{1}{2}})+\frac{1}{\beta}\exp(-\beta\delta\sqrt{n}) < \exp(n^{-r}) \quad \text{with any} \quad \delta>\varepsilon. \tag{6}$$

The above inequality is equivalent to

$$1+\frac{1}{\beta\exp(-n^{-\frac{1}{2}})\exp(\beta\delta\sqrt{n})} < \exp(n^{-r}).$$

Let $k_n = \exp(-n^{-\frac{1}{2}})\exp(\beta\delta\sqrt{n})$, which increases to infinity as n increases. The above inequality holds for every sufficiently large n because

$$\lim_{n\to\infty}\left(1+\frac{1}{\beta k_n}\right)^{k_n}=\exp(1/\beta)$$

and

$$\lim_{n\to\infty}\left(\exp(n^{-r})\right)^{k_n}=\lim_{n\to\infty}\exp(n^{-r}k_n)=\lim_{n\to\infty}\exp(n^{-r}\exp(-n^{-\frac{1}{2}}+\beta\delta n^{\frac{1}{2}}))=\infty.$$

The claim (5) follows.

Now, we have

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

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

$$\begin{aligned}
\sum_{j=1}^{n}\log\frac{1}{s_j} &= \sum_{j=1}^{n}\left(\left(\log\frac{1}{s}\right)+\int_{s-x_j}^{s}\frac{1}{z}dz\right)\\
&= n\log\frac{1}{s}+\sum_{j=1}^{n}\log\frac{s}{s-x_j}.
\end{aligned}$$

Note that in the last term,

$$\frac{s}{s-x_j}\leq\frac{s}{s-1}\leq\frac{\varepsilon n}{\varepsilon n-1}.$$

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 $c > \frac{1}{\varepsilon}$,

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

Therefore,

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

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.[21]

Lemma 12 [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*

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

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

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

By applying Lemma 12, we get

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

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

Lastly,

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

Thus far, r has been an arbitrary number between 0 and 1. By choosing $r > \frac{5}{7}$, we can guarantee that both of the last two terms converge to 0. This completes the proof. ■

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

We use Theorem 1 in Ashlagi, Kanoria, and Leshno (2013), which implies that for $0 < \lambda \le 1/2$,

$$
P_n \equiv P\left(\sum_{i=1}^{n} R^w_{\mu^f(i)}(\mu^f) \ge \frac{n}{-3\log\lambda}\right)
$$

converges to 1 as $n \to \infty$.

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

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

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. ■

10.5. Overall Efficiency – Proof of Proposition 7. 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 1, the claim follows.

2. We provide a bound on E_n for this environment. For each i, 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

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

We consider two random variable $\tilde{v}_{ij}^f$ 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

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

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

Therefore,

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

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

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

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

In addition, $H(\tilde{v}_{i(k)}^f)$ is equal to the k-th highest value of $\{H(\tilde{v}_{ij}^f)\}_{j=1}^n$, and $H(\tilde{v}_{ij}^f)$ 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 i has arcs to two workers with the highest realized values of $\tilde{v}^f_{ij}$. Similarly, each worker 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 denote maximum utilitarian efficiency achievable by matchings in G. We have

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

Each pair i, j matched in the utilitarian 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 $1 - \frac{\sqrt{2n}}{n+1}$.

Walkup (1979) illustrates that

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

Therefore, we have

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

Combining this bound with the bounds provided by Proposition 2 completes the claim. ∎

References

[1] Agarwal, Nikhil. 2013. "An Empirical Model of the Medical Match," mimeo.

[2] Anshelevich, Elliot, Sanmay Das, and Yonatan Naamad. 2009. "Anarchy, stability, and utopia: Creating better matchings," *Algorithmic Game Theory*, Springer Berlin Heidelberg, 159-170.

[3] Ashlagi, Itai, Yashodhan Kanoria, and Jacob D. Leshno. 2013. "Unbalanced random matching markets," mimeo.

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

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

[6] Boudreau, James W. and Vicki Knoblauch. 2010. "The price of stability in matching markets," University of Connecticut, Working Paper 2010-16.

[7] 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.

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

[9] Che, Yeon-Koo and Olivier Tercieux. 2013. "Efficiency and Stability in Large Matching Markets," mimeo.

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

[11] 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.

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

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

[14] 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.

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

[16] Dizdar, Deniz, and Benny Moldovanu. 2013. "Surplus Division and Efficient Matching," mimeo.

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

[18] Echenique, Federico and Alfred Galichon. 2013. "Ordinal and Cardinal Solution Concepts for Two-sided Matching," mimeo.

[19] Erdos, Paul and Alfred Renyi. 1964. "On Random Matrices," *Publications of the Mathematical Institute of the Hungarian Academy of Sciences,* 3(A), 455-460.

[20] 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.

[21] 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.

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

[23] Immorlica, Nicole and Mohammad Mahdian. 2005. "Marriage, Honesty, and Stability," *SODA*.

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

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

[26] 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*.

[27] Krishna, Vijay and John Morgan. 2012. "Majority Rule and Utilitarian Welfare," mimeo.

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

[29] Lee, SangMok. 2013. "Incentive compatibility of large centralized matching markets,"mimeo.

[30] Liu, Qingmin, and Marek Pycia. 2013. "Ordinal Efficiency, Fairness, and Incentives in Large Markets," mimeo.

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

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

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

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

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

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

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

[38] 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.

[39] 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.

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