informs
http://pubsonline.informs.org/journal/opre/

OPERATIONS RESEARCH
Vol. 66, No. 3, May–June 2018, pp. 620–636
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Convergence of the Core in Assignment Markets

Yash Kanoria,[a] Daniela Saban,[b] Jay Sethuraman[c]

[a] Graduate School of Business, Columbia University, New York, New York 10027; [b] Graduate School of Business, Stanford University, Stanford, California 94305; [c] IEOR Department, Columbia University, New York, New York 10027

Contact: ykanoria@columbia.edu, http://orcid.org/0000-0002-7221-357X (YK); dsaban@stanford.edu (DS); jay@ieor.columbia.edu, http://orcid.org/0000-0002-9985-0683 (JS)

Received: December 22, 2015
Revised: September 12, 2016; May 4, 2017
Accepted: July 31, 2017
Published Online in Articles in Advance: June 6, 2018

Subject Classifications: networks/graphs: matchings; games/group decisions: cooperative; probability: stochastic model applications
Area of Review: Games, Information, and Networks

https://doi.org/10.1287/opre.2017.1691

Abstract. We consider a two-sided assignment market with agent types and a stochastic structure, similar to models used in empirical studies. We characterize the size of the core in such markets. Each agent has a randomly drawn productivity with respect to each type of agent on the other side. The value generated from a match between a pair of agents is the sum of the two productivity terms, each of which depends only on the type (but not the identity) of one of the agents, and a third deterministic term driven by the pair of types. We prove, under reasonable assumptions, that when the number of agent types is kept fixed, the relative size of the core vanishes rapidly as the number of agents grows. Numerical experiments confirm that the core is typically small. Our results provide justification for the typical assumption of a unique core outcome in such markets, which is close to a limit point. Further, our results suggest that, given the market composition, wages are almost uniquely determined in equilibrium.

Funding: The research of J. Sethuraman was supported by National Science Foundation's Division of Civil, Mechanical and Manufacturing Innovation [Grant 1201045].
Supplemental Material: The e-companion is available at https://doi.org/10.1287/opre.2017.1691.

Keywords: assignment markets • matching • transferable utility • core • convergence • uniqueness of equilibrium • wages • stochastic model

1. Introduction

We study bilateral matching markets, such as marriage, labor, and housing markets, where participants form partnerships for mutual benefit. The two classical models of matching markets are the *non-transferable* utility (NTU) model of Gale and Shapley (1962), where payments are not allowed between the agents, and the Shapley–Shubik–Becker *transferable* utility (TU) model of Shapley and Shubik (1971) and Becker (1973), where monetary transfers are allowed between arbitrary sets of agents. For each of these models, a natural solution concept is that of a *stable* outcome in which there is no profitable pairwise deviation. In the TU matching market setting, it is well known that the notion of a stable outcome coincides with that of a competitive equilibrium. Furthermore, a stable outcome is guaranteed to exist in any two-sided market but is generically not unique (Shapley and Shubik 1971). Therefore, a fundamental question arises: How big is the core (i.e., the set of stable matchings) in a given market?[1]

The motivation for our study of core size in TU matching markets is twofold. First, it is of interest to know whether basic market primitives, such as the number of agents and the values of possible matches, are sufficient to determine the market outcome. As an example, can a labor market support higher wages for labor without adding jobs or improving productivity just by moving to a different equilibrium? To take another example, matching platforms such as Upwork, Airbnb, TaskRabbit, and others commonly have a "suggested wage/rate" feature, which is in the nature of a suggestion and not binding (Airbnb 2015). Does this feature help the platform designer (or participants) select between equilibria? A small core size would suggest negative answers to these questions.[2]

Second, several empirical studies in matching markets require or assume a nearly unique stable outcome to facilitate predictions, analyze comparative statics, etc. For instance, Choo and Siow (2006) and Galichon and Salanié (2010) make a continuum limit assumption in TU markets, and in the continuum limit there is a unique equilibrium. However, there is insufficient theoretical basis to justify such an assumption on the size of the core.[3] We ask when such an assumption is justified.

Our main contribution is to bound the size of the core as a function of the market primitives. In particular, we characterize the rate at which the core shrinks as a function of the size of the market. We find that the size of the core in TU matching markets is typically small. This suggests that online matching platforms have limited ability to redistribute welfare across agents, without changing the rules of the market. In addition, our findings justify the continuum limit assumption that is typically made in such markets. Under additional (mild) assumptions, we find that the core rapidly approaches the unique equilibrium in the continuum limit.

Model and Overview of Results. We consider the traditional assignment game model of Shapley and Shubik (1971), consisting of "workers" and "firms," who can each match to at most one agent on the other side. To model the different skills of the workers and the different requirements of the firms, we assume that there are K types of workers and Q types of firms. Matching worker i with firm j generates a value Φ_{ij} (this can be divided between i and j in an arbitrary manner since transfers are allowed), which we model as a sum of two terms: a term $u(\cdot,\cdot)$ that depends only on the *types* of i and j, and a term $\psi_{i,j}$ that represents the "idiosyncratic" contributions of worker i to firm j. In our model, the $u(\cdot,\cdot)$ is assumed to be fixed, but the ψ_{ij} is the sum of two random variables, the "productivity" of worker i with respect to the *type* of firm j and, symmetrically, the "productivity" of firm j with respect to the *type* of worker i. These productivities are assumed to be independently drawn from a distribution (satisfying some mild assumptions) for each (agent, type) pair. Such a generative model for the value of a match has been used in empirical studies of marriage markets, starting with Choo and Siow (2006) (see also Chiappori et al. 2015, Galichon and Salanié 2010).

We study the size of the set of stable outcomes for a random market constructed in this way. Shapley and Shubik (1971) showed that the set of stable outcomes (which is the same as the core) has a lattice structure, and thus has two extreme stable matchings: the worker-optimal stable match, where each worker earns the maximum possible and each firm the minimum possible in any stable matching; and the firm-optimal stable match, which is the symmetric counterpart. Also, in all stable outcomes, the matching between workers and firms must be such that social welfare is maximized. Given these structural properties, our metric for the relative size of the core is quite natural: we consider the difference between the maximum and minimum total utility of workers (equivalently, firms) in the core, scaled by the (maximum) social welfare. We show that if the number of types is fixed, then the core is small and we provide bounds on the size of the core under three different sets of assumptions.

First, we consider productivities drawn from a general distribution satisfying mild conditions.[4] Our first main result (Theorem 1) establishes a small core under some reasonable assumptions on market structure: specifically, the expected core size is $O^*(1/\sqrt[\ell]{n})$ in a market with n agents, and at most $\ell = \max(K, Q)$ types of agents on each side (with ℓ fixed).[5] We show that this bound is essentially tight by constructing a sequence of markets such that the core size is[6] $\Omega(1/\sqrt[\ell]{n})$. Thus, the core shrinks with market size, and this shrinking is faster when there are fewer types of agents. Additionally, we obtain a tighter upper bound in the special case of just one type of firm and more firms than workers (Theorem 2). Our upper bound in this case improves sharply as the number of additional firms m increases; we establish a bound of $O^*(1/(n^{1/\ell}m^{1-1/\ell}))$, where ℓ is the number of worker types.

Second, we again consider a fixed set of types, but assume that productivities are drawn from a distribution with unbounded support (Theorem 3), and obtain stronger results. We show an $O^*(1/n)$ bound on core size in this case. We also establish that the core solutions in the finite market converge to the unique equilibrium in the continuum limit market, bounding the distance between any core solution and the limit equilibrium by $O^*(1/\sqrt{n})$. In particular, this bounds the convergence rate of the (scaled) number of matched pairs belonging to each pair of types. We emphasize that this is the empirically observed quantity in many settings (transfers are often not observed, e.g., Choo and Siow 2006). This is also useful in bounding the error incurred when match utilities in empirical studies are estimated based on a continuum limit assumption. Both of our bounds are again tight.

To supplement our theoretical findings, we conduct computational experiments (Section 5). We run simulations with a variety of distributions for the idiosyncratic productivity terms and also go beyond our theoretical development by allowing the number of agent types to grow. In a broad range of settings, we find that the core is small, even in relatively small markets; the only exception we find to a small core is the case of productivities following a Pareto distribution, and a large number of agent types. In summary, our theoretical results and computational experiments strongly suggest that the core is small in practically relevant settings. In addition, our experiments also show that the (scaled) number of matched pairs belonging to any pair of types rapidly converges to the appropriate limiting value.

We conclude this overview with an observation that we found helpful in reasoning about the model and our proof techniques. Our model has the following property: there is a "price" associated with each (worker type, firm type) pair, such that for every matched pair of agents of these types, the utility of each agent is her productivity (with respect to the type on the other side) "corrected" additively on both sides of the market (in opposite directions) by the price. Each of our bounds on core size is proved by showing uniform bounds on variation in type-pair prices across core allocations. A key component of our analysis is to relate the combinatorial structure of the core to order statistics of certain independent, identically distributed (i.i.d.) random variables (RVs). These RVs are one-dimensional projections of point processes in (particular subregions of) the unit hypercube, where the point processes correspond to the market realization. An analytical challenge that we face is that the relevant projections as well as the relevant order statistics are themselves a random function of the market realization. We overcome this via appropriate union bounds. In our proof of Theorem 3, we further use an order-theoretic approach to control the set of candidate core solutions (this also yields a limiting characterization of core solutions as market size grows).

Related Literature. Most of the related literature focuses on the NTU model of Gale and Shapley (1962). Building on that model, a number of papers establish a small core under various assumptions such as short preference lists or many unmatched agents (Immorlica and Mahdian 2005, Kojima and Pathak 2009, Kojima et al. 2013, Menzel 2015, Peski 2015), and strongly correlated preferences (Holzman and Samet 2014, Azevedo and Leshno 2016).[7] In a recent paper, Ashlagi et al. (2017) show that in a random NTU matching market with long lists and uncorrelated preferences, even a slight imbalance results in a significant advantage for the short side of the market and that there is approximately a unique stable matching.[8] Furthermore, Ashlagi et al. (2017) find the near uniqueness of the stable matching to be robust to varying correlations in preferences and other features, suggesting that a small core may be generic in NTU matching markets.

There is an extensive literature on large assignment games that extends the many structural properties established by Shapley and Shubik for finite assignment games to a setting in which the agents form a continuum; see, for example, Gretsky et al. (1992, 1999). Those papers also show convergence of large finite markets to the continuum limit, including that the core shrinks to a point. However, unlike in our model, they model the productivity of each partnership as a deterministic function of the pair of types, with the only randomness being in the number of agents of each type. The work on assignment games that is most closely related to our work is a recent preprint of Hassidim and Romm (2015): in their model, all workers (firms) are a priori identical, and the value of matching worker i to firm j is a random draw from a bounded distribution, independently for every pair (i, j). For such a model, they establish an approximate "law of one price"—i.e., that workers are paid approximately identical salaries in any core allocation, and that the long side gets almost none of the surplus in unbalanced markets. By contrast, we work with multiple types of workers and firms, where the value of a match depends on the types of each agent, and on random variables that depend on the identity of one of the agents and the *type* (but not the identity) of the other agent. We believe that our model better captures features of real markets, as it allows for correlations in preferences across types.

There has been recent work in the operations literature characterizing equilibria in matching markets. Nguyen (2015) considers bargaining in a network en route to formation of coalitions (a generalization of matching) and characterizes the stationary equilibria of the game. Alaei et al. (2016) generalize the Shapley–Shubik model to the case of utilities that are not necessarily quasilinear in payments. They characterize equilibria and provide an algorithm for efficiently computing the extreme equilibria, and their work leads to a mechanism for ad auctions that has good properties even when there is inconsistency in click-through rate estimates. Overall, the goal of this line of work, including the present paper, is to obtain a refined understanding of equilibria to enable the design of better marketplaces for matching.

2. Model Formulation

2.1. The Standard Model

We consider a transferable utility matching market (or simply matching market) with a set of workers $\mathcal{L}$ and a set of firms $\mathcal{E}$. To keep things concrete, we cast the problem in terms of a labor market, but other interpretations are possible. There are $n_{\mathcal{L}}$ workers and $n_{\mathcal{E}}$ firms; we let $n := n_{\mathcal{L}} + n_{\mathcal{E}}$ denote the size of the market—i.e., the total number of agents. Firm i employing worker j creates a benefit $\Phi(i, j)$ that can be freely divided between i and j using monetary transfers, as all workers and firms are assumed to have preferences that are linear in money. We assume a one-to-one market so that each firm can employ at most one worker and each worker can work for at most one firm.

A *matching* M is a subset of $\mathcal{L} \times \mathcal{E}$. If $(i, j) \in M$, worker i is said to be matched to firm j in M; we sometimes abuse notation and say $M(i) = j$ or $M(j) = i$ whenever $(i, j) \in M$. If no element of M involves worker i, then i is said to be *unmatched*; likewise for the firms. An *outcome* for a matching market is a pair $(M, \boldsymbol{\gamma})$, where M is a matching and γ is a payoff vector such that $\gamma_i + \gamma_j = \Phi(i, j)$ for each $(i, j) \in M$, and $\gamma_a = 0$ for each unmatched agent a. That is, the vector $\boldsymbol{\gamma}$ indicates how the value of a match is divided between the agents involved in the match.

In this paper, we shall be concerned with outcomes that are in the *core*: in each such outcome, the division of the benefit to the agents is such that no subset of agents can produce a greater value among themselves than the sum of their gammas. Thus, a core outcome is one in which no coalition of agents has an incentive to deviate. Shapley and Shubik (1971) show that for this model, ruling out deviations by every worker–firm pair is sufficient for an outcome to be in the core. Such outcomes are said to be *stable*. That is, an outcome is stable if no agent prefers not to participate in the matching (because of a negative payoff), and if there is no *blocking pair* of agents who can both do better by matching with each other (because the value they generate by matching with each other exceeds the sum of their current payoffs). Thus, the stability condition can be mathematically described as $\gamma_i + \gamma_j \geqslant \Phi(i,j)$ for all $i \in \mathscr{L}$ and $j \in \mathscr{E}$, and $\gamma_a \geqslant 0$ for all $a \in \mathscr{L} \cup \mathscr{E}$. Furthermore, it is known that the matching M in any stable outcome must maximize total benefit; that is, $M \in \arg\max_{M' \in \mathscr{M}} \sum_{(i,j)\in M'} \Phi(i,j)$, where $\mathscr{M}$ is the set of all possible matchings. For a given matching M, we refer to $\sum_{(i,j)\in M} \Phi(i,j)$ as the *weight* of the matching.[9]

2.2. Our Model

We impose additional structure on the sets of workers and firms, and on the match benefit. We assume that the workers and firms are partitioned into "types," which are devices to group agents with similar characteristics. As an example, consider a labor market consisting of Ph.D. candidates (workers) and firms. Suppose that the type of a worker is the university and program she is graduating from, while the type of a firm corresponds to the industry it belongs to (e.g., tech, consulting, finance).[10] It is then natural to think that some portion of Φ will be determined by just the types of the matched agents alone: perhaps the core curriculum at one program is more relevant to the tech industry than to management consulting, while the opposite might be true at another program. Thus, we assume that matching two agents of a certain type will yield some baseline value. Additionally, each worker also has her own characteristics, which might make her more or less suitable for a particular industry, relative to other workers of the same type. Similarly, a particular firm's characteristics may be more or less attractive to workers of a given type, relative to other firms in the same industry.[11] These considerations lead us to formulate the following model.

The workers are partitioned into K types, and the firms are partitioned into Q types; the set of workers' and firms' types are denoted by $\mathscr{T}_{\mathscr{L}}$ and $\mathscr{T}_{\mathscr{E}}$, respectively. Let $\mathscr{T} = \mathscr{T}_{\mathscr{L}} \times \mathscr{T}_{\mathscr{E}}$ denote the set of pairs of types. For a given type $t \in \mathscr{T}_{\mathscr{L}} \cup \mathscr{T}_{\mathscr{E}}$, we denote by n_t the number of agents of type t. For each agent $a \in \mathscr{L} \cup \mathscr{E}$, let $\tau(a)$ denote the type of agent a; given a type t and an agent a, we say that $a \in t$ if $\tau(a) = t$. To capture the correlation in preferences between agents of the same type as well as individual variations, we model the value $\Phi(i,j)$ of matching i and j as the sum of two components: a type-type utility term $u(\tau(i),\tau(j))$, and a match-specific term $\psi_{i,j}^{\tau(i),\tau(j)}$. The utility term $u(\tau(i),\tau(j))$ depends only on the agents' types and allows us to express how well suited, in general, a worker of type $\tau(i)$ is to work in a firm of type $\tau(j)$. The match-specific term $\psi_{i,j}^{\tau(i),\tau(j)}$ potentially depends on both the identity of the agents and their types, and captures how useful are i's individual skills to perform job j. We further assume that $\Phi(i,j)$ is additively separable as follows.

Assumption (Separability). *We assume* $\Phi(i,j) = u(\tau(i),\tau(j)) + \epsilon_j^{\tau(i)} + \eta_i^{\tau(j)}$.

The separability assumption states that the match-specific component, $\psi_{i,j}^{\tau(i),\tau(j)}$, is further additively separable into two terms, $\epsilon_j^{\tau(i)}$ and $\eta_i^{\tau(j)}$; *each term depends on the identity of one agent and only the type of the other agent*. In particular, for any fixed employer j and two distinct workers $i, i' \in \mathscr{L}$, we have $\epsilon_j^{\tau(i)} = \epsilon_j^{\tau(i')}$ whenever $\tau(i) = \tau(i')$, as the term ϵ depends only on the *type* of the agents in $\mathscr{L}$. Analogously, the term η depends on the *individual* worker $i \in \mathscr{L}$ but only on the *type* of the firm $j \in \mathscr{E}$.

To illustrate the rationale behind the separability assumption, consider our previous example of a market consisting of Ph.D. candidates and firms. Students in a particular program have their own idiosyncratic skills; these are captured by the term $\eta_i^{\tau(j)}$, which indicates how an individual student i is valued by a certain type of firm $\tau(j)$, relative to other students of type $\tau(i)$. Note that all firms of type $\tau(j)$ value candidate i equally. Similarly, the productivity terms associated with firms ($\epsilon_j^{\tau(i)}$) capture how attractive that firm is to students emerging from a given program, relative to the other firms of type $\tau(j)$. Throughout the paper, we refer to the term $u(\tau(i),\tau(j))$ as the *type-type compatibility*, and to the terms η and ϵ as the *idiosyncratic productivity* terms.

Unless otherwise stated, we model the term $u(\tau(i),\tau(j))$ as a fixed constant, whereas the ϵ and η terms are modeled as random variables, independent across agent type pairs. We assume that the distributions of ϵ and η terms are supported on an interval (possibly unbounded), with density continuous and positive everywhere in the support. The continuum limit of such a model was first introduced by Choo and Siow (2006) to empirically estimate certain structural features of marriage markets. An attractive feature of this model is that it allows for reasonable correlation in agents' preferences, and also heterogeneity and idiosyncratic variation via the random productivities, while still remaining tractable.[12] The model has since been employed in other contexts as well (Galichon and Salanié 2010, Chiappori et al. 2015, Fox 2018).

2.3. Preliminaries

We now state some preliminary observations on the structure of the core under the separability assumption. For a matching M, let $M(t)$ be the set of agents who are matched to an agent of type $t \in \mathscr{T}_{\mathscr{L}} \cup \mathscr{T}_{\mathscr{E}}$, and let U be the unmatched agents. Observe that the weight of a matching depends only on the *type* of the partner that each agent is matched to, and so the maximum-weight matching is typically not unique. Nevertheless, as the idiosyncratic productivities are randomly drawn from non-atomic distributions, the following holds.

Observation 1. With probability 1, the sets U and $M(t)$ for $t \in \mathcal{T}_{\mathcal{L}} \cup \mathcal{T}_{\mathcal{E}}$ are the same under all maximum-weight matchings. We call this the type-matching.

We take this as the *definition* of the sets $M(t)$ and U for our purposes, since in any stable outcome, the matching is a maximum-weight matching; see Endnote 9. We next show that the payoffs can be expressed more conveniently.

Proposition 1. *For each $i \in \mathcal{L}$ and each type $q \in \mathcal{T}_{\mathcal{E}}$, let $\tilde{\eta}_i^q = u(\tau(i), q) + \eta_i^q$. Any core outcome $(M, \boldsymbol{\gamma})$ corresponds to a vector $\boldsymbol{\alpha} = \{\alpha_{kq}\} \in \mathbb{R}^{K \times Q'}$ such that*

- $\gamma_i = \tilde{\eta}_i^q - \alpha_{kq}$ *for all $i \in \mathcal{L}$ such that $\tau(i) = k$ and $i \in M(q)$.*
- $\gamma_j = \epsilon_j^k + \alpha_{kq}$ *for all $j \in \mathcal{E}$ such that $\tau(j) = q$ and $j \in M(k)$.*
- $\gamma_i = 0$ *for all $i \in U$.*

Proposition 1 follows directly from stability and formalizes the existence of a *single price* for every type-pair (k, q). In particular, if worker i and firm j of types k and q, respectively, are matched, the worker payoff is the type-type compatibility plus her own productivity with a type q firm minus the price α_{kq}; on the other hand, the firm's payoff is its productivity with a type k worker plus α_{kq}. As a consequence, if firms j' and j'' are of the same type and are matched to the workers of the same type, the difference in their utilities must be equal to the difference in their productivities with respect to that type of worker.

By Proposition 1, any core outcome can be thought of in terms of the maximum-weight matching M and the vector $\boldsymbol{\alpha}$. The following proposition states necessary and sufficient conditions for $(M, \boldsymbol{\alpha})$ to be a core outcome. (The maximum over an empty set is defined as $-\infty$.)

Proposition 2. *The following conditions are necessary and sufficient for $(M, \boldsymbol{\alpha})$ to be a core solution:*

(ST) *For every pair of types (k, q), $(k', q') \in \mathcal{T}$:*

$$\min_{i \in k' \cap M(q')} (\tilde{\eta}_i^{q'} - \tilde{\eta}_i^{q}) + \min_{j \in q \cap M(k)} (\epsilon_j^k - \epsilon_j^{k'}) \geqslant \alpha_{k'q'} - \alpha_{kq} \geqslant \max_{i \in k \cap M(q)} (\tilde{\eta}_i^{q'} - \tilde{\eta}_i^{q}) + \max_{j \in q' \cap M(k')} (\epsilon_j^k - \epsilon_j^{k'}).$$

(IR) *For every pair of types $(k, q) \in \mathcal{T}$:*

$$\min_{j \in q \cap M(k)} \epsilon_j^k \geqslant -\alpha_{kq} \geqslant \max_{j \in q \cap U} \epsilon_j^k, \quad \text{and}$$
$$\min_{i \in k \cap M(q)} \tilde{\eta}_i^q \geqslant \alpha_{kq} \geqslant \max_{i \in k \cap U} \tilde{\eta}_i^q.$$

We refer the reader to proposition 1 of Chiappori et al. (2015) for a proof. The first set of conditions (ST) follow from reexpressing the condition imposing the nonexistence of a blocking pair of matched agents; i.e., for every $(i, j) \in M$, we have $\gamma_i + \gamma_{j'} \leqslant \Phi(i, j)$ for all $j' \in \mathcal{E}$ and $\gamma_{i'} + \gamma_j \leqslant \Phi(i, j)$ for all $i' \in \mathcal{L}$. In particular, if the condition for pairs (k, q) and (k', q') fails to hold, it means that we can find either (1) a worker i of type k who is matched to a firm in q ($i \in k \cap M(q)$) and a firm j' of type q' matched to an agent in k' such that i and j' would rather be matched together, or (2) a worker i' of type k' who is matched to a firm in q' and a firm j in q matched to an agent in k such that i' and j prefer to be matched together. By considering the difference between two prices, these conditions state how much a type-pair price can vary relative to another type-pair price; therefore, these conditions impose bounds on the *relative variation* of prices.

The second set of conditions (IR) follow from combining two facts. The payoffs of matched agents are nonnegative, as otherwise they would rather be u-matched (see Proposition 1); this implies the left inequalities. The nonexistence of a blocking pair involving an unmatched agent implies the right inequalities. In particular, note that the right inequality for a type-pair price α_{kq} is meaningful only if an unmatched agent of type k or q exists. The (IR) conditions constitute bounds on the *absolute variation* of type-pair prices. Note that Proposition 2 highlights the lattice structure of the set of core solutions (Shapley and Shubik 1971).

We conclude with a definition of the *size of the core*, denoted by $\mathcal{C}$. We define $\mathcal{C}$ as the difference between the maximum and minimum total payoff of workers (or firms) among core outcomes, scaled by the overall social welfare (the total weight of M) in any core outcome. This can be equivalently stated in terms of the vector $\boldsymbol{\alpha}$. For each pair of types $(k, q) \in \mathcal{T}$, let $\alpha_{kq}^{\max}$ and $\alpha_{kq}^{\min}$ be the maximum and minimum possible values of α_{kq} in the core. Defining $\alpha^{\max} = (\alpha_{kq}^{\max})_{k \in \mathcal{T}_{\mathcal{L}}, q \in \mathcal{T}_{\mathcal{E}}}$, note that $(M, \alpha^{\max})$ is in the core and constitutes the firm-optimal stable outcome. Similarly, $(M, \alpha^{\min})$ is the worker-optimal stable outcome, where the definition of $\alpha^{\min}$ is analogous to that of $\alpha^{\max}$. The size of the core is defined in terms of $\alpha^{\max}$ and $\alpha^{\min}$ as follows.

Definition 1 (Size of the Core). Let M be the unique maximum-weight type-matching. For each pair of types $(k, q) \in \mathcal{T}$, let $N(k, q)$ denote the number of matches between agents of type k and agents of type q. Then, the size of the core is denoted by $\mathcal{C}$ and is defined as

$$\mathcal{C} = \frac{\sum_k \sum_q N(k, q) |\alpha_{kq}^{\max} - \alpha_{kq}^{\min}|}{\text{weight}(M)}.$$

It is worth noting that $\mathcal{C}$ is always between 0 and 1. This is because weight (M) is the total surplus produced by the match, while $|\alpha_{kq}^{\max} - \alpha_{kq}^{\min}|$ captures how much the surplus kept by one side can vary, which in turn is scaled by the number of matches involving agents of such a type-pair. The stability conditions imply that, for each match, the variation in surplus kept by each side must always be less than the value of the match.

We conclude this section by noting that the results in the paper can be further generalized.

Remark 1. All of our analysis and bounds extend to a model with taxation of transfers (Jaffe and Kominers 2014), which captures commissions charged by matching platforms. Think of the utility of each agent in a matched pair (i, j) as arising from both the sum of a base utility from participating in the match and the amount received/paid in a transfer payment between the matched partners. Suppose the base utility is $-c(\tau(i), \tau(j))$ for the worker, whereas the remaining match value (i.e., $\Phi(i, j) + c(\tau(i), \tau(j))$) accrues as the base utility to the firm (thus, all of the stochastic variation due to idiosyncratic productivities is in the base utility that accrues to the firm; the base utility/cost of the worker depends only on the pair of types). Suppose that the matching platform collects a fraction $\lambda \in [0, 1)$ of the transferred amount as taxes/commission. Following Jaffe and Kominers (2014), we assume that the base utility is (weakly) negative for workers—i.e., $c(k, q) \geqslant 0$ for all $k \in \mathcal{T}_{\mathscr{L}}, q \in \mathcal{T}_{\mathscr{E}}$, capturing the base cost to a worker of type k of matching with a firm of type q. This ensures that transfers are only from firms to workers, making this a so-called "wage market." Consider any such market $\mathscr{M}_{\text{tax}}$. There is a one-to-one correspondence between the core solutions in $\mathscr{M}_{\text{tax}}$ and those in a market $\tilde{\mathscr{M}}$ with no taxes, the same set of agents, and modified values from matching pairs of agents[13]

$$\tilde{\Phi}(i, j) = \Phi(i, j) - \frac{\lambda c(\tau(i), \tau(j))}{1 - \lambda} \quad \forall\, i \in \mathscr{L}, j \in \mathscr{E}.$$

The correspondence between core solutions of $\mathscr{M}_{\text{tax}}$ and $\tilde{\mathscr{M}}$ is as follows: the matching is identical in both markets, firm payoffs are identical in both markets, and the payoffs of workers are a factor $(1 - \lambda)$ smaller in $\mathscr{M}_{\text{tax}}$ relative to $\tilde{\mathscr{M}}$. It follows that the size of the core[14] in $\mathscr{M}_{\text{tax}}$ is identical to that in $\tilde{\mathscr{M}}$.

3. Results

We now turn our attention to the main objective of the paper, which is to understand how the size of the core scales as the market grows. Given the stochastic nature of our model, the size of the core $\mathscr{C}$ is itself a random variable. Therefore, the rest of this section is devoted to studying how the expected/typical value of $\mathscr{C}$ depends on the characteristics of the market. Throughout this section, we keep the number of agent types as well as $u(\cdot, \cdot)$ fixed and allow the number of agents to grow.

In markets with a finite number of agents, there is always a finite number of stability constraints. Thus, it is generically possible to marginally modify some payoffs in a core solution without violating stability, and therefore, the size of the core is strictly positive with probability one (Shapley and Shubik 1971).[15] However, as the size of the market increases (while the number of agent types stays the same), the number of stability constraints also increases, limiting the possible perturbations to the payoffs (cf. Proposition 2). Hence, one would expect the set of core vectors α to shrink as the market size increases. In this section, we characterize the *rate* at which the size of the core shrinks as the size of the market increases, as a function of the market primitives. In Section 3.1, we present results for the general case of productivities drawn independently from any bounded or unbounded distribution satisfying mild conditions. In Section 3.2, we present stronger results for distributions with (two-sided) unbounded support, including convergence of the core to a particular limiting point. The appendices, which have complete proofs of some results as well as some useful examples, appear in an online companion to the paper.

3.1. Idiosyncratic Productivities with a General Distribution

Throughout this subsection, we consider general distributions F for the idiosyncratic productivities, which satisfy the condition that F is supported on an interval of the form $[C_l, C_u]$ or $(-\infty, C_u]$ or $[C_l, \infty)$ for some finite C_l, C_u, and its density f is continuous and positive everywhere in the support.

One Type on Each Side of the Market. We start by considering the simple case of markets with one type on each side—that is, $K = Q = 1$. Given that there is only one type of agent on each side, the deterministic type-type utility term $u = u(\tau(i), \tau(j))$ will be the same for all matches, regardless of the identity of the agents. The value of a match between agents $i \in \mathscr{L}$ and $j \in \mathscr{E}$ is $\Phi(i, j) = u + \eta_i + \epsilon_j$, where one may think of η_i as representing the quality of worker i and ϵ_j as representing the quality of firm j. Suppose that $u > 0$ is a fixed constant.

Definition. We write $f(n) = O^*(g(n))$ if there exists $r < \infty$ such that $f(n) \leqslant r(\log n)^r g(n)$ for all $n \geqslant 2$.

Definition. A sequence of events $\mathscr{E}_n$ occurs with high probability if $\lim_{n\to\infty} \Pr(\mathscr{E}_n) = 1$.

Remark 2. The size of the core depends on the number of agents on each side of the market.

- In a balanced market (i.e., when $n_{\mathscr{L}} = n_{\mathscr{E}}$), the core is large if F is supported on positive values. For instance, if F is Uniform$(0, 1)$, the above market has $\mathscr{C} \geqslant u/(u + 2)$ with probability 1. In particular, we have $\mathrm{E}[\mathscr{C}] = \Omega(1)$.
- In any unbalanced market (i.e., $n_{\mathscr{L}} \neq n_{\mathscr{E}}$), there exists $f(n) = O^*(1/n)$ such that, with high probability, we have $\mathscr{C} \leqslant f(n)$. Also, $\mathrm{E}[\mathscr{C}] = \Omega(1/n)$.

To see why, note that in a balanced market where F is supported on positive values, all agents will be matched in a stable solution. By Proposition 1, we can describe the size of the core in terms of a single parameter α; by Proposition 2, the core consists of all $\alpha \in [-\min_j \epsilon_j, u + \min_i \eta_i]$. In other words, the value u that is part of $\Phi(i,j)$ for each (i,j) can be split in an arbitrary fashion between employers and workers. On the other hand, in any unbalanced market (i.e., $n_{\mathcal{L}} \neq n_{\mathcal{E}}$), the core is small and rapidly shrinks with market size. It turns out that the short side of the market has a significant advantage: if there are fewer workers than firms, the price α is always negative in the core, and is bounded as $\min_{j\in M} \epsilon_j \geqslant -\alpha \geqslant \max_{j \in U} \epsilon_j$. This observation agrees with some of the existing results for the non-transferable utility setting (Ashlagi et al. 2017).

Multiple Types on Both Sides of the Market. We now consider the general case of K types of labor and Q types of employers. The following condition generalizes the imbalance condition to the case of multiple types. The idea is to get rid of the cases that, for certain values of deterministic type-type utilities $u(\cdot,\cdot)$, may resemble a balanced problem.[16]

Assumption 1 (Generalized Imbalance). *For every pair of subsets of types $\mathcal{S} \subseteq \mathcal{T}_{\mathcal{L}}$ and $\mathcal{S}' \subseteq \mathcal{T}_{\mathcal{E}}$, we have $\sum_{t\in\mathcal{S}} n_t \neq \sum_{t\in\mathcal{S}'} n_t$. In words, there is no subset of types for which the induced submarket is balanced.*

We highlight that in our setting with fixed K and Q and growing n, "most" markets[17] satisfy Assumption 1.

Throughout this section, we allow the number of agents to grow, while keeping the number of types fixed. We limit the way in which the market grows by assuming that there is at least a linear number of agents of each type.

Assumption 2. *There exists $C > 0$ such that for all types $t \in \mathcal{T}_{\mathcal{L}} \cup \mathcal{T}_{\mathcal{E}}$, we have $n_t \geqslant Cn$.*

Under Assumption 2, each type has a comparable number of agents and no type vanishes as the market increases. We now present our main theorem.

Theorem 1. *Consider $K \geqslant 1$ types of labor, and $Q \geqslant 1$ types of employers. Let the idiosyncratic productivities be drawn i.i.d. from any fixed distribution F that is supported on an interval of the form $[C_l, C_u]$ or $(-\infty, C_u]$ or $[C_l, \infty)$ or $(-\infty,\infty)$ where $C_l > -\infty$ and $C_u < \infty$, and whose density f is strictly positive and continuous everywhere in the support.*[18,19] *Under Assumption 1 and Assumption 2, for a market with n agents, we have that*

- (Upper Bound) *There exists $f(n) = O^*(n^{-1/\max(K,Q)})$ such that, with high probability, we have*

$$\max_{\substack{(k,q)\in\mathcal{T} \\ N(k,q)>0}} |\alpha_{kq}^{\max} - \alpha_{kq}^{\min}| \leqslant f(n).$$

Also, we have $\mathrm{E}[\mathscr{C}] \leqslant f(n)$.

- (Lower Bound) *There exists a sequence of markets with K types of labor and Q types of employers such that* $\mathrm{E}[\mathscr{C}] = \Omega(n^{-1/\max(K,Q)})$.

In words, our main result says that under reasonable conditions, with high probability, the variation in the type-pair prices is uniformly bounded by $O^*(n^{-1/\max(K,Q)})$, which vanishes as n grows. The same bound holds for $\mathrm{E}[\mathscr{C}]$. In addition, this bound is tight in the worst case. Thus, the core size shrinks to zero as the market grows larger, at a rate that is faster (in the worst case) if there are fewer types of agents. The proof of Theorem 1 is sketched in Section 4.1; the complete proof is in Online Appendices C and D.

One implication of Theorem 1 is the following: as the type-matching is the same across all core solutions, the bound on the type-pair prices implies that the maximum difference between the utilities of *any* agent in any two core solutions vanishes at a rate of $n^{-1/\max(K,Q)}$, as the size of the market increases. Therefore, the fraction of welfare that can flexibly move from workers to firms (and vice versa) in core solutions vanishes. Hence, Theorem 1 shows that the vector of prices that support a competitive equilibrium (a core solution) is approximately unique, provided that the market is big enough.[20] As a result, the payoff of an agent is roughly the same in all stable solutions.

The upper bound in Theorem 1 can be improved if further constraints are imposed on the number of types and the imbalance. As an illuminating example, we show that for multiple worker types, a single employer type, and a type with more employers than workers, the size of the core can be bounded above by a function that depends on both the size of the market and the size of the imbalance in the market.

Theorem 2. *Let idiosyncratic productivities be drawn i.i.d. from any fixed distribution F that is supported on an interval of the form $[0, C_u]$ or $[0,\infty)$ where $C_u < \infty$, and whose density f is strictly positive and continuous everywhere in the support.*[21] *In addition, suppose that $u(k,1) \geqslant 0$ for all $k \in \mathcal{T}_{\mathcal{L}}$. Consider the setting in which $K \geqslant 2$, $Q = 1$, $n_{\mathcal{E}} > n_{\mathcal{L}}$, and let $m = n_{\mathcal{E}} - n_{\mathcal{L}}$. Under Assumption 2, we have, with high probability, that $\mathscr{C} \leqslant O^*(n^{-1/K} m^{-(K-1)/K})$. Also,* $\mathrm{E}[\mathscr{C}] \leqslant O^*(n^{-1/K} m^{-(K-1)/K})$.

For $m = O^*(1)$, the bound in Theorem 2 matches that in Theorem 1. However, the bound here becomes tighter as the imbalance m grows. In fact, for $m = \Theta(n)$, the core size is bounded as $O^*(1/n)$. It is noteworthy that, under a linear imbalance, the scaling behavior does not depend on the number of worker types.

A formal proof of Theorem 2 can be found in Online Appendix E. The idea behind the proof is to use the unmatched agents and condition (IR) in Proposition 2 (for the employers) to control absolute variation in one of the α's, and this control improves as m grows.

We separately control the relative variation of the α's in the core using condition (ST) in Proposition 2 under Assumption 2. Combining these, we obtain the stated bound on $\mathscr{C}$.

3.2. Idiosyncratic Productivities with Unbounded Support

For our second main result, we require the idiosyncratic terms (i.e., $\epsilon_j^{\tau(i)}, \eta_i^{\tau(j)}$) to be drawn independently from a distribution having *unbounded* support, with positive density everywhere (for instance, Gumbel and Gaussian distributions will be covered by our analysis here). While the distribution may depend on $\tau(i)$ and $\tau(j)$, to ease the burden of notation, we assume that all of these terms are drawn i.i.d. from a single distribution whose density $f\colon (-\infty,\infty) \to \mathbb{R}_+$ is positive everywhere and continuous.[22] Let $F\colon (-\infty,\infty) \to [0,1]$ be the corresponding cumulative distribution. Recall that $N(k,q)$ is the number of agents of type k matched to agents of type q, under the unique core type-matching M.

In the following theorem, we fix the fraction of agents of each type and scale the number of agents.[23] We show that the core prices and the fraction of matches corresponding to a type-pair each converge to a unique limit, and we further provide a bound on the rate of convergence. The second part of the theorem leverages these limit characterizations to bound the size of the core.

Theorem 3. *Fix K and Q and the distribution F. Also fix the fraction $\rho_t > 0$ of each agent type $t \in \mathcal{T}_{\mathcal{L}} \cup \mathcal{T}_{\mathcal{E}}$; that is, t can be a type of worker or a type of employer. Consider a market with n agents that includes $n\rho_t$ agents of each type[24] t, with idiosyncratic productivities drawn i.i.d. from the distribution F. We obtain the following bounds as a function of n.*

- *Limit characterization of α. There exists $\alpha^* = \alpha^*(K, Q, \rho, F)$ and corresponding $(\nu^*_{kq})_{(k,q)\in\mathcal{T}}$ such that as n to ∞, we have that both $\alpha^{\max}$ and $\alpha^{\min}$ converge (in probability and almost surely) to α^*, and $N(k,q)/n \to \nu^*_{kq}$ for all $(k,q) \in \mathcal{T}$. In fact, with high probability, all core outcomes (M, α) satisfy*

$$\|\alpha - \alpha^*\| \leqslant O^*(1/\sqrt{n}), \tag{1}$$

$$|N(k,q)/n - \nu^*_{kq}| \leqslant O^*(1/\sqrt{n}), \quad \text{for all } (k,q) \in \mathcal{T}. \tag{2}$$

- *Size of the core. With high probability, the size of the core is bounded as*

$$\mathscr{C} \leqslant O^*(1/n). \tag{3}$$

See Online Appendix F for a statement of Theorem 3 with additional technical details and a proof. The first part of our result (the limit characterization of α) is analogous to the one obtained in Azevedo and Leshno (2016) for NTU markets. Roughly, the limit point α^* is the unique price vector that clears the limit market, in the sense that for each type-pair (k,q), the "demand" of worker type k for employer type q is equal to the "supply" of employer type q for worker type k (both demand and supply take a value of ν^*_{kq}). Our analysis takes advantage of the fact that the limit market satisfies a strong gross substitutes condition on both sides (leading to uniqueness of the equilibrium price, and allowing it to be computed via tatonnement), whereas the finite n market satisfies a weak gross substitutes condition on both sides. We use an order-theoretic approach along with convergence of the empirical distribution of (type, productivity vector) among agents to the limiting distribution, to relate the finite n market to the limit market. This connection leads to the bounds that we obtain.

Theorem 3 has several noteworthy features:

1. The scaling of the bounds with n does not depend on the number of types K and Q. The stronger assumption of F having full support on $(-\infty,\infty)$ yields a stronger upper bound on core size than that in Theorem 1, without the need for Assumption 1 (generalized imbalance).[25]
2. The bound on core size (the second part of the theorem) is tight if F has a first moment, as demonstrated by Example 1 (Online Appendix A). The idea is that $1/n$ is the typical (expected) gap between consecutive order statistics of the idiosyncratic productivities.
3. The first part of the theorem states that the core converges to a limit point (α^*, ν^*), and provides a bound of $O^*(1/\sqrt{n})$ on the rate of convergence. Again, Example 1 (Online Appendix A) shows that this is tight. The idea is that the actual fraction of agents of a particular type who satisfy some fixed conditions on their idiosyncratic productivities has stochastic variations of order $1/\sqrt{n}$.

Past works that establish a small core and show convergence to a limit point (e.g., Gretsky et al. 1999, Kamecke 1992) are superficially similar. However, the novel feature of (random) idiosyncratic productivity terms in our model (as well as bounds on the rates of convergence) makes Theorem 3 significantly more powerful than previous results.

4. Theorem 3 provides a bound on the estimation error in empirical studies of matching. Often in empirical studies, only the $N(k,q)$'s are observed (cardinal utilities or transfers are typically unobserved), and the goal is to estimate type-type utilities (the $u(k,q)$'s) under suitable assumptions on the idiosyncratic productivities (e.g., that they follow a standard Gumbel distribution). The estimation is typically performed using a continuum limit assumption, as a result of which there is a tractable one-to-one mapping between observed quantities $(N(k,q))_{(k,q)\in\mathcal{T}}$, $(\rho_t)_{t\in\mathcal{T}_{\mathcal{L}}\cup\mathcal{T}_{\mathcal{E}}}$, and the estimated quantities $(u(k,q))_{(k,q)\in\mathcal{T}_{\mathcal{L}}\times\mathcal{T}_{\mathcal{E}}}$; cf. Choo and Siow (2006, equation (11)). Moreover, this mapping is

well behaved (e.g., it is Lipschitz continuous in both directions at interior points), and so we can immediately use Equation (2) to obtain a bound of $O^*(1/\sqrt{n})$ on the error in the estimate of $(u(k,q))_{(k,q)\in\mathcal{T}}$.

4. Overview of the Proof of Theorem 1

We now present an overview of our proof of Theorem 1. We first discuss the key steps in establishing the upper bound (the complete proof can be found in Online Appendix C), and then sketch the proof of the lower bound in Section 4.2 (completed in Online Appendix D). Throughout this section, we assume that there is a unique maximum-weight type-matching M as per Observation 1—i.e., unique $M(t)$ for all types $t\in\mathcal{T}_{\mathcal{L}}\cup\mathcal{T}_{\mathcal{E}}$ and unique U. Recall that $N(k,q)$ is defined as the number of matches between agents of type k and agents of type q in M.

We start by constructing a graph associated with matching M as follows. Let $G(M)$ be the bipartite graph whose vertex sets are the types in $\mathcal{L}$ and $\mathcal{E}$, and such that there is an edge between types $k\in\mathcal{T}_{\mathcal{L}}$, $q\in\mathcal{T}_{\mathcal{E}}$ if and only if there is an agent of type k matched to an agent of type q in M—i.e., $N(k,q)>0$. We say that $G(M)$ is the *type-adjacency graph* associated with matching M. The following lemma states a key fact regarding the structure of $G(M)$.

Lemma 1. *Let $G(M)$ be the associated type-adjacency graph associated with matching M. Suppose that we mark the vertex in $G(M)$ corresponding to type $t\in\mathcal{T}_{\mathcal{L}}\cup\mathcal{T}_{\mathcal{E}}$ if and only if at least one agent of type t is unmatched under M. Then, under Assumption 1 (generalized imbalance), every connected component in $G(M)$ must contain a marked vertex.*

4.1. Overview of the Upper Bound Proof

The intuition behind the proof of the result is as follows. We start by bounding the type-pair prices associated with types with at least one unmatched agent. By applying the bounds given by the (IR) conditions in Proposition 2, we show that the difference between the maximum and minimum $\alpha_{(k,q)}$ corresponding to a type-pair price associated with at least one such type (i.e., either k or q have at least one unmatched agent) is small with high probability. Using these bounds, we move on to bounding the prices associated with type-pairs with all agents matched. We use the (ST) conditions in Proposition 2, which bound the relative variation of a type-pair price with respect to other type-pair prices. The fact that such relative variation is also small with high probability, taken together with the absolute variation on the types with unmatched agents, will imply the result.

We first establish the result for the case where F is Uniform$[0,1]$. The two main challenges in the proof of this case are as follows. First, we must express the bounds in Proposition 2 as a function of the market primitives—i.e., market size and number of types. To do so, we reinterpret the stability conditions as geometric conditions over appropriately defined random regions in unit hypercubes, as explained in Section 4.1.1. Second, we need to relate the bounds given by the (IR) constraints to those given by the (ST) constraints, to obtain uniform absolute bounds on the variation of *all* type-pair prices. We address this issue in Section 4.1.2, by exploiting the combinatorial structure of the problem. Finally, we extend the proof to deal with general distributions. We show that under appropriate scaling and translation, the bounds obtained for the uniform case will also hold (within a constant factor) for any general distribution with positive density in an interval. This is shown in Section 4.1.3.

We provide a detailed overview of the proof next.

4.1.1. Geometric Interpretation of the Stability Conditions.

We now briefly describe the geometric interpretation of the (IR) and (ST) conditions. This will allow us to reexpress the constraints in Proposition 2 as bounds that depend on the market primitives.

The intuition is as follows. Once we focus on a single type $t\in\mathcal{T}_{\mathcal{L}}\cup\mathcal{T}_{\mathcal{E}}$, the random productivities of an agent of type t can be described by a $D(t)$-dimensional vector within the $[0,1]^{D(t)}$-hypercube, where $D(t)$ is the dimension of the productivity vector of agents of type t (i.e., $D(t)=K$ if $t\in\mathcal{T}_{\mathcal{E}}$ and $D(t)=Q$ if $t\in\mathcal{T}_{\mathcal{L}}$). Furthermore, the location of these points can be described by a point process in $[0,1]^{D(t)}$. Every such hypercube can be partitioned into (at most) $D(t)+1$ random regions: one region corresponding to agents matched to each of the $D(t)$ types on the other side, and one region corresponding to unmatched agents. For every α vector in the core, the regions are separated by hyperplanes as per the (ST) and (IR) conditions. Since there is a unique type-matching with probability one (by Observation 1), each agent (point) must fall into the same region for all core α's. Thus, the stability conditions can be interpreted as geometric conditions in the unitary hypercube that determine how far we can move the boundaries of each random region without changing the members of each region. Possible fluctuations in the boundaries are bounded by the distance between consecutive order statistics of the projections of points distributed independently in (subregions of) the hypercube. We show that the event where all such distances are bounded by functions in $O(f(n))$ ($f(n)$ is as described in the statement of Theorem 1) occurs with high probability and thus achieves a bound that depends only on the market primitives. Next, we formally define some of the regions, random sets, and random variables that will be useful in our analysis,[26] and we relate such definitions to the stability constraints as defined in Proposition 2.

To this end, consider a type $t \in \mathcal{T}_{\mathscr{E}}$. For each employer j with $\tau(j) = t$, there is a vector of idiosyncratic productivities ϵ_j distributed uniformly in $[0,1]^K$ and independently across employers. In this subsection, we consider these productivities for a given t. We suppress t in the definitions to simplify notation (i.e., n here corresponds to n_t, and so on). Analogous definitions can be made for $t \in \mathcal{T}_{\mathscr{L}}$.

Consider n i.i.d. points $(\epsilon_j)_{j=1}^n$, distributed uniformly in the $[0,1]^K$-hypercube. Here, $\epsilon_j = (\epsilon_j^1, \epsilon_j^2, \ldots, \epsilon_j^K)$. Let $\mathcal{T}_{\mathscr{L}} = \{1, 2, \ldots, K\}$ denote the set of dimension indices. Intuitively, for a fixed type $t \in \mathcal{T}_{\mathscr{E}}$ with unmatched agents, one can bound α_{kt} by using condition (IR) in Proposition 2: $\min_{j\in t\cap M(k)} \epsilon_j^k \geqslant -\alpha_{kt} \geqslant \max_{j \in t\cap U} \epsilon_j^k$. To apply this bound, we care just about the projection onto the kth coordinate of the points ϵ_j with $j \in M(k) \cup U$. The main analytical challenge we face is that these relevant subregions (containing sets such as $M(k) \cup U$) are themselves a random function of the market realization, as both $M(k)$ and U are themselves random sets. We overcome this difficulty by appropriately defining the region $\tilde{\mathscr{R}}^k(\delta) := \tilde{\mathscr{R}}^k(t, \delta)$ (throughout the rest of this section, the type t is suppressed in the definition of the associated regions, sets, and random variables[27]). Specifically, for $\delta \in (0, 1/2]$ and $k \in \mathscr{K}$, define[28]

$$\tilde{\mathscr{R}}^k(\delta) = \{x \in [0,1]^K \colon x^{k'} \leqslant \delta \;\forall\, k' \in \mathscr{K}, k' \neq k\}. \qquad (4)$$

By choosing k and δ appropriately, we can guarantee that $\tilde{\mathscr{R}}^k(\delta)$ contains only points corresponding to agents in $M(k) \cup U$. (Intuitively, we can choose δ so that the other types are not attractive enough with such a low productivity.) Once we have done that, it should be easy to see that $\min_{j \in t\cap M(k)} \epsilon_j^k - \max_{j\in t\cap U} \epsilon_j^k$ is upper-bounded by the maximum distance between two consecutive points in $\tilde{\mathscr{R}}^k(\delta)$, when they are projected onto their kth coordinate (the corner cases of all points in $M(k)$, or in U, turn out to be easy to handle). This becomes precise once we introduce the set $\tilde{\mathscr{V}}^k(\delta)$ and the random variable $\tilde{V}^k(\delta)$, defined as follows:

$$\tilde{\mathscr{V}}^k(\delta) = \{x \colon x = \epsilon_j^k \text{ for } \{j \colon \epsilon_j \in \tilde{\mathscr{R}}^k(\delta)\}\}, \qquad (5)$$

and $\tilde{V}^k(\delta) = \max\big($Difference between consecutive values in $\tilde{\mathscr{V}}^k(\delta) \cup \{0,1\}\big)$. (6)

Thus, $\tilde{\mathscr{V}}^k(\delta) \subset [0,1]$ is the set of values of the kth coordinate of the points lying in $\mathscr{R}^k(\delta)$, and $\tilde{V}^k(\delta) \in \mathbb{R}$ is the maximum difference between consecutive values in $\tilde{\mathscr{V}}^k(\delta) \cup \{0,1\}$. (As an example, if $\tilde{\mathscr{V}}^k(\delta) = \{0.3, 0.4, 0.8\}$, the differences between consecutive values in $\tilde{\mathscr{V}}^k(\delta) \cup \{0,1\}$ are $0.3, 0.1, 0.4, 0.2$, resulting in $\tilde{V}^k(\delta) = 0.4$. Note that $\tilde{\mathscr{V}}^k(\delta)$ is a random and finite set, and $\tilde{V}^k(\delta)$ is a random variable.) Therefore, the variation in α_{kt} is bounded by $\tilde{V}^k(\delta)$. One of our intermediate results is to show that with high probability $\{\max_k \tilde{V}^k(\delta) \leqslant f_1(n, K)\}$, for some $f_1(n, K) = O^*(1/n^{1/K})$.

To complete the proof overview, note that once the absolute variation of, say, $\alpha_{tt'}$ has been bounded, a bound on the absolute variation of $\alpha_{tt''}$ can be obtained by bounding the relative variation in prices (i.e., the variation in $\alpha_{tt'} - \alpha_{tt''}$) using the (ST) conditions; cf. Section 4.1.2. Therefore, we need to define additional regions, sets, and variables that will allow us to apply the (ST) conditions. Such regions will be analogous to those defined above. However, instead of being a function of a single parameter k, they will be defined as functions of two parameters k_1, k_2; this difference is due to the fact that, unlike the (IR) constraints that focus on a single type-pair, the (ST) constraints bound the variation between two type-pairs. Since the relationship between these random regions and the α's is more involved, their definition and analysis are relegated to Online Appendix B. However, analogously to the case explained above, we still care about the difference between (appropriately defined) consecutive order statistics in these regions.

4.1.2. Using the Combinatorial Structure of the Problem to Relate the Bounds Given by (IR) and (ST) Constraints. We consider some suitably defined "typical" events as discussed in Section 4.1.1, which occur with high probability in the markets being considered. When these events occur, the possible fluctuations in the boundaries of the random regions (see Section 4.1.1) are bounded by functions in $O(f(n))$, where $f(n) = O^*(1/n^{1/\max(K,Q)})$, and $f(n)$ agrees with that in the statement of Theorem 1. Under these events, we show that the variation in the type-pair prices is uniformly bounded as follows:

$$\max_{(k,q)\in\mathcal{T}, N(k,q)>0} |\alpha_{kq}^{\max} - \alpha_{kq}^{\min}| \leqslant f(n).$$

This bound on the type-pair prices, together with the fact that the defined events occur with a high probability, implies our result.

We divide the proof of the bound on the type-pair prices into two steps. First, note that types with at least one unmatched agent are "easier" to bound; we use the (IR) conditions in Proposition 2 to bound the absolute variation of prices α associated with such types. In particular, for each type t with at least one matched and one unmatched agent, we use the (IR) conditions to show that $\max_{t' \colon N(t,t')>0} (\alpha_{t,t'}^{\max} - \alpha_{t,t'}^{\min}) \leqslant O^*(1/(n^{1/\max(K,Q)}))$. This is done in Lemma 2.

It remains to show that the bound holds for prices between types such that all agents are matched. Unfortunately, as the (IR) conditions in Proposition 2 cannot be applied to such types, proving the result is not straightforward. Only the (ST) conditions can be used, which provide only relative bounds on the values of the associated type-prices.

To prove the bounds on the variations of the prices associated with fully matched types, we use the graph $G(M)$ as defined above. Given a type $t \in \mathcal{T}_{\mathcal{L}} \cup \mathcal{T}_{\mathcal{E}}$, let the *distance* $d(t)$ be defined as the minimum distance in $G(M)$ from t to *any* marked vertex. By Lemma 1, every unmarked vertex t must be at a finite distance from a marked one. Furthermore, $\max_{t \in \mathcal{T}_{\mathcal{L}} \cup \mathcal{T}_{\mathcal{E}}} d(t) \leqslant K + Q$, regardless of the realization of the graph. Our argument to control the variation in the α's is by induction on $d(t)$ starting with $d(t) = 0$. The induction base has already been established, as we showed that the variation in *all* of the relevant α's associated with marked types (these types have distance zero) is bounded. Note that these types are those with at least one matched and one unmatched agent.

In the inductive step, we assume that the bound holds for every α associated with a type whose distance is d or less—i.e., for every $(t, t') \in \mathcal{T}$ such that $\min(d(t), d(t')) \leqslant d$, we have $\alpha^{\max}_{t,t'} - \alpha^{\min}_{t,t'} \leqslant O^*(1/n^{1/\max(K,Q)})$. Then, we use the inductive hypothesis to show that the result must also hold for all types whose distance is $d+1$. By the definition of distance, for every type t such that $d(t) = d+1$, there must exist a type t^* such that $d(t^*) = d$ and $N(t, t^*) > 0$. Therefore, by our inductive hypothesis, we must have $\alpha^{\max}_{t,t^*} - \alpha^{\min}_{t,t^*} \leqslant O^*(1/n^{1/\max(K,Q)})$. Separately, we control the relative variation of the α's associated with type t in the core using the (ST) conditions—i.e., we show that $\alpha_{t,t_1} - \alpha_{t,t_2}$ for types t_1, t_2 with matches with type t can vary only within a range bounded by $O^*(1/n^{1/\max(K,Q)})$. Combining, we deduce the desired bound on the variation in all α's associated with type t. This is formally achieved in Lemma 5.

4.1.3. Extension to General Distributions. So far we have given an overview of the proof for the case in which F is Uniform(0, 1). We now argue that the same bounds can be obtained for any distribution satisfying the conditions in Theorem 1, via appropriate scalings and translations of productivities and base utilities.

We first show that in any core solution, each type-pair price must lie in a bounded interval whose extremes are independent of n.

Lemma 2. *Fix K, Q, u's, and F. Then, under Assumption 2, there exists $U = U(K, Q, (u_{kq})_{k,q}, F) < \infty$ such that, with high probability, any core outcome (M, α) satisfies*

$$-U \leqslant \alpha_{kq} \leqslant U, \quad \text{for all } k, q. \tag{7}$$

Lemma 2 is proved in Online Appendix G.

Recall that in our proof of Theorem 1 for the Uniform(0, 1) case, we were able to bound the variation in type-pair prices by bounding the distances between relevant consecutive order statistics of the productivities and their linear combinations. In particular, we followed an inductive approach where we first focused on agent types with an unmatched agent and bounded all of the prices corresponding to that type, and then we proceeded inductively to bound the prices corresponding to types where all agents are matched. Note that when F is an unbounded distribution, the distance between consecutive order statistics might be larger than any $O^*(1/n^{1/\max(K,Q)})$ function. However, using Lemma 2, it follows that the relevant interval where realized values matter is

$$I = \begin{cases} [\max(C_l, -U), \min(C_u, U)], & \text{for } \epsilon_i^{k}\text{'s}, \\ [\max(C_l, -U + \min_{k,q} u(k,q)), \\ \quad \min(C_u, U - \max_{k,q} u(k,q))], & \text{for } \eta_j^{q}\text{'s}. \end{cases} \tag{8}$$

In other words, if we fix a type $t \in \mathcal{T}_{\mathcal{E}}$, then we care about the consecutive order statistics of the ϵ's that occur in the I^K-hypercube, and analogously for $t \in \mathcal{T}_{\mathcal{L}}$. Now, $f_{\min} = \inf_{x \in I} f(x) > 0$; by assumption, f is positive and continuous, and I is a compact set. If we suitably translate the productivity (achieved via an equal and opposite change in u) and then scale utilities down by $|I|$, we can map I to $[0, 1]$, and the density in $[0, 1]$ is now lower-bounded by $f_{\min}|I|$. Thus, if the productivities of type $t \in \mathcal{T}_{\mathcal{E}}$ are being considered, we apply this translation to all $k \in \mathcal{T}_{\mathcal{L}}$, and then scale all utilities down by $|I|$. As a result, we obtain that the density of the productivity vector $(\epsilon_i^k)_{k \in \mathcal{T}_{\mathcal{L}}}$ for $\tau(i) = t$ is lower-bounded uniformly by $(f_{\min}|I|)^K$ everywhere in $[0, 1]^K$, and that this is now the relevant region where realized values matter. Hence, our uniform bounds on the gaps between consecutive order statistics for the case where F is Uniform(0, 1) suffice.[29] The corresponding bound on order statistics for the original problem is simply $|I|$ times (a constant factor) larger.

4.2. Proof of the Lower Bound

Our lower bound follows from the following proposition, proved in Online Appendix D.

Proposition 3. *Consider a sequence of markets (indexed by $\tilde{n}$) with $|\mathcal{T}_{\mathcal{L}}| = K$ types of labor, with $\tilde{n}$ workers of each type, and with a single type $1'$ of employer, and $(K-1)\tilde{n} + 1$ employers of this type. (Note that these markets satisfy Assumptions 1 and 2.) All productivities are drawn from a distribution with support in $[0, 1]$. Set $u(k_*, 1') = 0$ for some $k_* \in \mathcal{L}$, and $u(k, 1') = 3$ for all $k \in \mathcal{L} \backslash k_*$. For this market, we have $\mathrm{E}[\mathscr{C}] = \Omega^*(1/n^{1/K})$.*

The rough intuition for our construction in Proposition 3 is as follows. For our choice of u's, it is not hard to see that all workers of types different from k_* are always matched in the core. One employer j_* is matched to a worker of type k_*. Suppose that vector $(\alpha_k)_{k \in \mathcal{T}_{\mathcal{L}}}$ is in the core. Given that all types $k \neq k_*$ are a priori symmetric, we would expect that the α_k's for $k \neq k_*$ are close to each other (we formalize using Lemma 4 that they are typically no more than $\delta \sim 1/\sqrt{\tilde{n}}$ apart). Assuming this is the case, we can order employers based on $X_j = \max_{k \neq k_*} \epsilon_j^k - \epsilon_j^{k_*}$, and j_* should usually be the employer with the smallest X_j, since this employer has the largest productivity with respect to k_*, relative to the other types. Now, the X_j's are i.i.d., and a short calculation establishes that the distance between the first- and second-order statistics of $(X_j)_{j\in\mathscr{E}}$ is $\Theta(1/n^{1/K})$. This "large" gap between the first two order statistics allows $(\alpha_{k_*}, (\alpha_k + \theta)_{k\neq k_*})$ to remain within the core for a range of values of $\theta \in \mathbb{R}$ that has an expected length of $\Theta(1/n^{1/2})$ for $K = 2$ and $\Theta(1/n^{1/K}) - O(\delta) = \Theta(1/n^{1/K})$ for $K > 2$, which leads to the stated lower bound on $\mathscr{C}$.

We remark that the key quantity here—namely, the gap between the first two order statistics of $(X_j)_{j\in\mathscr{E}}$—is determined by the tail behavior of the distribution (both the left and right tails) of the ϵ's, along with the number of types K. See Section 6 for further discussion.

Note that the construction above can easily be adapted to accommodate $Q \leqslant K$ types of firms.[30] If $Q > K$, we simply swap the roles of workers and firms in our construction, leading to $\mathrm{E}[\mathscr{C}] = \Omega^*(1/(n^{1/Q}))$, as needed. Thus, the lower bound in Theorem 1 follows from Proposition 3.

5. Numerical Experiments

In the previous sections, we provided asymptotic bounds on the size of the core. The purpose of this section is twofold. First, in Section 5.1, we provide simulation results to illustrate our theoretical results, and study the core behavior in markets with a small number of agents. Second, in Section 5.2, we explore what happens when the number of types is allowed to grow. Our simulation results reveal that the core is indeed small across a wide range of settings, including relatively small markets, reinforcing our main conclusion—namely, that typical assignment markets are likely to have a small core.

5.1. Simulation of Theoretical Results

Our main results state that, if the number of types is fixed and bounded by K, then the core becomes small as the market size grows. In particular, the core size in a market with n agents and at most K types on each side is bounded as $O^*(1/n^{1/K})$ (Theorem 1). In addition, if the distribution of the idiosyncratic terms is supported on the interval $(-\infty, +\infty)$, the size of the core is bounded as $1/n$, and the core solutions converge to a limit point (Theorem 3). We now numerically investigate finite markets, including relatively small markets.

5.1.1. Illustration of Theorem 1. In Theorem 1, we show that our bound is tight in the worst case by carefully constructing a specific family of instances whose expected core size is $\Omega^*(1/n^{1/K})$. However, such markets are unlikely to occur in practice. We are now interested in understanding the size of the core in more "typical" settings. To that end, we define a *typical market* with a distribution over the number of agents of each type as follows.

Definition 2 (Typical Market). A typical market is a market of random size defined by the quadruple $(n_{\mathscr{L}}, n_{\mathscr{E}}, [p_1, \ldots, p_K], [r_1, \ldots, r_Q])$, where $[p_1, \ldots, p_K]$ (resp. $[r_1, \ldots, r_Q]$) is a list of positive reals of length K (resp. Q) such that $\sum_{i=1}^K p_i = 1$ and $\sum_{i=1}^Q r_i = 1$, and such that the number of agents of type $k \in K$ (resp. $q \in Q$) is given by an independent random draw from a Poisson distribution with mean $p_k n_{\mathscr{L}}$ (resp. $r_q n_{\mathscr{E}}$).

We highlight that the above definition of a typical market is valid regardless of the distribution used to determine the idiosyncratic productivities. If $n_{\mathscr{L}} = n_{\mathscr{E}}$, we say that the market is *balanced*. Otherwise, we say that the market is *unbalanced*. Note that even a balanced typical market is very unlikely to have exactly the same number of agents on each side; the likelihood of exact balance is $O(1/\sqrt{n})$, and such realizations have a negligible impact on our numerical results since we are focusing on the *median* core size across realizations.

We tested the intuition provided by Theorem 1 by simulating typical markets with a varying number of agents, types, and type-type utilities. We consider three types of distributions for the idiosyncratic productivities: uniform, exponential (right-unbounded with a lighter tail), and Pareto (right-unbounded with a heavier tail).[31] For each set of parameters, we ran 100 trials and reported the median core size over those markets.

To illustrate our findings, in Figure 1 we show the median core size as a function of the number of agents for balanced markets with a fixed number of types (two and five types on each side), the same expected number of agents of each type, and all type-type utilities equal to zero. The idiosyncratic terms were drawn independently from, respectively, a uniform(0,1) distribution, an exponential distribution with mean 1, and a Pareto distribution with shape parameter $\alpha = 2$ (the choice of scale parameter has no impact on core size). We also added the functions $1/n$ and $1/\sqrt{n}$ to the graph, and used logarithmic scales on both axes to facilitate the comparison with polynomials of n. In the figure, one observes that the dependence of decay rate on the number of types is weak, and the decay rate appears to be between $1/\sqrt{n}$ and $1/n$.

We also simulated the effect of imbalance and changing type-type utilities.[32] In our experiments, the core size tends to slightly decrease as the imbalance increases (see Theorem 2). However, the difference in the core size between balanced and unbalanced typical markets was not large, and the core size was found to always decay at a rate of between $1/\sqrt{n}$ and $1/n$. Furthermore, consistent with our theoretical insights, we find that the values of the type-type utilities do not play a significant role in the size of the core.

Figure 1. (Color online) (Left) Median core size as a function of number of agents, for balanced markets with two types of agents on each side and both the type-type productivities and idiosyncratic terms drawn from uniform (U), exponential (E), and Pareto (P) distributions, respectively. The graph uses logarithmic scales on both axes. The functions $1/n$ and $1/\sqrt{n}$ are added for comparison. (Right) Core size as a function of number of agents, for balanced markets with five types of agents on each side.

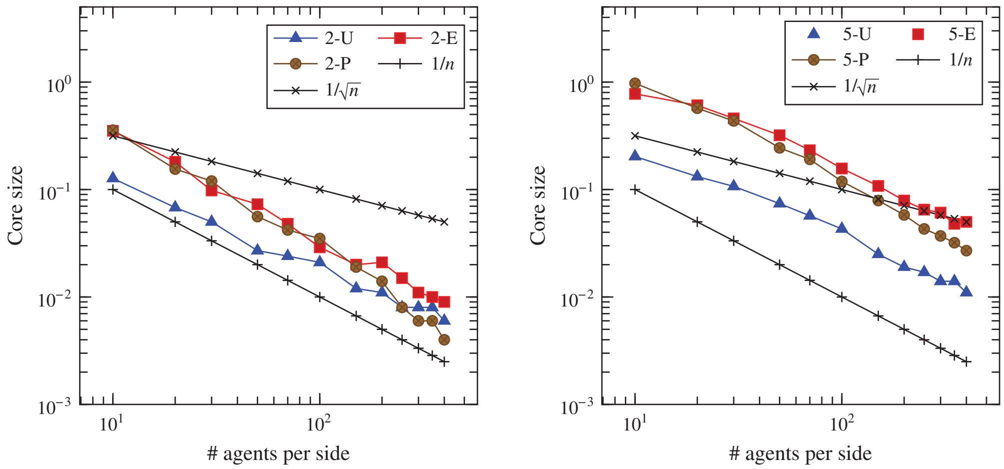

5.1.2. Illustration of Theorem 3. We now numerically illustrate our result in Theorem 3, using a standard Gumbel as the distribution for the idiosyncratic productivities. This case is particularly important: most empirical studies using this model assume a fixed number of types and idiosyncratic terms drawn from a Gumbel distribution (see, e.g., Choo and Siow 2006). For these experiments, we fixed the fraction ρ_t of agents of each type $t \in \mathcal{T}_{\mathcal{L}} \cup \mathcal{T}_{\mathcal{E}}$ and the type-type utilities, and increased the number of agents on each side of the market. For each set of parameters, we ran 100 trials and reported the average distance to the limiting type-type prices α^*, the average distance to the limiting (scaled) number of matches of each type-pair ν^*, and the median core size. As expected, the observed decrease in the core size occurs at a rate of approximately $1/n$, even when n is relatively small. We also found that the average distances from the realized α's and ν's to α^* and ν^*, respectively, decrease at a rate of $1/\sqrt{n}$, as predicted. Also, as predicted by Theorem 3, we found that the values of the type-type utilities do not alter the rate of convergence, and this holds even for small markets.[33]

We consider markets with two and five types on each side. Agents on each side are evenly distributed across types, but there is an imbalance across sides (we considered $1.2n$ workers and n firms). We also set the type-type utilities to be equal to 3.0, for each type pair. In this setting, all agents on each side are ex ante symmetric and thus, in the limit, the proportion of agents of type k that are matched to agents of type q, ν^*_{kq}, is the same for all type-pairs. Similarly, the limiting prices α^*_{kq} are equal across all type-pairs. In Figure 2 we illustrate this by plotting the median of $\text{Average}_{k,q}(|\alpha_{kq} - \alpha^*_{kq}|)$, the median of $\text{Average}_{k,q}(|N(k,q)/n - \nu^*_{kq}|)$, and the median size of the core as n increases. We observe that our asymptotic bounds capture the behavior of the core even in small markets.

5.2. Beyond Our Theory: Increasing the Number of Types

Our theoretical results assume that the maximum number of types (K) on each side remains fixed as the number of agents (n) increases. To conclude this section, we now numerically analyze what happens if we allow the number of types to increase.

In particular, we fix the number of agents n, and study the behavior of the core when the number of types K increases. We again show results for the case where the type-type utilities are drawn independently from the same distribution as the idiosyncratic terms. (We found that the core size is not sensitive to the distribution of the type-type utility term.) Clearly, the size of the core (see Definition 1) is bounded by 1. Therefore, it should not be surprising that the core size as a function of the number of types eventually flattens out. The question is, how much can the core grow if we allow the number of types to grow, and how many types of agents are needed before the core stops growing?

If the random productivities are drawn from a uniform distribution, then allowing the number of types to increase as n increases does not have a significant effect, and the core remains small. This is illustrated by Figure 3, where we consider balanced typical markets with 200 and 400 agents on each side, respectively, and plot the median core size as a function of the number of types for uniform, exponential, and Pareto distributions. Observe that under both uniform and exponential distributions, the curve flattens out when K is approximately 30 in each case. On the other hand, the

Figure 2. (Color online) (Left top) Median distance to α^* as a function of the number of workers for markets with two and five types of agents on each side, idiosyncratic terms drawn from standard Gumbel distribution, type-pair utilities set to zero, and 1.2 times as many workers as firms. for each simulation, we compute the distance to α^* as Average$_{(k,q)}|\alpha_{k,q} - \alpha^*|$. When $K = 2$, We have $\alpha^* = 0.132$ and, when $K = 5$, $\alpha^* = 1.026$. (Right top) Median distance to ν^*, where for each simulation, we compute the distance to ν^* as Average$_{(k,q)}|N(k,q)/n - \nu^*|$. When $K = 2$, we have $\nu^* = 0.106$ and, when $k = 5$, we have $\nu^* = 0.018$. (Bottom) Median core size as a function of the number of workers. All graphs use logarithmic scales on both axes. The functions $1/\sqrt{n}$ and $1/n$ are added respectively for comparison.

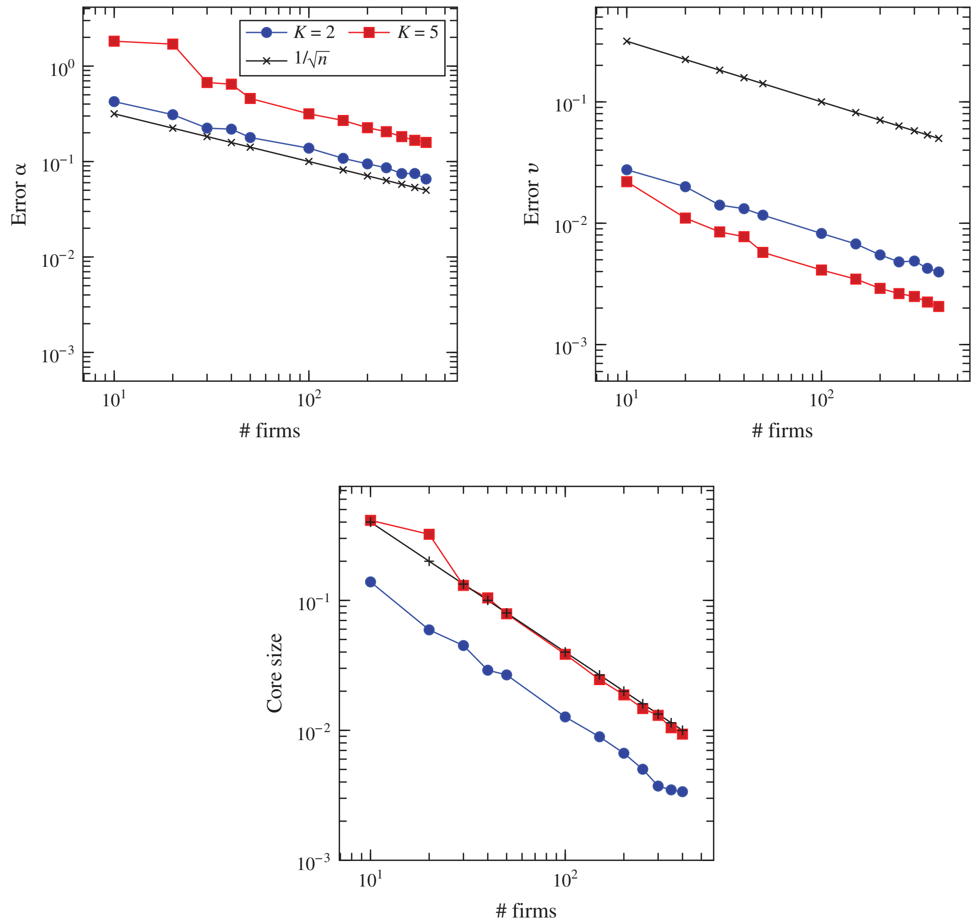

effect of increasing the number of types for Pareto distributions is significant, and a "large" core may arise as a result (with core size exceeding 0.1), albeit only when there is a large number of types (e.g., more than ~ 30 types with 200 agents on each side). As seen in Figure 3, the size of the core does not flatten out even with 100 types per side in this case.[34] However, such a large number of types would not be expected to occur in practice, and we consider our findings here as a confirmation that the core is indeed small under plausible values for n, K, and Q.

6. Discussion

This paper quantifies the size of the core in matching markets with transfers, as a function of market characteristics. We considered a model of an assignment market with a fixed number of types of workers and firms. We modeled the value of a match between a pair of agents as a sum of a deterministic term determined by the pair of types, and a random component that is the sum of two terms, each depending on the identity of one of the agents and the type of the other. Assuming a fixed number of types, we showed that the size of the core is bounded as $O^*(1/n^{1/\ell})$, where each side of the market contains no more than ℓ types, when the random terms are drawn from a distribution whose density is strictly positive and continuous everywhere in the (possibly unbounded) support. This bound holds in the worst case over the number of agents. These results imply that the vector of prices that support a competitive equilibrium (a core solution) is approximately uniquely determined, and thus the payoff of an agent is roughly the same in all stable solutions.

We provide additional results for the practically relevant case of distributions with support $(-\infty, \infty)$ (including Gumbel and normal): we show a tighter

Figure 3. (Color online) (Left) Median core size as a function of number of types, for balanced typical markets with 200 agents on each side and idiosyncratic terms drawn from a uniform distribution (U), an exponential distribution (E), and a Pareto distribution (P) with shape parameter $\alpha = 2$. The graph uses logarithmic scales on both axes. (Right) Median core size as a function of number of types, for typical markets with 400 agents on each side.

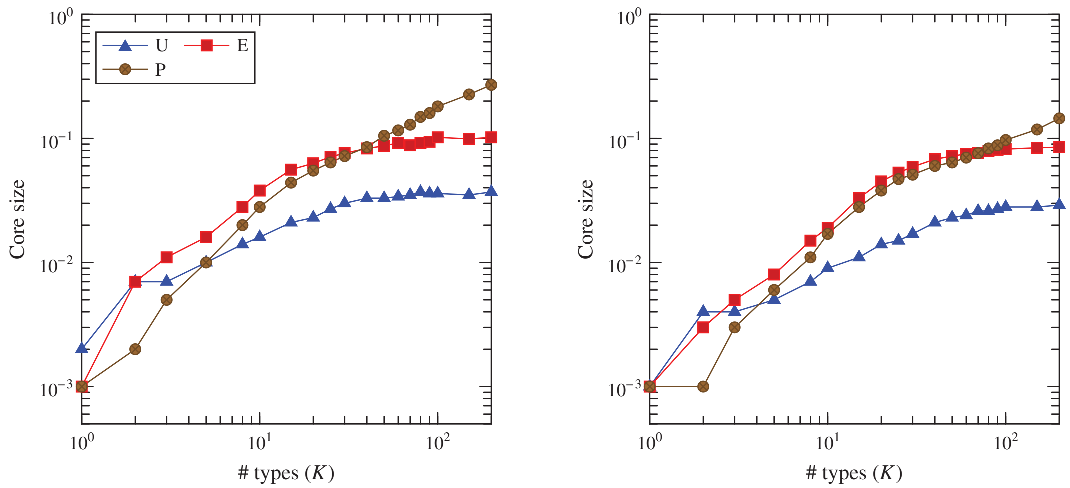

bound of $O^*(1/n)$ on the size of the core, as well as convergence to a unique limit of both the core prices and the fraction of matches of each type-pair, and bounds of $O^*(1/\sqrt{n})$ on these rates of convergence.

Using numerical experiments, we show that our small core finding holds across a range of practically relevant situations.

As future work, it would be interesting to extend our results to many-to-one markets, where employers can each have more than one opening. We expect that our results regarding the core (also our proofs) extend to the case where each employer's capacity is bounded by a constant, and each employer's utility is additive across matches. Another interesting direction would be to investigate the role of a marketplace operator (e.g., Airbnb, Upwork) in determining the wage/price levels in a decentralized market. Our work suggests that price recommendations alone may not be a good tool to control prices, since prices are uniquely determined in equilibrium. However, there are nontrivial search costs in many such marketplaces. Hence, search engine design and other aspects of the search/recommendation environment provided by a platform can play a role in modulating prices.

Acknowledgments

The authors thank the associate editor and the anonymous referees for their detailed comments and constructive criticism. In addition, they thank Itai Ashlagi, Eduardo Azevedo, Omar Besbes, Kostas Bimpikis, Yeon-Koo Che, Bernard Salanie, and Gabriel Weintraub for useful discussions and their comments on a first version of this paper. A preliminary version of this work appeared in SODA 2015 (Kanoria et al. 2015). That publication is an extended abstract containing only a subset of the current theoretical results, no proofs, and no numerical results.

Endnotes

[1] A small core has been found in special cases of the TU setting as in Gretsky et al. (1992, 1999) and Hassidim and Romm (2015), which we discuss below. In the NTU setting, real markets have almost always been found to contain a nearly unique stable outcome; see, e.g., Roth and Peranson (1999). A body of theory explains this phenomenon (Immorlica and Mahdian 2005, Kojima and Pathak 2009, Ashlagi et al. 2017, Holzman and Samet 2014, Azevedo and Leshno 2016).

[2] Our results extend to the case where transfers are taxed by the platform; cf. Remark 1.

[3] The core size and other features of real NTU markets can be readily studied as the data on ordinal preferences is often available. In TU matching markets, this becomes a harder task given that preferences (i.e., the value generated by a pair/match) are difficult to observe.

[4] Namely, we require the support of the distribution to be a (possibly unbounded) interval, and the density to be continuous and positive everywhere in the support. These requirements are satisfied by the normal, Gumbel, uniform, and exponential distributions, among others.

[5] We write $f(n) = O^*(g(n))$ if there exists $r < \infty$ such that $f(n) \leqslant r(\log n)^r g(n)$ for all $n \geqslant 2$. In words, this corresponds to the big-O notation where poly-logarithmic factors are also suppressed.

[6] Similarly, we write $f(n) = \Omega(g(n))$ if $\limsup_{n\to\infty} |f(n)/g(n)| > 0$. This corresponds to the standard big-O notation.

[7] Azevedo and Leshno (2016) consider a fixed number of schools and a growing number of students. One can think of each school as being replaced with a linearly growing number of agents (equal to the number of seats in the school) having identical preferences (identical to those of the school), without the core being affected. Hence, we see Azevedo and Leshno (2016) as implicitly considering preferences that are strongly correlated in a particular way.

[8] Our results on a small core in TU markets also hold under a similar "generalized imbalanced" requirement, which is generically satisfied; cf. Section 3.

[9] These results are formalized by Shapley and Shubik (1971). They show that the set of stable outcome utilities is the set of optima of the dual to the maximum-weight matching linear program, implying in particular that the matching M in a stable outcome must be a maximum-weight matching.

[10] More generally, several other relevant characteristics, such as location and experience, may be included to define a type.

[11] An alternative way of thinking about types is in the context of an empirical model. There, types represent the observable characteristics of the agents (age, sex, location, education level). In addition, agents also have some unobservable characteristics, captured by allowing for idiosyncratic variations.

[12] These features have also been important in facilitating identification (Choo and Siow 2006, Chiappori et al. 2015). A researcher is typically able to observe only the cross-section marriage/matching distribution—namely, the number of type t agents matched to type t' agents on the other side of the market.

[13] The modification can be absorbed in the $u(\tau(i), \tau(j))$ term, leaving the idiosyncratic productivities unaffected.

[14] One may define the size of the core in $\mathcal{M}_{\text{tax}}$ as the difference between the maximum and minimum total payoff of firms in core outcomes (this difference exceeds the corresponding difference for workers), divided by the weight of M (all stable matchings once again live on M, and weight (M) is again the total payoff of workers, employers, and platform combined).

[15] In the continuum limit markets as used in the empirical studies (e.g., Choo and Siow 2006, whose model is very similar to ours), there is a unique core solution. Real-world markets, on the other hand, are finite and thus always have multiple core solutions. Therefore, bounding the size of the core as a function of the market size and other market primitives may provide some justification for the continuum market assumption.

[16] Whenever Assumption 1 is not satisfied, one can construct a market for which $\mathrm{E}[\mathscr{C}] = \Omega(1)$, by extending the reasoning leading to Remark 2.

[17] Consider possible vectors $\mathcal{N} = \{(n_t)_{t\in\mathcal{T}_\mathcal{L}\cup\mathcal{T}_\mathcal{E}}: \sum_t n_t = n\}$ describing the number of agents of each type. Then, a fraction $O(1/n)$ of these vectors violate Assumption 1.

[18] Thus, allowed distributions include typical continuous distributions such as Gumbel, Gaussian, Pareto, and uniform.

[19] We remark that we in fact only need a weaker condition on f—namely, that it is bounded below by a function that is positive and continuous everywhere in the support, and is non-atomic. Also, F can be different for different type-pairs.

[20] This is similar in spirit to the result obtained by Hassidim and Romm (2015), albeit in a different model where a single marketwide price emerges.

[21] A positive lower limit in the support of F can be absorbed into the $u(\cdot,\cdot)$'s, and hence this case is covered.

[22] In fact, it is enough for f to be bounded below by a function that is positive everywhere and continuous.

[23] This scaling is a bit stronger than the one in Assumption 2, as we require the *fraction* to remain constant throughout. Note that this is necessary to ensure convergence to a limit point.

[24] In fact, it is sufficient that the fraction of agents of type t converge to ρ_t as $n to \infty$. If we have $\max_{t\in\mathcal{T}_\mathcal{L}\cup\mathcal{T}_\mathcal{E}} |\rho_t - (\text{\# agents of type t})/n| \leqslant \delta_n$ for some $\delta_n = o(1)$, then we obtain bounds of $O^*(\max(1/\sqrt{n}, \delta_n))$ in the first part of the theorem, while Equation (3) remains unchanged.

[25] The generalized imbalance assumption guarantees that at least one agent is unmatched, which is necessary for a small core to arise. Here, as F is unbounded on both sides, some agents will remain unmatched with high probability for a big enough market.

[26] The remaining definitions can be found in Online Appendix B.

[27] While for our purposes we consider the hypercube associated with a given type t, the results in this section are intended more generally as results for random point processes in a hypercube. Thus, types are omitted from the definitions.

[28] In our setting, $\mathcal{K}$ is equal to $\mathcal{T}_\mathcal{L}$. However, as previously explained, the results in this section are intended more generally as results for random point processes in a hypercube. Thus, we decided to use a more general notation.

[29] Formally, we divide realized productivity vectors into those that are considered and those that are not. The density of considered realizations is $(f_{\min}|I|)^K$ in $[0,1]^K$, and 0 elsewhere; overall, a fraction $(f_{\min}|I|)^K > 0$ of realizations are considered and hence, w.h.p., $\Omega(n)$ workers of type k are considered.

[30] Let each worker type $q \neq 1$ have $\tilde{n}$ agents each and $u(\cdot, q) = -2$. These workers are always unmatched, leaving the core unaffected.

[31] The Pareto distribution has two parameters: the shape parameter α and the minimum point in the support $x_{\min}$. Given those parameters, the CDF at $x \geqslant x_{\min}$ is given by $1 - (x_{\min}/x)^\alpha$.

[32] These results are not explicitly reported for the sake of brevity.

[33] Results for different type-type utilities are not reported for the sake of brevity.

[34] Our intuition for the underlying reason for a large core is as follows: When the number of types exceeds $\sqrt{n}$, then most of the matched pairs consist of agents who are especially well suited to each other in terms of their productivities for each other's type. This leaves significant leeway for each pair of agents to negotiate the division of the surplus they generate, even while keeping other agents' payoffs fixed and retaining stability.

References

Airbnb (2015) Using data to help set your price. *Airbnb Blog* (June 4), http://blog.atairbnb.com/using-data-to-help-set-your-price/.

Alaei S, Jain K, Malekian A (2016) Competitive equilibrium in two-sided matching markets with general utility functions. *Oper. Res.* 64(3):638–645.

Ashlagi I, Kanoria Y, Leshno JD (2017) Unbalanced random matching markets: The stark effect of competition. *J. Political Econom.* 125(1):69–98.

Azevedo E, Leshno JD (2016) A supply and demand framework for two-sided matching markets. *J. Political Econom.* 124(5): 1235–1268.

Becker GS (1973) A theory of marriage: Part I. *J. Political Econom.* 81(4):813–846.

Chiappori P-A, Salanié B, Weiss Y (2015) Partner choice and the marital college premium: Analyzing marital patterns over several decades. Working paper, Columbia University, New York.

Choo E, Siow A (2006) Who marries whom and why. *J. Political Econom.* 114(1):175–201.

Durrett R (2010) *Probability: Theory and Examples* (Cambridge University Press, New York).

Fox JT (2018) Estimating matching games with transfers. *Quant. Econom.* 9(1):1–38.

Gale D, Shapley LS (1962) College admissions and the stability of marriage. *Amer. Math. Monthly* 69(1):9–15.

Galichon A, Salanié B (2010) Matching with trade-offs: Revealed preferences over competing characteristics. Working paper, New York University, New York.

Gretsky NE, Ostroy JM, Zame WR (1992) The nonatomic assignment model. *Econom. Theory* 2(1):103–127.

Gretsky NE, Ostroy JM, Zame WR (1999) Perfect competition in the continuous assignment model. *J. Econom. Theory* 88(1): 60–118.

Hassidim A, Romm A (2015) An approximate law of one price in random assignment markets. *Proc.* 16*th Annual ACM Conf. Econom. Comput., EC* '15 (ACM, New York), 279–280.

Hatfield JW, Milgrom PR (2005) Matching with contracts. *Amer. Econom. Rev.* 95(4):913–935.

Holzman R, Samet D (2014) Matching of like rank and the size of the core in the marriage problem. *Games Econom. Behav.* 88(November):277–285.

Immorlica N, Mahdian M (2005) Marriage, honesty, and stability. *Proc. 16th Annual ACM-SIAM Sympos. Discrete Algorithms, SODA '05* (ACM, New York), 53–62.

Jaffe S, Kominers SD (2014) Matching markets with taxation of transfers. Working paper, Harvard University.

Kamecke U (1992) On the uniqueness of the solution to a large linear assignment problem. *J. Math. Econom.* 21(6):509–521.

Kanoria Y, Saban D, Jay S (2015) The size of the core in assignment markets. *Proc. 26th Annual ACM-SIAM Sympos. Discrete Algorithms, SODA '15* (SIAM, Philadelphia), 1916–1924.

Kelso AS Jr, Crawford VP (1982) Job matching, coalition formation, and gross substitutes. *Econometrica* 50(6):1483–1504.

Kojima F, Pathak PA (2009) Incentives and stability in large two-sided matching markets. *Amer. Econom. Rev.* 99(3):608–627.

Kojima F, Pathak PA, Roth AE (2013) Matching with couples: Stability and incentives in large markets. *Quart. J. Econom.* 128(4): 1585–1632.

Menzel K (2015) Large matching markets as two-sided demand systems. *Econometrica* 83(3):897–941.

Nguyen T (2015) Coalitional bargaining in networks. *Oper. Res.* 63(3):501–511.

Peski M (2015) Welfare and entropy in many-to-many matching models. Working paper, University of Toronto, Toronto.

Roth AE, Peranson E (1999) The redesign of the matching market for American physicians: Some engineering aspects of economic design. *Amer. Econom. Rev.* 89(4):748–780.

Shapley LS, Shubik M (1971) The assignment game I: The core. *Internat. J. Game Theory* 1(1):111–130.

Yash Kanoria is the Sidney Taurel Associate Professor of Business in the Graduate School of Business, Columbia University. His research focuses on the design and operations of marketplaces, especially matching markets and online platforms.

Daniela Saban is an assistant professor of operations, information, and technology at the Graduate School of Business, Stanford University. Her research interests include operations and revenue management, with a particular emphasis on the design of dynamic, online markets.

Jay Sethuraman is a professor in the Department of Industrial Engineering and Operations Research at Columbia University. His current research interests are in discrete optimization and its applications to matching and allocation problems, mechanism design, scheduling theory, and applied probability.