University of Pennsylvania
ScholarlyCommons

Business Economics and Public Policy Papers

Wharton Faculty Research

1-2014

Imperfect Competition in Two-sided Matching Markets

Eduardo M. Azevedo
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/bepp_papers

Part of the [Behavioral Economics Commons](#)

Recommended Citation

Azevedo, E. M. (2014). Imperfect Competition in Two-sided Matching Markets. *Games and Economic Behavior*, 83 (C), 207-223. <http://dx.doi.org/10.1016/j.geb.2013.11.009>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/bepp_papers/127
For more information, please contact repository@pobox.upenn.edu.

Imperfect Competition in Two-sided Matching Markets

Abstract

This paper considers a simple equilibrium model of an imperfectly competitive two-sided matching market. Firms and workers may have heterogeneous preferences over matches on the other side, and the model allows for both uniform and personalized wages or contracts. To make the model tractable, I use the [Azevedo and Leshno \(2013\)](#) framework, in which a finite number of firms is matched to a continuum of workers.

In equilibrium, even if wages are exogenous and fixed, firms have incentives to strategically reduce their capacity, to increase the quality of their worker pool. The intensity of incentives to reduce capacity is given by a simple formula, analogous to the classic Cournot model, but depends on different moments of the distribution of preferences. I compare markets with uniform and personalized wages. For fixed quantities, markets with personalized wages always yield higher efficiency than markets with uniform wages, but may be less efficient if firms reduce capacity to avoid bidding too much for star workers.

Keywords

matching markets, imperfect competition

Disciplines

Behavioral Economics | Economics

Imperfect Competition in Two-Sided Matching Markets*

Eduardo M. Azevedo[†]

January 10, 2011

Abstract

In two-sided matching markets participants care about who they interact with in the other side, such as colleges and students or employees and firms. This paper considers a simple equilibrium model of an imperfectly competitive matching market. A finite number of firms is matched to a continuum of workers. Firms and workers may have heterogeneous preferences over matches on the other side, and the model allows for both uniform and personalized wages or contracts.

In equilibrium, even if wages are exogenous and fixed, firms have incentives to strategically reduce their capacity, to increase the quality of their worker pool. The intensity of firms' incentives to reduce capacity is given by a simple formula, analogous to the classic Cournot model. If workers are paid uniform wages, firms distort capacity more with heterogeneous preferences, while if workers are paid personalized wages firms have more incentives to distort capacity with homogeneous preferences. We show that markets with personalized wages always yield higher efficiency given quantities, but may be less efficient if firms reduce capacity to avoid bidding too much for star workers. A new rationale for unravelling is also given, where firms hire early to obtain a first-mover advantage.

*Preliminary draft, comments welcome. I would like to thank Nikhil Agarwal, Itai Ashlagi, Susan Athey, Eric Budish, Carlos E. da Costa, Drew Fudenberg, Felipe Iachan, Scott Kominers, Jacob Leshno, Alex Nichifor, Alvin Roth, Rodrigo Wagner, Yosuke Yasuda and seminar participants at Harvard and USP for helpful comments and discussions. Finally, I would like to thank Jonathan Kluberg, Ruben Lobel, Alexander Rikun, and Gareth Williams for hosting me at the MIT Operations Research Center where most of this paper was written. All mistakes are my own.

[†]Harvard University, Department of Economics. azevedo@fas.harvard.edu.

1 Introduction

Two-sided matching markets are markets where participants on either side have preferences over who they interact with on the other side. Examples include matching CEOs to companies, husbands to wives, students to colleges, doctors to hospitals, lawyers to law firms, and advertisers to content providers.¹ A stable matching is an allocation where agents do not have incentives to break away from their matches and seek new ones. A well-known result by Roth (1985) shows that no mechanism that always produces a stable matching is strategyproof for the firms.² However, even though these markets are studied by a large literature originated by Becker (1973) and Gale and Shapley (1962), most contributions ignore strategic behavior by firms, assuming them to command insignificant market share, or to act naively. This is in contrast to the standard approach in industrial organization, which typically focuses on Nash equilibrium in imperfectly competitive markets. This paper considers the questions: What is the optimal behavior of a firm in an imperfectly competitive matching market?, e.g. What is the optimal admission policy for a college, how should a hospital play in a centralized clearinghouse?. What are the equilibrium consequences of strategic behavior?, What are the implications for the regulation and design of these markets?.

I consider markets where agents on one side (firms) may be matched to many agents on the other side (workers). Firms have nontrivial market share, and are assumed to act strategically. Both workers and firms may have heterogeneous preferences over matches in the other side. I consider both markets with uniform wages for each worker, and markets with personalized wages or contracts. The model is analogous to the classic Cournot model of imperfect competition in homogeneous goods markets. In the Cournot model, firms set quantities and prices are given by market clearing at those quantities. I follow Cournot (1838) and Konishi and Ünver (2006), and assume that firms set quantities, and the allocation is given by a stable matching given quantities. To make the analysis tractable, the paper uses the matching framework of Azevedo and Leshno (2010), where a finite number of firms is matched to a continuum of workers.³ Using the continuum model allows us to give simple first order conditions for the firms' optimal decisions.

A basic insight from the analysis is that some of the intuition of the Cournot model applies to matching markets. Market power causes the marginal revenue of a firm to be lower than the net productivity of a marginal worker. For that reason, firms shade capacities. A surprising finding is that firms may want to reduce quantities even if wages are fixed. The intuition is that by reducing capacity firms with market power can increase the quality of the worker pool

¹See, respectively, Gabaix and Landier (2008); Becker (1973); Gale and Shapley (1962); Roth and Peranson (1999); Ginsburg and Wolf (2003).

²More precisely, he shows that in the college admissions problem no stable mechanism is strategyproof for the colleges.

³Azevedo and Leshno (2010) show that this model corresponds to the limit of discrete economies with many workers per firm.

available to them. We provide a simple first-order condition that quantifies firms' incentives to reduce capacity. I provide applications of this idea, notably to the costs and benefits of flexible versus uniform wages, and also to unravelling.

The paper compares markets with personalized and uniform wages. For example, in the market for junior associates in elite New York law firms, most firms pay every incoming lawyer the same wage. In contrast, senior lawyers are often paid personalized wages (Ginsburg and Wolf (2003)). A series of papers have debated the desirability of using personalized wages, which is a key market design variable. Notably, Bulow and Levin (2006) have shown that uniform wages may reduce matching efficiency, and compress wages. I show however that if firms may shade capacities, this conclusion may be reversed. In the imperfect competition model, there is a tradeoff. Personalized wages always generate higher matching efficiency for a given level of capacity, but they may increase firms' incentives to distort capacity. If firms are very similar, personalized wages have little impact on matching efficiency, but may induce firms drastically reduce capacity to avoid entering in a bidding war for the best employees. In that case markets organized around uniform wages generate higher welfare. However, if firms are more heterogeneous, the loss from matching inefficiency dominates the loss from capacity shading, and personalized wages generate higher efficiency.

In addition, the paper proposes a new explanation for unravelling, the phenomenon of hiring dates moving earlier and earlier in labor markets. In the model, since there are strategic complementarities between firms' choices, unlike perfectly competitive models, it is advantageous for firms to move first. Therefore, firms might be willing to pay costs to hire early, to obtain a first mover advantage, or at least compete in the same par as its rivals. This rationale for unravelling does not rely on workers having incomplete information about their own abilities, and neither on firms hiring before revelation of information, as is the case of most traditional models of unravelling (Li and Rosen (1998)).

1.1 Related literature

The model is related to contributions in the industrial organization, matching, and market design literatures. In industrial organization, it extends the Cournot oligopoly model to matching markets. Classic themes such as strategic complementarities (Fudenberg and Tirole (1984); Bulow et al. (1985)) have important consequences to matching markets. Moreover, most of the industrial organization literature on two-sided markets has focused on the strategic competition among platforms that connect buyers and sellers (Rochet and Tirole (2006)).⁴ In contrast, we consider the strategic behavior of firms with nontrivial market share in one of the sides of the market.

The main departure of the present paper from the literature on assortative matching, introduced in the seminal Becker (1973) marriage model, is that we allow for non-transferable utility and heterogeneous preferences, so that match-

⁴See also Rochet and Tirole (2003); Armstrong (2006); Weyl (2007)

ing is not necessarily assortative. Some papers in this tradition have considered strategic behavior of firms. Most closely related to my model, Bulow and Levin (2006); Kojima (2007); Niederle (2007) consider the effect of personalized versus fixed wages in a matching market. Our model suggests a new tradeoff in this personalized versus uniform wages debate, that personalized wages may increase distortions caused by capacity manipulation. Gabaix and Landier (2008); Terzio (2008); Edmans et al. (2009) study the effect of assortative matching on wages in matching markets, and Kremer (1993) considers the same problem in the one-sided case.

The major difference between the model and most of the literature on matching with heterogeneous preferences (Gale and Shapley (1962); Roth and Sotomayor (1992)) is that we consider nontrivial equilibrium outcomes, as opposed to the strategyproofness of mechanisms. A stable mechanism is a mechanism that always produces a stable matching. Several papers in this literature explore conditions for the strategyproofness of stable mechanisms. Roth (1985) has shown that no stable mechanism is strategyproof for the colleges in the college admissions problem. Sönmez (1997) has shown that colleges may manipulate any stable mechanism by reducing capacity. Kesten (2008) shows that, if preferences are acyclic, it is not possible to gain by manipulating capacities. Roth and Peranson (1999); Immorlica and Mahdian (2005); Kojima and Pathak (2009) have given a series of results suggesting that in large markets the gains from manipulating stable mechanisms is small. The most closely related models to mine are in the literature on capacity manipulation games. Konishi and Ünver (2006) introduced these games, where a set of firms simultaneously choose capacities, and are assigned the corresponding stable matching. The most substantial differences of my model is that it considers a continuum framework, and allows for matching with contracts. Konishi and Ünver (2006) have shown that pure-strategy equilibria do not necessarily exist in these games. Subsequently, Kojima (2006) has shown that equilibria in mixed strategies exist, and that firms are weakly better off in such equilibria than under truthful reporting. Mumcu and Saglam (2009) studies games of capacity manipulation with an aftermarket, and Ehlers (2010) characterizes student-optimal mechanisms in the college admissions problem in terms of a strategyproofness condition.

Section 2.1 describes the baseline model, and characterizes optimality conditions for firms and equilibrium. Sections 2.2 and 2.3 then extend the model to include endogenous wages and matching with contracts, and Section 3 discusses applications, comparing uniform and personalized wages, and providing a new rationale for unravelling. Section 4 concludes. Proofs of all results are given in the Appendix.

2 Model

Section 2.1 considers the baseline model, of matching with uniform and exogenous wages. Sections 2.2 and 2.3 then incorporate endogenous wages and matching with contracts into the analysis. A casual reader may want to read

Section 2.1 to grasp the basic model, skim over Sections 2.2 and 2.3 and proceed to the applications in Section 3.

2.1 Baseline: matching with uniform exogenous wages

2.1.1 Firms, workers, and stable matchings

Firms $i = 1, \dots, I$ compete for a continuum mass of workers. Abusing notation, we also denote by I the set of firms $\{1, \dots, I\}$. Worker of type θ has productivity e_i^θ in $[0, 1]$ at firm i . Note that a worker's productivity may differ in different firms. We denote by e^θ the I -dimensional vector of worker productivity. Each worker has a complete preference ordering $\succ^\theta$ over all firms, and over being unmatched. Let $\mathcal{P}$ be the set of all strict preference relations over the set of firms and being unmatched. The set of worker types is $\Theta = [0, 1]^I \times \mathcal{P}$. The distribution of workers is given by a finite measure η in Θ . We make the following assumption on the measure η .⁵

Assumption 2.1. (*Strict preferences*) *The measure of all indifference curves $\eta(\{\theta \in \Theta | e_i^\theta = x\})$ is 0, for all firms i and productivity levels x .*

The model builds upon the continuum matching model of Azevedo and Leshno (2010). A matching is a function $\mu : \Theta \times I \rightarrow I \times 2^\Theta$ such that each worker is assigned either to a firm or to itself and each firm is assigned to a set of workers. Moreover, if a worker is matched to a firm, the firm is matched to the worker, and vice versa.⁶ Consider now a vector of capacities q . We say that a matching μ is stable (with respect to $[\eta, q]$) if⁷

1. Firms' capacity constraints are respected.
2. No worker or firm receives an unacceptable match.
3. No firm-worker pair could be made better off by matching to each other.
4. μ is right-continuous.

⁵For the sake of rigour, we make standard measurability assumptions. Note that Θ is the union of a finite number of I dimensional cubes. The measure η is assumed to be defined over the Borelians of Θ . Moreover, a matching (defined below) has to be measurable with respect to the σ -algebra generated by the Borelians.

⁶Mathematically, for all $\theta \in \Theta$, $\mu(\theta) \in I \cup \{\theta\}$ and for all $i \in I$, $\mu(i) \in 2^\Theta$. In addition, $i = \mu(\theta)$ iff $\theta \in \mu(i)$. In addition, we require the matching to be measurable with respect to the σ -algebra generated by the Borelians of Θ (see footnote 5).

⁷Formally,

1. For all i we have $\eta(\mu(i)) \leq q_i$,
2. If $\mu(\theta) = i$, then $i \succ^\theta \theta$.
3. If $i \succ^\theta \mu(\theta)$, then $\eta(\mu(i)) = q_i$, and for all $\theta' \in \mu(i)$ we have $e_i^{\theta'} \geq e_i^\theta$.
4. For all sequences θ^k , with $\theta = \lim_{k \rightarrow \infty} \theta^k$, and all $e^{\theta^k} \geq e^\theta$, we have $\mu(\theta) = \lim_{k \rightarrow \infty} \mu(\theta^k)$.

Conditions 1-3 are analogous to those in the definition of a stable matching in the discrete college admissions model of Gale and Shapley (1962). A stable matching is a resting point for the market, where agents cannot gain by breaking off their matches and seeking new ones. Markets organized around centralized clearinghouses often use algorithms that produce stable matchings. Roth (1991) gives evidence that market clearinghouses organized around stable allocations tend to thrive, while those that are not tend to unravel and eventually collapse. In decentralized markets, there is also evidence that sometimes outcomes correspond to stable matchings. Hitsch et al. (2010) and Sorensen (2007) give evidence that matchings in an online dating website, and among venture capitalists and biotechnology firms are stable. Some theoretical models give conditions under which decentralized matching processes result in stable matchings (Adachi (2003); Niederle and Yariv (2009)), at least when search frictions are small. Moreover, Roth and Vate (1990) show that a stable matching can be achieved by starting with any matching, and having blocking pairs randomly break their current connections and match to better partners. These findings motivate using stability as a solution concept in a first pass at modelling imperfect competition in matching markets. As for condition 4, it is a technical condition, which avoids multiplicity of stable matchings which coincide up to measure 0, and simplifies the analysis without sacrificing generality. Note that condition 3 must hold pointwise. That is, no firm-worker pair can be better off matching outside the allocation, not even a measure 0 set of workers.

The following result gives sufficient conditions for a unique stable matching to exist. Let $T = \{\theta \in \Theta \mid i \succ^\theta \theta \forall i \in I\}$.

Proposition 2.2. *If all subsets of T of the form*

$$(\{e^\theta < p\} \setminus \{e^\theta < p'\}) \cap T$$

with $p' \leq p$ and $p' \neq p$ have strictly positive η measure, then for all q there exists a unique stable matching with respect to $[\eta, q]$. This holds, in particular, if η has full support.

Henceforth, we will simply take η such that for all q there exists a unique stable matching (with respect to $[\eta, q]$). The Proposition guarantees that there is no great loss of generality in doing so, as we could for example have assumed that η has full support.⁸ Moreover, Azevedo and Leshno (2010) show that, when the continuum model has a unique stable matching, it equals the limit of the set of stable matchings of any approximating sequence of discrete economies. This justifies using the continuum model as an approximation of discrete economies with a large number of workers per firm.

Assumption 2.3. (Uniqueness) *For all q , there exists a unique stable matching with respect to $[\eta, q]$.*

⁸In addition, AL show that as long as η satisfies minimal smoothness requirements, for almost every q there is a unique matching stable with respect to $[\eta, q]$. Therefore, uniqueness is the typical case in this model.

2.1.2 The game

We now lay out the oligopoly game considered, where the players are the firms I . The primitives are $I, \Theta, \eta, c(\cdot), w_i, Q_i$.

1. Firms simultaneously choose capacities q_i in compact intervals Q_i .
2. After capacity choices q , workers are hired according to the unique matching μ_q stable with respect to $[\eta, q]$.
3. Each firm's payoff is given by

$$\Pi_i = \int_{\mu_q(i)} [e_i^\theta - w_i] d\eta(\theta) - c_i(q_i), \quad (1)$$

where w_i is firm i 's wage, and the continuously differentiable function $c_i(\cdot)$ is the cost of investing in capacity. That is, profits are the integral of the productivity of the workers hired, net of the wage, minus the costs of investing in capacity.

Therefore, the game corresponds to a situation where firms simultaneously commit to capacity choices. Workers are then assigned according to a stable matching, holding capacities fixed. The model is the matching analogue of the Cournot (1838) model in homogeneous good markets. In the next sections, we will generalize the model, to consider the case where firms may set wages w_i , and also the case where wages are personalized for each worker. Because stable matchings are used to model centralized and decentralized outcomes, this can correspond to a variety of situations, which need not even be labor markets. For concreteness, we will maintain the firm-worker terminology, until we discuss applications, but the following examples should also be kept in mind.

1. In the United States, every year 25,000 new medical school graduates apply to residency positions in hospitals through a centralized clearinghouse, the National Resident Matching Program (NRMP).⁹ In several other countries, doctors are matched to residency positions in similar clearinghouses.¹⁰ The algorithm used by the NRMP, and many other clearinghouses, is a variant of the deferred acceptance mechanism, proposed by Gale and Shapley (1962), and always produces a stable matching. Although the deferred acceptance mechanism is strategy-proof for the doctors, it is not strategy-proof for the hospitals. One interpretation of the model is that $c(\cdot) \equiv 0$, and hospitals with real capacity $\bar{q}_i$ are strategically deciding whether to misrepresent their capacities.

⁹See Roth and Peranson (1999). One aspect in which application does not fit the model well, is that the number of doctors matched to each program is small, and the continuum model represents a large number of doctors. In addition, the mechanism used is not exactly equal to the deferred acceptance mechanism, as it includes special provisions for couples, for example.

¹⁰Roth (1991)

2. In Hungary and Turkey, college admissions are coordinated by centralized clearinghouses, which use the deferred acceptance mechanism. Similar mechanisms are used for assigning students to public high schools in Hungary, and in some American cities.¹¹ The model can be interpreted as an equilibrium model of how universities or schools behave in the mechanism.
3. In a hypothetical country, the majority of the best and brightest students of each generation attend one of a handful of elite universities. Therefore, each one of these has a relatively large share of the market, and some market power. Capacity investment, in dormitories and other facilities, is made a year in advance. Although the admissions process is decentralized, it approximately corresponds to a stable matching.
4. Two hypothetical strategic consulting companies, B and M, control most of the market for new college graduates going into the business. In each firm, all entering employees are paid the same wage. They choose the number of hires (and possibly wages) a month in advance of the market, which is decentralized but produces a stable matching.
5. In the United States, many of the graduates of the nation's best law schools join one of the top law firms in the country (Ginsburg and Wolf (2003)). The entry level position in this market is referred to as associate. An interesting feature of this market is that inside each firm the vast majority of associates are paid the same wage. Moreover, across firms, wages are mostly the same, with even the end of year bonuses being equal. Despite compensation being uniform across firms, candidates have strong preferences over firms, and pay close attention to prestige rankings in the industry. Most firms hire a large number of associates every year, sometimes over 100. Moreover, since firms only compete directly with other firms with a similar prestige ranking, it is reasonable to assume that they have some market power.
6. In markets for highly differentiated services, buyers and sellers may have (possibly heterogeneous) preferences over trade partners. For example, entrepreneurs and venture capitalists (Sorensen (2007)), advertisers and content providers (newspapers, content websites, search engines, or television channels). Moreover, in some of these markets, firms on one of the sides have large market share, and therefore have some market power (e.g. Google, Yahoo!, Bing and Baidu's large share of the search engine advertisement market, or CBS, NBC, ABC, Fox and Univision's large share of the American broadcast TV market).

The solution concept we adopt for most of the analysis is pure-strategy Nash equilibrium. Let $Q = \times_i Q_i$. A profile $q^* \in Q$ is an equilibrium if

$$\Pi_i(q^*) \geq \Pi_i(q_i, q_{-i}^*),$$

¹¹See Balinski and Sönmez (1999); Biró (2007); Abdulkadiroglu et al. (2009).

for all i and $q_i \in Q_i$. In some cases, pure strategy equilibria may fail to exist. We define a mixed strategy as a probability distribution σ_i over Q_i . We will abuse notation and denote expected payoffs given a mixed strategy profile by $\Pi(\sigma)$. A profile of mixed strategies σ^* is a mixed strategy Nash equilibrium if it is a pure strategy Nash equilibrium of the game where action spaces are ΔQ_i . Azevedo and Leshno (2010) show that, in our setting, the unique stable matching with respect to $[\eta, q]$ varies continuously in q . This guarantees that a mixed strategy equilibrium always exists.

Proposition 2.4. *At least one mixed strategy equilibrium exists.*

2.1.3 Cutoffs

In this section we make a preliminary observation, which is a centerpiece of the analysis. We note that any stable matching must have a very particular structure. Moreover, stable matchings can be defined in a decentralized way, by considering the productivity of a marginal hired worker at each firm. A cutoff is simply a threshold p_i in $[0, 1]$, such that firm i accepts workers with a productivity higher than p_i , and rejects workers with lower productivity. A cutoff vector is a vector p specifying cutoffs for each firm. Given cutoffs, we may define a worker's demand as her favorite firm that would accept her. That is,

$$D^\theta(p) = \arg \max_{i \in \mathcal{I}} \{i | p_i \leq e_i^\theta\}.$$

To simplify notation, we assume that $D^\theta(p)$ is an I dimensional vector, with value 1 in the coordinate corresponding to the chosen firm, and 0 in the other coordinates. We can then define aggregate demand as the I dimensional vector

$$D(p) = \int D^\theta(p) d\eta(\theta).$$

We now define a market clearing cutoff

Definition 2.5. A cutoff vector p is a market clearing cutoff with respect to $[\eta, q]$ if for all i

$$D_i(p) \leq q_i$$

with equality if $p_i > 0$.

The first observation we make is that there is a natural bijection between market clearing cutoffs and stable matchings. Given a stable matching μ , consider the operator $p = \mathcal{P}\mu$, where

$$p_i = \min_{\mu(i)} e_i^\theta,$$

if $\eta(\mu(i)) = q_i$, and $p_i = 0$ otherwise. And given a market clearing cutoff p let $\mu = \mathcal{M}p$ where

$$\mu(\theta) = D^\theta(p).$$

We then have

Lemma 2.6. (*Cutoff Lemma - Azevedo and Leshno (2010)*) Given $[\eta, q]$, there exists a unique market clearing cutoff p . Let μ be the unique stable matching. Then $p = \mathcal{P}\mu$ and $\mu = \mathcal{M}p$.

Azevedo and Leshno (2010) give a more general statement of this Lemma, in the case where there may be multiple stable matchings. The intuition is that there is a parametrization of stable matchings by the admission thresholds in each firm. The possibility of such a parametrization is suggested by a result by Roth and Sotomayor (1989), who show that the pools of workers hired by a firm in different stable matchings are always ordered by first order stochastic dominance. The cutoff Lemma generalizes a similar characterization by Abdulkadiroglu et al. (2008), who introduced the term cutoffs, in a setting where all firms have the same preferences. Note that, in the discrete college admissions problem, a similar result holds, with the caveat that a single stable matching may correspond to several market clearing cutoffs. The fact that this relationship is bijective in the continuum model makes it more tractable. In the discrete setting, this alternative definition of stability was stated in a different form by Biró (2007). We refer the interested reader to Azevedo and Leshno (2010) for proofs and a discussion of the result in both discrete and continuous settings.

In addition to being an admission threshold, a cutoff p_i is the productivity of a marginal hired (or rejected) worker at firm i . Therefore, we may think of cutoffs as being the marginal value of capacity for firm i ,¹² holding the pool of applicants fixed. As the shadow price of capacity, cutoffs share many properties with prices. Indeed, this analogy and basic price theory will be a central part of our argument, and we will analyze distortions caused by market power as wedges between marginal revenues and cutoffs faced by firms with nontrivial market share. However, we must caution that cutoffs are not prices. Therefore, several price theoretical insights do not carry over to cutoffs. For example, unlike prices, an agent may only demand a firm if her productivity *at that firm* is greater than the cutoff. Consequently, ratios between cutoffs are not informative of marginal rates of substitution. This dual interpretation of cutoffs is in the same spirit of a characterization of stable matchings due to Adachi (2000), who shows that stable matchings can be parametrized by the utilities of agents on both sides of the matching. This influential idea has been fruitfully applied to several problems in matching theory (Echenique and Oviedo (2004, 2006); Hatfield and Milgrom (2005); Ostrovsky (2008)). However, it is different from cutoffs, which only require keeping track of one threshold for each firm, as opposed to a reservation value for each firm and a reservation value for each worker.

2.1.4 Equilibrium

Cutoffs play a major role in the analysis of the oligopoly game. As explained in the definition of the game, each strategy profile q induces a unique stable matching μ_q . Henceforth, we will denote by $P(q)$ the vector of market clearing cutoffs associated with μ_q . That is, the function P gives the productivity $P_i(q)$

¹²That is, net of wages and investment costs.

Figure 1: The relevant portion of the set of agent types in example 2.8. The left square represents agents who prefer firm 1, and the right square agents who prefer firm 2. Coordinates correspond to productivities. The numbers denote the firm to which agents in each region are matched, for a cutoff vector (p_1, p_2) . The red lines represent the regions Δ_{12} and Δ_{21} of agents which are in the margin of switching firms, which are used to define the quantities M_{ij} and $\bar{P}_{ij}$ in Proposition 2.11.

of a marginal hired worker at firm i when the quantities played are q . First note that, when a firm increases its quantity, it lowers cutoffs facing all firms.

Lemma 2.7. *If $q' \geq q$, then $P(q') \leq P(q)$.*

This is the continuum analogue of the comparative statics results of Gale and Sotomayor (1985a,b). To clarify the definitions, and develop intuition, we now consider a simple example.

Example 2.8. There are two firms, which have a maximum capacity of 1, so that $Q_i = [0, 1]$. For simplicity, assume that wages and the cost of capacity are $w_i \equiv c_i(\cdot) \equiv 0$. There is a mass 1 of agents with preference list $\succ^\theta = 1, 2, \theta$ and a mass 1 with preferences $2, 1, \theta$. Productivity vectors e^θ are uniformly distributed in $[0, 1]^2$, and independent of preferences. Figure 1 depicts the relevant portion of the set of agent types. Note that, since all workers are productive, the most efficient allocation would be for all of them to be employed. If both firms set $q_i = 1$, all workers are hired by their favorite firm, implying cutoffs $p_i = 0$.

Consider now the market clearing equations for this economy. If capacities are given by q , market clearing equations are

$$q_i = (1 + p_{-i})(1 - p_i).$$

That is, for each firm i , a measure $1 - p_i$ of the agents with preferences $i, -i, \theta$ are accepted. Plus, a measure $p_{-i}(1 - p_i)$ of the agents with preferences $-i, i, \theta$,

which were rejected by firm $-i$. Figure 1 illustrates the allocation given cut-offs. These equations have a unique solution, which defines the unique market clearing cutoff $P(q)$ as a function of capacities. For example, assume firm 1 sets a quantity of $q_1 = 1/2$, while firm 2 sets $q_2 = 1$. Solving the system then yields $P_1 = (\sqrt{17} + 1)/8 \approx .64$, and $P_2 = (\sqrt{17} - 1)/8 \approx .39$. Therefore, when firm 1 reduces its capacity to $1/2$, it becomes more selective, and raises its cut-off, from 0 to .64. Even though firm 2 is still supplying full capacity $q_2 = 1$, its cutoff also goes up, albeit only to .39. Even though we computed the stable matching using the market clearing equations, Azevedo and Leshno (2010) show that it could also be computed as the outcome of the continuum version of the deferred acceptance mechanism. Since there is a unique stable matching, it does not matter whether the worker-proposing or firm-proposing version of the algorithm is used, the outcome is the same. However, cutoffs provide an analytically convenient way to compute the stable matching.

The market clearing equations can also be used to calculate optimal strategies by each firm. Due to the uniform distribution, profits are given by $\Pi_i = q_i \cdot (1 + p_i)/2$. To find the optimal q , we use the implicit function theorem to calculate the marginal revenue to firm i increasing its capacity. Straightforward algebra shows that

$$MR_i = P_i - (1 - P_i) \frac{1 - P_i}{2} \left(-\frac{dP_{-i}}{dq_i} \right).$$

In a symmetric equilibrium, we must have $MR_i(q^*, q^*) = 0$. Solving this equation we get $q^* = 4\sqrt{5} - 8 \approx .94$. Therefore, in equilibrium, some workers remain unemployed, even though firms cannot affect wages $w_i = 0$, and there are no costs of providing more capacity.

The puzzling feature of the example is that firms do not hire some workers with positive net productivity, even though they have enough capacity to do so. The reason why firms choose to reduce capacity in equilibrium is the possibility of rejection chains (Sönmez (1997)). That is, by rejecting a worker, the firm sends him back to the worker pool. The rejected worker may then be hired by a competing firm, which will in turn reject another worker. Possibly, this newly rejected worker will then apply to the original firm, and be more productive than the original rejected worker. By reducing capacity, firms are shedding marginal workers, but they may gain workers who are marginal to the other firms. If preferences are not perfectly correlated, these workers may be better than the rejected workers.

To gain some insight into this mechanism, we now consider an expression for the marginal revenue, net of wages and investment costs of each firm. That is, let

$$\begin{aligned} R_i(q) &= \int_{\mu_q(i)} e_i^\theta d\eta(\theta), \text{ and} \\ MR_i(q) &= \partial_{q_i} R_i(q), \end{aligned}$$

in the case where this derivative exists. We now restrict attention to interior points.

Definition 2.9. A vector of strategies $\tilde{q}$ is interior if it is in the interior of the set $\{q \in \times_i Q_i | \eta(\mu_q(i)) = q_i\}$.

In any interior equilibrium q^* , if profits are differentiable, then firm i 's quantity choice must satisfy the first order condition:

Proposition 2.10. *In any equilibrium q^* where firm i 's profits $\Pi_i(q)$ are differentiable in q_i and q^* is an interior point we have*

$$MR_i(q^*) = w_i + c'(q_i).$$

This expression is analogous to the first order condition facing a firm in the Cournot oligopoly model (Figure 3). It says that the marginal productivity gain from increasing quantity must equal the marginal cost in wages and investment in capacity. If firm i has no market power, increasing the quantity of hires by a small amount dq would add dq marginal workers, of productivity P_i to its worker pool. Therefore marginal revenue would be $MR_i = P_i$. However, when firms have market power, rejection chains induce a wedge between marginal revenue MR_i and cutoffs P_i . In the case where η admits a continuous density f , there is a simple intuitive expression for this wedge. Denote the set of workers which would be accepted at firm i , but are marginally accepted by a firm j that the worker prefers as

$$\Delta_{ij}(q) = \{e_i^\theta \geq p_i, e_j^\theta = p_j, j \succ^\theta i \succ^\theta \theta, k \succ^\theta i \Rightarrow e_k^\theta < p_k\},$$

where $p = P(q)$. Roughly speaking, these are the workers which firms i and j compete for, and firm i may hope to poach them from firm j (see Figure 1). Let

$$\begin{aligned} M_{ij} &= \int_{\Delta_{ij}} f(\theta) d\theta, \\ \bar{P}_{ij} &= \int_{\Delta_{ij}} e_i^\theta f(\theta) d\theta / M_{ij} \text{ if } M_{ij} \neq 0 \\ &= P_i \text{ if } M_{ij} = 0. \end{aligned}$$

That is, M_{ij} is the $I - 1$ dimensional mass of the set Δ_{ij} of disputed agents, and $\bar{P}_{ij}$ is their average productivity. If firm i has some market power, its quantity decisions affect the cutoff P_j . Therefore, by reducing quantity, firm i increases the cutoffs of firm j , and gains some of the agents in the disputed set Δ_{ij} . That is, by rejecting a small mass of agents dq , firm i can cause firm j to reject some of the agents in Δ_{ij} , which will then apply to firm i via a rejection chain. We have the following expression for the wedge between marginal revenue and productivity of a marginal worker.

Theorem 2.11. *Assume η admits a continuous density f . Then $P(q)$ is continuously differentiable at almost every interior point q , and*

$$MR_i(q) = P_i(q) - \sum_{j \neq i} [\bar{P}_{ij}(q) - P_i(q)] \cdot M_{ij}(q) \cdot \left(-\frac{dP_j(q)}{dq_i}\right).$$

Consequently, $MR_i(q) \leq P_i(q)$.

The intuition for this formula is as follows. When firm i reduces capacity by dq , it loses a measure dq of workers. If firm i had no market power, those workers would have productivity equal to the cutoff, P_i , and hence the first term. The second term measures the distortions caused by market power. If firm i may affect the cutoffs of other firms, by hiring more workers it loses some workers in the set Δ_{ij} , which are marginal to firm j , but may be better than marginal for firm i . The difference $\bar{P}_{ij} - P_i$ measures this difference in productivity, and the term $M_{ij} \cdot (-\frac{dP_j}{dq_i})$ measures the mass of workers in this marginal set that are displaced. The intuition can be further clarified by considering the particular case of example 2.8. Figure 2 displays the effect of a small increase in quantity dq for firm 1. This increase in q_1 leads both p_1 and p_2 to decrease, so that cutoffs move from p to p' . The set of workers gained by firm 1 is highlighted as the two blue rectangles, while the set of workers that are lost is highlighted as the yellow rectangle. Notice that the workers gained have productivities close to p_1 . Hence, if firm has negligible market power, and its increase in quantity has a small effect in the cutoff of firm 2, the mass of workers lost would be small, and marginal revenue would be close to p_1 . However, if that is not the case, firm 1 also has to take into account that it loses the mass of workers in the yellow rectangle. Those workers have average quality $\bar{P}_{12}$. In addition, the total mass of workers lost must be approximately $M_{12} \cdot (-dP_2/dq_1) \cdot dq$. Therefore, the change in revenue is approximately $\{p_1 - (\bar{P}_{12} - p) \cdot M_{12} \cdot (-dP_2/dq_1)\}dq$.

The wedge between the productivity of a marginal worker and marginal revenue is analogous to the wedge between prices and marginal quantities in the Cournot model. Figure 3 plots marginal revenue and cutoffs, as a function of the quantity chosen by firm i . Note that the marginal revenue curve MR_i is lower than the cutoff curve P_i . Moreover, when $q_i = 0$, we have that $\bar{P}_{ij} = P_i = 1$, so that $MR_i = P_i$. Note that the first order condition for firm i is to provide quantity up to the point where the marginal revenue curve crosses the marginal cost plus wages curve, $w_i + c'_i$. Therefore, the equilibrium quantity is given by the point $q^* < 1$. At this point, we have $P_i > MR_i = w_i + c'_i$. Therefore, in equilibrium the firms do not hire some workers with strictly positive net productivity.

Equilibrium does not rely on firms reasoning through the rejection chains they set off. All that is necessary is that firms set their quantities optimally given the strategies of other firms. In several markets, firms (which may represent colleges, hospitals, or television networks) do seem to spend a lot of time deciding the quantities to be supplied. A college, for example, faces a quantity versus quality tradeoff when deciding on the size of each year's entering class. The

Figure 2: The effects of a small increase in quantity by firm 1.

Figure 3: Marginal revenue $MR_i(q)$, and the cutoff $P_i(q)$, as a function of q_i keeping q_{-i} fixed. The best response is denoted q_i^* .

equilibrium assumption is that the college gets this decision right, by any mix of trial and error, experience, or abstraction. However, it does not depend on each college fully understanding its impact on the rest of the market.

Notice that, if all n firms were acting in unison, as a monopolist maximizing profits, they would have even greater incentives for reducing capacity. Under suitable differentiability assumptions, at an interior point q we may write the first order condition with respect to q_i as

$$MR_i(q) + \sum_{j \neq i} \frac{d\Pi_j}{dq_i} \leq MR_i(q).$$

where we used the fact that all $d\Pi_j/dq_i \leq 0$.¹³ Therefore, a cartel has more incentives for quantity reduction than the oligopolists.

An immediate consequence of the previous discussion is that quantity distortions in markets with fixed wages are driven by preference heterogeneity. The reason why firm i can profit from rejection chains is that it may reject a worker θ that is accepted by firm j , and that leads firm j to reject a better worker θ' . However, if firms i and j have the same preferences this is not possible. It can also be shown that when workers have the same preferences, such rejection chains are not profitable either.

Proposition 2.12. *Assume that q is interior and either*

- (1) *(Homogeneous ordinal worker preferences) All worker types in the support of θ have the same preference ordering $\succ^\theta$, or*
- (2) *(Homogeneous ordinal firm preferences) For any two worker types $\theta \neq \theta'$ in the support of η , we have that either $e_i^\theta > e_i^{\theta'}$ for all i or $e_i^\theta < e_i^{\theta'}$ for all i .*

Then $MR_i = P_i$.

This Proposition is in line with results from the discrete matching literature, which show that stable mechanisms are strategyproof when preferences are acyclic (Kesten (2008)).¹⁴ It implies that, in settings where agents are ordered by some vertical measure of quality, there is no quantity reduction in matching with fixed wages. We will see that the opposite result holds in matching with flexible wages. Therefore, a major part of our comparison of flexible and fixed wages depends on whether preferences are heterogeneous.

There is another way to frame firms' first order condition, which yields an analogue of the classic Lerner index formula. Let the average productivity of

¹³This follows from the monotonicity result in Lemma 2.7.

¹⁴It is possible to prove such a result in our setting. That is, that under a condition analogous to the condition used by Kesten in the discrete setting firms have no incentives to reduce capacity. This can be derived for example using the Kesten result and Theorem 2 of Azevedo and Leshno (2010), which shows that the stable matching of a continuum economy correspond to the limits of stable matchings of discrete economies. We consider the less general form of the Proposition above for simplicity, and focus instead on the equilibrium consequences of firms manipulating capacity, as opposed to conditions for strategyproofness.

accepted workers be given by

$$\begin{aligned}\bar{P}_i(q) &= \int_{\mu(i)} e_i^\theta d\eta/\eta(\mu(i)) \text{ if } \eta(\mu(i)) \neq 0, \\ &= 0 \text{ otherwise.}\end{aligned}$$

At an interior point q , we may write firm i 's profits as

$$\Pi_i = (\bar{P}_i(q) - w_i) \cdot q_i - c_i(q_i).$$

Assuming $\bar{P}_i$ is differentiable, we can denote the elasticity of quantity $\bar{q}_i$ with respect to average productivity $\bar{P}_i$ as $\epsilon_i = -\bar{P}_i/q \partial_{q_i} \bar{P}_i$. Then, at any interior equilibrium q^* firm i 's first order condition can be written

$$\frac{\bar{P}_i(q^*) - w_i - c'_i(q_i^*)}{\bar{P}(q^*)} = \frac{1}{\epsilon_i}.$$

The term on the left is the fraction of the marginal worker's productivity that is above wages plus investment costs. The optimality condition says that this markup must equal the inverse of the elasticity of quantity with respect to average worker productivity. Therefore, the quantity setting game with fixed wages is equivalent to a Cournot oligopoly game, where the demand curve is replaced by the average productivity curve.¹⁵

2.2 Matching with uniform endogenous wages

In the previous section, we have taken wages w_i at each firm to be fixed and exogenous. While this may be realistic for some applications (e.g. admissions to public schools), in other settings it is more natural to consider wages as an endogenous strategic variable, which firms control (e.g. two companies competing for employees). This section shows that the arguments in the previous section extend to the case where the wage is endogenous. Moreover, the model yields predictions on the level of equilibrium wages.

Denote a cardinal agent type by θ_u . Each cardinal type derives utility $u_i^{\theta_u} + w_i$ from being employed at firm i receiving wage w_i . The utility of being unemployed is normalized to 0. Utilities belong to a compact interval U . Moreover, a cardinal type θ_u has a productivity vector $e^{\theta_u} \in [0, 1]^I$ which specifies her productivity at each firm. The set of cardinal types is $\Theta_u = [0, 1]^I \times U^I$. Let η_u be a measure over cardinal types. We maintain assumption 2.1 that firms have strict preferences in the cardinal type space. Moreover, we assume that there is no positive mass of agents that has exactly the same certainty equivalent between two choices.

¹⁵The analogy with the Cournot model illuminates the issue of equilibrium existence. The Cournot model may fail to have pure strategy equilibria, even under well behaved parameters (Roberts and Sonnenschein (1976)). However, in many concrete example it does have pure strategy equilibria. In particular examples of the present model, it is often the case that the functions $\Pi_i(q)$ are quasiconcave in q_i , and therefore that a pure strategy equilibrium exists (Fudenberg and Tirole (1991); Glicksberg (1952)). In addition, when $c_i \equiv w_i \equiv 0$, then log concavity of the function $\bar{P}_i(q)$ in q_i implies that $\Pi_i(q)$ is quasiconcave, and therefore that a pure strategy Nash equilibrium exists.

Assumption 2.13. (*Strict preferences*) For all x, i, j the sets

$$\{u_i^{\theta_u} = x\}$$

and

$$\{u_i^{\theta_u} - u_j^{\theta_u} = x\}$$

have η_u measure 0.

Under this assumption, for any vector of wages w , the set of workers who are indifferent between two firms has measure 0. Therefore, any pair $[\eta_u, w]$ induces a distribution $\eta([\eta_u, w])$ over the (ordinal) set of worker types Θ . We make the assumption that there is a unique stable matching given a measure η and vector of quantities, as in the case of exogenous wages.

Assumption 2.14. (*Uniqueness*) For all q, w in space $Q \times W$ of strategies there is a unique matching stable with respect to $[\eta([\eta_u, w]), q]$.

We define the oligopoly game with endogenous wages as follows.

1. Firms simultaneously chose quantities q_i and wages w_i from compact intervals Q_i and W_i .
2. After choices q, w , workers are hired according to the unique matching $\mu_{q,w}$ stable with respect to $[\eta([\eta_u, w]), q]$.
3. Payoffs are given as before by equation 1.

We will maintain the solution concepts, pure and mixed strategy Nash equilibrium. As in the case of exogenous wages, the results from Azevedo and Leshno (2010) guarantee that payoffs depend continuously on q, w . Consequently, at least one mixed strategy equilibrium exists.

Proposition 2.15. *The oligopoly game with endogenous wages has at least one mixed strategy equilibrium.*

Denote by $P(q, w)$ and $\bar{P}(q, w)$ the vectors of market clearing cutoffs and of average productivities given a quantity wage pair.

Theorem 2.16. *Assume η_u admits a continuous density, and that $\bar{P}$ is continuously differentiable at an interior point (q, w) . Then marginal revenue with respect to quantity is given by the formula in Proposition 2.11, and the derivative of profits with respect to wages is*

$$\partial_{w_i} \Pi_i(q, w) = \partial_{w_i} \bar{P}(q, w) - 1.$$

The Theorem guarantees that the basic first order condition of the optimal quantity choices is preserved in the model with endogenous wages. In a pure-strategy equilibrium, even if wages are endogenous, whatever the level of wages the same tradeoff analysed in the previous section must hold. Consequently, it is still the case that firms choose to reduce quantities, and do not hire workers

with strictly positive productivity net of wages and investment costs, when preferences are not perfectly correlated.

As for wage setting, the Theorem shows that in any interior pure-strategy equilibrium, the derivative of average productivity with respect to wages has to be equal to 1. That is, at the optimum, an increase of \$1 in wages increases average productivity of the hired workers by exactly \$1, holding quantity fixed. If the gain in productivity were greater, firms would be willing to increase wages, and were it smaller they would gain by reducing wages. Moreover, the number of workers whose ordinal preferences would change depends on how many workers are close to indifferent between two firms, or between working for a firm and being unemployed (which can be viewed as pursuing some alternative employment, in home production or in an unmodeled nonstrategic set of firms). Therefore, wages may be very low, and perhaps pushed towards the boundary condition $w_i = \min W_i$, if workers' decisions are relatively inelastic with respect to wages. On the other hand, if there is a large mass of workers that can potentially be poached from other firms, wages can be high, and even in the right boundary $w_i = \max W_i$ of the set of possible wages.

2.3 Matching with contracts

2.3.1 Firms, workers, and stable matchings with contracts

So far we have assumed that, within each firm, all workers are paid the same wage. However, this is not a realistic assumption in some markets. In many cases, firms and workers may be able to negotiate not only personalized salaries, but also terms of employment. For example, an economics department hiring an assistant professor may negotiate her salary, teaching load, and research budget. Moreover, different job candidates may receive different offers from the same department. Not only are flexible contracts common, but the choice of how much amplitude agents should have to personalize contracts is an important market design variable, and the subject of some debate (Bulow and Levin (2006)). In this section, we consider quantity manipulation games when wages are personalized. We will use the model of matching with contracts introduced by Azevedo and Leshno (2010). This framework is similar to models of matching with contracts introduced by Kelso Jr and Crawford (1982) and Hatfield and Milgrom (2005).¹⁶ The main differences are that it considers a continuum of workers and simpler preference structures.

A contract $x = (i, \theta, w)$ specifies a firm i , a worker θ , and other terms of employment w . The set of all available contracts is denoted by X . X also includes a null contract $\emptyset$, which corresponds to being unmatched. Firms $i \in I$ have a profit function $\pi_i(x)$ that prescribes its profit for each contract which it may be part of. Workers $\theta \in \Theta_X$ have utility functions $u^\theta(x)$ over contracts that include them. We normalize the profits and utility of being unmatched to 0. Workers are distributed according to a measure η_X over Θ_X . A matching, $\mu : \Theta \cup I \rightarrow X \times 2^X$ assigns each worker to either the empty contract or to a

¹⁶See also Hatfield et al. (2010); Hatfield and Kojima (2009, 2010)

contract that contains her, and each firm to a set of contracts that contain her (possibly the empty set). Consider now a vector of capacities q . A matching is said to be stable (with respect to $[\eta_X, q]$) if¹⁷

1. Firms' capacity constraints are respected.
2. No worker is assigned to a firm she considers unacceptable.
3. No firm-worker pair could be made better off by matching to each other.

Note that the conditions are analogous to those of section 2. The main difference is the omission of the right continuity condition. That condition was useful in section 2 to avoid some multiplicities in the set of stable matchings up to measure 0, which were not consequential to the game. In the case of matching with contracts, it is somewhat more intricate to formulate a similar condition, and the expositional gains are small, so we omit it.

For simplicity, we now assume that the only term which agents may vary in contracts are wages.¹⁸ That is

$$X = I \times \Theta_X \times \mathfrak{R}.$$

Moreover, we assume that agents have quasilinear preferences, so that for any contract $x = (i, \theta, w)$ we have $\pi_i(x) = e_i^\theta - w$, and $u^\theta(x) = u_i^\theta + w$.¹⁹ While the Azevedo and Leshno (2010) model does not depend on this assumption, and it is straightforward to extend the model without it, we maintain it for concreteness, as it clarifies the role of flexible wages in quantity competition. Given a worker-firm pair, we may define the total surplus from the match as $s_i^\theta = u_i^\theta + e_i^\theta$. Henceforth, we will assume that the set of surplus vectors s^θ of all worker types equals $[0, 1]^I$. We will make the following assumption, which plays a similar role to strict preferences in the case of fixed wages.

Assumption 2.17. (*Strict preferences*) All level sets of the surplus function of firm i , and of the difference in surplus between two firms have η measure 0. That is, for all $y \in \mathfrak{R}$ and $i \neq j \in I$ we have

$$\eta(\{\theta | s_i^\theta = y\}) = \eta(\{\theta | s_i^\theta - s_j^\theta = y\}) = 0.$$

¹⁷Mathematically, these conditions can be state as follows.

1. For all i $\eta(\{\theta | \text{a contract } x = (i, \theta, w) \in \mu(i)\}) \leq q_i$.
2. For all θ , if $\mu(\theta) = x$, then $u^\theta(x) \geq 0$.
3. There is no firm-worker pair i, θ , and contract $x = (i, \theta, w)$ such that $u^\theta(x) > u^\theta(\mu(\theta))$, and either $\eta(\{\theta | \text{a contract } x = (i, \theta, w) \in \mu(i)\}) < q_i$ or there is $x' \in \mu(i)$ such that $\pi_i(x') < \pi_i(x)$.

¹⁸The loss of generality in this assumption is actually smaller than it seems. See Echenique (2010)

¹⁹So far, I have postponed making measurability assumptions. This is for simplicity, and because measurability does not play a big role in the analysis. However, with the quasilinearity assumptions, it is easy to make the appropriate requirements. We assume that Θ_X is the set $[0, 1]^{2I}$ of all (u, e) vectors, with the σ algebra generated by the Borelians. Therefore X is just the union of a finite number of cylinders in Euclidean space, and we also endow it with the Borelian σ algebra. Matchings are assumed to be measurable.

Azevedo and Leshno (2010) show that, under this assumption, stable matchings correspond to market clearing cutoffs. A cutoff $p_i \in [0, 1]$ specifies a threshold for firm i , such that it will only accept contracts that yield profits in excess of p_i . Equivalently, it may be thought as the profits derived from a marginal selected (or rejected) contract, or as the shadow price of capacity. Given a stable matching μ , the associated cutoffs are given by $p = \mathcal{P}\mu$, where

$$p_i = \inf_{x \in \mu(i)} \pi_i(x),$$

if $\eta_X(\mu(i)) = q_i$ or 0 otherwise.²⁰ Given cutoffs, demand for each firm is

$$D_i(p) = \eta_X(\theta | s_i^\theta - p_i \geq s_j^\theta - p_j \text{ for all } j \text{ in } I \text{ and } s_j^\theta - p_j \geq 0).$$

Similarly to matching with fixed wages, there are mild conditions that guarantee that an economy has a unique market clearing cutoff.

Proposition 2.18. *If all subsets of Θ_X of the form*

$$\{s^\theta < p\} \setminus \{s^\theta < p'\}$$

with $p' \leq p$ and $p' \neq p$ have strictly positive η_X measure, then for all q there is a unique market clearing cutoff with respect to $[\eta_X, q]$. This holds, in particular, if η_X has full support.

Based on this Proposition, henceforth we will simply assume uniqueness.

Assumption 2.19. (Uniqueness) *For all q there exists a unique market clearing cutoff with respect to $[\eta_X, q]$.*

Under these assumptions, the set of stable matchings has the following simple characterization.

Lemma 2.20. *There exists a unique market clearing cutoff p , solving*

$$D_i(p) \leq q_i,$$

with equality if $p_i > 0$. There exist (possibly more than one) stable matchings. For every stable matching μ , $\mathcal{P}\mu = p$. Moreover,

$$\mu(\theta) \in \arg \max_{i \in I \cup \{\emptyset\}} s_i^\theta - p_i,$$

where we define $s_\emptyset^\theta - p_\emptyset = 0$.

The Lemma guarantees the existence of stable matchings, and of a market clearing cutoff. In addition, any stable matching corresponds to the unique market clearing cutoff. Moreover, the Lemma says that each agent is matched to the firm where $s_i^\theta - p_i$ is the highest. Intuitively, this is the firm which is willing to bid more for the agent, in a combination of wages and match specific

²⁰Here we have abused notation by denoting the set of workers matched to i as $\mu(i)$.

utility u_i^θ . An important difference between matching with contracts and with fixed wages, is that in matching with contracts a single market clearing cutoff may correspond to several stable matchings. The basic reason is that, given an allocation of workers to firms, wages are not uniquely determined. Consider a worker θ that is assigned to firm i . Firm i would be willing to pay the worker a share of up to $s_i^\theta - p_i$ of the total surplus s_i^θ of the relationship. If the worker's second best choice is firm j , the worker would demand receiving at least a share $s_j^\theta - p_j$ of the surplus, or else she would rather match with firm j . Therefore, the worker's utility from the match may be any point in the interval

$$s_j^\theta - p_j \leq u_i^\theta + \text{wage} \leq s_i^\theta - p_i.$$

Since for given quantities there are multiple stable matchings, in defining the oligopoly game we must specify how to select one of them. For simplicity, we will assume that the selected matching is that in which the firm has all the bargaining power, and captures all of the gains of the relationship vis a vis the worker's second best offer. That is, wages are such that each worker is just indifferent between working for the firm that hires it, or joining the second best alternative. Formally, given cutoffs p , consider a worker θ . If $i^* = \arg \max_{i \in I \cup \{\emptyset\}} s_i^\theta - p_i$, let

$$\bar{u}^\theta = \max_{I \cup \{\emptyset\} \setminus \{i^*\}} s_i^\theta - p_i.$$

We will assume that the worker's wage is such that $u_{i^*}^\theta + \text{wage} = \bar{u}^\theta$, so that the firm i^* derives surplus

$$\pi(x) = e_i^\theta - \text{wage} = p_i + (s_i^\theta - p_i) - \bar{u}^\theta = s_i^\theta - \bar{u}^\theta$$

from hiring worker θ . We will say that a stable matching that satisfies this equation is a stable matching with minimal wages.

2.3.2 The game

We define the oligopoly game with flexible wages as follows. The primitives are $X, \Theta_X, \eta_X, c(\cdot), Q_i$ and the set of players I .

1. Firms simultaneously choose quantities q_i in the compact intervals Q_i .
2. After capacity choices q , workers are hired according to a matching stable with respect to $[\eta_X, q]$. Let p denote the unique vector of market clearing cutoffs, and let μ be a stable matching with minimal wages.
3. Each firm's payoff is given by

$$\Pi_i = \int_{\mu(i)} \pi_i(\mu(\theta)) d\eta(\theta) - c_i(q_i)$$

where the continuously differentiable function $c_i(\cdot)$ is the cost of investing in capacity. That is, profits are the integral of the productivity of all

contracts signed, minus the costs of investing in capacity. Because the matching has minimal wages, this must equal

$$\Pi_i = \int_{\mu(i)} s_i^\theta - \bar{u}^\theta d\eta(\theta) - c_i(q_i),$$

that is, total surplus from all matched workers, minus the worker's utility from her second best option.

The game corresponds to a situation where firms first invest in capacities. The matching and wages are then determined by the firm-optimal stable allocation given those capacities. The interpretation is that first firms invest in capacities, and then straightforwardly compete bidding up workers wages. Another plausible specification of the game would have been to assume that firms commit both quantities and to a schedule of wages, conditional on worker types, and then matching takes place. That is, firms can commit to paying wages according to a formula, that takes all observable worker characteristics into account. Even though the latter model seems reasonable in many cases, I chose to analyze the first because it typically seems more relevant, given the limited level of commitment that firms have in a market with flexible wages. In addition, it more closely resembles the equilibrium concepts used in previous work in the assortative case (Bulow and Levin (2006); Gabaix and Landier (2008); Tervio (2008); Mailath et al. (2010)). Moreover, it relates more closely to the Cournot model, where firms commit to capacities, and then prices are given by competitive equilibrium. However, it would be interesting to examine other strategic variables. In the homogeneous goods case, for example, games with commitment variables such as prices (Bertrand), prices and quantities (Kreps and Scheinkman (1983)), or supply functions (supply function equilibria of Grossman (1981); Hart (1985); Klemperer and Meyer (1989)) yield interesting insights. These possibilities may also be interesting in the two-sided matching case, but in the interest of space we leave them for future research.

2.3.3 Equilibrium

Let $P(q)$ be the unique vector of market clearing cutoffs with respect to $[\eta_X, q]$. We will first consider a simple example of the model.

Example 2.21. There two firms, with no cost of investing in capacity, $c_i \equiv 0$, and $Q_i = [0, 1]$. There is a mass 2 of workers, with surplus vectors $(s_1^\theta, s_2^\theta) = s^\theta = e^\theta + u^\theta$ uniformly distributed in $[0, 1]^2$. Figure 4 illustrates a typical stable matching, with cutoffs (p_1, p_2) . Workers are always assigned to the firm where $s_i^\theta - p_i$ is the highest, provided it is positive. Therefore, all workers with surplus vectors in region $H_1 \cup H_{12}$ are assigned to firm 1, and the workers with surplus vector in region $H_2 \cup H_{21}$ are assigned to firm 2 (Figure 4). To illustrate how cutoffs are determined, consider the case where firm 1 sets $q_1 = 1/2$, while firm 2 supplies maximum capacity $q_2 = 1$. Since the mass of unemployed workers is $1/2$, we must have

$$2 \cdot p_1 p_2 = 1/2.$$

Figure 4: A stable matching, plotted over the set of possible surplus vectors $s^\theta = (s_1^\theta, s_2^\theta) = u^\theta + e^\theta$.

Moreover, the market clearing equation for firm 1 yields

$$q_1 = 1/2 = 2(1 - p_1)p_2 + 2(1 - p_1)^2/2.$$

Solving these equations yields $p_1 \approx .60$ and $p_2 \approx .42$.²¹ Note that not all workers are assigned to the firm where surplus s_i^θ is the highest. Since firms have limited capacity, workers are instead assigned to where $s_i^\theta - p_i$ is the highest. In the example, because firm 1 has smaller capacity, it has a higher cutoff, and therefore some worker with higher surplus at firm 1 ($s_1^\theta > s_2^\theta$) are assigned to firm 2.

Consider now a symmetric pure strategy equilibrium, where $q_1 = q_2 = q^*$. By market clearing, cutoffs are given by $q^* = 1 - p^{*2}$, where $p^* = P_i(q^*, q^*)$. Some algebra shows that the derivative of Π_i with respect to q_i is given by

$$MR_i(q^*, q^*) = p^* - (1 - p^*)^2 \cdot \left(-\frac{dP_j}{dq_i}\right).$$

Solving the first order condition then yields that in equilibrium $p^* \approx .36$ and $q^* \approx .87$.²²

The example shows that, unsurprisingly, firms have incentives to reduce quantity when wages are personalized. The interesting point is that, although

²¹The exact formulas are complicated, as p_1 is the solution to a cubic equation. Some algebra yields

$$p_1 = -2/3 \cos\left(1/3 \arctan\left(3/5 \sqrt{3}\sqrt{37}\right)\right) + 2/3 + 2/3 \sin\left(1/3 \arctan\left(3/5 \sqrt{3}\sqrt{37}\right)\right) \sqrt{3}$$

and $p_2 = 1/4p_1$.

²²The exact value is the solution to a cubic equation.

$$p^* = -1/6 \sqrt{3}\sqrt{7} \cos\left(1/3 \arctan\left(2/9 \sqrt{3}\sqrt{79}\right)\right) + 1/2 + 1/2 \sqrt{7} \sin\left(1/3 \arctan\left(2/9 \sqrt{3}\sqrt{79}\right)\right),$$

and $q^* = 1 - p^{*2}$.

firms may reduce quantity with either fixed or flexible wages, the reasons to do so are very different. When wages are uniform, firms choose to reduce quantity when their preferences are heterogeneous, so that firing a worker may set off a profitable rejection chain. With personalized wages, we will see that firms choose to reduce quantity exactly when their incentives coincide. Given quantities q , let the set of agents that are hired by firm i but are not coveted by any other firms as

$$H_i(q) = \{\theta | s_i^\theta \geq p_i, s_j^\theta < p_j \text{ for all } j \neq i\}.$$

Moreover, define as $H_{ij}(q)$ the set of agents who are hired by firm i , and whose second best choice is firm j .

$$H_{ij}(q) = \{\theta | s_i^\theta - p_i \geq s_j^\theta - p_j \geq s_k^\theta - p_k \text{ for all } k \neq i, j \text{ and } s_j^\theta - p_j \geq 0\}.$$

Consider now the expression for the marginal revenue of firm i . Denote by the revenue of firm i its profits net of investment costs. That is, $R_i = \Pi_i(q) + c_i(q_i)$. We may write revenue as

$$R_i(q) = \int_{H_i(q)} s_i^\theta d\eta(\theta) + \sum_{j \neq i} \int_{H_{ij}(q)} s_i^\theta - s_j^\theta + P_j(q) d\eta(\theta). \quad (2)$$

The first term is the revenue from hiring workers for whom firm i is the sole bidder. Since firm i captures all the surplus of the relationship, revenues in this region are simply the integral of the surplus of each employment relationship. The other terms are the sum of the profits from hiring workers for whom firm i competes with firm j . In these region, revenue is the integral of surplus, but net of the cost of outbidding firm j , which must be $s_j^\theta - P_j(q)$ (that is, the surplus firm j could obtain, minus firm j 's shadow price of capacity).

Let MR_i be the derivative of revenue with respect to q_i . The following Theorem gives an expression for the marginal revenue of firm i .

Theorem 2.22. *Assume the distribution of surplus vectors s^θ has admits a continuous density. Then $P(q)$ is continuously differentiable at almost every interior point q , and*

$$MR_i(q) = P_i(q) - \sum_{j \neq i} \eta(H_{ij}(q)) \cdot \left(-\frac{dP_j(q)}{dq_i}\right).$$

Consequently, $MR_i(q) \leq P_i(q)$.

The intuition behind the Theorem is as follows. If firm i raises its quantity by a small amount dq , it hires some additional workers. These come from small changes in the sets H_i, H_{ij} of hires. Since all these new workers were on the margin of being hired (because wages are flexible), the average profit from hiring them must be P_i . This means a gain of $P_i dq$ in productivity. On the other hand, by raising its quantity firm i drives up the cutoffs P_j of other firms. This means that firm j will bid more aggressively for workers. Therefore firm i must pay an extra $-\frac{dP_j}{dq_i} \cdot dq$ for the mass $\eta(H_{ij})$ of workers in the set H_{ij} .

As in the case of fixed wages, the marginal revenue curve MR_i is lower than the cutoff curve P_i . Therefore, in equilibrium, the profitability of a marginal hired worker is again higher than the cost of investing in capacity (as in Figure 3, replacing the $c'_i + w_i$ curve for c'_i). Consequently, firms with market power have incentives to reduce their capacities. The reason why is that a firm i with nontrivial market share may affect the cutoffs P_j of other firms, and therefore how much they are willing to bid for workers.

Note that firm i has more incentives to reduce capacity the larger the mass of contested workers on the sets H_{ij} . That is, the greater the number of workers that both firms are interested in, the greater the incentives to reduce capacity. This is the opposite of the conclusion from the analysis of markets with fixed wages. With fixed wages, firms further distort their capacity choices the more they disagree on who the best workers are. With flexible wages, distortions are greater the more firms' preferences agree.

3 Applications

3.1 Comparison between uniform and personalized wages

While in some markets workers are paid uniform wages in each firm, in others wages are personalized. For example, most top American law firms pay all graduating law students joining the firm as an associate the same wage. (Ginsburg and Wolf (2003)). In contrast, senior lawyers receive personalized wages. When a firm hires a partner from a rival firm, what practitioners call a lateral transaction, the offer is often personalized, and lawyers of comparable seniority receive very different wages. Likewise, major Web search engines charge different prices from different advertisers. Google for example assigns each advertiser a “quality score”, calculated according to a secret formula, which confers advantages in its auctions for advertisement slots (Edelman et al. (2007)). Moreover, in markets organized around a centralized clearinghouse, a key design variable is whether the matching mechanism should use uniform wages, or allow for flexible wages, as in the mechanism proposed by Crawford (2008). Several papers have looked at the question of which form of pricing is more efficient. However, in most models in this literature, personalized prices always yield higher efficiency. In contrast, the present model implies that personalized prices have benefits (through higher matching efficiency) and costs (by possibly exacerbating quantity distortions). We have the following result.

Theorem 3.1. *For a given vector of capacities q , matching with personalized wages is always more efficient. However, if firms are allowed to set quantities, uniform wages may dominate personalized wages, as the latter may induce more capacity manipulation.*

Mailath et al. (2010) have proposed a continuum model of one-to-one assortative matching in these comparing flexible and uniform wages. In their model, due to the simple preferences, and each firm only hiring one worker, matching

is always efficient. However they show that when prices are not personalized, workers may not have the right ex-ante incentives to invest in skills, as the market does not fully compensate them for it, which causes uniform wages to be less efficient.²³ More closely related to our model, Bulow and Levin (2006) propose a discrete model of one-to-one assortative matching where firms must make wage offers simultaneously. Because firms must make offers simultaneously, if wages are uniform they employ mixed strategies. Therefore, the uniform wages generate a small degree of inefficiency vis a vis personalized wages. The model also implies that wages are lower, and more compressed than in the flexible case.²⁴

To see how uniform wages may yield higher efficiency than personalized wages, we consider an example of the Bulow and Levin (2006) model (albeit with each firm hiring a continuum of workers). We will show that, varying how similar firms are, we can reverse the Bulow and Levin (2006) result, and make either uniform or personalized wages more efficient. There are two firms $i = 1, 2$, and a mass 1 of workers. Worker types are indexed by a productivity parameter θ uniformly distributed in $[0, 1]$. At firm 1, a worker of type θ generates $e_1^\theta = \theta$. However, firm two is more productive, and multiplies the output of a worker $e_2^\theta = A\theta = (1 + \epsilon)\theta$, $0 < \epsilon < (\sqrt{5} - 1)/2$. Workers are assumed to have utility 0 from matching with either firm, and in case they offer equal wages they choose randomly between them.²⁵ Capacity at each firm is constrained to the interval $Q_i = [0, 1/2]$, and there are no costs of investing in capacity $c_i \equiv 0$.

Consider first the case of uniform wages, where firms play the game in Section 2.2. In this case, equilibrium coincides with the solution in Bulow and Levin (2006). Firms offer random wages, since a deterministic wage offer could be undercut by a rival. However, the most efficient firm 2 pays on average higher wages, and $\Pr(w_2 > w_1) = (1 + 2\epsilon)/2(1 + \epsilon) > 1/2$. Firms always supply the maximum quantity $q_i = 1/2$. While this is a mixed equilibrium, this is consistent with the findings of Section 2.2, where we saw that the incentives to reduce capacity when wages are uniform and preferences are homogeneous is small. Therefore, uniform wages imply allocative inefficiency with positive probability, but no quantity reduction.

The opposite happens with personalized wages. Given quantities, matching is efficient. However, aligned preferences and flexible wages give both firms strong incentives to reduce capacity. Applying our formulae for marginal revenue, we obtain $MR_1 = p_1 - q_1/A$, while $MR_2 = p_2 - q_2$. Equilibrium quantities are therefore highly depressed, as depicted in Figure 3. We have

$$q_1^* = \frac{1 + \epsilon}{3 + 3\epsilon + \epsilon^2}, q_2^* = \frac{1 + 2\epsilon + \epsilon^2}{3 + 3\epsilon + \epsilon^2}.$$

²³The inefficiency due to workers not investing in skills because they do not receive the full value of the surplus they generate is similar to the inefficiency in the classic analysis of the hold up problem due to incomplete contracting, albeit here it arises in a market setting (Klein et al. (1978); Grossman and Hart (1984)).

²⁴See also Kojima (2007); Niederle (2007) for some responses to this model. Crawford (2008) proposes a modification that increases the wage flexibility of the NRMP.

²⁵This violates assumption 2.13, which was made to ensure that a stable matching exists. Nevertheless, a stable matching exists in this example.

If $\epsilon \approx 0$, we have $q_i \approx 1/3$.

Calculating surplus, we find that uniform wages are more efficient for $\epsilon \leq .41$, while flexible wages are more efficient for $\epsilon > .41$. The intuition is that, when both firms have similar productivities, the cost of allocative inefficiency is small. Therefore the quantity distortion dominates the allocative inefficiency. However, when one firm is much more productive than the other, the allocative inefficiency of uniform wages dominates, and personalized wages yield higher efficiency.

3.1.1 College Admissions and the Overlap

A well-known anti-trust case in matching markets is *U.S. v. Brown University, et al.*, where the U.S. Department of Justice sued the eight the Ivy League schools and MIT.²⁶ Since the 1950s, the schools in the Ivy League had agreed that it was in their best interest not to bid for top students that had been admitted to multiple schools. The main reasoning was that, were they to spend money bidding for star students, less funds would be available to provide support for low income students. The schools then started meeting to share information on students which were admitted to multiple schools. In the “Overlap meetings”, they agreed to what was the necessary financial aid to the commonly admitted students, and not to provide more aid than what was justified on a need basis. The number of schools participating on the Overlap meetings grew over the years, and in the 1970s it included MIT and 14 other schools.

In 1991, the Antitrust Division of the DOJ sued the Ivies and MIT for violating Section 1 of the Sherman act, by engaging in a conspiracy to fix prices. The schools did not deny they were engaging in cooperative behavior. Instead, their defense was that Antitrust law should not apply to them, for they are not for-profit institutions. Moreover, they argued that the purpose of the Overlap was not to raise revenues, but to guarantee that financial aid resources would be allocated to students in need, and not to wealthy meritorious students. Eventually, all the schools except for MIT agreed to cease the practice. MIT was then found guilty of price fixing in 1992. However, Congress passed a new law in 1992 that allowed the schools to engage in some cooperative practices. The 1992 MIT conviction was then overturned in 1993, and a subsequent trial was ordered. Finally, the government reached a 1993 settlement with MIT that allowed it to engage in most of the challenged conduct.

Carlton et al. (1995) provide some empirical evidence that the practice of the overlap had no effect on the average prices paid by the students. Moreover, they argue that, since schools are not for-profit enterprises, and instead have a complex objective function, the Antitrust laws should not apply. In addition, that the “conduct prevented the flow of school resources to high-income students at the expense of needy students”. On the other side of the debate, Grossman (1995) argues that since money is fungible, a reduction in merit aid does not necessarily increase need-based aid, and may reduce efficiency.

A normative analysis of this issue depends on assumptions on the objectives

²⁶See Carlton et al. (1995).

of the Ivy schools, and is beyond the scope of this paper. However, from a purely positive perspective, it seems reasonable to assume that not-for-profit Universities prefer some students from others, and also that they attach some positive value to financial resources (which may be used to maximize other goals they may have, however complex those may be). Therefore, the present framework may be used as a descriptive model of this market, with firms representing the schools, and workers representing the students. A coarse approximation is to consider the overlap as offering uniform wages (Section 2.2), and the non-overlap market having flexible wages and schools bidding for the best workers (Section 2.3). The model then predicts that the overlap changes the tradeoff faced by the Universities when determining optimal class size q_i . If we assume that preferences are homogeneous, so that schools mostly agree on which the best students are, then marginal revenue curves are lower under personalized rather than uniform wages. Therefore, firms would have more incentives to reduce class size, as depicted in Figure 3. The schools could argue, that ending the overlap would cause them to bid for star students, but without significantly changing the matching between students and universities. However although the bidding would imply a small gain in matching efficiency, it could give universities large incentives to reduce capacity, and make the final allocation less efficient. Even though Carlton et al. (1995) give evidence that the overlap did not change the average tuition payments, to my knowledge the effects of the overlap on class size have not been investigated.

3.2 Unravelling

A common market failure in entry level job markets is unravelling - the tendency of firms to hire workers earlier each year, trying to move ahead of each other. Unravelling of hiring dates has been observed in field (Roth and Xing (1994); Niederle and Roth (2003); McKinney et al. (2005)) and experimental (McKinney et al. (2005); Niederle et al. (2008)) settings, and has been the key market failure associated with the collapse of several markets (Roth (1991); Roth and Xing (1994)).

The standard rationale for unravelling is that there is some imperfect information about preferences. By accepting an early offer, a worker is in effect buying insurance against the possibility that she will receive a low quality match (Roth and Xing (1994); Li and Rosen (1998); Suen (2000); Niederle et al. (2008); Halaburda (2009); Fainmesser (n.d.); Ostrovsky and Schwarz (2010)). In particular, these theories explain why a firm may choose to hire at a date where less information is available. However it does not explain why firms would want to act first, but at a point in time where the same information about preferences is known.²⁷

²⁷Recently, these ideas have been applied to early action and early admission programs at American universities (Avery et al. (2001, 2004); Avery and Levin (2009); Chade et al. (n.d.); Lee (2009)). Avery and Levin (2009) propose a model where early action programs, which give students the option to enroll early, are used to let students signal their preferences over schools. Moreover, early decision programs, where students commit to accept or reject an

The imperfect competition model suggests a complementary rationale for unravelling, which does not rely on the arrival of new information. Simply put, strategic complementarities give firms incentives to act first. That is, early hiring can be viewed as a commitment device, yielding a first mover advantage. If, for example, the quantities played by firms are strategic substitutes, a firm has incentives to hire a large number of workers early on, effectively committing to be more aggressive, and therefore inducing the other firm to be less aggressive.

To see how firms have incentives hire early in a simple setting, we return to the example of Section 2.1.4, where two firms with $Q_i = [0, 1]$ compete for a mass 2 of workers with uniformly distributed productivities and preferences. We assume, however, that each firm may pay a cost c_i to hire a day early. The c_i are private information of each firm, and drawn independently from a distribution $F(\cdot)$, with a continuous density $f(\cdot)$ with support $[0, \infty)$. Timing is as follows. First, firms observe the c_i . They then decide simultaneously whether to incur the cost c_i and hire early. Next, firms hiring early simultaneously supply capacity q_i . Firms hiring late then choose capacity. Finally, workers are matched according to the unique stable matching given reported capacities. We make the assumption that the outcome is the unique stable matching with respect to (q_1, q_2) for starkness, to completely shut down the insurance channel considered in the literature. In effect this assumption implies that early hiring simply changes the order in which the q_i are chosen, but not the final matching given q_i . The motivation is that, since there is perfect information, rational workers will anticipate the quantity choices of firms that act late. Therefore, no worker would accept an early offer that is worse than what she could receive in the second day of hiring.

Conditional on the decision to enter early, firm behavior is simple. Denote by $\Pi_{tt'}$ the payoff net of c_i of a firm hiring at date t , while its opponent hires at date t' , for t, t' equal to E, L standing for early and late. If both firms hire at the same date, equilibrium is as in Section 2.1.4, yielding payoffs of $\Pi_{EE} = \Pi_{LL} = 14 - 6\sqrt{5}$. However, if only one of the firms hires early, it is a Stackelberg leader, and gets to set its quantity first. The optimum is for the leader to set quantity equal to the cap of 1. The intuition is that the leader acts more aggressively, inducing the other firm to reduce its capacity. Payoffs in this case are given by $\Pi_{EL} = 3/5$ and $\Pi_{LE} = 9/16$.

Since the expected payoff of entering early is strictly decreasing in c , in equilibrium there must be a value of $c^* > 0$ such that a firm enters iff $c_i \leq c^*$. In a symmetric equilibrium, a firm with cost of exactly c^* must be indifferent between entering or not. Therefore, c^* must solve

$$F(c^*) \cdot [\Pi_{EE} - \Pi_{LE}] + [1 - F(c^*)] \cdot [\Pi_{EL} - \Pi_{LL}] = c^*.$$

For $c^* = 0$, the left-hand side of the expression is positive, and the right-

early offer, are used by low ranked schools in a manner analogous to early offers in unravelling models, offering insurance to students who are unsure of their qualifications. While their model relies on (at least some level of) incomplete information about one's own skills, the present model suggests that early action policies could also be used due to imperfect competition, and universities seeking a first mover advantage.

Figure 5: Probability of unravelling given the mean λ of the cost of early hiring.

hand side is zero. As c^* approaches $+\infty$, the left-hand side is bounded, while the right-hand side goes to infinity. Therefore, at least one equilibrium exists. To illustrate the result, consider the case where $F(\cdot)$ is an exponential distribution with mean λ . Figure 5 plots the probability that unravelling occurs (at least one firm hires early), given λ . To gauge the size of the cost c , note that the value to either firm of hiring the entire worker pool would be 1. Therefore, if the mean costs of hiring one day in advance are at most 1% of the value of the entire worker pool, the unravelling probability is close to 100%. Moreover, even if the mean costs are approximately 10% of the value of the entire worker pool, the unravelling probability is still above 30%. This suggests that strategic complementarities may be an important channel generating unravelling in matching markets.

4 Conclusion

The large literature in two-sided matching markets leaves strategic interactions between participants to the sidelines. In contrast, the industrial organization literature puts strategic interactions, imperfect competition, and Nash equilibrium in the center stage, but typically does not consider settings with rich heterogeneous preferences over trading partners. This paper contributes to bridging this gap, by offering a simple, tractable equilibrium model of imperfect competition in many to one matching markets. This is a first step towards understanding firm behavior, and its implications for the design and regulation of matching markets.

The main contribution of the analysis is to extend standard price-theoretic insights of the Cournot model to matching markets. Market power induces a wedge between the marginal revenue of a firm, and the net productivity of a marginal worker it may hire. Interestingly, the determinants of the size of this wedge are very different when wages are uniform (as in the market junior law associates) or personalized (as in the market for senior lawyers). With uniform wages, the wedge exists due to heterogeneous preferences between firms, which mean that rejecting a worker may create a beneficial rejection chain. When wages are personalized, though, the wedge exists because of aligned preferences. Firms that reduce capacity increase the pool of available workers, which induces its competitors to bid less aggressively for workers that both firms covet.

These insights help to inform the debate over the desirability of uniform versus flexible wages. We have seen that taking strategic capacity setting into account qualifies the Bulow and Levin (2006) conclusion that flexible wages always generate more efficiency. In their model, if firms are allowed to choose capacity, flexible wages do produce higher matching efficiency given quantities, but they also give more incentives for firms to reduce capacities. Flexible wages are still more efficient if firms are sufficiently heterogeneous. However, if they are very similar, so that the loss from matching inefficiencies is small, uniform wages produce higher welfare, as they cause less capacity reduction. The model also yields a new rationale for unravelling, that firms may wish to hire early to get a first mover advantage, due to strategic complementarities.

An important omission of the model is the focus on quantity competition. While in industrial organization quantity competition models figure prominently, they are by no means the only models available. It would be interesting to explore matching markets where firms have a different set of strategic variables at their disposal. For example, how much firms try to differentiate themselves, target specific market segments, or allowing firms to misrepresent their preferences. Another shortcoming is that, at the level of generality that we consider here, many interesting phenomena from real matching markets cannot be explained. A hallmark of matching markets is the importance of institutional details. Therefore, it would be interesting in future research to specialize the model to study features of specific markets, and inform the market design in real-world cases. Hopefully, the present model will provide a useful basic framework to pursue this line of research.

References

- Abdulkadiroglu, A., P.A. Pathak, and A.E. Roth, “Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match,” *The American Economic Review*, 2009, 99 (5), 1954–1978.
- , Y.K. Che, and Y. Yasuda, “Expanding “Choice” in School Choice,” *Unpublished manuscript*, 2008.

- Adachi, H.**, “On a characterization of stable matchings,” *Economics Letters*, 2000, 68 (1), 43–49.
- , “A search model of two-sided matching under nontransferable utility,” *Journal of Economic Theory*, 2003, 113 (2), 182–198.
- Armstrong, M.**, “Competition in two-sided markets,” *The RAND Journal of Economics*, 2006, 37 (3), 668–691.
- Avery, C., A. Fairbanks, and R. Zeckhauser**, “What worms for the early bird? Early admissions at selective colleges,” Technical Report, working paper, Kennedy School of Government, Harvard University 2001.
- , – , and – , *The early admissions game: Joining the elite*, Harvard Univ Pr, 2004.
- and **J. Levin**, “Early admissions at selective colleges,” *NBER Working Paper*, 2009.
- Azevedo, E.M. and J. Leshno**, “The college admissions problem with a continuum of students,” 2010.
- Balinski, M. and T. Sönmez**, “A Tale of Two Mechanisms: Student Placement,” *Journal of Economic Theory*, 1999, 84 (1), 73–94.
- Becker, G.S.**, “A theory of marriage: Part I,” *Journal of Political economy*, 1973, 81 (4), 813.
- Biró, P.**, “Higher education admission in Hungary by a score-limit algorithm,” in “The 18th International Conference on Game Theory at Stony Brook” 2007.
- Bulow, J. and J. Levin**, “Matching and price competition,” *American Economic Review*, 2006, 96 (3), 652–668.
- Bulow, J.I., J.D. Geanakoplos, and P.D. Klemperer**, “Multimarket oligopoly: Strategic substitutes and complements,” *The Journal of Political Economy*, 1985, 93 (3).
- Carlton, D.W., G.E. Bamberger, and R.J. Epstein**, “Antitrust and higher education: was there a conspiracy to restrict financial aid?,” *The RAND Journal of Economics*, 1995, 26 (1), 131–147.
- Chade, H., G. Lewis, and L. Smith**, “A Supply and Demand Model of the College Admissions Problem.”
- Cournot, A.**, *Recherches sur les Principes Mathématiques de la Théorie de la Richesse*, Hachette, Paris, 1838.
- Crawford, V.P.**, “The flexible-salary match: a proposal to increase the salary flexibility of the national resident matching program,” *Journal of Economic Behavior & Organization*, 2008, 66 (2), 149–160.

- Echenique, F.**, “Contracts vs. Salaries in Matching,” 2010.
- **and J. Oviedo**, “Core many-to-one matchings by fixed-point methods,” *Journal of Economic Theory*, 2004, *115* (2), 358–376.
 - **and –**, “A theory of stability in many-to-many matching markets,” *Theoretical Economics*, 2006, *1* (2), 233–273.
- Edelman, B., M. Ostrovsky, and M. Schwarz**, “Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords,” *The American Economic Review*, 2007, *97* (1), 242–259.
- Edmans, A., X. Gabaix, and A. Landier**, “A multiplicative model of optimal CEO incentives in market equilibrium,” *Review of Financial Studies*, 2009, *22* (12), 4881.
- Ehlers, L.**, “Manipulation via capacities revisited,” *Games and Economic Behavior*, 2010.
- Fainmesser, I.P.**, “Social Networks and Unraveling in Labor Markets.”
- Fudenberg, D. and J. Tirole**, “The fat-cat effect, the puppy-dog ploy, and the lean and hungry look,” *The American Economic Review*, 1984, *74* (2), 361–366.
- **and –**, *Game theory*, MIT Press, 1991.
- Gabaix, X. and A. Landier**, “Why Has CEO Pay Increased So Much?*,” *Quarterly Journal of Economics*, 2008, *123* (1), 49–100.
- Gale, D. and L. Shapley**, “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, 1962, *69*, 9015.
- **and M. Sotomayor**, “Ms. Machiavelli and the stable matching problem,” *American Mathematical Monthly*, 1985, pp. 261–268.
 - **and –**, “Some remarks on the stable marriage problem,” *Discrete Applied Mathematics*, 1985, *11*, 223–232.
- Ginsburg, T. and J.A. Wolf**, “Market for Elite Law Firm Associates, The,” *Florida State University Law Review*, 2003, *31*, 909.
- Glicksberg, I.L.**, “A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points,” *Proceedings of the American Mathematical Society*, 1952, *3* (1), 170–174.
- Grossman, H.**, “Scholarships: Need or Merit?,” *Cato Journal*, 1995, *14*, 519–526.
- Grossman, S.J.**, “Nash equilibrium and the industrial organization of markets with large fixed costs,” *Econometrica: Journal of the Econometric Society*, 1981, pp. 1149–1172.

- **and O. Hart**, “The costs and benefits of ownership A theory of vertical integration,” 1984.
- Halaburda, H.**, “Unravelling in two-sided matching markets and similarity of preferences,” *Games and Economic Behavior*, 2009.
- Hart, O.D.**, “Imperfect competition in general equilibrium: an overview of recent work,” *Frontiers of economics*, 1985, pp. 100–149.
- Hatfield, J.W. and F. Kojima**, “Group incentive compatibility for matching with contracts,” *Games and Economic Behavior*, 2009, 67 (2), 745–749.
- **and –**, “Substitutes and stability for matching with contracts,” *Journal of Economic Theory*, 2010.
- **and P.R. Milgrom**, “Matching with contracts,” *The American Economic Review*, 2005, 95 (4), 913–935.
- **, S.D. Kominers, A. Nichifor, M. Ostrovsky, and A. Westkamp**, “Stability and Competitive Equilibrium in Trading Networks,” 2010.
- Hitsch, G.J., A. Hortacsu, and D. Ariely**, “Matching and sorting in online dating,” *The American Economic Review*, 2010, 100 (1), 130–163.
- Immorlica, N. and M. Mahdian**, “Marriage, honesty, and stability,” in “Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms” Society for Industrial and Applied Mathematics 2005, p. 62.
- Jr, A.S. Kelso and V.P. Crawford**, “Job matching, coalition formation, and gross substitutes,” *Econometrica: Journal of the Econometric Society*, 1982, 50 (6), 1483–1504.
- Kesten, O.**, “On Two Kinds of Manipulation for School Choice Problems,” *Tepper School of Business*, 2008, p. 224.
- Klein, B., R.G. Crawford, and A.A. Alchian**, “Vertical integration, appropriable rents, and the competitive contracting process,” *The Journal of Law and Economics*, 1978, 21.
- Klemperer, P.D. and M.A. Meyer**, “Supply function equilibria in oligopoly under uncertainty,” *Econometrica*, 1989, 57 (6), 1243–1277.
- Kojima, F.**, “Mixed strategies in games of capacity manipulation in hospital–intern markets,” *Social choice and Welfare*, 2006, 27 (1), 25–28.
- **,** “Matching and price competition: comment,” *The American economic review*, 2007, 97 (3), 1027–1031.
- **and P.A. Pathak**, “Incentives and stability in large two-sided matching markets,” *American Economic Review*, 2009, 99 (3), 608–627.

- Konishi, H. and M.U. Ünver**, “Games of Capacity Manipulation in Hospital-Intern Markets,” *Social choice and Welfare*, 2006, 27 (1), 3–24.
- Kremer, M.**, “The O-ring theory of economic development,” *The Quarterly Journal of Economics*, 1993, 108 (3), 551–575.
- Kreps, D.M. and J.A. Scheinkman**, “Quantity precommitment and Bertrand competition yield Cournot outcomes,” *The Bell Journal of Economics*, 1983, pp. 326–337.
- Lee, S.H.**, “JUMPING THE CURSE: EARLY CONTRACTING WITH PRIVATE INFORMATION IN UNIVERSITY ADMISSIONS*,” *International Economic Review*, 2009, 50 (1), 1–38.
- Li, H. and S. Rosen**, “Unraveling in matching markets,” *American Economic Review*, 1998, 88 (3), 371–387.
- Luenberger, D.G.**, *Optimization by vector space methods*, Wiley-IEEE, 1969.
- Mailath, G.J., A. Postlewaite, and L. Samuelson**, “Pricing in Matching Markets,” 2010.
- McKinney, C.N., M. Niederle, and A.E. Roth**, “The collapse of a medical labor clearinghouse (and why such failures are rare),” *American Economic Review*, 2005, 95 (3), 878–889.
- Milnor, J.W.**, *Topology from the differentiable viewpoint*, Princeton Univ Pr, 1997.
- Mumcu, A. and I. Saglam**, “Games of capacity allocation in many-to-one matching with an aftermarket,” *Social Choice and Welfare*, 2009, 33 (3), 383–403.
- Niederle, M.**, “Competitive wages in a match with ordered contracts,” *The American Economic Review*, 2007, 97 (5), 1957–1969.
- , **A.E. Roth, and U. "Unver**, “Unraveling Results from Comparable Demand and Supply,” 2008.
- **and –**, “Unraveling reduces mobility in a labor market: Gastroenterology with and without a centralized match,” *Journal of Political Economy*, 2003, 111 (6), 1342–1352.
- **and L. Yariv**, “Decentralized Matching with Aligned Preferences,” *NBER Working Paper*, 2009.
- Ostrovsky, M.**, “Stability in supply chain networks,” *The American Economic Review*, 2008, 98 (3), 897–923.
- **and M. Schwarz**, “Information Disclosure and Unraveling in Matching Markets,” *American Economic Journal: Microeconomics*, 2010, 2 (2), 34–63.

- Roberts, J. and H.F. Sonnenschein**, “On the existence of Cournot equilibrium without concave profit functions,” *Journal of Economic Theory*, 1976, 13 (1), 112–117.
- Rochet, J.C. and J. Tirole**, “Platform competition in two-sided markets,” *Journal of the European Economic Association*, 2003, 1 (4), 990–1029.
- and – , “Two-sided markets: A progress report,” *The RAND Journal of Economics*, 2006, 37 (3), 645–667.
- Roth, A.E.**, “The college admissions problem is not equivalent to the marriage problem,” *Journal of Economic Theory*, 1985, 36 (2), 277–288.
- , “A natural experiment in the organization of entry-level labor markets: Regional markets for new physicians and surgeons in the United Kingdom,” *The American economic review*, 1991, 81 (3), 415–440.
- and **E. Peranson**, “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design,” *The American Economic Review*, 1999, 89 (4), 748–780.
- and **J.H.V. Vate**, “Random paths to stability in two-sided matching,” *Econometrica*, 1990, 58 (6), 1475–1480.
- and **M. Sotomayor**, “The college admissions problem revisited,” *Econometrica*, 1989, 57 (3), 559–570.
- and **M.A.O. Sotomayor**, *Two-sided matching: A study in game-theoretic modeling and analysis*, Cambridge Univ Press, 1992.
- and **X. Xing**, “Jumping the gun: imperfections and institutions related to the timing of market transactions,” *The American Economic Review*, 1994, 84 (4), 992–1044.
- Sönmez, T.**, “Manipulation via Capacities in Two-Sided Matching Markets,” *Journal of Economic Theory*, 1997, 77 (1), 197–204.
- Sorensen, M.**, “How smart is smart money? A two-sided matching model of venture capital,” *Journal of Finance*, 2007, p. 2725.
- Suen, W.**, “A competitive theory of equilibrium and disequilibrium unravelling in two-sided matching,” *The Rand journal of economics*, 2000, pp. 101–120.
- Tervio, M.**, “The difference that CEOs make: An assignment model approach,” *The American Economic Review*, 2008, 98 (3), 642–668.
- Weyl, E.G.**, “The Price Theory of Two-Sided Markets,” 2007.

A Proofs of Propositions and Lemmas

A.1 Uniform Wages

A.1.1 Stable Matchings

Before proving results relating to the game per se, we have to develop some preliminaries on stable matchings. We consider a set of agents Θ and firms I , as in section 2.1.1. Cutoffs will play an important role in the analysis, therefore the definitions of section 2.1.3 will also be used. The first result guarantees that with strict preferences there is a bijection between stable matchings and market clearing cutoffs.

Lemma A.1. (*Cutoff Lemma - Azevedo and Leshno (2010)*) *Assume η satisfies the strict preferences assumption 2.1. Then stable matchings exist. Moreover, if μ is a stable matching, then $\mathcal{P}\mu$ is a market clearing cutoff. If p is a market clearing cutoff, then $\mu = \mathcal{M}p$ is a stable matching. In addition, $\mathcal{P}$ and $\mathcal{M}$ are inverses of each other in the sets of market clearing cutoffs and stable matchings.*

Proof. This is a restatement of the Cutoff Lemma of Azevedo and Leshno (2010), which we refer to for a proof. $\square$

Note that this result does not depend on the assumption of a unique stable matching existing. We now use it to prove Proposition 2.2, which gives sufficient conditions for the set of stable matchings to have a unique element. Moreover, the particular version of the cutoff Lemma used in the text, Lemma 2.6, follows directly from this result, and our assumption in the text that there is a unique market clearing cutoff associated with $[\eta, q]$.

Proof. (Proposition 2.2) By the previous Lemma, we only have to show that there is a unique market clearing cutoff. By the Lattice Theorem (Azevedo and Leshno (2010)), there exists a minimum and a maximum market clearing cutoffs $p' \leq p$. Any worker that is employed under p must be employed over p' . Moreover, all workers in the set

$$A = T \cap (\{e^\theta < p\} \setminus \{e^\theta < p'\})$$

must be employed under p' but not under p . Therefore, the difference in the mass of unemployed workers under p and p' must be at least as large as $\eta(A)$. However, by the Rural Hospitals Theorem (Azevedo and Leshno (2010)), the mass of unemployed workers is constant in all stable matchings. Consequently, $\eta(A) = 0$. By the assumption made in the statement of the Proposition, this implies $p = p'$, and therefore there is a unique stable matching. $\square$

In the text, we denote by $P(q)$ the unique market clearing cutoff associated with a vector of capacities q , and implicitly a measure η . In the appendix, it will be useful to sometimes make the dependence on η explicit. We will abuse notation and write $P([\eta, q])$. Moreover, whenever there is more than one market clearing cutoff associated with η , we will use the same notation

for the correspondence of market clearing cutoffs. Some of our results depend on topological properties of this correspondence. Henceforth, we will use the standard Euclidean topology for the set of cutoffs. For the set of measures η , we will use the weak-* topology, sometimes also referred to as weak convergence of measures.

Proposition A.2. *(Continuity) Assume η satisfies the strict preferences assumption 2.1, and that $P([\eta, q])$ has a unique element. Then P is continuous in a neighborhood of $[\eta, q]$.*

Proof. This is part of the statement of Azevedo and Leshno (2010) Theorem 2. $\square$

Finally, we provide a proof of the monotonicity result in Lemma 2.7.

Proof. (Lemma 2.7) Let $p = P(q)$, $p' = P(q')$, and $\hat{p}$ be the sup of the vectors p and p' . For every i it is always the case that either $\hat{p}_i = p_i$ or p'_i . If $\hat{p}_i = p_i$, then because for all $j \neq i$ we have $\hat{p}_j \geq p_j$ we must have $D_i(\hat{p}) \geq D_i(p)$. If $\hat{p}_i = p'_i$ then by the same logic $D_i(\hat{p}) \geq D_i(p') \geq D_i(p)$. Therefore $\hat{p} \geq p$ and $D(\hat{p}) \geq D(p)$. Since $\sum_i D_i(\hat{p}) \leq \sum_i D_i(p)$, we must have $D(\hat{p}) = D(p)$. If $p_i > 0$, then $D_i(\hat{p}) = D_i(p) = q_i$. If $p_i = 0$ and $\hat{p}_i = 0$ then $D_i(\hat{p}_i) \leq q_i$. Finally, if $p_i = 0$ and $\hat{p}_i > 0$ then $D_i(\hat{p}_i) \geq D_i(p'_i) = q'_i \geq q_i$, so that $D_i(\hat{p}_i) = q_i$. Therefore, $\hat{p} = P(q)$. By the uniqueness assumption, $\hat{p} = p$, and therefore $p' \leq p$. $\square$

A.1.2 The Oligopoly Game with Exogenous Wages

We may now prove Proposition 2.4, which guarantees the existence of a mixed strategy equilibrium.

Proof. (Proposition 2.4) Since payoffs $\Pi(q)$ can be written as a continuous function of $P(q)$, and by Proposition A.2 $P(q)$ depends continuously on q , then payoffs must depend continuously on q . Moreover, the strategy space of each agent is a compact interval of the real line. Therefore, it follows from Glicksberg's 1952 theorem (FT pp 35 theorem 1.3) that a mixed-strategy equilibrium exists. $\square$

We may also prove Proposition 2.11, which guarantees that the profit functions are continuously differentiable almost everywhere, and provides an expression for marginal revenues.

Proof. (Proposition 2.11)

First, note that since η admits a continuous density, the demand function $D(p)$ may be written as

$$D(p) = \int_{(\mathcal{M}_p)(i)} f(\theta) d\theta,$$

where

$$\mathcal{M}p(i) = \{\theta \in \Theta | e_i^\theta \geq p_i, e_j^\theta < p_j \text{ for all } j \text{ such that } j \succ^\theta i, i \succ^\theta \theta\}. \quad (3)$$

Therefore, by Leibniz's rule for differentiation under the integral sign, $D(p)$ is continuously differentiable.

Let Q^* be the set of interior points of Q . That is, the interior of the set of points q such that all $\eta(\mu_q(i)) = q_i$. Note that, in Q^* , market clearing cutoffs $P(q)$ are the single root of the equation $D(P(q)) = q$. By Sard's theorem,²⁸ for almost every point $q \in Q^*$, $D(\cdot)$ is continuously differentiable at $P(q)$, and its derivative is nonsingular. Therefore, by the inverse function theorem, $P(q)$ is continuously differentiable in a neighborhood of q .

Given a quantity vector q , and cutoffs $P(q)$, the revenue of firm i may be written as

$$R_i(q) = \int_{(\mathcal{M}P(q))(i)} e_i^\theta \cdot f(\theta) d\theta. \quad (4)$$

If P is continuously differentiable at q , the formula for the marginal revenue in the Proposition follows directly from an application of Leibniz's rule. $\square$

We now prove Proposition 2.12, which guarantees that when either side of the market has homogeneous cardinal preferences, then firms have no incentives to reduce capacity.

Proof. (Proposition 2.12) Consider first the case where all workers have the same preferences. Without loss of generality, assume that all of them have preference ordering $1, 2, \dots, I$.²⁹ Note that, by the market clearing equations, cutoffs $P_i(q)$ do not depend on q_j for $j > i$. In addition, by the formula for the set of matched students $\mathcal{M}p(i)$ in equation 3 in the proof of Proposition 2.11, we have that revenue $R_i(q)$ may be written as a function $\tilde{R}_i(P_1(q), P_2(q), \dots, P_{i-1}(q), P_i(q))$, that does not depend on $P_k(q)$ for $k > i$. If we consider an interior point q and q' with $q'_i = q_i + \epsilon$, $\epsilon > 0$ and $q'_j = q_j$ for all other coordinates, we have

$$\begin{aligned} R_i(q') &= \tilde{R}_i(P_1(q'), P_2(q'), \dots, P_{i-1}(q'), P_i(q)) \\ &= \tilde{R}_i(P_1(q), P_2(q), \dots, P_{i-1}(q), P_i(q')). \end{aligned}$$

Therefore, using again equation 3, $R_i(q') - R_i(q)$ may be written as

$$\int_A e_i^\theta d\eta(\theta),$$

where

$$A = \{e_j^\theta < P_j(q) \text{ for all } j < i, P_i(q') \leq e_i^\theta < P_i(q)\}.$$

²⁸See Milnor (1997).

²⁹We are assuming that all firms are considered acceptable by the workers, as firms which no worker finds acceptable play no role in the proof.

Because $P(\cdot)$ is continuous, the productivity e_i^θ of all workers in the set A is approximately $P_i(q)$. Moreover, as the measure of A is ϵ , we have that

$$R_i(q') - R_i(q) = P_i(q) \cdot \epsilon + o(\epsilon).$$

Therefore, $R_i(\cdot)$ is differentiable at q with derivative $P_i(q)$, completing the proof.

Now consider the second case, where firms have the same ordinal preferences in the support of η . First, note that there must exist increasing continuous functions $f_2, f_3, \dots, f_I$ such that the support of η equals the set

$$\{(e_1^\theta, f_2(e_1^\theta), \dots, f_I(e_1^\theta)) | e_1^\theta \in [0, 1]\}.$$

To see this, note first that the support must include points with all possible values of $e_1^\theta \in [0, 1]$. If it did not, the unique stable matching assumption would be violated, as there would exist $\bar{p}$ with $\{p \in [0, 1]^I : \|p_1 - \bar{p}_1\| < \delta\}$ outside the support of η . Therefore, $D(p)$ would be constant in p_1 in a neighborhood of $\bar{p}$, and so there would be more than one market clearing cutoff associated with η and $q = D(\bar{p})$. In addition, for a given value of e_1^θ , by the homogeneous ordinal preferences assumption, the support may only contain one point. We denote this point as $(e_1^\theta, f_2(e_1^\theta), \dots, f_I(e_1^\theta))$, which defined the functions f_i . Again, by the homogeneous ordinal preferences assumption, the f_i are strictly increasing. Moreover, they must be continuous, as otherwise the support of θ would not include points with some value of $e_i^\theta \in [0, 1]$, which would violate the unique stable matching assumption, by the argument used before.

With this observation in hand, the rest of the proof is simple, and similar to the first part. Let f_1^θ be the identity map. Note that, for all i , $e_i^\theta \leq P_i$ iff $e_1^\theta \leq f_i^{-1}(P_i)$. Therefore, given an interior point q , we may denote the firms which are more selective than firm 1 as

$$I_+(q) = \{i : f_i^{-1}(P_i(q)) > P_1(q)\}$$

Note that, by the market clearing equations, a small change in q_i does not affect $P_i(q)$ for $i \in I_+$. Now take an interior point q and q' with $q'_i = q_i + \epsilon$, $\epsilon > 0$ and $q'_j = q_j$ for all other coordinates. By the definition of the demand function, all points in the symmetric difference $\mu_q(1) \triangle \mu_{q'}(1)$ must satisfy

$$P_i(q') \leq e_i^\theta < P_i(q)$$

for some $i \in I$. For small ϵ , we have $P_i(q') = P_i(q)$ for $i \in I_+(q)$, and therefore the equation has to hold for some $i \in I \setminus I_+(q)$. Therefore we must have $e_1^\theta \leq P_1(q)$. Since, by the definition of the demand function, every point in the symmetric difference must also satisfy $P_1(q') \leq e_1^\theta$, we must have

$$P_1(q') \leq e_1^\theta < P_1(q).$$

Therefore, using the same argument as before on the continuity of P , we must have that the revenue of firm 1 is differentiable at q , and $MR_1(q) = P_1(q)$. $\square$

A.1.3 Endogenous Uniform Wages

We begin with the Proposition that guarantees existence of a mixed strategy equilibrium.

Proof. (Proposition 2.15) Note that the under our assumptions the measure $\eta([\eta_u, w])$ varies continuously on w . Therefore the continuity result in Proposition A.2 implies that P varies continuously in (q, w) . The same argument used in the proof of Proposition 2.4 then guarantees existence of a mixed-strategy Nash equilibrium. $\square$

We now turn to the characterization of marginal revenue.

Proof. Since $\bar{P}$ is continuously differentiable, so is revenue, and $P_i(q) = q_i \cdot \bar{P}(q)$. Having established the differentiability of P , the formula for the derivative with respect to q follows from the argument in the proof of Proposition 2.11. The formula for the derivative with respect to w is a straightforward calculation. $\square$

A.1.4 Matching With Contracts

The matching framework used in Section 2.3 is a particular case of the model in the Appendix D of Azevedo and Leshno (2010). As observed in Azevedo and Leshno (2010)'s Appendix D.3, the existence of stable matchings, Rural Hospitals Theorem, and Lattice Theorem hold in this setting. Therefore, the proof of Proposition 2.18, which guarantees that there exists a unique market clearing cutoff, is identical to the proof of Proposition 2.2.

Proof. (Proposition 2.20) The Proposition simply collects some results from Azevedo and Leshno (2010). The existence of stable matchings and market clearing cutoffs is established in Azevedo and Leshno (2010) Appendix D.3. The fact that $\mathcal{P}$ and $\mathcal{M}$ take stable matchings into market clearing cutoffs and vice versa follows from the cutoff lemma for matching with contracts in Appendix D.2. As for the fact that agents are always matched one of the firms i with highest $s_i^\theta - p_i$, this is proved in Azevedo and Leshno (2010) Appendix D.4 which considers the transferable utility case. $\square$

It only remains to prove Proposition 2.22, which characterizes marginal revenue.

Proof. (Proposition 2.22) The proof that $P(q)$ is differentiable for almost every interior point q is exactly the same as in the case with exogenous wages given in the proof of Proposition 2.11. We therefore take an interior point q where $P(\cdot)$ is differentiable, and derive the formula for the marginal revenue. The formula for marginal revenue then follows directly from the formula for the revenue $R_i(q)$ of firm i in equation 2, and a direct application of Leibniz's formula for differentiation under the integral sign. $\square$

A.2 Applications

A.2.1 Comparison between uniform and personalized wages

Proof. (Theorem 3.1) The proof that, with endogenous capacities, uniform wages may dominate follows from the example given in the text. Therefore it only remains to prove that, for a fixed capacity vector q , personalized wages generate at least as much welfare as matching with uniform wages. Consider a set of workers Θ_X satisfying the requirements of both the model with uniform and the model with flexible wages. Let a generalized allocation be a measurable map

$$x : \Theta_X \rightarrow [0, 1]^{I+1}$$

designating a distribution $x(\theta)$ of each worker type over firms, with $I + 1$ representing being unemployed. Therefore both a stable matching with uniform wages and a stable matching with contracts induce an allocation, and one that only takes values in the extreme points of the simplex. Given a generalized allocation, we can calculate social welfare as

$$\int_{\Theta_X} s^\theta \cdot x(\theta) d\eta(\theta),$$

as social welfare does not depend on the specific contracts nor on wages. Consider now the problem of finding a generalized allocation that maximizes social welfare subject to feasibility constraints

$$\begin{aligned} \max \quad & \int_{\Theta_X} s^\theta \cdot x(\theta) d\eta(\theta) \\ \text{s.t.} \quad & \int_{\Theta_X} x_i(\theta) d\eta(\theta) \leq q_i \quad \text{for } i = 1, \dots, I. \end{aligned}$$

A standard compacity argument implies that such a maximum value is attained by at least one generalized allocation x^* . Moreover, because the problem has allocations x where all constraints are strictly at slack, $\int_{\Theta_X} x_i(\theta) d\eta(\theta) < q_i$ for all i , strong Lagrange duality holds. By Theorem 1 pp. 217 from Luenberger (1969), there exist numbers $\lambda_i \geq 0$ such that x^* maximizes

$$\int_{\Theta_X} s^\theta \cdot x(\theta) d\eta(\theta) + \sum_{i=1}^I \lambda_i \cdot [q_i - \int_{\Theta_X} x_i(\theta) d\eta(\theta)]$$

over all generalized allocations. Moreover, if $\lambda_i > 0$, then $\int_{\Theta_X} x_i(\theta) d\eta(\theta) = q_i$. Note that we can rewrite the expression above as

$$\sum_{i=1}^I \int_{\Theta_X} (s_i^\theta - \lambda_i) \cdot x_i(\theta) d\eta(\theta).$$

Therefore, any maximizer x^* satisfies that almost every type θ is matched with probability 1 to a firm that maximizes $s_i^\theta - \lambda_i$. So the measure of workers x^* allocates to each firm equals the demand for each firm when market clearing

cutoffs are equal to λ , in the matching with contracts model. Therefore λ is a vector of market clearing cutoffs in the matching with contracts model. Since the market clearing cutoffs are unique, we have $\lambda = P(q)$. Therefore, the generalized allocation x^* coincides almost everywhere with the generalized allocation induced by any stable matching with contracts. Therefore any stable matching with contracts maximizes social welfare. In particular, any stable matching with uniform wages generates weakly lower welfare, completing the proof. $\square$

Details of the calculation from the example comparing uniform and personalized wages follow below.

For the case of uniform wages, the reasoning follows Bulow and Levin (2006). Although in their model each firm hires a single worker, the argument is similar in the discrete case. Assume for now that both firms set $q_i = 1/2$. The firm that offers higher wages attracts workers with an average quality of $3/4$, while the firm with lower wages attracts workers with an average quality of $1/4$. Following their algorithm to characterize equilibrium, firms must offer random wages, with distributions $G_i(\cdot)$ with the same interval as support. To find this support, we consider each firm's first order condition with respect to wage in an interior point of the interval, as each firm must be indifferent between offering any wage in the support. We have

$$\begin{aligned} g_1(w) \cdot A \cdot 1/2 &= 1 \\ g_2(w) \cdot 1/2 &= 1. \end{aligned}$$

Therefore the density of firm 2's offer is $g_2(w) = 2$. As argued in Bulow and Levin (2006) pp. 659, the lowest wage offered must be zero. Therefore the support of the distributions is $[0, 1/2]$. Firm 2 offers a wage uniformly in this interval. Firm 1 has a density of only $2/A$. According to the Bulow and Levin (2006) algorithm, with probability $1/A$ it offers a wage uniformly at random in the interval $[0, 1/2]$, and offers 0 otherwise. Consequently, the probability that firm 2 offers a higher wage is $1/2A$.

We now show that it is in the interest of both firms to set $q_i = 1/2$. We will do the calculation for firm 2, as firm 1's case is analogous. If firm 1 plays $q_1 = 1$ and w_1 uniformly distributed in $[0, 1/2]$, firm 2's gain from offering q_2, w_2 with $w_2 \leq 1/2$ is

$$2w_2 \cdot \left(\frac{2-q_2}{2} - w_2\right) \cdot q_2 + (1-2w_2) \cdot \left(\frac{1-q_2}{2} - w_2\right) \cdot q_2.$$

We can see that the maximum of this expression in $q_2 \in [0, 1/2]$ is attained with $q_2 = 1/2$.

Consider now the case of personalized wages. Workers have surplus vectors $s^\theta = (s_1^\theta, s_2^\theta)$ uniformly distributed in the segment $[(0, 0), (1, A)]$ in $\mathbb{R}^2$. Therefore firm 2 always hires the q_2 best workers, and firm 1 hires the next q_1 best ones. The market clearing equations imply $p_1 = 1 - q_1 - q_2$, and $p_2 = \epsilon(1 - q_2) + p_1$.

We have $\eta(H_{21}) = q_2$, whereas $\eta(H_{12}) = 1 - q_2 - p_2/A = q_1/A$. Therefore the marginal revenue formula yields

$$\begin{aligned} MR_2 &= p_2 - q_2 \\ MR_1 &= p_1 - q_1/A. \end{aligned}$$

To solve for equilibrium we only have to set $MR_1 = MR_2 = 0$, and the formula in the text for q_1^* and q_2^* obtains.