March 2021

Big Names, Bigger Barriers:

Firm Reputation and its Role as a Barrier to Entry

Hannah Ellery

We investigate the potential for barriers to entry for new firms and an implicit tuition in worker wages stemming from incumbent reputation. To do so, we introduce two models of worker preferences within a framework of matching with externalities and transferable utility based on the work of Kelso and Crawford (1982) and Pycia and Yenmez (2017). Our first preference model allows workers to value the opinions of others by incorporating these opinions into their utility functions. We then build a new fixed point algorithm which increases the efficiency of finding equilibria in this context. Our second preference model introduces explicit preferences for reputation and prestige among workers. We then show that when workers make employment decisions based on firm prestige, we expect to see more concentrated markets and barriers to entry for new firms. Stylized facts from a unique dataset on hiring in large law firms present preliminary empirical evidence supporting the model, indicating that workers in a low-entry industry do have a preference for reputation and prestige in their employers.

1. Introduction

As markets have become increasingly concentrated in the United States (Grullon et al., 2018), antitrust enforcement has become more important than ever. The verdict in such antitrust cases often hinges on findings that the incumbent has erected barriers to entry in an industry.[1] These barriers are usually assumed to stem from restricted access to physical or financial resources, or from explicit regulatory barriers.

However, as economists become more aware of the role that intangible capital plays in other areas of economics,[2] we should broaden the definition of barriers to entry to include intangible barriers. This paper suggests that a lack of established firm reputation, a type of intangible capital, can also prevent entry. In short, reputation acts on the market by making the process of worker acquisition

[1] See as two recent examples *U.S. and Plaintiff States v. Google LLC* and *U.S. v. Microsoft* 2001.

[2] See Corrado and Hulten (2010), Corrado et al. (2009), and Chen (2018) for the roles that intangible forms of capital such as organizational skill can play in growth.

less successful and more costly for the entrant than for the incumbent. When choosing a job, it is common for prospective employees to consider culture, hours, pay, and benefits. However, reputation, defined in this paper as *public perception of prestige, success, and desirability*, is frequently an important metric for decision as well,[3] and one that may have negative consequences for the nascent firm. Providing theoretical and empirical foundations for reputation as a barrier to entry indicates that any study of market structure is incomplete without considerations of intangible capital, such as prestige.

To begin to demonstrate the existence of such a barrier, we study a framework of worker-firm matching. We introduce two models of preferences into this framework. In the first model, we allow workers' utility to be determined endogenously by the preferences and matches of their fellow employees. In the second, we allow for workers' preferences to be shaped both by their fellow workers' preferences, as in the first model, and by the success of the firm in the market, thereby introducing a preference for reputation. These models have at their roots the classic many-to-one matching model of Kelso and Crawford (1982), but both incorporate significant matching externalities: namely, workers care about others' opinions of the firm that they choose to work for, as well as about its performance, a statistic which is affected by the quality of workers at that firm.

The first contribution of this work, our model of labor market matching, adapts and extends existing models of matching with externalities and introduces a fixed point algorithm that more easily solves transferable utility matching problems. Building a fixed point algorithm based on previous work by Kelso and Crawford (1982) and Pycia and Yenmez (2017) that we term the *Stealing Process*, we demonstrate that it is indeed possible to find stable equilibria in this context.

We then introduce our first preference model within this framework, allowing workers' preferences to affect those of their fellow employees while maintaining substitutability. This interaction of worker preferences has implications for the salaries of the workers in the model, who are willing to give up some pay to stay with a firm rated more highly in the eyes of others. However, as the second contribution of this paper demonstrates, any matching model built on substitutable preferences will be unable to show the existence of a reputational barrier to entry in the context of a Kelso-Crawford style labor matching model. Though externalities that do not violate substitutability are sufficient to create an implicit tuition in salaries when workers go to firms which others rank highly, we find that when workers base their choices on utility functions and always prefer employment to

[3]See recent examples of popular literature on the topic, including Burgess (2016) and the annual Opinium survey on firm reputation, in particular the survey from 2019.

unemployment, substitutability implies that workers' choices of employment cannot depend in any way on the broader matching, and therefore cannot depend on a firm's success.

We thus introduce our second preference model, enriching our framework by building complementarities into workers' preferences. We construct a second novel fixed-point algorithm, the *Preference Updating* process, under which we iteratively update workers' preferences as a function of firms' success in the labor market and find equilibria which correspond to these preferences, simultaneously determining employment allocations and workers' views of firm prestige. We then explore what equilibria look like—when they exist—when workers have such dynamic preferences over prestige. In particular, we demonstrate how preferences that depend upon perception of firm prestige can lead to barriers to entry.

Finally, we build a unique dataset, giving preliminary evidence for workers having reputation-based preferences and exploring the relationship between reputation and entry. We collect evidence from elite law firms and demonstrate a strong and significant positive association between the reputation of the firm and the quality of workers that it is able to hire, as measured by the average ranking of new hires' alma mater law schools. These results indicate that law firms which are rated more highly tend to have access to higher quality workers, though more work is needed before any causal claims can be made.

Together, the results from the model and empirical work suggest that reputational barriers are worthy of addition to the list of phenomena and circumstances we currently consider to limit firm entry.

The rest of the paper is laid out as follows. Section 2 begins by discussing related literature and the challenges of matching with externalities. Section 3 presents the formal model and definitions. Section 4 discusses the Stealing Process for finding stable equilibria. Section 5 describes our first model of preferences. Section 6 presents the Preference Updating process, simultaneously introducing our second fixed point algorithm and our second preference model. Section 7 analyzes economic results indicated by the two preferences models. In Section 8, we present the results of our empirical work on large law firms and their employees. Finally, Section 9 concludes.

2. Review of Related Literature

The very phrase "barrier to entry" is one of the more controversial in economics. The list of possible barriers to entry has grown steadily since the term was first coined by Joe S. Bain in the 1950's (Bain, 1956), though the specific items included under the label of "barrier" have shifted as different models of entry have gained popularity (McAfee et al., 2004). The first proposed barriers, including economies of scale and the power of incumbent brands in directing and capturing consumer tastes, have been joined by a myriad of others in the intervening half-century. Today, economists have considered evidence for a much wider variety of barriers to entry, ranging from breakup fees (Bedre-Defolie and Biglaiser, 2017) to incumbent firm discounts (Ide et al., 2016).

To demonstrate the theoretical existence of a barrier to entry stemming from reputation, we rely on one method of analyzing the operation of labor markets: the two-sided matching model. The first forms of this model were pioneered in Gale and Shapley (1962) with a one-to-one model of the marriage and college admission markets, as well as in Becker's 1974 application to marriage markets of earlier work by Koopmans and Beckmann (1957). The two-sided matching model has since been expanded to cases with both non-transferable utility, in which each participant's match utility is limited to what they themselves reap from the pairing (well represented by the 1962 Gale-Shapley model) and to cases of transferable utility, such as the model of goods transfer built in Shapley and Shubik (1972), in which transfers between agents are without restrictions.

The Gale-Shapley model began what is now a flourishing literature, and similar ideas have since been applied to various labor markets. In particular, the well-known general labor market application of Kelso and Crawford (1982) was foundational for its elegant use of heterogeneity on both sides of the match as well as transferable utility. More general frameworks have since been built, including Hatfield and Milgrom (2005), which presented a model of matching with contracts and embeds both Kelso and Crawford (1982) and Gale and Shapley (1962) through the use of generalized contracts, allowing for both transferable and non-transferable utility. The model presented below takes part of its base from Kelso and Crawford, utilizing some parts of their framework for a many-to-one market with transferable utility.

Such matching models allow for a broad range of heterogeneity in agents' preferences, as well as offering the ability to analyze equilibrium outcomes that fall within what is known as the core of the game. Under these models, equilibria are defined in a manner closely related to the definition of competitive equilibria using two base conditions, each formulated in different notation for different models. First, the concept of individual rationality states that all agents prefer their match to remaining unmatched. Second, agents must be unable to form coalitions such that all agents do weakly better and at least one does strictly better under the new coalition than under their old match; this is known as non-blocking. The union of individual rationality and non-blocking in an allocation is called stability. Stability is frequently guaranteed by a notion of substitutability, first discussed in the aforementioned Kelso and Crawford paper and broadly defined as a condition on models in which agents preferences' over their matches do not have complementarities. The set of stable matchings is then called the core of the game.

While both individual rationality and non-blocking are well-studied and well-defined in markets without externalities, the discussion of non-blocking in particular in models with externalities is relatively new. Though early papers like Shapley and Shubik (1969) demonstrated that stable equilibria need not exist in a model with externalities, recent literature has given us a deeper understanding of how to think about stability in such markets. An early paper in the literature on matching with externalities, Sasaki and Toda (1996), found a result akin to that of Shapley and Shubik, but also found that equilibria may exist if we consider stability to hold as long as there is no coalition that benefits from re-matching under all possible results after the change in the match. Subsequent papers which explored this topic continued to place restrictions on either the definition of stability or the definition of preferences. Two examples of the latter include requiring stronger versions of substitutability (Pycia and Yenmez, 2017), or requiring concavity in agent's valuation functions (Hatfield and Kominers, 2015). Recent work has additionally branched into complementary externalities, showing that environments with complementarities across contracts signed by agents can still have stable equilibria in certain many-to-many settings (Rostek and Yoder, 2020).

Reputation considered as an externality that impacts the hiring of workers has been explored only tangentially in either the matching literature or in the more general economics literature. The most closely related paper is Agarwal (2015), which considers the willingness of workers in the context of the hospital residency match to pay for more desirable programs by taking reduced salaries. The paper finds both theoretical and empirical evidence for the existence of this implicit tuition, though it relates this implicit tuition to the presence of on-the-job training.

The externality of reputation lies somewhere in between individual externalities, such as workers' preferences over coworkers, explored in Dutta and Masso (1997) and Echenique and Yenmez (2007), or firms' preferences over rivals' workers, discussed in Bando (2012, 2014), and the idea of aggregate externalities, recently discussed in Fisher and Hafalir (2016). One piece of recent work that also lies in the space between individual and aggregate externalities is Jaffe and Morris (2017), which discusses agent preferences over locations with positive externalities when other workers choose the same location. Indeed, the externality of reputation is one that depends on the individual workers at a firm, but is largely aggregate in nature; a single bad coworker does little to tarnish the worker's vision of a firm as a whole. At the same time, however, a workforce made entirely of low-quality workers does indeed seem likely to lead to a loss of reputation.

In constructing a model with this reputation externality, then, we rely on insights both from basic matching theory, and on work done in the space of matching with externalities.

3. The Model

We consider an economy with W workers indexed by $i = 1, \ldots, w$ and F firms indexed by $j = 1, \ldots, f$.

Each firm may hire as many workers as it wishes, paying each worker some salary and seeking to maximize profit. Each worker may match with only one firm, and seeks to maximize utility from this match. A match between a worker and a firm at a salary is called a *contract* x, where $\mathsf{w}(x)$ denotes the worker, $\mathsf{f}(x)$ the firm, and $\mathsf{s}(x)$ the salary associated with the contract. Call the set of all possible contracts $\mathbb{X}$. For any set of contracts $X \subseteq \mathbb{X}$, define $X_i = \{x \in X : \mathsf{w}(x) = i\}$ and $X_j = \{x \in X : \mathsf{f}(x) = j\}$; these are the sets of all contracts in X that contain worker i and firm j respectively. Additionally, define $\mathsf{W}(X)$ to be the set of all workers such that if $x \in X$, then $\mathsf{w}(x) \in \mathsf{W}(X)$.

A *matching* is an assignment function $\mu : \{1, \ldots, w\} \to \mathbb{X}$ such that each worker is assigned to one contract $\mu(i)$, and if $\mu(i) = x$, then we have $\mathsf{w}(x) = i$. A worker i is employed by a firm j at salary $s_{i,j}$ if and only if $\mu(i) = x$, $\mathsf{f}(x) = j$ and $\mathsf{s}(x) = s_{i,j}$. If i is unmatched, i is unemployed and has contract denoted $\mu(i) = \emptyset$. If j is the firm associated with contract x and j employs $\mathsf{w}(x)$ under μ, we then have $x \in Y$, where Y is the set of contracts held by firm j under μ.

Given that each firm j holds a set of contracts Y, employing a set of workers $\mathsf{W}(Y)$, it earns revenue from a production function $y_j(\mathsf{W}(Y))$, which may be different across firms. Each worker $\mathsf{w}(x) \in \mathsf{W}(Y)$ with contract x receives compensation from the firm in the form of salary $\mathsf{s}(x)$ as given in the contract. Firms thus have profit:

$$\pi_j(Y) = y_j(\mathsf{W}(Y)) - \sum_{x \in Y} \mathsf{s}(x).$$

Worker utility may be a function of individual qualities like firm characteristics, worker preferences, and salary transfers, as well as the overall matching between firms and workers in the economy; in other words, workers may derive utility from the overall placement of other workers in the economy. Worker i thus recieves utility $u_i(x|\mu)$ for accepting contract x conditional on reference matching μ.

The reference matching determines the value of any externalities present in the utility function, and represents the beliefs the workers currently hold about where other workers are placed; note that this may include where they believe their own current placement is. In Section 5, we discuss the implications of specific utility functions.

Finally, we define *choice sets*. First, we consider workers. For a set of contracts X, define $\omega_i(X|\mu)$ to be the contract such that:

$$\omega_i(X|\mu) = \text{argmax}_{x \in X_i} \{u_i(x|\mu)\}.$$

In words, the contract $\omega_i(X|\mu)$ is the contract which maximizes i's utility for the subset of this set of contracts of which i is a part, and is thus i's choice set for that set of contracts X and a given matching μ. If there is a tie, workers may break this tie however they wish. If no contract in X_i offers the worker utility higher than $u_i(\emptyset|\mu)$, the worker may remain unemployed with $\omega_i(X|\mu) = \emptyset$ and receive unemployment utility. Call the union of all such individual choice sets $\Omega(X|\mu)$. Note that this union defines a matching, as all workers choose one contract; in the following sections, we use it as such.

For all firms j, the choice set of an individual firm from a set of contracts X is given as:

$$C_j(X) = \text{argmax}_{C \subseteq X_j} \{\pi_j(C)\}.$$

As is the case for workers, if firms have a tie between sets, they may break this tie however they wish, with the caveat that they must choose in favor of hiring a worker if the worker's salary is equal to their marginal product, and they would desire to hire the worker at any lower salary. Define C_j^μ to be the set of workers chosen by firm j in matching μ. Finally, call the union of all individual firm choice sets $\Gamma(X)$.

We note that, as defined, given the stipulation that each agent is utility or profit maximizing, both firms' and workers' choices satisfy the irrelevance of rejected contracts. This condition, first discussed in Aygun and Sonmez (2013), is crucial for finding stable outcomes in environments with weakened substitutability, which includes contexts like that discussed in this paper.

Definition 1 (Irrelevance of Rejected Contracts)**.** A choice function H satisfies the *irrelevance of rejected contracts* if for any sets of contracts X, $X' \subseteq X$ and any matching μ, then:

$$H(X|\mu) \subseteq X' \subseteq X \Rightarrow H(X|\mu) = H(X'|\mu).$$

We discuss the irrelevance of rejected contracts for both sides of the market. Firms derive profit only from workers that they hire, and earn the same profit from hiring the same workers regardless of who is rejected; thus, the the irrelevance of rejected contracts must be satisfied. This means that if firms choose a subset of workers from a larger set, they will choose the same subset when offered only some of the workers in the original set, if the subset contains their choice set. Note that, as the firm's profits are independent of the broader matching, we omit the argument μ from their choice functions.

Likewise, workers offered a menu of salaries and firms will choose the firm that maximizes their utility, and will still be happy to work for this firm at the given salary even when presented only with that firm-salary contract and some subset of the other rejected firms. Analogous to firms, workers derive utility only from their matched firm and the overall economy's matching, and thus derive the same utility from a given firm conditional on a given matching regardless of who else they have rejected.

3.1. Assumptions

First, we consider firms, and place three assumptions on their choices and production functions: a fixed base rate of production, substitutability, and a finite marginal product for all workers.

We assume that firms produce at a **fixed non-negative base rate**[4] when their choice set of workers is the empty set:

$$\pi_j(\emptyset) = y_j(\emptyset) \geq 0.$$

We additionally assume that each worker has a **finite marginal product at each firm**, though this marginal product may vary across firms.

Finally, we assume that firm production functions satisfy the **substitutes** rule first described in Kelso and Crawford (1982). Updated for the contract space, this rule says the following:

Definition 2 (Substitutibility for Firms)**.** We say firm choice functions *are substitutable* if for all $j \in F$ and for all X, X' such that $X' \subseteq X$:

$$\forall x \ s.t. \ x \in X, x \in X', \ x \in C_j(X) \Rightarrow x \in C_j(X').$$

Thus, if a contract is chosen from a larger set X, it must also be chosen from a smaller subset containing that contract. Firm profits display no externalities, implying that this definition of substitutability will be enough to ensure that firm-side considerations do not harm our ability to find stable equilibria.

[4]A note on formatting: we use bold for assumptions and theorems, while special terms are written in italics.

Next, we describe our conditions for workers. We have three assumptions about the utility functions of workers: first, we assume that **workers prefer higher salaries**, all else being equal. Second and third, we assume that worker utilities satisfy both consistency with a defined preorder, as well as a more complex substitutes assumption than what we require of firms. This substitutability condition is based off of the definition of substitutability used in Pycia and Yenmez (2017) in dealing with externalities.

To define these conditions, we must first define a preorder, a relationship that is both reflexive and transitive, over states of the world. Let $\succeq^W$ be a preorder over the states of the world. This ordering is defined as:

$$u_i(\mu(i)|\bar{\mu}) \geq u_i(\mu'(i)|\bar{\mu}) \ \forall i \in W, \bar{\mu} \subseteq \mathbb{X} \iff \mu \succeq^W \mu'.$$

In words, one matching, μ, is preferred to another, μ', under $\succeq^W$, if all workers receive weakly higher utility under the firm that they are matched to under μ, given that all workers consider the same reference matching across the two states, and that this is true for all reference matchings. Thus, we say $\mu \succeq^W \mu'$.[5] Next, define:

Definition 3 (Consistency)**.** A preorder $\succeq^W$ is *consistent* if utility functions satisfy the property that:

$$X \supseteq X', \mu \succeq^W \mu' \Rightarrow \Omega(X|\mu) \succeq^W \Omega(X'|\mu').$$

We assume that all worker utility functions satisfy **consistency**. In addition, a similar condition can be set for firms: if all firms have weakly higher profits under μ as compared to μ', we say $\mu \succeq^F \mu'$.

Next, for any reference matching μ and any set of contracts X, which may include the same firm-worker pair in multiple contracts with different salaries, define the rejection set:

$$R_i(X|\mu) = X \setminus \omega_i(X|\mu).$$

If workers are indifferent between multiple firms, they may break ties however they wish. $R_i(X|\mu)$ is thus the set of contracts that worker i rejects given the set X of choices and the broader matching μ. Note that we always have $|R_i(X|\mu)| \geq |X| - 1$, as a worker may choose at most only one firm and one contract with that firm. Workers may also reject all contracts and take unemployment.

[5] Note that different definitions of a preordering are possible; for example, we could define an ordering $\hat{\succeq}^W$ such that $\mu \hat{\succeq}^W \mu'$ would imply that all workers simply had higher salaries under μ.

Using this definition, we can now define our **substitutability** condition, which we assume all worker utility functions satisfy.

Definition 4 (Substitutability for Workers)**.** Given matchings $\mu \succeq \mu'$ and two sets of firm-salary pairs $X' \subseteq X$, workers' choices *are substitutable* if all of the following hold:

$$X' \subseteq X \Rightarrow R_i(X|\mu) \supseteq R_i(X'|\mu) \tag{1}$$

$$\mu \succeq \mu' \Rightarrow R_i(X|\mu) \supseteq R_i(X|\mu') \tag{2}$$

$$X' \subseteq X, \mu \succeq \mu' \Rightarrow R_i(X|\mu) \supseteq R_i(X'|\mu'). \tag{3}$$

We will now consider each of these properties individually. First, consider (1): for any matching μ, we must have:

$$X' \subseteq X \Rightarrow R_i(X|\mu) \supseteq R_i(X'|\mu).$$

This equation gives us the standard substitutability construction, as presented in the firm assumptions above; for any matching μ, workers reject more contracts from a superset of contracts than from a subset, and any contracts rejected from the subset must also be rejected in the superset.

Next, consider (2) in order for substitutibility to hold, we must have for any set of contracts X:

$$\mu \succeq \mu' \Rightarrow R_i(X|\mu) \supseteq R_i(X|\mu').$$

This equation tells us that workers reject the same contracts from a given set X conditional on a more preferred matching as they do conditional on a less preferred matching. Thus, workers must become more picky about which firm they match to, given that other workers have matched to better firms. Note here that either the sets are equal, or the worker has rejected all contracts in X under μ when they would have accepted one of these contracts under μ'.

Finally, if (1) and (2) hold, then we know that the conjunction of these two properties, (3), holds true, defining our substitutability condition:

$$X' \subseteq X, \mu \succeq \mu' \Rightarrow R_i(X|\mu) \supseteq R_i(X'|\mu').$$

Given more choices and a more preferred overall matching, workers must reject more contracts. The union of all workers' rejection functions is given as $R^W(X)$.

Again, we can describe a similar rejection function for firms. For a set of contracts X, which again may include the same firm-worker pair in multiple contracts with different salaries, firms have rejection functions:

$$R_j(X) = X \setminus C_j(X).$$

Firms reject all those worker-salary pairs that are not in a profit-maximizing subset of the available labor contracts. The union of all firms' rejection functions is given as $R^F(X)$.

3.2. Individual Rationality and Stability

Definition 5 (Individual Rationality)**.** A matching μ is *individually rational* if for all workers i and firms j, we have:

$$u_i(\mu(i)|\mu) \geq u_i(\emptyset|\mu)$$
$$\pi_j(C_j^\mu) \geq \pi_j(\emptyset) \text{ and } \forall x \in C_j^\mu,\ \pi_j(x) \geq 0.$$

A matching is individually rational if all workers are either unemployed or prefer being employed at their current job to being unemployed, if all firms prefer their current choice set of contracts to hiring no workers at all, and all firms are either indifferent to or prefer to hire each worker individually, as opposed to not hiring that worker.

In order to define stability, we must first define a blocking coalition. We adopt a short-sighted definition of blocking, assuming that workers do not consider their own impact on the broader economy when they move.

Definition 6 (Blocking Coalition)**.** A *blocking coalition* is some set of contracts X, a group of workers $\mathsf{W}(X)$, and one firm j, where each contract x in X has $\mathsf{f}(x) = j$ and each $\mathsf{w}(x)$ is present in $\mathsf{W}(X)$ exactly once, such that:

$$u_i(X_i|\mu) \geq u_i(\mu(i)|\mu)\ \forall i \in \mathsf{W}(X)$$
$$\pi_j(X) \geq \pi_j(C_j^\mu)$$

with strict inequality holding for at least one agent in $\{j\} \cup \mathsf{W}(X)$.

A blocking coalition is thus a group of workers and a firm that can profitably deviate from a given matching. Note that any blocking coalition with more than one firm can be broken down into component coalitions with a single firm each, at least one of which must block. Finally, we define stability.

Definition 7 (Stability)**.** A matching μ is *stable* if it is individually rational and there exist no blocking coalitions.

4. Finding Equilibrium Outcomes with a "Stealing Process"

In order to analyze outcomes for market concentration and determine whether or not barriers to entry may be present when workers have preferences for firm reputation, we must first find an equilibrium in the labor market in this context. We can think of reputation as a form of externality imposed on workers; how the broader set of workers think of a firm may change how an individual worker views that same firm, and thereby impact their potential job satisfaction when working there. In general, finding stable equilibria in matching settings often relies on the use of a fixed point algorithm, often with many similarities to the deferred acceptance algorithm of Gale and Shapley (1962). In spite of the presence of externalities, a similar technique can be employed here, with some modifications which ensure that the model works well for transferable utility. We call this fixed point algorithm the *Stealing Process*.

Previous work in the space of matching with externalities has largely dealt with non-transferable utility.[6] One notable exception is Pycia (2012), which deals with transferrable utility under a strict set of bargaining rules. Pycia and Yenmez (2017) can be extended to a transferable-utility scenario with a lowest monetary unit, however, and we do so here.

One of the challenges of transferable utility is that we have a very large contract space to search over. Any contract must stipulate not only which firm and worker are making the agreement, but also a value for the salary that the firm will pay to the worker. Contract spaces can thus become very large with a small base unit—say, one cent, \$0.01—and reasonably large marginal products for workers.[7]

One way around this is to search over our contract space in a more strategic manner. We begin with a smaller subset of the possible contracts in our space, namely, the set of marginal product contracts and contracts that make workers indifferent between working at the firm and taking unemployment.[8]

[6] See Sasaki and Toda (1996)), Alcalde and Revilla (2004), Echenique and Yenmez (2007), Bando (2012), Fisher and Hafalir (2016), and most of Pycia and Yenmez (2017), to name a few.

[7] Given that we place no restrictions that force the salary to be positive, this contract space is conceivably infinite, even with a lowest monetary unit. In practice, this lower unboundedness will never to be an issue where workers have finite utility from working for a firm; workers will always choose unemployment before an infinitely negative salary, and firms will thus never be able to match to such a contract. Furthermore, in cases where workers do not actively harm firm productivity, firms will be willing to offer a salary of 0 to workers, ensuring that employment at some firm will likely supersede unemployment regardless. In general, as long as workers are not infinitely harmful to firm profits, this infinite lower bound will never raise an issue.

[8] For further discussion of why we chose this initial set, see Section 4.2.

Next, we find a stable equilibrium for this subset of the contract space. The algorithm for doing so proceeds in two parts, both based on Pycia and Yenmez (2017). First, we construct an auxiliary matching such that all workers are happier with their contract under the auxiliary match than they would be with any other contract, given the same matching.

We then use this auxiliary matching as workers' initial assumption about the matching of the economy and run a form of deferred acceptance, the algorithm introduced by Gale and Shapley (1962). Workers choose contracts from the total set while assuming that the auxiliary matching defines the market, while firms choose from an initially empty set. The contracts chosen by the workers are then added to the set available to the firm, and the workers form a new belief about the state of the matching in the market based on the choices of their fellows. We then begin the process again, with firms now choosing from an increasingly large set of contracts; furthermore, any contract that firms have rejected is removed from the workers' set. When the two sets stabilize and the reference matchings of the workers cease to move, we have found an equilibrium for the subset of the contract space we began with.

However, it may not be the case that this equilibrium is stable for the entire contract space; in particular, a firm may wish to offer a worker a salary that is between two marginal products. Thus, we expand the space to include contracts between firms and workers that the firms do not currently employ that (i) are in the possible set of discrete contracts, (ii) the firm is willing to offer in some state of the world and (iii) are more attractive under the current matching than the worker's current job. We call these new contracts *stealing contracts*, intended to take a worker from their previous match. With this new, expanded set of contracts, we re-calculate our stable equilibrium. If nothing changes from the last equilibrium, we are done; if we have new matches, we check to see if more contracts are needed. This proceeds until we have finished our matching, or until we have added all possible contracts to our set.[9]

The detailed process for a set of firms F and a set of workers W is as follows.

S0: Construct a set of contracts, X_0, which includes each worker-firm pair once at some salary $s^0_{i,j}$ such that for contract x with $\mathsf{s}(x) = s^0_{i,j}, \mathsf{f}(x) = j$:

$$u_i(x|\emptyset) = u_i(\emptyset|\emptyset),$$

where utility conditional on the empty set implies utility conditional on all workers being unmatched.

[9]Note that the repeated addition of stealing contracts, in essence a gauranteed blocking coalition, to the contract space bears some similarity to work showing that there are paths to stability when, starting from an arbitrary matching, we repeatedly allow blocking pairs to match. See Roth and Vande Vate (1990).

We call these our *lower bound salaries.*

Next, for all sets $C \subseteq W$ and for all firms $j \in F$, construct a contract x for each worker $i \notin C$ such that $\mathsf{f}(x) = j$, $\mathsf{w}(x) = i$. For some lowest salary increment c, define the operator $[s]_c$ to mean the largest multiple of c less than s. Then set the salary of the contract x equal to:

$$\mathsf{s}(x) = [y_j(C \cup \{i\}) - y_j(C)]_c.$$

If this contract is not already present in the current set of contracts, add it. Thus, for every possible marginal product that a worker i may have at a firm j, a contract is offered with the closest discrete salary that the firm is willing to pay. Note that this set is finite; call it set X^0.

We now begin to cycle through the following steps, **S1-S3**; call the current cycle k.

S1: We first construct the auxiliary matching for cycle k. Using the current set of contracts X^k, we construct an auxiliary matching $\hat{\mu}$ such that for all workers $i \in W$:

$$u_i(\hat{\mu}(i)|\emptyset) \geq u_i(\omega_i(X^k|\hat{\mu})|\emptyset).$$

The matching $\hat{\mu}$ is such that all workers weakly prefer their match to any other contract they may choose from X^k given the matching $\hat{\mu}$. We construct this by setting $\mu_0 = \emptyset$, and then defining the matching recursively; for each round t of this definition, set $\mu_t = \Omega_t(X^k|\mu_{t-1})$ for every $t \geq 1$. Since the number of contracts considered in a given set is finite, there are a finite number of possible matchings, and thus the auxiliary matching algorithm will eventually return to a previous matching with some $\mu_t = \mu_{t+q}$ for some $q \geq 1$. Take the minimum t such that this is true and set $\hat{\mu} = \mu_t$. This satisfies the property that:

$$\hat{\mu} \succeq^W \Omega(X^k|\hat{\mu}).$$

By the definition of our preorder $\succeq^W$, then, for all $i \in W$, $u_i(\hat{\mu}(i)|\emptyset) \geq u_i(\omega_i(X^k|\hat{\mu})|\emptyset)$. The proof that such a matching exists is **Result 1** in Appendix A, based directly on Pycia and Yenmez (2017).

S2: We now construct the full matching for the current set of contracts X^k. We do this iteratively, with the current iteration denoted by t. First, construct the sets $X^W(1) = X^k$, $X^F(1) = \emptyset$, where $X^W(1)$ is the set of contracts available to the workers in the first iteration, and $X^F(1)$ the set available to the firms in the first iteration. Next, construct the first reference matching, $\mu(1) = \hat{\mu}$.

In each iteration $t = 1, 2, \ldots$, construct new sets:

$$X^W(t+1) = X^k \setminus R^F(X^F(t))$$
$$X^F(t+1) = X^k \setminus R^W(X^W(t)|\mu(t))$$
$$\mu(t+1) = \Omega(X^W(t)|\mu(t)).$$

The intuition is as follows: in each iteration, the firms reject some contracts from those available to them; the workers then have progressively fewer contracts available, while the firms, receiving new opportunities to sign new contracts, have progressively more. Workers' reference matching will appear progressively worse in terms of $\succeq^W$ as firms reject the workers' favorite offers. This monotonicity will allow the two sides to reach a stable equilibrium.

For a full proof of convergence, see **Result 5** in Appendix A, again based directly on Pycia and Yenmez (2017).

We stop this phase of the algorithm when we have $X^W(t+1) = X^W(t)$, $X^F(t+1) = X^F(t)$, and $\mu(t+1) = \mu(t)$. The outcome is then $X^F \cap X^W$. Call this matching μ^k. Matching μ^k now represents a stable equilibrium for the set of contracts X^k.

S3: We now find *stealing salaries* for the current matching and set of contracts, under which firms may attempt to entice workers from their current match. For matching μ^k and contract set X^k, we construct the set S of new contracts with all firm and worker pairs i, j such that $i \notin \mathsf{W}(C_j^{\mu^k})$. For x_0 with $\mathsf{f}(x_0) = j$, $\mathsf{w}(x_0) = i$, and $\mathsf{s}(x_0) = y_j(\{i\})$, if it is the case that:

$$u_i(\mu^k(i)|\mu^k) < u_i(x_0|\mu^k).$$

Then we add a new contract x to S with $\mathsf{w}(x) = i$ and $\mathsf{f}(x) = j$. Call s^* the salary such that if we were to set $\mathsf{s}(x) = s^*$, then we would have:

$$u_i(\mu^k(i)|\mu^k) = u_i(x|\mu^k).$$

For each such contract, we then define the contract's salary to be:

$$\mathsf{s}(x) = [s^*]_c + c.$$

Where c is as above the positive increment defined as the smallest possible change in salary.

Intuitively, S is the set of contracts such that, all else equal, worker i would take the contract x from j over their current match $\mu(i)$, and firm j finds it profitable or at least profit-neutral to offer

x under some possible state of the market. Finally, to avoid duplicating any existing contracts, set:

$$X^{k+1} = S \cup X^k.$$

Termination: At the end of each run of **S3**, return to **S1** and begin the next cycle. Stop the algorithm when no new contracts are added in **S3**; that is, when $X^{k+1} = X^k$. The final matching, μ^k, is a stable equilibrium for the set of contracts defined with salaries in increments of c.

4.1. Properties of the Stealing Process

Before turning to particular utility functions and the analysis of equilibria found by the Stealing Process, we first establish some basic facts about the process that are independent of functional forms. First, we establish convergence and the stability of the outcomes of the process. Second, we lay the groundwork for the issue of modeling a barrier to entry by showing that a worker's choice of firm will always be independent of the reference matching if worker preferences are substitutable.

Lemma 1. *The Stealing Process converges in finitely many iterations.*

Proof. First, we note that by **Result 5**, for a finite number of contracts, each cycle **S1-S2** converges in finite time. Thus, it remains only to show that there exists a cycle k such that:

$$X^{k+1} = S \cup X^k$$

and thus that the full algorithm terminates.

Because we add only new contracts to the set X, the algorithm will terminate if either no new contracts are added, or all possible stealing contracts have been added. The former may happen at any time; the latter is certain to happen. Because worker marginal products are finite and there are a finite number of salaries such that s is a multiple of c below the values of these marginal products, eventually either no firms will desire to bid, or all contracts will be added, implying that we will have

$$X^{k+1} = S \cup X^k$$

for some round k, and the process will terminate. □

Lemma 2. *Let W be a set of workers and F be a set of firms whose choices satisfy the assumptions made in Section 3.1. Then the outcome of the Stealing Process μ is individually rational.*

Proof. By **Result 5**, the outcome of each cycle **S1-S2** is individually rational for the market it is defined in. Because any outcome of the Stealing Process is necessarily an outcome of **S1-S2**, it must also be individually rational. □

Theorem 1. *Let W be a set of workers and F be a set of firms whose choices satisfy the assumptions made in Section 3.1. Then the Stealing Process converges to a stable allocation μ.*

Proof. By **Lemma 1**, the process terminates, and by **Lemma 2**, the outcome is individually rational. Thus, it remains only to show that the outcome has no blocking coalitions. Thus, if the result of the Stealing Process, μ, is not a stable, then there must exist a blocking coalition of workers and a firm j together with a set of contracts X such that $\mathsf{f}(x) = j \; \forall \; x \in X$, and $\mathsf{W}(X)$ defines the set of workers in the coalition. This implies that:

$$u_i(X_i|\mu) \geq u_i(\mu(i)|\mu) \; \forall i \in \mathsf{W}(X)$$

$$\pi_j(X) \geq \pi_j(C_j^{\mu}).$$

We observe that because the outcome of each cycle **S1-S2** is stable for the set of contracts it is defined over, any blocking coalition defined by contract set X must include at least one new contract not previously available in any round of the Stealing Process.

The set of workers that j employs under μ is $\mathsf{W}(C_j^{\mu})$. In order for the coalition to block, it must be the case that there exists a worker i such that $i \in \mathsf{W}(X)$, but $i \notin \mathsf{W}(C_j^{\mu})$. This is so because if $\mathsf{W}(X) = \mathsf{W}(C_j^{\mu})$, no worker could be made better off without the firm being made worse off and vice versa, as the only change in worker utilities or firm profits would come from changes in salary.

We have established that the firm wishes to add a currently unavailable contract; we will now show that any new contracts the firm wishes to add to the current set must be with workers not already with the firm. To prove this, we will first show that no current worker may have their salary lowered, and if a current worker's salary is raised, there exists a blocking coalition in which the worker's salary remains the same. The reasoning is as follows.

First, no worker both already matched to the firm and in the blocking coalition may have their salary lowered without decreasing their utility when the two are compared under a fixed reference matching. Thus, a set of contracts X with a worker i in $\mathsf{W}(C_j^{\mu})$ employed at a salary less than $\mathsf{s}(\mu(i))$ cannot block.

Second, assume that a worker already at the firm has their salary raised under X. Then it must be the case that the same coalition of agents also blocks with contracts X', in which all workers currently at the firm see no raises, and the firm makes strictly more profit than under X. Thus, any blocking coalition with contracts X with raises for current workers means that there exists a blocking coalition with contracts X' with no raises for current workers and higher firm profits.

However, if X' then includes no workers that are not currently employed at j with currently unavailable contracts, it cannot block μ by the stability of each cycle. Thus, because X' is different from X only in the salaries of currently employed workers, any blocking coalition must include a worker i not currently at firm j with a salary not currently in the available set of contracts.

But, if there exists a contract with a salary s less than i's maximum marginal product at j such that i would leave $\mu(i)$ for j, and this contract has an offerable salary for lowest monetary unit c, then j would be able to offer it in **S3**, and the algorithm cannot have terminated. □

We next turn to a general limitation of the assumptions we rely on. Because of the strict restrictions placed upon worker utility by the substitutability condition, we can show that while workers' utility may depend on the reference matching, their choices cannot. Our proof proceeds by showing that any case in which a worker's choice of firm depends on the matching may violate substitutability.

Theorem 2. *If workers always prefer employment to unemployment and preferences are substitutable, then for any set of contracts X and any two reference matchings μ, μ', it must be the case that:*

$$\omega_i(X|\mu) = \omega_i(X|\mu').$$

In words, under substitutability, a worker's choice of firm is independent of the reference matching.

Proof. Consider any economy with workers W and firms F such that $\theta, \lambda \in F$ are two such firms. Suppose that there exists a worker i with choices that do depend on the matching; that is, for a set of contracts X containing a contract for each of θ and λ, denoted x_θ and x_λ respectively, there exists a matching μ_θ such that $\omega_i(X|\mu_\theta) = x_\theta$, and a matching μ_λ such that $\omega_i(X|\mu_\lambda) = x_\lambda$.

It is either the case that $\mu_\theta \succeq^W \mu_\lambda$ or vice versa; the states may also be equal under the ordering. Assume without loss of generality that $\mu_\theta \succeq^W \mu_\lambda$. Note that the method of ordering the world can be the one given above in Section 3, or any other; the specific structure of this ordering does not impact the result.

By the structure of the worker's preferences, we must have:

$$R_i(X|\mu_\theta) = \{x_\lambda\}$$

$$R_i(X|\mu_\lambda) = \{x_\theta\}.$$

However, by substitutability, we must have:

$$\mu_\theta \succeq^W \mu_\lambda \Rightarrow R_i(X|\mu_\theta) \supseteq R_i(X|\mu_\lambda).$$

However, here, we have disjoint sets, and neither is a superset of the other. Thus, here, $R_i(X|\mu_\theta) \not\supseteq R_i(X|\mu_\lambda)$, and the worker's preferences are not substitutable. Thus, worker choices may not depend on the matching. □

Note that with substitutability alone and without the assumption that workers always prefer employment, workers may choose to reject all contracts in some set X under a certain reference matching μ, when they would have chosen one of these under a different reference matching μ'. Thus the above holds only wherever workers always prefer being employed to being unemployed, or vice versa.

4.2. Initial Contracts and the Need for Stealing Salaries

Before we continue, we permit ourselves a brief digression about the reasoning behind the choice of the initial set of contracts in the Stealing Process. Any initial set of contracts will result in a stable equilibrium; any firm can always bid on any worker in the stealing phase **S3**, even an unmatched worker, as long as there are contracts that the firm wishes to offer and has not yet bid. For example, both the empty set and the set of all contracts will result in a stable equilibrium when used as the initial contract set. Choosing a certain set of initial contracts can, however, help reduce the time we need to solve.

Our choice is the set of marginal product contracts. Why is this better than, for example, all possible contracts, or all integer-salary contracts? For small values of c, such as $c \leq 0.01$, or many workers, enumerating and searching over all possible contracts can be very time consuming. Integer salaries can also result in large initial contract spaces, but may do little to prevent bidding wars among firms where the marginal product of a worker is fractional. This connects to the reason for using a larger initial space than just the lower bound salaries; using only these salaries, we may have to run many rounds of the Stealing Process as firms gradually push up the salaries of a set of workers.

Concretely, the marginal product contract space has two advantages. First, it is always dominated in size by the total contract space, and the number of contracts can be quite manageable for small to medium sets of workers. Second, and more specifically, it ensures that there are never any stealing salaries when all firms have constant marginal products for each worker (though these marginal products may differ across firms). For an example, see Section 5.2.

Theorem 3. *If all firms have constant marginal products for all workers, the Stealing Process will terminate after one round with no contracts added to the initial set.*

Proof. Suppose that all firms have constant marginal products for each worker; for example, their production functions are linearly separable in workers. Suppose for a contradiction that after the first cycle **S1-S2**, there exists a stealing salary s such that firm j could attract worker i away from their current firm, $\mu(i)$. This implies that there exists a salary that both pushes i's utility above their current match and that firm j is willing to offer that is unavailable to be offered from a set of contracts containing only lower bound salaries and marginal product salaries.

However, because the worker has constant marginal product at j, it must be the case that j is always willing to offer the maximum salary to i, namely i's marginal product. Because this is the highest salary that the firm is willing to offer, it must be higher than or equal to s. This implies that the worker would be willing to move to the firm at the marginal product salary, which the firm is willing to offer by constant marginal products. Thus, the marginal product contract with i and j blocks the outcome of **S1-S2**. However, this outcome must be stable for the current set of contracts, establishing a contradiction. $\square$

Thus, the marginal product initial set provides quick results for common substitutable production functions, with the prime example being additively separable production functions.

5. Utility Functional Forms

We now specify an initial functional form for worker utility. In this first preference model, we allow workers' preferences to impact each others' utility.

We have utility functions for worker i accepting a contract x with firm j and salary $\mathsf{s}(x) = s_{i,j}$ of the form:

$$u_i(x|\mu) = s_{i,j} + p_{i,j} + \gamma_i r_j(\mu, \bar{r}),$$

where $p_{i,j}$ is the worker's fixed a priori preference for a given firm, and r_j is the externality representing others' opinions of the firm, which may depend on the matching μ. Each worker i has additional parameter γ_i, which is a weight which determines how much value the worker places on others' opinions.

When a worker is unemployed, their utility takes the form:

$$u_i(\emptyset|\mu) = \gamma_i r_\emptyset(\mu, \bar{r}).$$

Thus, our externalities are represented by the parameter r_j. One choice for this parameter is as follows.

We assume that each worker i has a fixed human capital value h_i which serves as an input to firm production functions. Let $H = \sum_{i\in M} h_i$, the sum of total human capital values for all workers. Define $m_{i,j}$ to be the contract with $\mathsf{f}(m_{i,j}) = j$, $\mathsf{w}(m_{i,j}) = i$, and $\mathsf{s}(m_{i,j}) = y_j(\{i\}) = \bar{s}_{i,j}$, the marginal product of worker i being the only worker at firm j. Then we have:

$$r_j(\mu, \bar{r}) = \sum_{i\in M} \left(\frac{h_i}{H}\right) \left(u_i(m_{i,j}|\emptyset, r_j = \bar{r}) - u_i(\mu(i)|\emptyset, r_{\mu(i)} = \bar{r})\right).$$

Thus, r_j is determined by the worker-quality-weighted average of how the utility of a given firm j compares to workers' matched firms under μ, denoted $\mu(i)$. This allows a worker to take into account both the opinions of others about a given firm via the first, fixed term, $u_i(m_{i,j}|\emptyset, r_j = \bar{r})$, as well as the changing fortunes of other workers in evaluating their own outcome, via the variable second term $u_i(\mu(i)|\emptyset, r_{\mu(i)} = \bar{r})$ which depends upon the current reference matching μ. To deal with the recursive nature of this definition, we define a $\bar{r} \geq 0$ which is a fixed proxy for all r_j for the purposes of calculating this utility. When a worker is unmatched, we set the second term equal to $\gamma_i \bar{r}$.

We can then write workers' utility of being unemployed as:

$$u_i(\emptyset|\mu) = \gamma_i r_\emptyset(\mu, \bar{r}) = \gamma_i \sum_{w\in M} \left(\frac{h_w}{H}\right) \left(\gamma_w \bar{r} - u_w(\mu(w)|\emptyset, r_{\mu(w)} = \bar{r})\right).$$

Note that this implies that unemployment is always less attractive than any contract from which a worker derives positive utility under an empty reference matching; in other words, if a worker likes a firm independent of how others are matched, that firm will always be more attractive to the worker than unemployment. In addition, we will show below that this functional form for utility meets substitutability requirements. Thus, our utility functions meet the all of the requirements of **Theorem 2**.

5.1. Properties of the Utility Function

First, we show that this utility function satisfies consistency as described in definition 3.

We recall that if worker choice functions satisfy consistency, then:

$$X \supseteq X', \mu \succeq^W \mu' \Rightarrow \Omega(X|\mu) \succeq^W \Omega(X'|\mu').$$

In order to show consistency, then, we must show that workers will choose more highly ranked contracts under $\succeq^W$ when offered more contracts or when facing a higher-ranked reference matching.

First, we consider the size of the contract set. Each firm's r_j is independent of the contracts which are offered to the worker, given a fixed reference matching. Thus, when the workers are offered more contracts, they will choose those that raise their utility for the given reference matching, and will necessarily choose weakly better contracts than they did under the smaller set. This means that consistency is satisfied with respect to the size of the set of contracts offered.

We next establish consistency with respect to the reference matching by showing that workers choose weakly better contracts given a better reference matching. For two matchings μ', μ, if $\mu \succeq^W \mu'$, by the definition of our preorder:

$$u_i(\mu(i)|\bar{\mu}) \geq u_i(\mu'(i)|\bar{\mu}) \ \forall i \in W, \bar{\mu} \subseteq \mathbb{X}.$$

Then we know that r_j will be lower for all firms under μ: each worker will earn a weakly higher utility from their match under μ as opposed to their μ' match, raising $u_i(\mu(i)|\emptyset, r_{\mu(i)} = \bar{r})$ relative to $u_i(\mu'(i)|\emptyset, r_{\mu'(i)} = \bar{r})$ and thereby implying a lower relative r_j for all firms under μ.[10] Given that all of the same contracts are available, workers will choose contracts with weakly higher $p_{i,j}, s_{i,j}$ and fixed $r_j(\emptyset, \bar{r})$ in response. Thus, they will choose a contract such that their choice under μ delivers weakly higher utility for a fixed reference matching than their contract under μ'.

The conjunction of these two conditions, choosing better contracts when there are more options or when there is a better reference matching, provides consistency.

Next, we show that the utility functions as described above satisfy substitutability as introduced in definition 4. We recall that substitutability is defined as the conjunction of the following properties of choice functions:

$$X' \subseteq X \Rightarrow R_i(X|\mu) \supseteq R_i(X'|\mu)$$
$$\mu \succeq \mu' \Rightarrow R_i(X|\mu) \supseteq R_i(X|\mu').$$

[10]However, note that relative rankings in r_j between firms will be the same under both reference matchings; indeed, all firms' values of r_j will fall by the same amount, including the utility of being unemployed.

First, note again that adding new contracts does not change the externality when the reference matching is fixed (the value of r_j remains fixed for all j regardless of the contract set), implying that all contracts rejected from a smaller set will be rejected from a larger set.

Second, for two matchings μ, μ', if we have for all i:

$$u_i(\mu(i)|\emptyset) \geq u_i(\mu'(i)|\emptyset).$$

Then we have $\mu \succeq^W \mu'$. It will then be the case that the value of r_j of all firms must weakly decline by the argument made above for consistency, and furthermore must decline by the same amount. Workers will then seek contracts with weakly higher values of the fixed parameters $p_{i,j}, s_{i,j}$ and fixed $r_j(\emptyset, \bar{r})$. This implies that we satisfy substitutability because no firm rejected under μ' can have improved relative to other firms under μ, and thus no firm rejected under μ' will be accepted under μ.

5.2. A Simple Illustration of the Stealing Process

In order to illustrate both the Stealing Process and the use of the above utility functions, we now present a simple example of finding an equilibrium when workers have preferences as described in this section.

Consider an economy with two firms, θ and λ, and two workers, a and b. For a set of workers C, each firm has production:

$$y_j(C) = \sum_{i \in \mathsf{W}(C)} h_i.$$

We set $\bar{r} = 0$. Each worker is defined by the following initial conditions:

Worker	h_i	γ_i	$p_{i,\theta}$	$p_{i,\lambda}$
a	2	1	2	1
b	1	1	1	2

We will proceed through each step in the Stealing Process.

S0: First, we construct our set of contracts; here, these are a set of tuples defined entirely by a firm, a worker, and a salary. Because firms have linear production functions, this set is simply:

$$X^0 = \{(a, \theta, 0), (a, \theta, 2), (b, \theta, 0), (b, \theta, 1), (a, \lambda, 0), (a, \lambda, 2), (b, \lambda, 0), (b, \lambda, 1)\}.$$

S1: Next, we build our auxiliary matching, which will give our workers initial priors about the quality of match they can expect. We allow each worker to successively choose contracts based on their choices in the previous iteration. Begin with $\mu_0 = \emptyset$. This means that each firm has initial r_j values:

$$r_\theta = \frac{2}{3}(2+2) + \frac{1}{3}(1+1) = \frac{10}{3}$$
$$r_\lambda = \frac{2}{3}(2+1) + \frac{1}{3}(1+2) = 3.$$

In the course of building the auxiliary matching, the workers will always choose the contracts with the highest salaries for a given firm. For example, a chooses between $(a, \theta, 2)$ and $(a, \lambda, 2)$, ignoring the zero salary contracts. Denote these highest-salary contracts by θ_i and λ_i for each worker i. The workers thus compare:

$$u_a(\theta_a) = 2 + 2 + \frac{10}{3} > 2 + 1 + 3 = u_a(\lambda_a)$$

$$u_b(\lambda_b) = 1 + 2 + 3 > 1 + 1 + \frac{10}{3} = u_b(\theta_b).$$

Thus, worker a chooses θ_a, worker b chooses λ_b, and we have the result of our first iteration of the auxiliary matching: $\mu_1 = \{(a, \theta, 2), (b, \lambda, 1)\}$.

Now that each worker receives positive utility from their match under μ_1, the value of the parameter r_j falls by definition. This results in r_j values:

$$r_\theta = \frac{2}{3}(2+2) + \frac{1}{3}(1+1) - \frac{2}{3}(2+2) - \frac{1}{3}(1+2) = \frac{-1}{3}$$
$$r_\lambda = \frac{2}{3}(2+1) + \frac{1}{3}(1+2) - \frac{2}{3}(2+2) - \frac{1}{3}(1+2) = \frac{-2}{3}.$$

Again, the workers compare:

$$u_a(\theta_a) = 2 + 2 - \frac{1}{3} > 2 + 1 - \frac{2}{3} = u_a(\lambda_a)$$

$$u_b(\lambda_b) = 1 + 2 - \frac{2}{3} > 1 + 1 - \frac{1}{3} = u_b(\theta_b).$$

Again, worker a chooses θ_a, while worker b chooses λ_b. Thus, the cycle is complete, and we have auxiliary matching $\hat{\mu} = \{(a, \theta, 2), (b, \lambda, 1)\}$.

This auxiliary matching demonstrates the results found in **Theorem 2**; namely, we see that the workers make the same choice regardless of the reference matching being μ_0 or μ_1. Though both workers have less utility under the second reference matching as the r_j values fall, these values fall equally for both firms, resulting in identical choices.

S2: We now consider the process of creating an actual matching for our initial set of contracts, taking into account both firms' and workers' choices as we run a form of deferred acceptance. We can describe the full match in the following table:

Round	X^W	X^F	μ	Ω	Γ
1	$\mathbb{X}$	$\emptyset$	$\{(a,\theta,2),(b,\lambda,1)\}$	$\{(a,\theta,2),(b,\lambda,1)\}$	$\emptyset$
2	$\mathbb{X}$	$\{(a,\theta,2),(b,\lambda,1)\}$	$\{(a,\theta,2),(b,\lambda,1)\}$	$\{(a,\theta,2),(b,\lambda,1)\}$	$\{(a,\theta,2),(b,\lambda,1)\}$
3	$\mathbb{X}$	$\{(a,\theta,2),(b,\lambda,1)\}$	$\{(a,\theta,2),(b,\lambda,1)\}$	$\{(a,\theta,2),(b,\lambda,1)\}$	$\{(a,\theta,2),(b,\lambda,1)\}$

In the initial round, the workers have the full set of contracts to choose from, and choose the same contracts as they did in the auxiliary matching. The firms have no contracts available to choose among, and therefore reject none. Moving to round 2, the firms now have the workers' non-rejected contracts to pick from; here, this is the workers' round-1 choice set. Because the workers earn their marginal product in each contract, the firms find them acceptable, and each is chosen by its respective firms. The same choices are repeated in round 3, and the algorithm comes to an end. Thus, the auxiliary matching becomes the full matching; call this μ. This is a stable matching for the initial set of contracts, X^0.

S3/Termination: Now we consider stealing salaries, salaries that a firm may wish to offer to attract a worker currently matched elsewhere. Given the current matching $\{(a,\theta,2),(b,\lambda,1)\}$, we must check if θ can add a stealing contract to X^0 with b, and if λ can add a stealing contract with worker a. Because utilities are linear in all arguments, we can subtract each not-currently-matched firm's best possible offer from each worker's current contract, and check if the result is negative in order to determine if the firm would be able to offer stealing salaries. Recall that the marginal salary contract for firm j and worker i is denoted $m_{i,j}$. We thus consider:

$$u_a(\mu(a)|\mu) - u_a(m_{a,\lambda}|\mu) = 2 + 2 - \frac{1}{3} - 2 - 1 + \frac{-2}{3} = \frac{1}{3} > 0$$

$$u_b(\mu(b)|\mu) - u_b(m_{b,\theta}|\mu) = 1 + 2 - \frac{2}{3} - 1 - 1 + \frac{1}{3} = \frac{2}{3} > 0.$$

Thus, there are no salaries that either firm is willing to offer that would entice the workers away from their current match, and the algorithm ends here. If such salaries had existed, we would have added them to the original set of contracts and re-run the algorithm from the beginning.

There are no stealing salaries in this example; this is a direct consequence of the result proved in **Theorem 3**, which tells us that for situations in which the marginal product of workers is constant, as it is in the example above, we will find a stable equilibrium without needing to add stealing salary contracts.

When, then, do we need stealing salaries? One place the need may arise is if a firm wants to swap one worker for another, but is unwilling to pay the new worker their entire marginal product. Consider the following example. We have two firms, θ and λ. The firm θ has $y_\theta(C)$ such that the revenue is equal to the human capital value h_i of the first worker in the set $\mathsf{W}(C)$. The firm λ is similar, except that it has revenue equal to twice the human capital value h_i of the first worker in $\mathsf{W}(C)$. For example, if θ accepts two workers, the first with $h_i = 4$, and the second with $h_i = 5$, it will have $y_\theta(C) = 4$. We furthermore set $\bar{r} = 0$.

We have two workers, a and b. The workers have initial conditions as follows:

Worker	h_i	γ_i	$p_{i,\theta}$	$p_{i,\lambda}$
a	10	1	1	0
b	3	1	0	1

The full marginal product contract set has four contracts for each firm, and within each subset, each worker has two contracts: one with salary 0, and one with a salary equal to their marginal product of being the first worker at the firm. Thus, the initial contract set is:

$$X^0 = \{(a, \theta, 0), (a, \theta, 10), (b, \theta, 0), (b, \theta, 3), (a, \lambda, 0), (a, \lambda, 20), (b, \lambda, 0), (b, \lambda, 6)\}.$$

Running a single round of the Stealing Process, ending when an allocation is reached for the initial eight contracts, we find a matched to θ for a salary of 10, and b matched to λ for a salary of 0. Call this matching μ. Firm θ has profit of 0, and firm λ has profit of 6. The firms have final r_j values:

$$r_\theta = \frac{10}{13}(10 + 1 - 10 - 1) + \frac{3}{13}(3 + 0 - 0 - 1) = \frac{6}{13}$$
$$r_\lambda = \frac{10}{13}(20 + 0 - 10 - 1) + \frac{3}{13}(6 + 1 - 0 - 1) = \frac{108}{13}.$$

Workers thus have final utility:

$$u_a(\mu(a)) = 10 + 1 + \frac{6}{13} = 11\frac{6}{13}$$

$$u_b(\mu(b)) = 0 + 1 + \frac{108}{13} = 9\frac{4}{13}.$$

However, we can see that this allocation is blocked by the contract x with $\mathsf{w}(x) = a, \mathsf{f}(x) = \lambda$, and $\mathsf{s}(x) \in (3\frac{2}{13}, 14)$. For example, with $\mathsf{s}(x) = 10$, firm λ increases its profit to 10, while worker a now has utility:

$$u_a(x) = 10 + 0 + \frac{108}{13} = 18\frac{4}{13} > 11\frac{6}{13}.$$

Thus, both parties would benefit from such an arrangement. This illustrates one of the major roles that stealing salaries play; for firms already making positive profit, they allow for beneficial arrangements between workers and firms when all possible beneficial contracts lie in between two marginal products.

6. Incorporating Firm Reputation Through Preference Updating

Throughout the above analysis, we have relied heavily on substitutability as a tool to ensure that we can find stable equilibria. However, this substitutability restriction is extremely limiting, even when it has been extended to deal with some classes of externalities, as in the work above. We cannot, for example, have workers explicitly care about the quality of their coworkers at a firm and still hope to reliably find stable matchings, as the core of such a game may be empty (Echenique and Yenmez, 2007).

Indeed, given that changes in our externalities are derived from changes in our matching, any externalities that impact the utility that certain firms provide more than the utility provided by other firms may result in workers' preferences changing with the matching. **Theorem 2** then tells us that such externalities violate substitutability and may thus result in an empty core. What this means is that although the utility functions described above are able to capture, through the use of the term r_j, workers' preferences for a well-regarded firm, we have not yet been able to capture a preference for reputation per se, because firms' relative images in the eyes of workers are essentially fixed, and we are thus unable to capture changes in a firm's fortune and prestige.

Therefore, in order to determine the influence of reputation on the market for workers in a labor matching model, we will have to utilize a more complex and complimentary structure for worker utility, and to allow for a possibly empty core. This allows us to, at last, explicitly introduce a preference for prestige-based reputation into workers' utility. We do so via our second preference model, which modifies and extends the utility functions in Section 5.

One possible model of prestige, and the one which we will employ below, is that workers observe a noisy signal of the amenities of the firm $\alpha_{i,j}$, and a concrete signal of firm efficiency and success, e_j, and take preferences to be a convex combination of the two:

$$p_{i,j} = \psi_i e_j + (1 - \psi_i)\alpha_{i,j}$$

This captures the fact that different workers may have different preferences over amenities and efficiency, as well as the fact that different workers may have different assessments of firm culture and amenities based on promotional material or interactions with current and former employees. We can then utilize an updating algorithm to generate e_j based on the success of the firm in the labor market.

The algorithm, which we call *Preference Updating*, begins by setting worker's initial preferences $p_{i,j}$ equal to their initial assessment of firm amenities, ignoring for the moment common observations of success. Thus, each worker begins with exclusively idiosyncratic, possibly unique preferences over the set of firms. We then run the Stealing Process for the current set of preferences, resulting in a stable equilibrium for these parameters.

From this equilibrium, workers update their preferences according to the new allocation of employees to firms. Taking a convex combination of shared efficiency signals and idiosyncratic opinions, workers change their preferences to account for the new information. In particular, we utilize two signals of firm success. The first, average worker quality, provides prospective employees with a sense of the exclusivity of the firm and the desirability of their future coworkers. The second, firm market share, as determined by the firm's share of total industry revenue, works as a proxy for firm performance. We then re-run the Stealing Process until either preferences converge or we return repeatedly to a previously-held set of preferences, implying that cycling occurs.

This algorithm terminates in one of two outcomes. First, if the algorithm converges, it gives us a stable set of preferences and an outcome of the Stealing Process which is stable for a market in which workers have these preferences. Convergence in this case is thus the same as producing a stable outcome. Second, if the algorithm does not converge, the preferences will cycle, as we will discuss further below. This indicates that the algorithm fails to reach a stable outcome for preferences; however, for a given (unstable) set of preferences, it does produce a stable outcome of the Stealing Process.

Explicitly, for our case, Preference Updating works as follows.

P0: Set initial preferences equal to a noisy signal of firm amenities. For example, one possible way to do so for all $i \in W$ and $j \in F$ is to set:

$$p^0_{i,j} = (1 - \psi_i)\alpha_{i,j}.$$

Where $\alpha_{i,j} \sim \mathcal{N}(a_j, \sigma)$, and a_j is the true value of the firm's amenities. This induces heterogeneity in worker's perception of firms, implying a difference in worker contact with firm recruitment materials or job descriptions, or a difference in workers' perceptions of various amenities—say, weighing vacation time against dress code. Run the Stealing Process once with these preferences.

P1: Now, begin to cycle through **P1-P2**. Call the current cycle t. Using the outcome of the most recent Stealing Process, μ, for each firm j set:

$$\bar{h}_j = \frac{1}{|C_j^\mu|} \sum_{i \in \mathsf{W}(C_j^\mu)} h_i$$

$$\tau_j = \frac{y_j(\mathsf{W}(C_j^\mu))}{\sum_{k \in F} y_k(\mathsf{W}(C_k^\mu))}.$$

Thus, $\bar{h}_j$ is the average worker quality at firm j, and τ_j is the firm's market share. Using weight w for the firm's market share, update worker preferences to:

$$p_{i,j}^t = \psi_i(\tau_j w + \bar{h}_j) + (1 - \psi_i)\alpha_{i,j}.$$

P2: Re-run the Stealing Process with updated preferences.

P3: Stop the algorithm when $p_{i,j}^t = p_{i,j}^{t-q}$ for some $q \geq 0$.

Though this process is not guaranteed to converge, it will either converge or cycle. To see this, note that the possible values each preference can take on are finite. For a given set of initial conditions, the possible values of w, ψ_i, h_i, and $\alpha_{i,j}$ are fixed, as are the revenues a firm earns from each coalition of workers. Thus, preferences change only with the allocation of workers. Because the number of workers and firms are finite, the number of allocations are finite, and the market must eventually return to the same allocation and thus the same preferences.[11]

The Preference Updating process introduces a second layer to our Section 5 utility functions. While the original utility functions allowed for worker's preferences to influence each other, we now add to this the ability for workers to have preferences which are based upon the changing fortunes of firms in the market. In essence, we have added a preference for reputation which more closely mirrors our own reality. In the work that follows, we will investigate the character of the equilibria found through the Preference Updating process and the difficulty of firm entry into the equilibria it generates.

[11] We note here that if the Preference Updating process converges, it produces a fixed point at which preferences remain unchanged even if we allow workers to recontract; however, if it does not converge, it does not imply anything about the existence or non-existence of the core. In the following section, which includes our economic results, all proofs which hold for results of the Preference Generation process also hold more generally for stable equilibria with the preference structure described in this section.

7. Economic Results

In the previous three sections, we have presented two different methods of equilibrium construction and specified two models of worker preferences. The first method of equilibrium construction, the Stealing Process, in conjunction with the utility functions described in Section 5, gives us equilibria when workers take into account the preferences of other workers. The second method and preference model, Preference Updating, described in Section 6, allows workers to both take into account the preferences of others, and for these preferences to be based on dynamic indicators of success—namely, worker quality and market share. Thus, through Preference Updating, we model an explicit preference for firm prestige and success, amplified by the agreement of fellow workers; in other words, an explicit preference for reputation.

In the following section, we explore the theoretical insights which can be gained from each method of equilibrium construction by investigating both the impact that the preferences introduced in the Stealing Process have on the labor market, and the influence of adding a preference for reputation via the Preference Updating process.

We begin by characterizing equilibrium salaries following matching via the Stealing Process with utility functions as defined in Section 5 and via converged runs of the Preference Updating process. We find that the addition of other workers' opinions into workers' preferences, as in a Stealing Process equilibrium, has an ambiguous impact on workers' salaries. Preference Updating has similarly ambiguous results, though in cases in which a monopsony has dominated the labor market, we see a definitive reduction in salaries after the introduction of a preference for reputation.

Next, we investigate vacancy chain dynamics for Stealing Process equilibria, asking how a firm closure or worker retirement impacts remaining firms and workers in the market. As in most standard matching models of the labor market, we see that a firm closure benefits the remaining firms at the expense of the workers, and vice versa for the retirement of a worker. This reflects a crucial notion of market power. Firms facing less competition for workers are able to offer weakly less attractive contracts, yet are still able to get employees.

This then leads us to a broader notion of equilibrium characterization: when we begin with a certain number of possible firms, how many will be able to acquire workers at the end of each of the two equilibrium generation processes? Though any number of firms is feasibly supported in either process for certain initial characterizations of workers, simulations of the two algorithms with randomized initializations show that including an explicit preference for prestige (as in Preference Updating) leads to more concentrated markets, even above the impact of workers taking into account the preferences of others.

Finally, we consider the core issue of the paper: how does workers' explicit preference for reputation limit firm entry, if at all? Here, we first show that substitutable preferences in a labor matching model cannot lead to a reputation-based barrier to entry, motivating the need for Preference Updating and an explicit, complementary model of preferences. Next, we demonstrate that the necessary conditions for firms to acquire workers become more difficult when workers have a preference for reputation, as in the Preference Updating process. Finally, we present several numerical examples of reputation acting as a barrier to entry, and consider this barrier's impact on the agents in the market.

7.1. Bounding the Equilibrium Salary

Given the utility functions described in Section 5, we can construct bounds for a worker's salary in a stable equilibrium. Consider worker i who has taken contract x at firm $\mathsf{f}(x) = j$ with salary $\mathsf{s}(x) = s_{i,j}$, with the broader economy matched according to μ.

First, we know that an equilibrium salary will be individually rational for the firm paying it; thus, the absolute highest the firm would be willing to pay is a worker's maximum marginal product. Because we assume that there are no complementarities between workers, it must be the case that a worker's (not necessarily unique) maximum marginal product is achieved when they are the only worker at the firm. Thus, any worker i at firm j in equilibrium must have a salary less than or equal to:

$$y_j(\{i\}) - y_j(\emptyset).$$

Next, we consider the lower bound of a worker's salary. If the equilibrium is stable, it must be the case that there is no salary that any firm is willing to pay such that, at the current matching, i would prefer that firm at that salary to j. Find firm k such that:

$$k = \text{argmax}_{f \in F}[(y_f(\mathsf{W}(C_f^\mu) \cup \{i\}) - y_f(\mathsf{W}(C_f^\mu))) + p_{i,f} + \gamma_i r_f(\mu, \bar{r})].$$

This firm is the worker's best possible outside option; the firm which could, when paying the worker their marginal product, offer the worker their highest possible utility outside of their current match. However, their current match must by stability be giving the worker weakly more utility. Thus, we know that:

$$s_{i,j} + p_{i,j} + \gamma_i r_j(\mu, \bar{r}) \geq (y_k(\mathsf{W}(C_k^\mu) \cup \{i\}) - y_k(\mathsf{W}(C_k^\mu))) + p_{i,k} + \gamma_i r_k(\mu, \bar{r}).$$

Rearranging, we know that:

$$s_{i,j} \geq (y_k(\mathsf{W}(C_k^\mu) \cup \{i\}) - y_k(\mathsf{W}(C_k^\mu))) + (p_{i,k} - p_{i,j}) + \gamma_i(r_k(\mu, \bar{r}) - r_j(\mu, \bar{r})).$$

We can simplify this further by utilizing our formula for r_j:

$$\begin{aligned} r_k(\mu, \bar{r}) - r_j(\mu, \bar{r}) &= \sum_{w \in W} \left(\frac{h_w}{H}\right) [u_w(m_{w,k}|\emptyset, r_k = \bar{r}) - u_w(m_{w,j}|\emptyset, r_j = \bar{r})] \\ &= \sum_{w \in W} \left(\frac{h_w}{H}\right) [(\bar{s}_{w,k} - \bar{s}_{w,j}) + (p_{w,k} - p_{w,j})] . \end{aligned}$$

This leaves us with lower bound:

$$s_{i,j} \geq \underbrace{(y_k(\mathsf{W}(C_k^\mu) \cup \{i\}) - y_k(\mathsf{W}(C_k^\mu)))}_{\text{Marginal salary from outside option}} + \underbrace{(p_{i,k} - p_{i,j})}_{\text{Preference differential}} + \gamma_i \underbrace{\sum_{w \in W} \left(\frac{h_w}{H}\right) [(\bar{s}_{w,k} - \bar{s}_{w,j}) + (p_{w,k} - p_{w,j})]}_{\text{Difference in other's opinions of } j \text{ and } k} .$$

The worker's salary must be at least greater than the sum of the three terms. This includes the first term, which is the marginal salary from their best outside option; the second term, which is the amount that they prefer k over j, though this may be negative; and the third term, the sum of others' opinion differential between the two firms, including both the marginal salary differential and the preference differential.

Thus, we see salary is weakly increasing in the worker's productivity at the outside option $(y_k(\mathsf{W}(C_k^\mu) \cup \{i\}) - y_k(\mathsf{W}(C_k^\mu)))$ and their preference for that firm $p_{i,k}$, as well as the outside option's overall productivity and other workers' preferences for it, via r_k. It is falling in the currently matched firm's overall productivity and attractiveness both to i and to others in the economy, both effects coming via r_j. Note that it is independent of the broader matching beyond firms j and k; the difference between the r values of two firms is constant for fixed preferences.

Only the third term, consideration for the preference of others, has effects beyond what we would expect from workers making choices based their outside options. We see that this term can be positive or negative. It can push up the salary of i at j when j is less well-regarded than its peer k, forcing the firm to pay higher salaries. At the same time, it can also allow a very well-regarded firm to pay lower salaries, creating an implicit tuition at the firm based on the opinions of other workers about the firm.

What impact does adding an explicit preference for reputation via Preference Updating have on this bound? Even for a fixed matching, it may raise or lower the salary lower bound; if a worker's outside option gains more market share and higher quality workers, more so than their own matched firm, it will push up the lower bound. The opposite may also be true, however, if the worker's matched firm gains a high market share and better workers than competitors.

There is one situation in which we can say definitively that Preference Updating has a negative impact on salaries: the case of an immediate monopsony. Suppose that in the first round of Preference Updating, a single firm j acquires all workers. Compared to the initial preferences, all workers with ψ_i higher than 0 have a higher estimation of this firm j than when the matching process began:

$$\psi_i > 0 \Rightarrow \psi_i(w + H/N) + (1 - \psi_i)\alpha_{i,j} > (1 - \psi_i)\alpha_{i,j}.$$

Thus, because j is the only firm with preferences rising, it will acquire all workers again in the following round, and the algorithm will reach completion. With all else fixed, the higher preferences for j implies possible lower salaries for all workers.

Note that this applies only to the lower bound of the salaries: as competitors become more desperate for workers, firms will be willing to offer salaries that represent a worker's marginal product at the firm when they work alone, which by substitutability must be their highest salary that the competitor is willing to offer in any state of the world. This may push up the monopsony's salaries, implying that workers may still get upper-bound salaries in a single employer world, if there is competition from other firms.

In sum, we see that the addition of consideration for other's preferences and of an explicit preference for reputation, though they make no definitive change to the salaries all firms must offer, will tend to lead to higher salaries for less-reputable firms, and lower salaries for those whose stature is higher.

7.2. Vacancy Chain Dynamics

We next turn to the effects on remaining workers and firms if a worker retires or a firm closes. This is of interest because it can instruct us on, most centrally to this paper, the benefits or drawbacks to workers of having more or fewer firms to choose from. In this section we will deal exclusively with the basic Stealing Process, independent of the utility functions used, provided that they meet the requirements of Section 3. Changes in the equilibrium following Preference Updating will be discussed further in Section 7.4.

In order to tackle this question, we need to formulate two additional assumptions. Namely, in the vein of Kelso and Crawford (1982), we need to assume that there are no ties between firms or sets of workers in any actor's preferences. We can describe this as follows.

First, we assume that workers have a **strict a priori ordering** over all possible contracts. This means that for any two contracts x, x', either $u_i(x|\mu) > u_i(x'|\mu)$ or vice versa for all workers $i \in W$ and all matchings μ. Note that this assumption is only a minor extension of the fact that worker choices are independent of the matching, assuming strictness where we might otherwise have equality. Second, we assume that **firms have no ties**; no two sets of workers have the same output at any firm. These two properties will allow us to ensure that trivial changes to equilibria do not occur, whereby a firm or worker switches one equally attractive contract for another. If the initial conditions of a market do not conform to these assumptions, it is easy to construct a market which does by perturbing by some epsilon the preferences or quality of certain workers.

Now that we have our additional assumptions, we begin by asking what happens when a firm leaves the market, a process known as vacancy chain dynamics. *Vacancy chain dynamics* is thus the process of dynamic recontracting from a stable equilibrium after an agent leaves the market.

In practice, this works as follows. Consider a set of firms F and a set of workers W. For this set of firms, we run the Stealing Process with firms offering until we have an equilibrium μ with final set of contracts X. This is identical to the worker-offering process described above, with the exception of using a firm auxiliary matching, and having firms start with the full set of contracts and workers with no contracts as options.

We then choose one firm, j, and assume that j refuses all workers. Define a new preference ordering $\hat{\succeq}^F$ such that all worker sets are less attractive than the empty set to firm j; in other words, a preference ranking that is consistent when j rejects all workers. Starting from the end of the most recent cycle of **S2**, where workers have available contracts X^W, firms have X^F, and the reference matching is μ, the economy begins recontacting. All workers previously matched to j are now rejected by the firm, and must find new matches.

Theorem 4. *Suppose that we have a set of workers W and a set of firms F, each with substitutable preferences. Let μ be the outcome of the Stealing Process for this set of agents. Then the vacancy chain dynamics following the closure of a firm converges to a new stable equilibrium μ^* such that $\mu \succeq^W \mu^*$, and $\mu^* \hat{\succeq}^F \mu$.*

Proof. We can begin by using **Result 6** in Appendix A, drawn from Pycia and Yenmez (2017), to see that the end of the initial **S2** cycle after the removal of a firm will result in a stable matching $\bar{\mu}$ such that $\mu \succeq^W \bar{\mu}$, and $\bar{\mu} \hat{\succeq}^F \mu$ for all firms still in the market. In words, removing a firm will result in a matching that remaining workers rate less highly under $\succeq^W$, and which gives each remaining firm weakly more profit. The logic of this is simple; any worker that was matched to the firm will now have to choose a contract they did not find as attractive as the old one, and any firm that has remained open has a chance to win new workers which may bring additional profit. Furthermore, any worker displaced by a worker from the closed firm will have to choose a contract they find less attractive than their previous contract. In other words, the vacancy chains caused by the firm's closure will negatively impact all workers.

Next, we consider the outcome of the full Stealing Process after a firm has closed. Above, we saw that a firm's closure will result in a match in that round such that workers are worse off per $\succeq^W$ and firms better off. Following this, firms have a chance to add new contracts with stealing salaries. If no new contracts are added, the process terminates in the stable equilibrium discussed above.

If, however, a new contract is added, this implies that some firm k is now able to offer a stealing salary to a worker that they could not offer before. Note that because worker's choices do not depend on the broader matching, each worker has a complete ordinal ranking over all possible contracts in this market, with each contract having the same rank regardless of matching. If firm k is now able to offer a stealing salary that they could not offer before, the worker to whom they are offering the salary must have taken a new contract; furthermore, that contract must be worse per $\succeq^W$ than their match under μ due to the properties of $\bar{\mu}$, and thus the new stealing contract must have a lower ordinal ranking than the μ match as well.

Starting a new cycle after the removal of j and the addition of any new salaries, workers face a set of contracts in which their maximum utility, minus externalities, is either unchanged or has fallen if their first choice was the closed firm. Call the matching for this new set of contracts μ'. Because choices are independent of the broader matching, there are no ties between contracts, and μ is stable for the market it is defined in, any worker that was willing to go to a firm under μ for some salary must still be willing to go to that firm for that salary; this implies that all open firms will have weakly higher profit under μ' as opposed to μ.

The firms, however, will offer the lowest salaries first, and no worker who went to a firm for some salary under μ will see a higher salary at that firm under μ'. Thus any worker matched to their μ firm has weakly lower utility. Any worker who is not matched to their firm under μ must also be matched to a less attractive contract per their ordering, for two reasons. First, we have shown above that no more attractive contracts than the μ matches have been added as stealing contracts. Second, any pre-existing contract that a worker orders above their μ match and that the firm corresponding to the contract would be willing to offer under μ' would block μ, as the firm still has access to its μ match, and this would imply that the firm prefers to add the contract to its chosen set.

All subsequent rounds must also add only stealing contracts lower in rank for the workers than their μ matches, and the above is thus true for all subsequent rounds. Because the number of possible stealing contracts is finite, the algorithm must terminate in an allocation such that firms still open have weakly higher profit, and workers have weakly lower ranked contracts; it must be the case that, if the final outcome is μ^*, then $\mu \succeq^W \mu^*$, and $\mu^* \hat{\succeq}^F \mu$. $\square$

Note that while the workers must choose weakly lower ranked contracts, it is not necessarily true that their utility falls. For the case of the utility functions considered in this work in particular, if the workers at the firm which closes are particularly fond of the firm (having very high $p_{i,j}$ values), the closure of the firm and their move to other, less-attractive firms may actually raise the r_j values of other firms enough to compensate for any decrease in utility experienced by other workers. The above applies exclusively to a consistent ordering over states of the world and thus, in our case, to the fixed portions of the utility function. In general, however, these results indicate to us that market concentration is likely to be good for firms, and negative for workers.

We now move on to considering what happens when a worker retires. For now, we define this as a worker refusing to accept any contracts, while still remaining a part of the preferences considered by other workers. Again, consider a set of firms F and a set of workers W. We run the worker-offering Stealing Process with these sets, reaching an equilibrium μ. At the end of the final run of **S2**, one worker, i, refuses all contracts, and matching continues from the original equilibrium matching and contract sets. Once again, define the new preorder $\hat{\succeq}^W$ to be consistent when i refuses all contracts.

Theorem 5. *Suppose that we have a set of workers W and a set of firms F, each with substitutable preferences. Let μ be the outcome of the Stealing Process for this set of agents. Then the vacancy chain dynamics following the retirement of a worker converges to a new stable equilibrium μ^* such that $\mu \succeq^F \mu^*$, and $\mu^* \hat{\succeq}^W \mu$.*

The proof of this theorem is analogous to the proof of **Theorem** 4; the complete proof is available in Appendix B.

A note specific to the utility functions used in this work is appropriate. Here, if we truly retire a worker and remove their preferences from the calculation of r_j, the ordinal ranking of contracts each worker has may shift dramatically, making statements comparing the pre- and post-retirement equilibria difficult. For example, if we remove a worker with very positive opinions about all firms, all workers may lose a great deal of utility. Another issue is removing a worker with very strong and diverse opinions about firms; this may result in a large shifting of these ordinal rankings in such a way that workers end up at both much higher individually ranked and much lower individually ranked firms. In this case, a retired worker can be thought of as an advisor; their experience still shapes the opinions of others, even though they do not seek a job themselves.

7.3. How Many Firms are there in Equilibrium?

Under both the Stealing Process and Preference Updating, any number of firms can theoretically be supported in equilibrium by assigning initial preferences and other starting conditions in particular ways. However, it is important to understand how the explicit preference for reputation can influence market concentration. While we cannot reliably analytically compare the equilibria produced by the two methods, we can run simulations which produce equilibrium outcomes for small sets of workers and firms, enabling us to compare the distribution of outcomes under each.

Thus, in order to get a better sense of how each algorithm systematically produces matches, we run each multiple times with randomized initializations and consider the distribution of firms acquiring workers under each. We call size of the set of firms which manage to hire at least one worker in equilibrium the *number of firms*. The number of firms in the first iteration of the Stealing Process within the Preference Updating algorithm, prior to the addition of prestige metrics to preferences, represents the number of firms at the end of a run of the basic Stealing Process; the final number of firms represents the full post-Preference Updating number of firms.

The general framework is as such. We consider an economy with 16 workers and 8 firms. Workers have utility functions as described in Section 5, which will be modified as described in Section 6 in the course of Preference Updating. Each firm is a priori identical across all runs, with revenue equal to the sum of the human capital values of its workers:

$$y_j(C) = \sum_{i \in \mathsf{W}(C)} h_i \ \forall j \in F.$$

Each worker has $\gamma_i = 1$ and $\psi_i = 0.5$ for all runs. For each worker in each run, h_i takes a random integer value between 1 and 5, and each worker's preferences for each firm are randomly distributed according to a Gaussian distribution with mean 2 and standard deviation 3. We have $w = 1$ and $\bar{r} = 0$ for all runs. For 1000 such randomly generated initial conditions, we collect the number of firms following the initial run of the Stealing Process and following the completion of the Preference Updating process.

7.3.1. Convergence

Before looking at the results of the simulations, we consider the question of convergence. While the Stealing Process is always guaranteed to converge to a stable equilibrium, due to the complementarities in preferences introduced by the Preference Updating process, the latter algorithm may cycle and thus not reach stability. We therefore ask as an initial question if the runs of the Preference Updating process on which we do reach stable equilibria are representative of the whole for our desired object of study, namely the number of firms.

In each run, we allow 200 iterations of Preference Updating before terminating; most runs of the algorithm which reach stability and find fixed preferences do so in less than 10 iterations. For 1000 random generations, about two-fifths reach a stable equilibrium; we have exactly 373 observations from runs which reached stability. There appears to be some slight bias in which rounds converge based on the number of final round firms. We display a percentage comparison in figure 1.

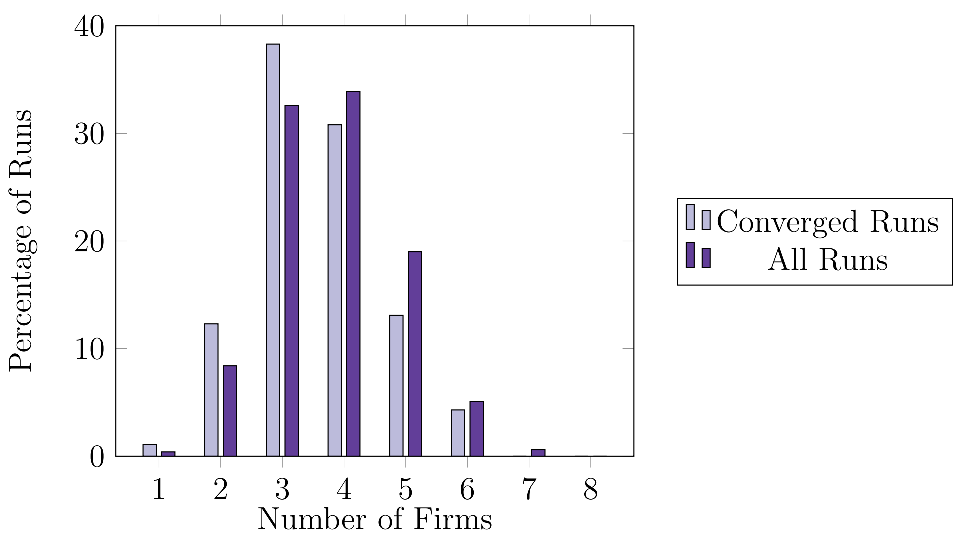

Figure 1: Comparison Between Convergent and Non-Convergent Runs

We see that the converged runs tend to have slightly lower number of firms; indeed, converged runs stochastically dominate non-converged runs. At first glance, this may be due to one of two causes. First, we may simply have preferences that are less likely to cause cycling be those which lead to fewer firms; or we may have the cycle itself probabilistically outputting a result that is at the high end of an oscillation. A trial run of 50 randomly generated initializations shows that the second possibility is not the case, however, as no run requires more than 6 rounds to reach its final number of firms. A closer look at the rounds that cycle shows that after reaching a fixed number of firms n, these same n firms simply trade workers back and forth across rounds. This indicates either that the algorithm has found a local cycle, or that the core of such runs may be empty. Characterizing why this is the case and in general why so few runs converge presents an interesting avenue for future research.

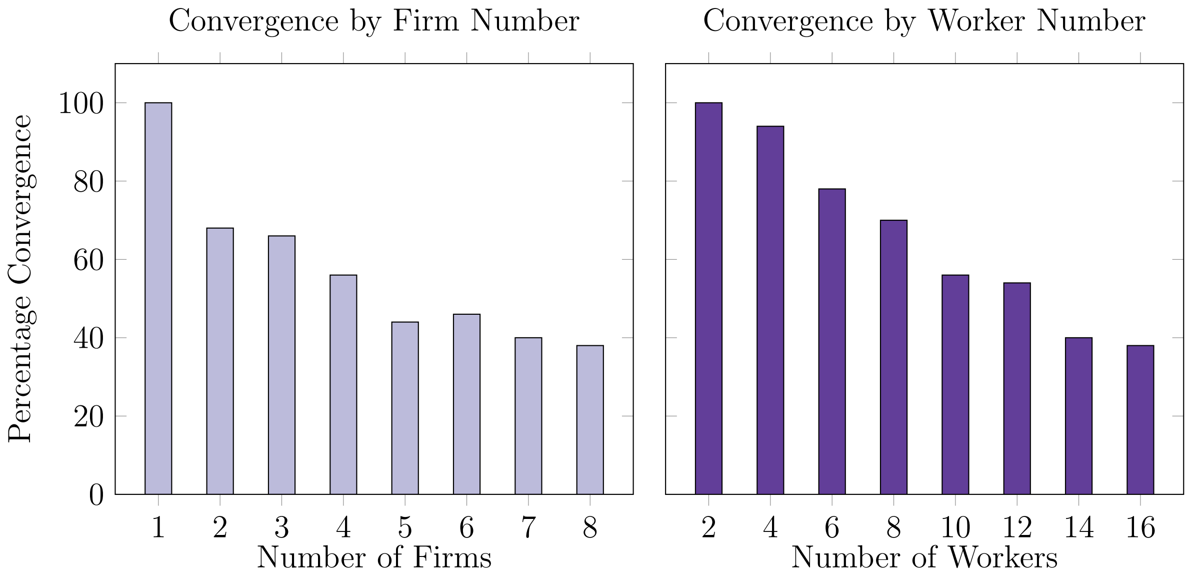

Figure 2: Convergence by Number of Agents

Lastly, we also consider that the probability of different numbers of firms and workers at initialization inducing more or less convergence. We plot convergence probability by worker and firm number for trial runs of size 50, fixing the size of the other agent coalition at 8 and 16 respectively, in figure 2.

As we would expect, fewer actors on either side leads to a higher rate of convergence, reflecting a decreased likelihood of getting the types of preferences needed to cause cycling; this may suggest that the major issue is with the algorithm being stuck in local cycles, as there would necessarily be fewer of these with fewer actors. This data also suggests that the workers play a larger role in determining convergence than the firms. For example, the convergence rate with only two firms and sixteen workers, 68%, is below that of eight firms and eight workers, 70%.

7.3.2. Basic Comparison

We compare the number of firms acquiring workers in the end of the Stealing Process to the number in the end of the Preference Updating process, reporting the number of firms in each of the converged observations in figure 3. Our y-axis measures the number of runs which display a given number of firms at convergence.

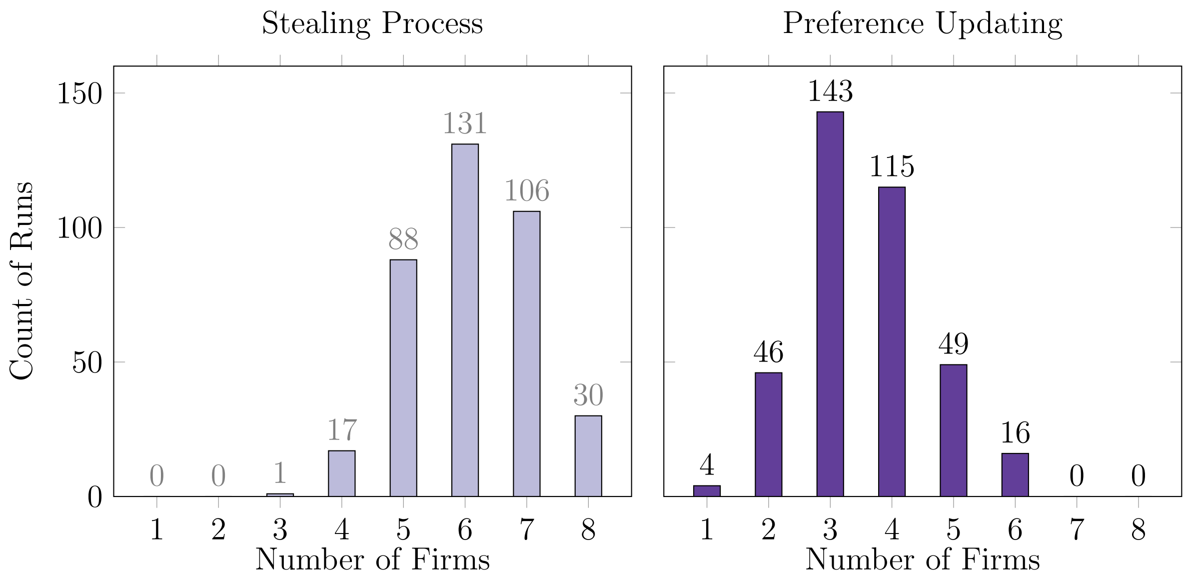

Figure 3: Number of Firms Acquiring Workers in Equilibrium

We see that the introduction of a preference for reputation via the Preference Updating process definitively reduces the average number of firms in the market: if we consider more firms to be better, the Stealing Process results have first-order stochastic dominance over the Preference Updating results. This aligns with expectations well. When workers have only their own idiosyncratic preferences and the similarly idiosyncratic opinions of those around them, they are more likely to follow their own hearts than to all agree on a single, best firm. In a way, this can be thought of as a form of horizontal product differentiation. Each worker has different overall rankings for each firm, and thus may choose differently than their fellows.

However, once we allow for preferences to depend on measures of prestige that make up reputation, in this case coworker quality and market share, worker preferences become more akin to vertical differentiation. Each worker adds $\psi_i(\tau_j w + \bar{h}_j)$ to their original value of $p_{i,j}$, and their preferences thereby become more aligned. Workers may differ in their willingness to pay for a certain firm—directly analogous to the salary bounds discussed above—but many more will agree on which firm is the best. This implies that more workers will choose the same firm within a given iteration of the Preference Generation process (with a temporarily fixed set of preferences) by the independence of choice from the reference matching, a fact which is established in **Theorem 2** above.

Stepping back from the model, this indicates to us that we may see a broader pattern in labor markets whereby industries in which workers have a preference for reputation tend to have larger, more entrenched incumbents. We present preliminary evidence for the existence of such a preference for reputation in an industry with few entrants—the world of large corporate law firms—in Section 8.3.

With an eye to the fact that our other parameters may greatly change the outcome, we can try the same exercise for different, randomized values of γ_i, ψ_i, and w; for each, we again consider 1000 runs.

7.3.3. Randomized γ

Now, we consider workers with initial conditions which are the same as in our basic comparison, with the exception that we assign each worker a random γ_i, i.e. a random weight on the parameter encompassing others' opinions in their utility. In the Stealing Process, this represents only the weight on others' preferences over firms. In Preference Updating, however, this both weights the opinions of others, and provides a multiplier on the prestige-based portions of the worker's preference, which will be a combination of the identical prestige indicators in the worker's own $p_{i,j}$ and in the $p_{w,j}$ of others.

First, in order to ensure a wide distribution of average γ_i across runs, we choose a random $\bar{\gamma} \sim \text{Unif}[0, 2]$. From here, we assign each worker a $\gamma_i \sim \mathcal{N}(\bar{\gamma}, 0.25)$. This allows for uniqueness among workers, while still enabling us to see the impact that the value of reputation on an industry level has on the number of firms. The results of this set of simulations can be found in figure 4.

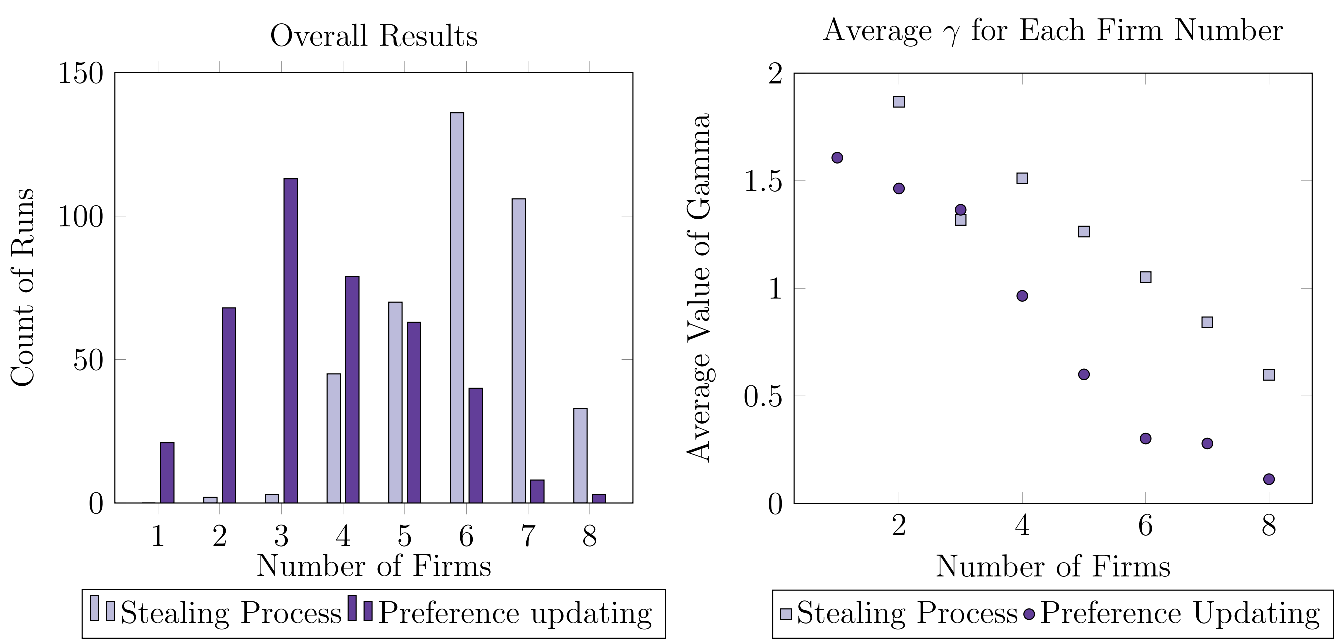

Figure 4: Number of Firms Acquiring Workers in Equilibrium, Variable γ

Again, we saw about two-fifths of our runs converge, for a total of 395 observations. As before, very little appears different between those runs that converged and those that did not.

Here, we see a reprisal of the same divergence between the Stealing Process and Preference Updating that we saw in the initial runs above. The market has more firms acquiring workers when workers care about the idiosyncratic preferences of others, and fewer firms when workers' preferences begin to align on concrete measures of prestige.

Interestingly, the distributions for both processes have flattened, with Preference Updating now occasionally supporting seven or even eight firms, and the basic Stealing Process supporting as few as two. We see indications for the reason why in the average γ for each number of firms. Not only does a lower weight on other's preferences increase the number of firms post-Preference Updating, it also does so for the Stealing Process, if at a lower intensity. The results from varying γ in the Stealing Process suggests that for very high values of γ, we begin to see a similar level of verticality to preferences as is present in Preference Updating; the common preference term, r_j, begins to outweigh the individual's preferences. This corroborates our earlier conjecture that the number of firms directly correlates to the level of horizontal differentiation among firms. This may indicate that we would expect to see more concentration in industries where workers can easily communicate their preferences to others, and where they value the preferences of those around them.

We also see that the value of γ needed to produce a certain level of market concentration in the Stealing Process equilibrium is overall about .5 higher than that for the Preference Updating equilibrium. This indicates that when workers care more about the prestige and reputation of the firm they are going to, they are more easily swayed by the opinions of others and are more likely to choose the firm which the general consensus finds most preferable. This leads to even fewer firms than we would expect if workers simply take their fellow employee's opinions into account. It may also indicate that we would expect to see more monopolization in industries where workers care about the brand of their firm and less where firms are interchangeable along this prestige dimension, further corroborating the results of the basic comparison above. This theory will be discussed in the context of law firms in Section 8.

7.3.4. Randomized ψ

Next, we vary values for ψ_i. This explicitly puts more weight on workers' preferences for a firm's prestige and reputation.

Again, we first randomize the mean value of ψ for each run, drawing a value for each run such that $\bar{\psi} \sim$ Unif$[0,1]$. Workers then have $\psi_i \sim \mathcal{N}(\bar{\psi}, 0.1)$. The results of these runs are displayed in figure 5.

Surprisingly, we find that close to two thirds of these runs converge; we thus have 628 observations. It appears that runs with higher ψ are more likely to converge, indicating that the below runs may be slightly more biased towards more concentrated markets than would be present had all runs converged. However, it is clear that the major trends present in the set of runs which converge are also present in the set of runs which do not.

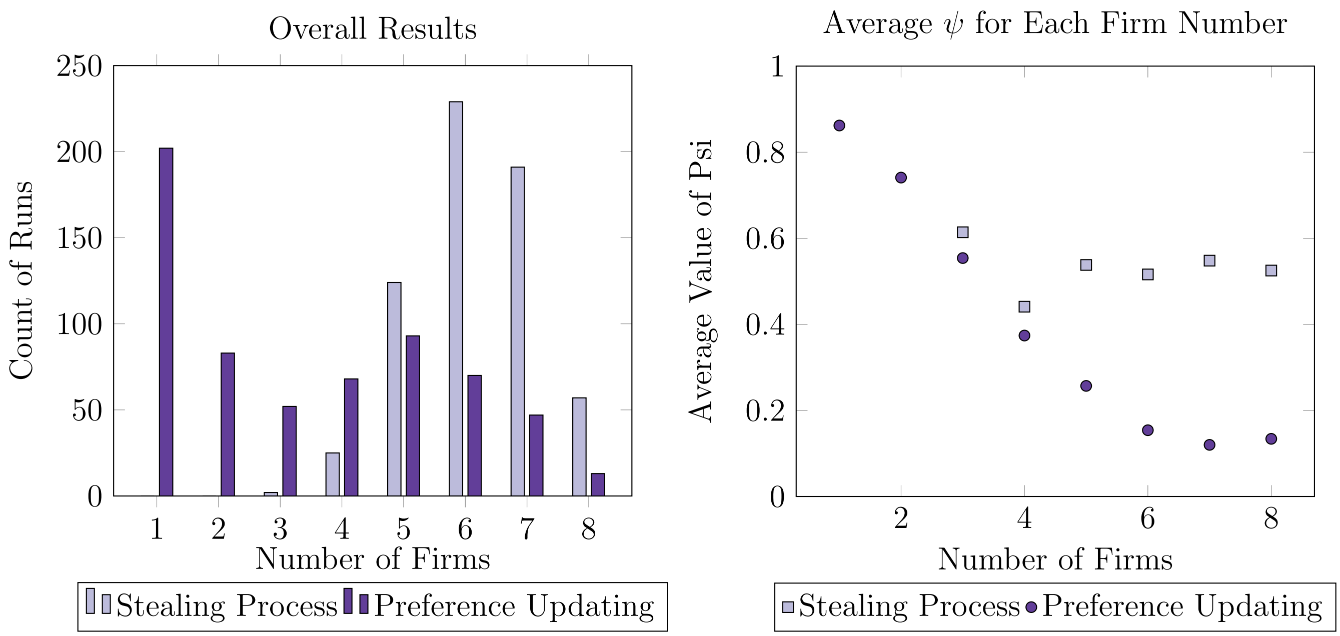

Figure 5: Number of Firms Acquiring Workers in Equilibrium, Variable ψ

First, we note that, as expected, there is no impact on the outcome of the original Stealing Process. Initial decisions of workers under basic preferences are independent of ψ, which we see reflected in the right graph of figure 5. The value of ψ hovers around .5 for all firm numbers, indicating that any firm number supported under the original runs is also supported with varying values of ψ.

This contrasts directly with the results for Preference Updating, again as expected. When workers give more weight to the vertical differentiation between firms, we see fewer firms supported in equilibrium. On the other hand, when ψ is very low, worker's idiosyncratic preferences dominate, and even as many as eight firms can find workers. Interestingly, randomized ψ gives us our first departure from a standard bell curve in the overall results of our simulations. It seems that high values of ψ, which lead to homogenous preferences, often result in monopoly, while low values may lead to a broader range of firms; note the similarity in average ψ across $6, 7$ and 8 firms.

As presented here, it is easy to see the continuation of the story in the data about horizontal and vertical differentiation, and the balance between objective and subjective measures of firm quality. When more objective measures exist, and it is easier for workers to see clearly which firms are higher in objective quality—say, offering better benefits, having better coworkers, or sending alumni on to higher positions—we may be more likely to see a more concentrated industry. However, when workers must rely on their subjective experiences, in interviews, with current employees, and otherwise, and firms are difficult to compare vertically, we may be more likely to see a higher number of firms. An avenue for further research will be to compare the concentration of industry with the availability of information that vertically differentiates firms and the importance workers in the industry place on these markers.

7.3.5. Randomized w

Finally, we consider variations of w, the parameter which controls the weight on market share.

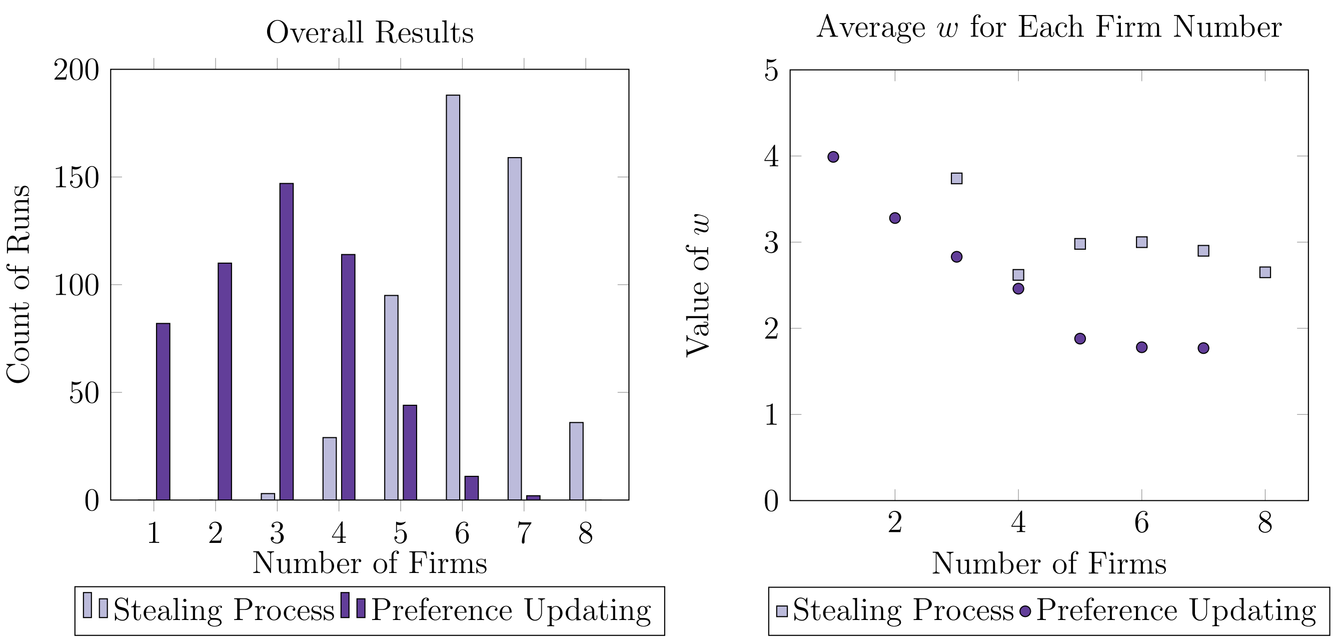

Figure 6: Number of Firms Acquiring Workers in Equilibrium, Variable w

Raising this parameter gives workers an explicit preference for larger firms. As w is common to all workers, we randomize by drawing w from a uniform distribution: $w \sim \text{Unif}(0, 5)$. Interestingly, as with varying ψ, more runs converge; and, again, this is biased toward a higher w. However, much like with ψ, the same trends are present in both the converged and non-converged runs. We have 510 converged runs, the results of which are displayed in figure 6.

First, note that, much like the results from the Stealing Process for the variation of ψ, the variation of w appears to have no impact prior to Preference Updating, as we would expect. The final result is a version of what we saw by varying ψ; as workers care more about going to a large company, they are more likely to choose a larger firm, leading to a consensus choice and fewer overall firms in the market. This effect is reduced, however, by the relatively low value of ψ in the above runs, fixed at .5. Again, we see a continuation of the story told by other simulations: when workers have more vertical preferences, often as a result of concern for prestige and reputation, we see more market concentration.

7.4. Entry and its Effects on Welfare

In determining the presence of barriers to entry in this market, we can consider the ability of a firm to enter after the completion of the Stealing Process, where workers care about their fellows' preferences, and after the completion of Preference Updating, where workers place an explicit premium on prestige and reputation. Below, we model a new firm attempting to enter a market with one or several incumbents. This model can then provide intuition into how we should think about entry in a world with reputation-concerned workers. Here, rather than being primarily concerned with side-level evaluations of matchings, we are concerned with the ability of a new firm to acquire workers.

What does it mean for a firm to enter in this context? We define two terms to help us distinguish between entry into our two different kinds of equilibria. First, we consider entry into a Stealing Process equilibrium. Because workers' base preferences remain unchanged throughout this process, we term it *Unbiased Entry*:

Definition 8 (Unbiased Entry)**.** Define μ_F to be an outcome of the Stealing Process for the set of firms F and workers W. *Unbiased Entry* into μ_F is the process by which the set $\{j\} \cup F$, where j is a firm such that $j \notin F$, is matched with workers via the Stealing Process, taking μ_F as the initial auxiliary matching.

We then consider a firm attempting to enter after workers' preferences have been shaped by the Preference Updating algorithm, which we term *Biased Entry*:

Definition 9 (Biased Entry)**.** Define μ_F to be an outcome of the Preference Updating process for the set of firms F and workers W, and $\mathbf{P}$ to be the equilibrium preference matrix for the workers W. *Biased Entry* into μ_F is the process by which the set of firms $\{j\} \cup F$ are then matched to the W workers, taking $\mathbf{P}$ as the workers' initial preferences and μ_F as the initial auxiliary matching.

This is different from dynamics typically studied in matching models, including those discussed in Section 7.2 above, which usually proceed from the removal of a firm at stable equilibrium, and typically begin with the original stable equilibrium being used both for worker priors and for the available contract set, as well. This more standard framework is dealt with above, and encompasses a form of 'entry' as the reverse of exit. However, here, we attempt to model the ability of an untested firm, competing against incumbents, to hire from some pool of workers with prior opinions, thereby perturbing our stable equilibrium.

First, the more general question: does the externality present in the basic Stealing Process prevent a firm from entering after the process has completed? The answer is a definitive no. We can show that any firm which can always acquire workers when it starts in the initial set of firms can also acquire workers during unbiased entry into a Stealing Process equilibrium. This is true for all utility functions such that the assumptions made in Section 3 hold, namely, that utility is substitutable and the assumption of consistency as defined in Section 3 is met, and which meet the conditions of **Theorem 2**.

Theorem 6. *Suppose we have a set of workers W and firms F which satisfy the assumptions made in Section 3.1, and that workers always prefer employment to unemployment. Let M be the set of all possible outcomes of the Stealing Process in this market. Consider any firm $j \in F$ such that $|C_j^\mu| > 0$ for all $\mu \in M$, where C_j^μ is the set of workers hired by firm j under matching μ. Let $\hat{\mu}$ be any outcome of the Stealing Process for the market defined by workers W and firms $F \setminus \{j\}$. Then j will be able to hire a non-empty set of workers following unbiased entry into $\hat{\mu}$.*

Proof. To begin, we show that unbiased entry as defined above still results in a stable equilibrium over the initial set of contracts within the first post-entry run of **S2**. This follows from the fact that worker choices are independent of the reference matching, by **Theorem 2**. When we re-run the algorithm with the added firm, workers face the same possible contracts X^W, which here includes firm j, in this first iteration of **S2** under either an auxiliary matching for the firm set $F \cup \{j\}$ or the μ from the previous run of the Stealing Process. Because worker choices (though not their utilities) are independent of the matching, and firm choices are entirely independent, any equilibrium that can be reached beginning with a constructed auxiliary matching can be reached beginning with the previous μ, and vice versa; the two processes will play out identically. Thus, the new process will converge, and because all equilibria of the Stealing Process with auxiliary matchings must be stable, all equilibria of the new process will be stable. This implies that the result of the first post-entry run of **S2** is stable for the initial set of contracts used.

Because any equilibrium reached in the initial iteration of **S2** under an auxiliary matching must also be possible under entry and vice versa, all equilibria of future rounds must be possible under entry and when all firms are present from the beginning. In future rounds, only the information about the current equilibrium and the current set of contracts will be carried forward, both of which are indistinguishable between the two initial conditions.

Thus, because the only difference between the outcome of the Stealing Process when j was present at the beginning of the process and the outcome when j entered at the end is the auxiliary matching used in the initial iteration of **S2** and choices are independent of the auxiliary matching, if all matchings resulting from the former case result in j acquiring a set of workers with cardinality greater than zero, then all matchings resulting from the latter will also see the firm acquiring a set of workers with cardinality greater than zero. □

Sections 7.1 and 7.3 have suggested that interdependencies among workers' preferences may affect salaries and market concentration, respectively; however, here we see that interdependencies alone, like those present in the Stealing Process, are not enough to result in a barrier to entry in this context. This motivates our explicit introduction of reputation into preferences. As such, we now consider the case of biased entry.

As with the case of any number of firms being technically supportable both with or without explicit reputation preferences, a sufficiently attractive firm will be able to enter into any given market. However, we can also show that the necessary conditions for a firm to be able to enter are more difficult to meet under biased entry, as opposed to unbiased entry.

First we must define a notion of the *necessary conditions for entry.* In order for a firm j to enter, it must necessarily be the case that there exists some worker i who would prefer to work for j for their highest marginal product salary at j over any other firm for their lowest marginal product salary at that other firm. If this is not the case, and every worker has at least one firm k they would prefer to work for at their lowest marginal product salary at k over their best contract with j, those firms will always be willing to pay the workers that salary, and thus no worker will ever go to j. In other words, j offering a maximal salary must be able to attract some worker if all other firms offer minimal salaries. If this is not the case, no worker will ever join j.

We define an equation to capture this effect as follows. First, define worker i's lowest marginal product salary at a firm j to be $\underline{s}_{i,j}$. We recall that the worker i's highest marginal product salary at firm j is given by $\bar{s}_{i,j}$. Then, in order for firm j to enter when preferences are described by $\mathbf{P}$, it is a necessary (but not sufficient) condition that there exists a worker i such that for all firms k, the following is positive:

$$d(j,i,k,\mathbf{P}) = (\bar{s}_{ij} - \underline{s}_{i,k}) + (p_{i,j} - p_{i,k}) + \gamma_i(r_j(\mu,\bar{r}) - r_k(\mu,\bar{r}))$$

Note that the difference between the two terms r_j and r_k is independent of the reference matching, as discussed in Section 5. However, the difference does depend on the underlying preferences of workers. Thus, we define the following function:

$$R(j,k,\mathbf{P}) = r_j(\mu,\bar{r}) - r_k(\mu,\bar{r}) = \sum_{w \in W} \frac{h_w}{H}(m_{w,j} - m_{w,k} + p_{w,j} - p_{w,k}).$$

With this background complete, we show that allowing reputation and prestige to play a role in preferences may lead to more difficult entry.

Theorem 7. *Consider a market with firms F and workers W, with utility functions defined as in Section 5. Let μ_S be a stable equilibrium of the Stealing Process with Section 5 utilities in this market, with preferences described by $\mathbf{P}_S$. Next, let μ_P be a stable outcome of the Preference Updating process, with updated preferences described by $\mathbf{P}_P$. Then a firm that meets the necessary conditions for biased entry into μ_P (entry with worker preferences described by $\mathbf{P}_P$) will meet the necessary conditions for entry into μ_S (entry with worker preferences described by $\mathbf{P}_S$), but the reverse may not hold.*

Proof. First, we note that in the case of unbiased entry, the necessary conditions for entry are fixed: there is no change throughout the algorithm in any component of $d(j,i,k,\mathbf{P})$. This corresponds to the results of **Theorem 6**, which shows that entry cannot be blocked in the Stealing Process.

In the case of biased entry, the firm must meet the necessary conditions for entry in the first iteration of Preference Updating in order to have a chance at gaining workers in equilibrium. This is so because if the firm fails to meet these conditions in the first iteration, it will then acquire no workers at any point in the Preference Updating process. Preferences toward the entrant will then remain unchanged after the first iteration, implying that there will be no change in any incumbent's contracts, and that the original equilibrium will then continue to be stable.

Therefore, it suffices to compare $d(j,i,k,\mathbf{P}_S)$ and $d(j,i,k,\mathbf{P}_P)$, the necessary conditions for unbiased and biased entry respectively, where the latter corresponds to the necessary conditions for entry on the first round of the Preference Updating process. We then want to show that $d(j,i,k,\mathbf{P}_S)$ is always higher than $d(j,i,k,\mathbf{P}_P)$, implying that some firms may be able to enter when facing idiosyncratic preferences, and yet be unable to enter when incumbents' reputation and prestige are preferred by workers.

Consider $d(j, i, k, \mathbf{P}_S)$, written out with preferences decomposed into idiosyncratic and common portions:

$$d(j, i, k, \mathbf{P}_S) = (\bar{s}_{ij} - \underline{s}_{i,k}) + (1 - \psi_i)(\alpha_{i,j} - \alpha_{i,k}) + \gamma_i R(j, k, \mathbf{P}_S).$$

Next, consider $d(j, i, k, \mathbf{P}_P)$:

$$d(j, i, k, \mathbf{P}_P) = (\bar{s}_{ij} - \underline{s}_{i,k}) - \psi_i(\tau_k w + \bar{h}_k) + (1 - \psi_i)(\alpha_{i,j} - \alpha_{i,k}) + \gamma_i R(j, k, \mathbf{P}_P).$$

We can then show that for all workers i and firms k, $d(j, i, k, \mathbf{P}_S) \geq d(j, i, k, \mathbf{P}_P)$. Compare the difference between the two:

$$d(j, i, k, \mathbf{P}_S) - d(j, i, k, \mathbf{P}_P) = \psi_i(\tau_k w + \bar{h}_k) + \gamma_i(R(j, k, \mathbf{P}_S) - R(j, k, \mathbf{P}_P)).$$

The latter two terms are:

$$R(j, k, \mathbf{P}_S) = \sum_{w \in W} \frac{h_w}{H}(m_{w,j} - m_{w,k} + (1 - \psi_i)(a_{w,j} - a_{w,k}))$$

$$R(j, k, \mathbf{P}_P) = \sum_{w \in W} \frac{h_w}{H}(m_{w,j} - m_{w,k} - \psi_w(\tau_k w + \bar{h}_k) + (1 - \psi_w)(a_{w,j} - a_{w,k})).$$

Taking the difference between these two, we can rewrite:

$$d(j, i, k, \mathbf{P}_S) - d(j, i, k, \mathbf{P}_P) = \psi_i(\tau_k w + \bar{h}_k) + \sum_{w \in W} \frac{h_w}{H}(\psi_w(\tau_k w + \bar{h}_k)).$$

As all terms are weakly positive, we have $d(j, i, k, \mathbf{P}_S) \geq d(j, i, k, \mathbf{P}_P)$. This implies that any firm which meets the necessary conditions for biased entry will meet them for unbiased entry, however, the reverse may not hold; it may be the case that $d(j, i, k, \mathbf{P}_S) > 0 > d(j, i, k, \mathbf{P}_P)$. Therefore, a firm which is able to enter when facing preferences $\mathbf{P}_S$ may fail to gain workers in biased entry facing preferences $\mathbf{P}_P$. □

In economic terms, this result indicates to us that we may indeed expect to see less entry in industries where firm prestige and reputation play a role in worker decisions. The implications of this barrier for market efficiency and antirust policy will be areas for future research. In the following numerical examples, we attempt to gain some intuition as to how this barrier to entry may impact different agents in a market.

Consider the case of two firms, denoted respectively as θ and λ, with linear production functions:

$$y_\theta(C) = \sum_{i \in \mathsf{W}(C)} h_i$$

$$y_\lambda(C) = 1.1 \left[\sum_{i \in \mathsf{W}(C)} h_i \right].$$

Thus, λ is more efficient than θ, as it can produce more given any coalition of workers than θ can with the same coalition. The two firms compete over four workers, a, b, c, and d, who have Section 5 utility functions and are characterized as follows:

Worker	h_i	γ_i	ψ_i	$\alpha_{i,\theta}$	$\alpha_{i,\lambda}$
a	4	1	0.5	0	5
b	3	1	0.5	5	0
c	2	1	0.5	5	0
d	1	1	0.5	5	0

Worker a is the most productive, and prefers the amenities at firm λ, while the other workers prefer the amenities at firm θ. When we run the Preference Updating process with both firms present, we have the following final contracts:

Worker	Firm	Salary	Final Utility
a	λ	4.4	8.31
b	θ	3	6.79
c	θ	2	5.79
d	θ	1	4.79

Each worker matches with the firm for which they have the highest original estimation of amenities, and each receives their marginal product in salaries. Note that both firms get workers in equilibrium.

We can next consider what happens when λ attempts to enter after θ has collected workers as a monopolist. Running the Preference Updating process for θ alone, all workers match to θ and receive their marginal product. We can then allow λ to enter, taking the preferences and reputations established in the θ-only run as initial conditions. We find new final matching:

Worker	Firm	Salary	Final Utility
a	θ	4	5.75
b	θ	3	7.25
c	θ	2	6.25
d	θ	1	5.25

We see that λ now fails to enter; the preferences established via the earlier run of the Preference Updating process outweigh a's high valuation of λ's amenities. Thus, Preference Updating may prove a barrier to entry for some initial conditions. This attempts to model what may happen in an economy over time; as certain businesses cement themselves in workers' minds as being good places to work, it may be more difficult for entrants to attract workers, unless those workers have very high preference for the entrant, or don't care much for reputation (say, $\gamma_i = 0$).

Given the analysis above, we can attempt to make a statement about the welfare effects of entry after Preference Updating. For the entrant, entry is beneficial, given that all equilibria of the Preference Updating process must be individually rational. For entrants that are more efficient than incumbents, their entry may benefit the overall productivity of the economy, with the natural caveat that we only model the labor market here. In the example above, for instance, we see that the sum of production of the workers increases when λ acquires workers, which may benefit consumers as opposed to the monopoly economic configuration of θ alone.

When we consider workers, however, the story becomes more complex. As we see in the above example, workers b, c, and d have lower utility when λ is present than when it fails to enter; even if it is able to enter at the end—for example, if $\alpha_{a,\lambda} = 7$, enough to overcome the reputation barrier—utilities for the other three workers fall. This suggests that workers who have a vested interest in the incumbent, for example if they like certain benefits it provides, or value stability, may actually lose utility when the entrant gains workers, thereby decreasing their own firm's relative success and reputation. That being said, it is also possible for an entrant to be both preferred by all workers in terms of base amenities and to be more productive, and still fail to enter, if these advantages do not beat the generated preferences for the incumbent; in that case, entry would benefit the workers. For example, we consider again θ and λ as defined above, but set workers' initial conditions as follows:

Worker	h_i	γ_i	ψ_i	$\alpha_{i,\theta}$	$\alpha_{i,\lambda}$
a	4	1	0.5	1	2
b	3	1	0.5	0	1
c	2	1	0.5	1	2
d	1	1	0.5	0	1

All workers prefer λ, which is also more productive, and will go to λ if offered the choice prior to Preference Updating; however, if λ attempts to enter after θ has already gained workers as a monopoly, it will fail, and θ will hold all workers.

Finally, we consider incumbent firms. Like the utility of workers, the profit of incumbents may be affected both positively and negatively by entry. At a surface level, firms may lose workers, profit, and productivity to the entrant. However, an incumbent may actually benefit from the entry of a new firm if that firm hurts their competitors more than it hurts them. For example, if a certain subset of high quality workers wishes to go to one incumbent over another, they may bring with them some other less-decisive workers, drawn to reputation, that the other incumbent would nonetheless prefer to have. However, if an entrant takes the high-quality workers, and the workers with weaker preference differentiation aren't attracted by the start up, they may choose to go to the other incumbent instead.

We can see this dynamic play out in a numerical example. Consider three identical firms θ, ζ, and λ, each with linear production functions. There are four workers in the economy, a, b, c, and d, with initial conditions:

Worker	h_i	γ_i	ψ_i	$\alpha_{i,\theta}$	$\alpha_{i,\zeta}$	$\alpha_{i,\lambda}$
a	3	1	0.5	4	0	10
b	2	1	0.5	2.8	3	0
c	2	1	0.5	0	4	0
d	2	1	0.5	0	4	0

If θ and ζ compete against each other, θ attracts both a and b, while ζ has c and d, in spite of b's slight a priori preference for ζ. However, if we allow λ to attempt to enter after Preference Updating, worker a will leave θ for λ, and, seeing this fall in market share and coworker quality for θ, b will instead move to ζ, who now gets all three workers b, c, and d.

We can take several lessons from this example. First, it is a testament to the limitations of the rules of substitutability in accurately modeling reality, and also highlights the flexibility of a model able to take some form of complementarity into account (though, naturally, we will always be concerned about the downside of non-convergence). Second, it highlights the ways in which the story of entry can be extremely complex. It may be the case that entry hurts all incumbents; but it also may be the case that entry by small firms that chip away largely at a competitor's business may be good news for certain incumbents.

Indeed, as shown above, entry can be both harmful and helpful for nearly every player in the market. Importantly, however, the reputational barrier to entry described above does not discriminate in blocking by the overall efficiency of the entrant, and this is unlikely to be the case in the real labor marketplace either. Future research may attempt to determine the extent to which this barrier is present in various industries, and how it impacts their development.

8. Empirical Application: Law Firms

Finally, we present evidence for a preference for reputation in a real-world industry, thereby motivating the assumptions of our model. In order to do so, we require both a credible measure of firm reputation, and a credible account of the quality of the workers at that firm or hired by that firm in a given period. For most industries, finding this information is difficult if not impossible without inside connections to several firms or schools. However, one industry which relies heavily on the reputation of both firms and workers and publicizes data about such is the legal industry; in particular, so-called "Big Law" firms serving corporations worldwide.

These firms serve big-name clients and recruit from highly selective schools, putting them in direct competition with one another to hire the top tier of students. As such, these firms tend to pay highly similar salaries, being locked in something of a stalemate with their opponents; when one firm raises salaries, the others must raise them to remain competitive (Holt, 2020; Moody, 2018). Furthermore, these firms are known for their grueling hours and poor work-life balance. In an industry with such homogeneity in salary and intensity, there can be little doubt that potential workers will choose their firm in part based on its reputation.

Indeed, informal interviews with a handful of law students at two highly-ranked law schools suggest that this is the case. Most discussed using reputation as an initial screener, looking at online rankings and discussing with peers in order to build an initial pool of firms, often the most highly ranked firms, and from there choosing based on fit. The students also suggested that reputation was extremely important for career movement following graduation, and that going to a prestigious firm offers an easy stepping stone to future career success. Therefore, in analyzing the relationship between firm reputation and worker quality, law firms and their employees provide an excellent dataset from which to work, not only in terms of data availability, but also in terms of relationship isolation.

In addition to being an industry with a high value for reputation, large corporate law firms tend to be old, and entry into the space is low. Recent analysis has shown that the average age of an S&P 500 company has fallen below 20 years (Sheetz, 2017), while the firms of the Vault Law 100, a ranking of the prestige of large law firms which we will use in our analysis below, have an average age of 117 years.[12] In addition, no firm on the Vault list was formed in the last 20 years, and only 35% were formed in the last century. Determining the existence of an explicit relationship between entry and the value placed on reputation by workers in an industry will be an important place for future work; the analysis in this paper, which shows that there is a link between reputation and worker quality in law, a low-entry industry, begins that work.

Finally, it is important to note that this empirical application is extremely limited in scope, and serves only to provide stylized evidence for the theoretical claims made above. One area for further research is the extent of the importance of reputation in other industries; another is the empirical validity of the connection between entry and the importance of reputation, as indicated by the low entry in the legal industry and by the theoretical work in Section 7.4. We will further discuss the connection between the entry and reputation in this setting below.

8.1. Data and Limitations

To analyze the relationship between firm prestige and worker quality, we build a novel dataset of firm reputations and worker education levels and tenure with the firm.

The top law firms are ranked annually by several companies, including Vault, a firm that provides industry insights to job seekers. This includes ranking companies within industries by their attractiveness. The Vault Law 100 is particularly interesting. To create this ranking, a survey is distributed annually to American law firms, which is then filled out by more than 15,000 associates (in some years, such as 2021, the number can reach 20,000). The associates are asked to rank firms based explicitly on how prestigious it is to work for the firm in question, omitting their own firm and any other firm with which they are unfamiliar. The average of the associates' rankings of a given firm then becomes the firm's score. Firms are then ranked in order of score. The Vault system

[12] This average was constructed using the oldest age of any company that forms a portion of a current company; several firms in the Vault 100 are composed of the mergers of other large law firms.

also includes headquarters, size and salaries for the latest year of their survey.

This ranking in particular is useful for two reasons. First and most importantly, it directly measures reputation as we have defined it, as community sentiment about prestige and success. Second, the rankings lead reality by approximately six months; thus, any worker hired by a firm in 2020 is likely to have already seen the 2020 ranking of that firm. This means that we can make a natural assumption that the present year's ranking is most salient for workers hiring in.

Given this credible measure of reputation, we must then ask how it relates to the quality of workers hired by a firm in a given year. Fortunately, law firms are uniquely transparent about the education and backgrounds of their attorneys. Thus, utilizing data from company websites and LinkedIn, we can find the education, experience and hire date of a large subset of these employees. By comparing this data to the firm reputation scores, we can begin to see the correlation between worker quality and firm reputation.

For each worker in each of five firms, we collect from the firm's website the worker's job title (partner, counsel, or associate), whether or not they had a previous judicial clerkship, their undergraduate degree school, and their highest degree school. If the worker has multiple higher degrees, we collect the highest law degree; for example, if the worker has a PhD in chemistry, a JD, and an LLM, we collect the school at which they acquired their LLM. If the worker has multiple LLMs, we take either the most recent, or, if only one was done in the United States, the domestic (and therefore ranked) LLM.

We then pair this with LinkedIn information about their graduation year and the year they joined the firm. Finally, we add the current 2021 ranking of the law school the employee attended, sourced from US News and World Report. While there is some debate over the usefulness of these rankings, they do provide a rough estimate of the quality of students, including a mix of average undergraduate GPA, LSAT and GRE scores, and placement success, all of which very directly reflect upon the quality of the average candidate from that school (Morse et al., 2021). Though there are other rankings which may be more accurate and helpful for potential students, for our purposes, US News and World Report also provides one of the few complete rankings over all US law schools, enabling us to compare schools in a quantitative way.

Using this data, we can create a comparison of Vault scores with the average school quality or average clerkship likelihood for each firm and year. While the field of law provides uniquely useful data for this analysis, there are also several obstacles we must consider when drawing conclusions from the data.

One possible issue is that we have limited information on the rankings of schools in past years, and are forced to utilize the most recent rankings rather than those that correspond to the year the student was hired. This may be problematic if there are large differences across years in the rankings of individual law schools. This may bias the results if law school rankings have tended to reverse in the past ten years; for example, if the highly ranked law schools of 2010 are now lower and the low high, or vice versa. This would pose a problem if a firm always draws from the same tier of school, but the firm's rankings have moved with those of the schools they have drawn from. However, this concern seems unlikely to impact our results, as we can show trend matching between firm and school ranking across firms with very different trends in Vault ranking.

A second and more intractable problem is that we have access only to those workers who are still at the firm. Not only does this mean that worker numbers become extremely sparse before 2010, limiting usefulness, it also implies that there may be bias in the school quality of who stays at the firm. For example, if a firm recruits a set of mostly high quality workers in a low-rank year, but most of the high-quality workers change jobs, leaving only the lower quality workers, our data would assume that the firm hired only low quality employees in that year. Thus, we are concerned with a systematic relationship between the ranking in the hiring year, school quality of the individual worker, and the likelihood of leaving the firm. Though the firms display different ranking trends, this relationship may still exist if any worker who comes from a more or less highly ranked school than the firm's reputation in that year has earned tends to leave. For example, if such workers are more likely to feel "out of place" at the firm, it may result in a bias in the data.

Another issue to consider is the fact that data shows only who has been hired, not the pool of workers that the firms have to choose from. While it may be a natural assumption that firms choose the highest quality workers from their pool, which we would assume would tend to relate to school quality, this may not necessarily be the case. For example, the firm's hiring team may be more likely to choose lawyers who come from their own alma mater, which may not conform to a strict quality based hiring scheme.

With these issues in mind, we can still note the marked trends in this comparison, and draw some stylized conclusions. More detailed data, likely from non-public sources, will be instrumental in investigating these claims more thoroughly in the future.

8.2. Summary Statistics and By-Firm Breakdown

We begin by analyzing our data at the firm level. We have data for five large law firms: Boies Schiller Flexner, Cadwalader, Proskauer, Quinn Emanuel, and Shearman & Sterling, across the years from 2009 to 2020. These five firms were chosen in order to get a wide range of patterns in Vault rankings: for example, while Quinn Emanuel rises in ranking consistently over the years surveyed, Shearman & Sterling falls reliably.

Table 1: Summary Statistics by Firm

Variable	Boies	Cadwalader	Proskauer	Quinn	Shearman
Total Lawyer Count	203	377	799	818	735
Vault Rank	21.58	40.67	35.83	17.58	29.58
	(9.26)	(9.13)	(6.87)	(6.26)	(8.61)
Average School Rank	24.55	28.94	30.81	16.19	21.56
	(19.92)	(11.30)	(7.58)	(2.25)	(9.96)
% Hires Former Clerks	51.6%	6.8%	9.5 %	32.7 %	4.7%
	(18.9%)	(9.4%)	(6.0 %)	(12.2%)	(3.8%)
Average At Hire Experience	6.63	5.69	6.19	5.65	4.65
	(8.69)	(7.83)	(8.21)	(7.59)	(7.02)
Worker Tenure	9.87	7.95	7.88	7.58	8.63
	(7.33)	(8.25)	(7.81)	(5.98)	(8.24)

Notes. Standard deviations are provided in parentheses. Vault Ranking corresponds to the firm's average ranking in the Vault Law 100 from the years 2009-2020. Average School Rank corresponds to the average of the average rankings of the alma maters of the lawyers hired in a given year from 2009-2020. % Hires Former Clerks represents the average percentage of the workers hired in that year who held a clerkship prior to being hired from 2009-2020. Average At Hire Experience represents the average years since graduation from law school of a worker hired in a given year from 2009-2020. Worker Tenure is an average over all workers, including all years of hire, and represents how long workers have been with the firm.

We first look at some summary statistics from the five firms, presented in table 1. We have a broad range of sizes, from Boies, our smallest, to Quinn Emanuel, more than four times its size. As we would expect, with smaller groups of lawyers hired annually, the two smallest firms Boies and Cadwalader have larger standard deviations for measures of school rank. Interestingly, however, they also have somewhat higher standard deviations in Vault ranking, something that is harder to explain; an interesting avenue for further research is an investigation of the relationship between reputation volatility and firm size. Overall, the five firms chosen provide a good range of average Vault rankings, ranging from Quinn's high average to Cadwalader's low.

A few other pieces of information are worth noting. First, firms differ widely in their propensity to hire former clerks. Boies averages 50% former clerks in a class of new hires, while Quinn, though hiring a lower percentage of clerks, hires a far larger absolute number, given larger recruitment. These two firms are also the firms with the best average Vault ranking, suggesting that lawyers capable of getting clerkships may also be more drawn to prestigious firms. This relationship will be discussed in more detail in the Section 8.3.

Second, we note that average experience differs relatively little among firms, though Shearman does seem to recruit more young workers directly out of law school than the other four. Indeed, workers recruited with no years of experience at Shearman make up more than 30% of the total body of employees, while no other firm broaches 27%. The number is as low as 13% for Boies.

Third, in addition to having the highest average experience level for new hires, Boies also has by far the longest tenure for employees. This may be due to any number of factors, from firm culture to investment in a smaller class of recruited workers, and remains a matter for another paper.

Next, we discuss the time trends of Vault ranking and average school ranking of each firm in turn below. Strikingly, not only do the trends in the rankings of the firm and the schools it hires from mirror each other, the raw numbers are tightly related as well.

8.2.1. Shearman & Sterling

We begin with Shearman & Sterling, a firm which provides some of the most clear evidence for the relationship between firm reputation and worker quality, then continue with the other four firms in alphabetical order.

Shearman was founded in 1873, and is based in New York. They have 24 offices across four continents, and 736 total partners, counsels and associates, of which partners make up 28%. A great many of these lawyers attended only international schools, so we use data from 332 employees for this firm's school ranking.

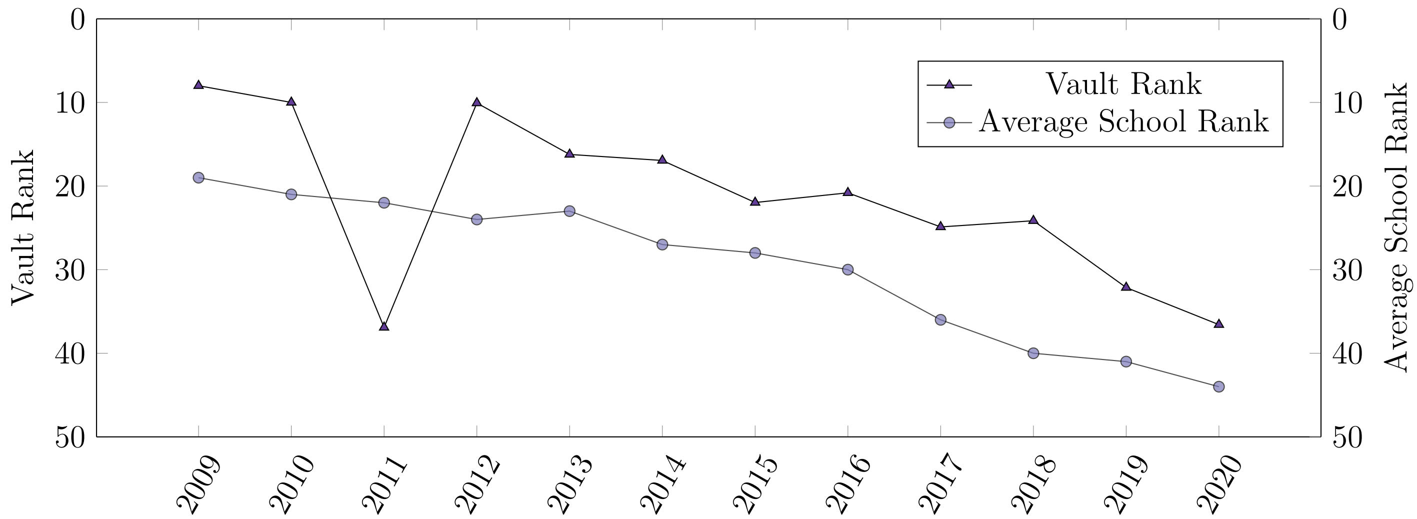

Figure 7: Shearman & Sterling, Vault Rank and School Rank by Year

As we see in figure 7, Shearman has been falling steadily in both ranking and hire quality since 2009. With the exception of 2011, the school and Vault rankings fall together through the period studied. The progression of this trend is remarkably similar in school ranking and in Vault ranking, a relationship which will continue to be apparent as we look at other firms.

8.2.2. Boies Schiller Flexner

Boies Schiller Flexner, headquartered in New York, is by far the youngest firm in our group of five. Founded in 1997, Boies has only 203 partners, counsels and associates, of which partners make up 49%, a far larger percentage than any other firm in the set. Boies has 13 total offices, only one of which is outside of the United States, in London. However, in part due to the long average tenure of employees, only 91 of their lawyers were hired after 2009, educated in the United States, and had useable data for our purposes.

Though the graph presented in figure 8 is more noisy than for other firms with more workers, Boies presents an interesting case study because it is the only one of our five firms to have a distinctly U-shaped distribution of Vault rankings. Though the school quality data jumps up and down—likely because some year buckets have as few as four observations—it appears to dance around this same U-shape. Thus, Boies presents an interesting alternative to the flat lines or simple slopes of other firms' rankings.

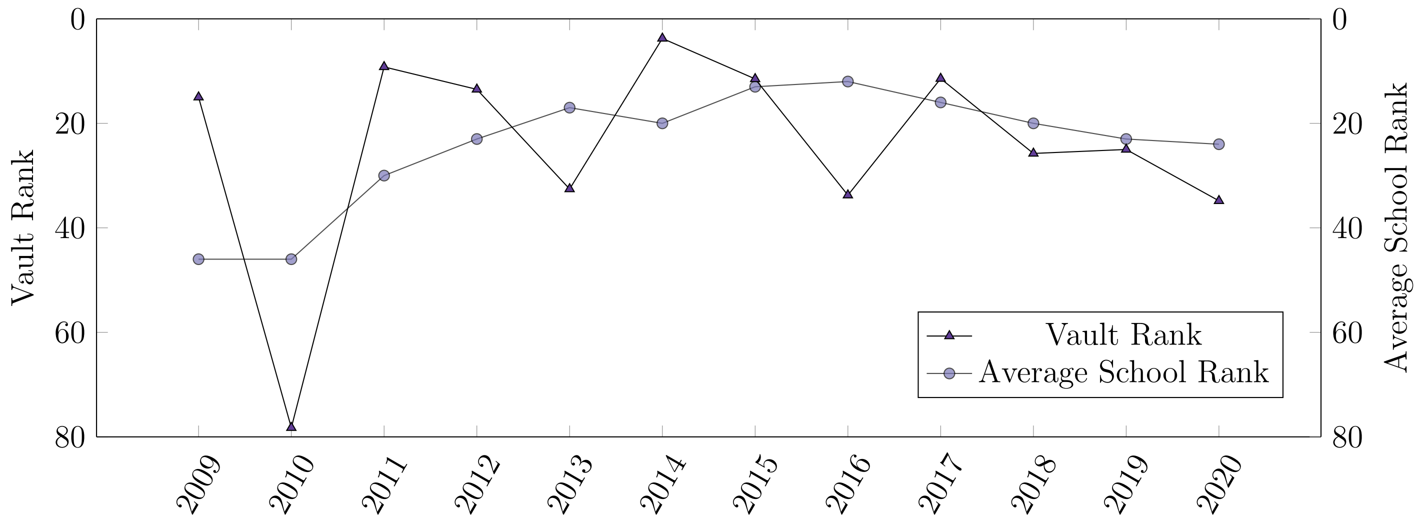

Figure 8: Boies Schiller Flexner, Vault Rank and School Rank by Year

8.2.3. Cadwalader

While Boies Schiller Flexner is the youngest firm in our data set, Cadwalader is the oldest. Founded in New York in 1792, the firm is the oldest Wall Street firm still in existence, according to Vault. It has 377 partners, counsels and associates, of which 33% are partners. Its employees are spread across four total offices, three in the United States and one in London. Of the 377 total lawyers, 188 had usable data.

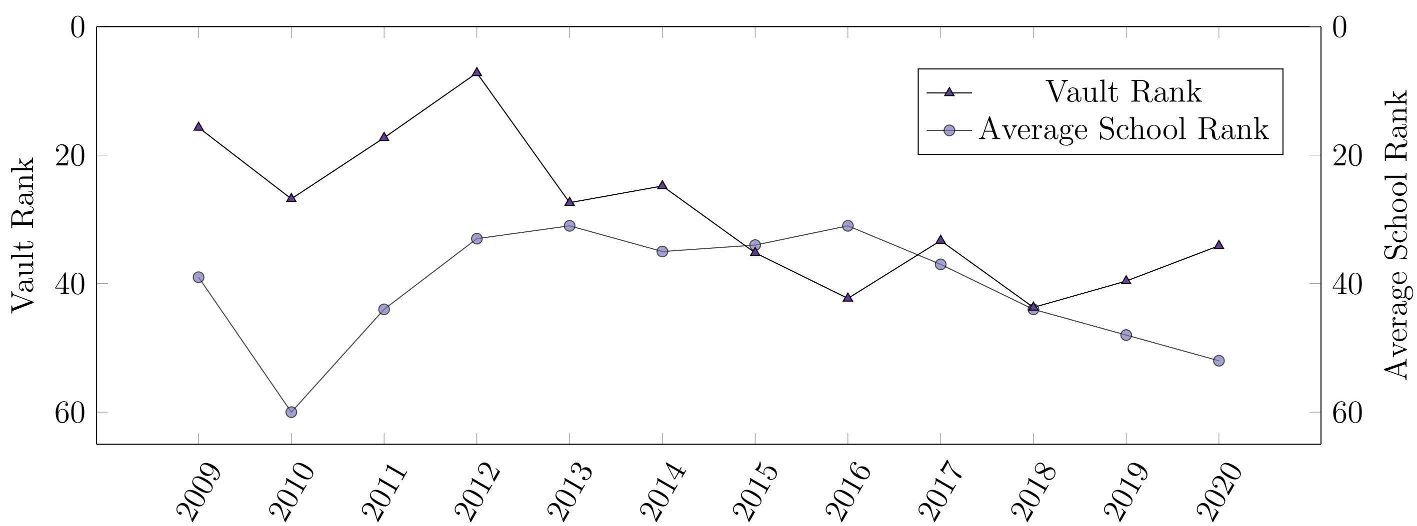

Figure 9: Cadwalader, Vault Rank and School Rank by Year

Among our firms, Cadwalader averages the lowest on the Vault rankings, and only slightly outperforms Proskauer in overall average school ranking. Figure 9 plots the year-over-year change in each ranking for Cadwalader.

Cadwalader's pattern of rankings appears mostly flat, with one interesting exception: a brief fall in Vault rank of some 20 points from 2009 to 2010, with a rise of another 20 points the next year. While this does correspond to a sharp drop in the quality of hires, we also see that it is not until 2014 that the school rankings become as closely knit to the Vault rankings as they are in other firms.

8.2.4. Proskauer

We next consider Proskauer, founded in New York in 1875. The firm has eight domestic offices and five international offices in Asia, Europe and South America. Proskauer employed 799 partners, counsels and associates at the time of data collection, of whom 39% are partners. A total of 421 employees provided usable data.

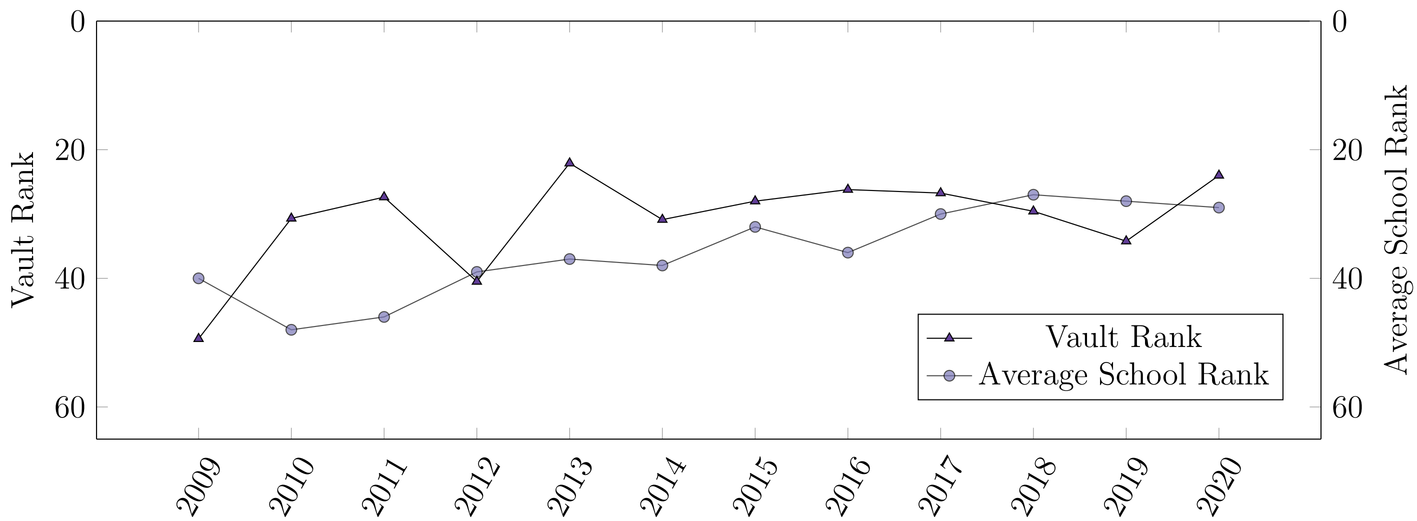

Figure 10: Proskauer, Vault Rank and School Rank by Year

As we can see in figure 11, Proskauer provides us with an example of a firm which slopes generally upward (toward lower, better values) in both Vault ranking and average school quality.

8.2.5. Quinn Emanuel

Quinn Emanuel Urquhart & Sullivan, a relatively new law firm founded in 1986, is based in Los Angeles. They have 23 offices spread around the world, including spaces in London, Tokyo

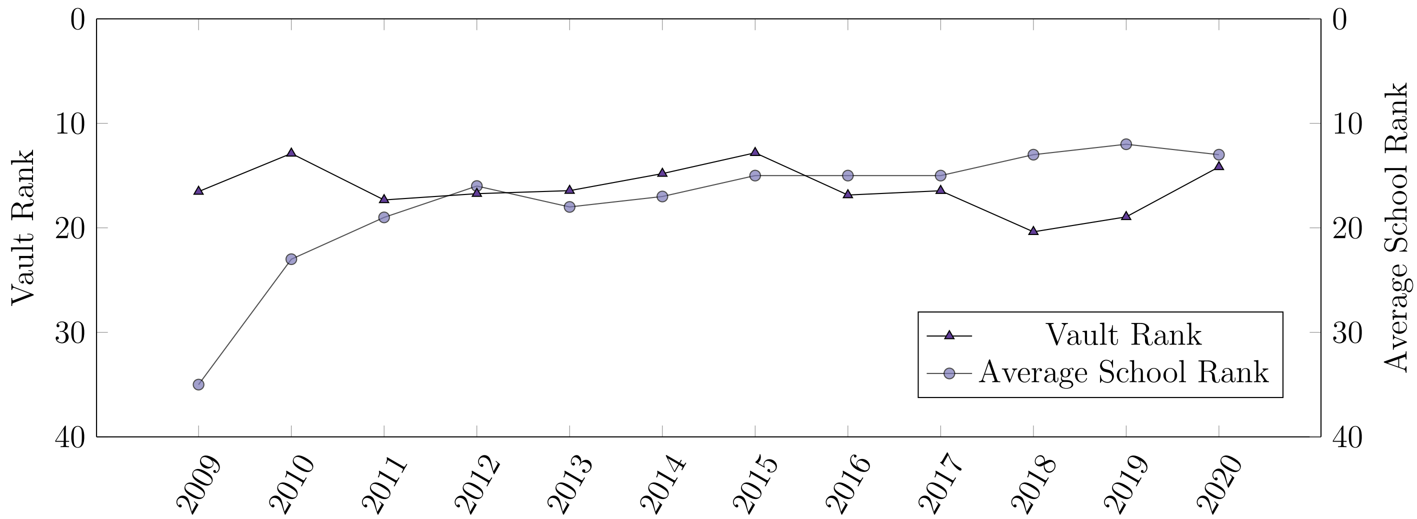

Figure 11: Quinn Emanuel, Vault Rank and School Rank by Year

and Paris. The firm is fairly large, boasting 818 total partners, counsels and associates as of data collection, of which 34% are partners. 443 of these total lawyers attended American law schools, were hired after 2009, and had enough public data to be used in our school quality estimates.

Quinn is the most stable of the firms in our dataset. Plotting firm ranking beside school quality, the data from Quinn Emanuel paints a picture of a new firm that rose quickly through the ranks of large law institutions before settling in the early 2010's into a slower drift upward. This is reflected in the data on school quality as well. In this respect, Quinn Emanuel is one of the firms with the least change over the period in question: as is visible in figure 11, it has the lowest standard deviation in school rank among the five firms.[13]

8.3. Results

In order to better understand how reputation impacts worker choices, we look at differences in worker quality both across firms and within a given firm over time. We first consider evidence of a connection between reputation and worker quality across different firms. In figure 12, we plot the firm Vault rank versus the average school rank for each firm-year pair.[14]

[13]However, we should note that these five firms were chosen explicitly for their high variability in Vault rankings; many firms in the Vault top 15 persist in nearly the same spot over more than a decade. Thus, Quinn is not so much notable for its stability as the other four are for their variability.

[14]We omit a single outlier firm-year: 2010 for Boies. This year has only four observations in our dataset, resulting in an average school rank of 78.25. It is omitted here solely for clarity of visualization.

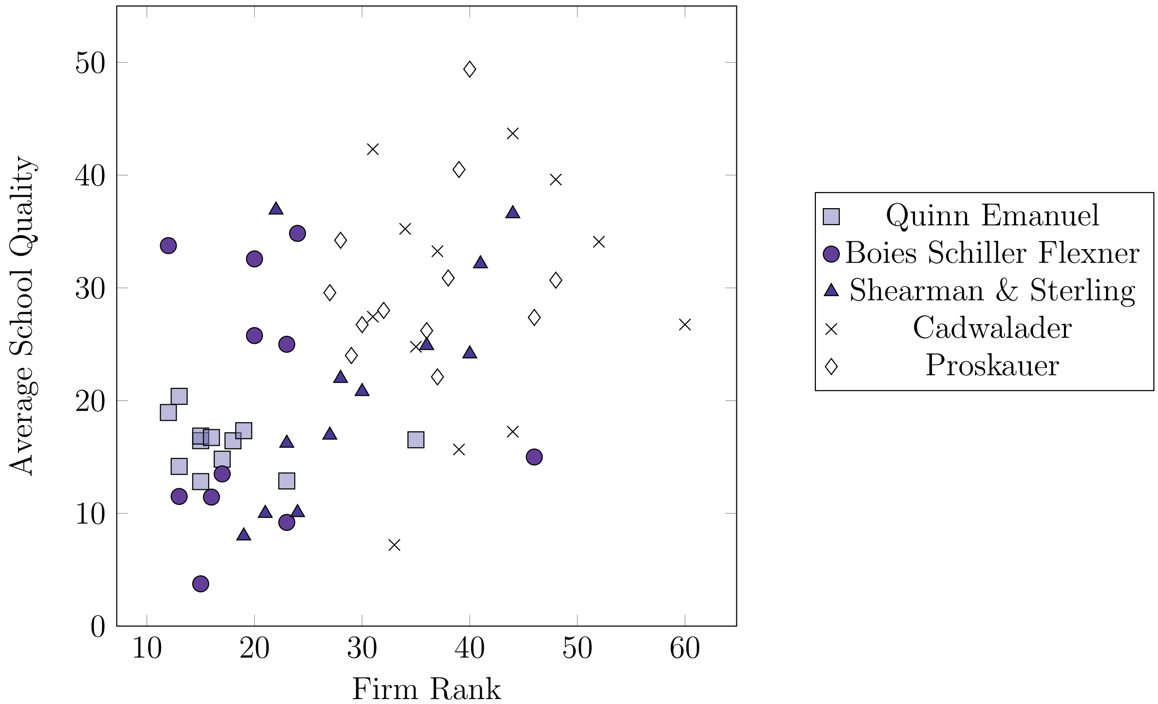

Figure 12: Vault Rank vs. Average School Rank, All Firms

In this graph, we see further confirmation of the general trends present in the individual firm data; firms with higher reputation in a given year tend to hire higher quality workers that year, as measured by school quality, an effect that is apparent both across and within firms. We can also consider the raw correlation between Vault rank and school rank, as presented in table 2, along with the correlations between firms and rank.

Table 2: Correlations Between School Rank, Vault Rank and Firm

	Vault	Av. School	Boies	Cadwalader	Proskauer	Quinn	Shearman
Vault	1.000	0.419	−0.322	0.501	0.292	−0.495	0.023
Av. School	0.419	1.00	0.005	0.183	0.258	−0.331	−0.115

Notes. The unit of observation is the firm-year. “Av. School” denotes the average rankings of the alma maters of the lawyers hired in a given year from 2009-2020. Correlation between a firm and a rank is the correlation between the rank and the indicator variable representing a firm.

We note that there is a moderate to high correlation between a firm’s Vault rank and school rank, a numerical representation of the strong, if messy, trend we see in figure 12. We also see that correlations tend to move in the same direction for both Vault rank and school rank within firms, though both Boies and Shearman remain nearly flat in school rank and Vault rank, respectively.

The correlation between Vault rank and school rank indicates that the relationship we saw in the graphs of each firm's rank extends to a cross-firm relation, showing that firms which are more highly ranked on average also hire more highly ranked workers on average. Indeed, a simple regression of school rank on Vault rank achieves significance at the 1% level, as we can see in table 3.

Table 3: Regression of School Rank on Vault Rank

	(1) Average School Rank
Vault Rank	0.448^{***}
	(0.127)
Constant	11.38^{**}
	(3.984)
N	60

Notes. Standard errors are given in parentheses. The unit of observation is the firm-year. Average School Rank is the average rank of the law school from which workers hired in a year received their highest degree.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We now turn to within-firm effects of reputation on worker quality. In order to formalize our intuition from Section 8.2, which indicated that a firm's reputation and hire quality will fall and rise together, we consider a set of simple regressions with measures of worker quality as the dependent variable. We estimate several equations of the from:

$$\text{Worker Quality}_{f,y} = \beta_1(\text{Vault Rank}_{f,y}) + \beta_2(\text{Firm FE}_f) + \beta_3(\text{Controls}_{f,y}).$$

Our three measures of worker quality to be tested are the average alma mater ranking of the workers hired by a given firm in a given year, the number of workers hired in a given year that are former clerks, and the average experience years of the hired workers. We then regress worker quality on the firm's Vault rank, a firm fixed effect, and controls made up of the other two potential dependent variables as well as the number of workers hired in a given year.

The results are presented in table 4.

Table 4: The Effects of Firm Vault Ranking on Measures of Worker Quality

	(1) Sch. Rank	(2) Sch. Rank	(3) Sch. Rank	(4) Sch. Rank	(5) Clerks	(6) Av. Exp.
Vault Rank	0.425^{**}	0.416^{**}	0.425^{**}	0.416^{**}	0.00543	0.0778
	(0.191)	(0.198)	(0.193)	(0.200)	(0.0624)	(0.0515)
Num.	0.113	0.114	0.110	0.111	0.143^{***}	−0.0105
	(0.0783)	(0.0793)	(0.103)	(0.104)	(0.0245)	(0.0266)
Av. Exp.		0.106		0.106	−0.0116	
		(0.548)		(0.553)	(0.166)	
Clerks			0.0231	0.0239		−0.00822
			(0.463)	(0.467)		(0.118)
Sch. Rank					0.00215	0.00680
					(0.0420)	(0.0354)
Firm Fixed Effects	YES	YES	YES	YES	YES	YES
N	60	60	60	60	60	60

Notes. Standard errors are in parentheses. The unit of observation is the firm-year. Num. is the number of total workers in our dataset hired by a given firm in a given year. Sch. Rank is the average rank of the law school from which workers hired in a year received their highest degree. Av. Exp. is the average number of years since a worker hired in a given year graduated from law school. Clerks is the number of workers hired in a year with a former clerkship.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

In each of our first four regressions, we have a positive, nearly identical coefficient on Vault rank as a predictor of school rank, indicating that as firms lose prestige, they tend to hire from less highly-ranked schools. These coefficients are also quite significant, with $p < 0.05$, confirming what we intuitively see in the graphs in Section 8.2. This indicates that workers may be attracted to prestige, providing support for the assumptions of our model; moreover, it indicates that higher ranked firms have access to a larger set of workers, likely due to more potential candidates being attracted by their reputation. In reality, this effect seems to imply that the most highly ranked firms are able to take the highest quality employees from the pool of those interested in large law firms.

Interestingly, however, we also see positive coefficients on the Vault rank variable in regressions of both Clerks and Average Experience, possibly signaling that when firms fall in prestige, they may choose to hire workers with more experience or with a clerkship under their belt, even when they may come from less prestigious universities. This may indicate a tradeoff between the three measures of quality; it might also indicate that those workers most recently out of college or who attended a prestigious law school are precisely those most concerned with reputation, leaving less prestigious firms to look for other signals of quality in workers. That being said, these coefficients are not significant, and both are very near zero, so more research will be required to establish a definitive connection or tradeoff.

One additional point on clerkships is worth discussion. The relationship shown here between firm rank and clerkships, indicating that firms hire more clerks when they have lower rank, stands in marked contrast to the summary data in table 1, which had our two firms with highest average rank recruiting a higher percentage of clerks. This indicates that while clerks may on average go to higher ranked firms, for a given firm, a drop in ranking may lead to a larger number of clerks hired. However, as noted above, the relationship here is near zero and non-significant.

Though the results are mixed across different measures of worker quality, within school rank as a measure of employee desirability, the evidence indicates a connection between worker quality and firm rank. In such a prestige-driven industry as law, this is to be expected. Law also happens to be an industry with very little firm entry. Establishing an empirical connection between these two facts, however, will require more work. Future avenues for empirical research may include combining survey data weighing workers' value for prestige and reputation, and comparing this with the rate of entry in a given industry. This more comprehensive work may allow us to determine whether such preferences for prestige truly do translate into a barrier to entry.

9. Concluding Discussion

Understanding barriers to entry is necessary as we seek to establish fair markets that promote competition. Broadening this definition to include prestige-based barriers can allow us to have a more realistic view of how our markets operate and come to better decisions in antitrust prosecution. Our analysis of the two theoretical models of preferences presented above provide a foundation for how a firm's reputation may serve as a barrier to entry when workers place value on prestige, and a unique dataset derived from large law firms provides initial evidence for such preferences. We show that when workers place value on others' preferences, we see an implicit tuition in salaries.

Furthermore, when workers value reputation and prestige, we are likely to see more concentrated markets and to see less entry by new firms. Future research can help us to determine what industries feel the effects of this barrier the most, and how reputation impacts entry in practice.

While much more research is needed before conclusions can be drawn about the extent of repetitional barriers' effect on certain markets, investigation of entry has never been more important as we see a sustained decrease in market competition across many sectors (Philippon, 2019). Thus, considering whether and how reputation has contributed to the slowdown in competition in the United States is a worthwhile project as we work to build an economy that is both efficient and fair.

References

U.S. and Plaintiff States v. Google LLC. URL https://www.justice.gov/atr/case/us-and-plaintiff-states-v-google-llc.

U.S. V. Microsoft, June 2001. URL https://www.justice.gov/atr/complaint-us-v-microsoft-corp.

Everyones business reputation tracker. Technical Report 3, Opinium, Aug. 2019. URL https://workplaceinsight.net/wp-content/uploads/2019/09/CBI_report_vFinal.pdf.

N. Agarwal. An empirical model of the medical match. *American Economic Review*, 105(7):1939–1978, July 2015.

J. Alcalde and P. Revilla. Researching with whom? Stability and manipulation. *Journal of Mathematical Economics*, 40(8):869–887, Dec. 2004.

O. Aygun and T. Sonmez. Matching with contracts: comment. *American Economic Review*, 103 (5):2050–2051, Aug. 2013.

J. S. Bain. *Barriers to new competition, their character and consequences in manufacturing industries*. Cambridge: Harvard University Press, 1956.

K. Bando. Many-to-one matching markets with externalities among firms. *Journal of Mathematical Economics*, 48(1):14–20, Jan. 2012.

K. Bando. A modified deferred acceptance algorithm for many-to-one matching markets with externalities among firms. *Journal of Mathematical Economics*, 52:173–181, May 2014.

G. S. Becker. A theory of marriage. In T. W. Schultz, N. B. of Economic Research, and P. Council, editors, *Economics of the family: marriage, children, and human capital*. Published for the National Bureau of Economic Research by the University of Chicago Press, Chicago, 1974.

O. Bedre-Defolie and G. Biglaiser. Contracts as a barrier to entry in markets with nonpivotal buyers. *The American Economic Review*, 107(7):2041–2071, July 2017.

W. Burgess. A bad reputation costs a company at least 10% more per hire. *Harvard Business Review*, Mar. 2016. URL https://hbr.org/2016/03/a-bad-reputation-costs-company-at-least-10-more-per-hire.

W. Chen. Cross-country income differences revisited: accounting for the role of intangible capital. *Review of Income and Wealth*, 64(3):626–648, Sept. 2018.

C. Corrado, C. Hulten, and D. Sichel. Intangible capital and U.S. economic growth. *Review of Income and Wealth*, 55(3):661–685, Sept. 2009.

C. A. Corrado and C. R. Hulten. How do you measure a "technological revolution"? *The American Economic Review*, 100(2):99–104, 2010.

B. Dutta and J. Masso. Stability of matchings when individuals have preferences over colleagues. *Journal of Economic Theory*, 75(2):464–475, Aug. 1997.

F. Echenique and M. B. Yenmez. A solution to matching with preferences over colleagues. *Games and Economic Behavior*, 59(1):46–71, Apr. 2007.

J. C. Fisher and I. E. Hafalir. Matching with aggregate externalities. *Mathematical Social Sciences*, 81:1–7, May 2016.

D. Gale and L. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, Jan. 1962.

G. Grullon, Y. Larkin, and R. Michaely. Are U.S. industries becoming more concentrated? Technical report, Social Science Research Network, Rochester, NY, Oct. 2018.

J. W. Hatfield and S. D. Kominers. Multilateral matching. *Journal of Economic Theory*, 156: 175–206, Mar. 2015.

J. W. Hatfield and P. R. Milgrom. Matching with contracts. *American Economic Review*, 95(4): 913–935, Sept. 2005.

J. Holt. Biglaw salary scale with bonuses: 1999 - 2020, Feb. 2020. URL https://www.legalreader.com/biglaw-salary-scale-with-bonuses-1999-2020/.

E. Ide, J.-P. Montero, and N. Figueroa. Discounts as a barrier to entry. *The American Economic Review*, 106(7):1849–1877, July 2016.

S. Jaffe and S. Morris. Efficient location choice and the returns to agglomeration. 2017.

A. S. Kelso and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.

T. C. Koopmans and M. Beckmann. Assignment problems and the location of economic activities. *Econometrica*, 25(1):53, Jan. 1957.

R. P. McAfee, H. M. Mialon, and M. A. Williams. What is a barrier to entry? *The American Economic Review*, 94(2):461–465, May 2004.

M. Moody. BigLaw to $190K- Who? What? When? Where? Why?", June 2018. URL https://www.vault.com/blogs/vaults-law-blog-legal-careers-and-industry-news/biglaw-to-190k-who-what-when-where-why.

R. Morse, A. Castonguay, and J. Vega-Rodriguez. Methodology: 2021 best law schools rankings, Mar. 2021. URL https://www.usnews.com/education/best-graduate-schools/articles/law-schools-methodology.

T. Philippon. The economics and politics of market concentration. *NBER Reporter*, (4), Dec. 2019.

M. Pycia. Stability and preference alignment in matching and coalition formation. *Econometrica*, 80(1):323–362, 2012.

M. Pycia and M. B. Yenmez. Matching with externalities. 2017.

M. Rostek and N. Yoder. Matching With complementary contracts. *Econometrica*, 88(5):1793–1827, 2020.

A. E. Roth and J. H. Vande Vate. Random paths to stability in two-sided matching. *Econometrica*, 58(6):1475–1480, Nov. 1990. URL https://web.stanford.edu/~alroth/papers/1990_E_Random_Paths.pdf.

H. Sasaki and M. Toda. Two-sided matching problems with externalities. *Journal of Economic Theory*, 70(1):93–108, July 1996.

L. Shapley and M. Shubik. On the core of an economic system with externalities. *The American Economic Review*, 59(4):678–684, 1969.

L. Shapley and M. Shubik. The assignment game I: the core. *International journal of game theory*, 1(2):111–130, 1972.

M. Sheetz. Technology killing off corporate America: Average life span of companies under 20 years. *CNBC*, Aug. 2017. URL https://www.cnbc.com/2017/08/24/technology-killing-off-corporations-average-lifespan-of-company-under-20-years.html.

10. Appendix A: Proofs in Pycia and Yenmez (2017)

Below we present for convenience a selection of proofs first given in Pycia and Yenmez (2017) which are utilized above, modified slightly for this context. If the proof has a labeled counterpart in the paper, it is given beside the result number.

10.1. Proof of the Convergence of S1

Result 1 (Unlabeled)**. S1** results in a matching $\hat{\mu}$ such that $\hat{\mu} \succeq^W \Omega(X|\hat{\mu})$ for a given set of contracts X.

We know that because $\mu_t = \Omega(X|\mu_{t-1})$, by the irrelevance of rejected contracts, we have $\Omega(\mu_t|\mu_{t-1}) = \mu_t$ for every $t \geq 1$. We then use induction to find our result.

For the base case where $t = 1$, we consider $\mu_1 = \Omega(X|\emptyset)$. Here, the consistency of the ordering $\succeq^W$ and the fact that $X \supseteq \emptyset$ imply that:

$$\mu_1 = \Omega(X|\emptyset) \succeq^W \Omega(\emptyset|\emptyset) = \mu_0.$$

In the general case, if we assume that $\mu_t \succeq^W \mu_{t-1}$ and know that $X \supseteq \mu_t$, we know by consistency that:

$$\mu_{t+1} = \Omega(X|\mu_t) \succeq^W \Omega(\mu_t|\mu_{t-1}) = \mu_t.$$

Thus, the sequence $\{\mu_t\}_{t\geq 1}$ is monotone with respect to preorder $\succeq^W$. Because the set X is finite, there exists some iterations n and $m \geq n$ such that $\mu_{m+1} = \mu_n$; if not, it would imply that there are infinite contracts. Choosing the minimum m such that this property is satisfied, we call $\hat{\mu} = \mu_m$. It then follows that, because μ_m comes after μ_n in the sequence:

$$\Omega(X|\mu_m) = \mu_{m+1} = \mu_n \preceq^W \mu_m.$$

Thus, **S1** succeeds in finding such a matching.

10.2. Proof of the Convergence and Stability of S2

In order to reproduce the proof for this stage, we need to introduce some of the machinery used in Pycia and Yenmez (2017) which is omitted above. Call the current set of contracts X.

First, define the function f to be the following:

$$f(X^W, X^F, \mu) = (X \setminus R^F(X^F), X \setminus R^W(X^W|\mu), \Omega(X^W|\mu)).$$

This function represents the output of each step of the algorithm in **S2**. Two properties of this function, namely the monotonicity of the function and the stability of the fixed points of the function, will be required to show that **S2** reaches a stable equilibrium. First, we show monotonicity.

Result 2 (Lemma 2)**.** Suppose that the choice functions of workers and firms satisfy substitutability. Then the function f is monotone increasing with respect to the preorder $\sqsubseteq$ defined as follows:

$$(X^W, X^F, \mu) \sqsubseteq (\tilde{X}^W, \tilde{X}^F, \tilde{\mu}) \iff X^W \subseteq \tilde{X}^W, X^F \supseteq \tilde{X}^F, \mu \preceq^W \tilde{\mu}.$$

Proof. We know that function f is monotonic under $\sqsubseteq$ because for any $X^W \subseteq \tilde{X}^W, X^F \supseteq \tilde{X}^F, \mu \preceq^W \tilde{\mu}$, substitutability of choice functions implies that:

$$X \setminus R^F(X^F) \subseteq X \setminus R^F(\tilde{X}^F)$$

$$X \setminus R^W(X^W|\mu) \supseteq X \setminus R^W(\tilde{X}^W|\tilde{\mu}).$$

Furthermore, consistency implies that:

$$\Omega(X^W|\mu) \preceq \Omega(\tilde{X}^W|\tilde{\mu}).$$

Thus, if $(X^W, X^F, \mu) \sqsubseteq (\tilde{X}^W, \tilde{X}^F, \tilde{\mu})$, then:

$$f(X^W, X^F, \mu) = (X \setminus R^F(X^F), X \setminus R^W(X^W|\mu), \Omega(X^W, \mu)) \sqsubseteq$$

$$(X \setminus R^F(\tilde{X}^F), X \setminus R^W(\tilde{X}^W|\tilde{\mu}), \Omega(\tilde{X}^W, \tilde{\mu})) = f(\tilde{X}^W, \tilde{X}^F, \tilde{\mu}).$$

Thus, if one input to f is ranked higher by $\sqsubseteq$ than another, the output will be as well, and the function is monotone. $\square$

Result 3 (Lemma 3)**.** Let (X^W, X^F, μ) be a fixed point of function f. Then $X^W \cup X^F = X$ and:

$$\mu = X^F \cap X^W = \Omega(X^W|\mu) = \Gamma(X^F).$$

Proof. Because (X^W, X^F, μ) is a fixed point of f, it must be the case that $X^W = X \setminus R^F(X^F)$ and $X^F = X \setminus R^W(X^W|\mu)$. Thus, we know that:

$$X^W \cup X^F = X^W \cup [X \setminus R^W(X^W|\mu)] \supseteq X^W \cup [X \setminus X^W] = X.$$

Therefore, $X^W \cup X^F = X$.

By the same logic, we can see that:

$$X^W \cap X^F = X^W \cap [X \setminus R^W(X^W|\mu)] = X^W \setminus R^W(X^W|\mu) = \Omega(X^W|\mu).$$

This implies that $X^W \cap X^F = \Omega(X^W|\mu)$, and analogously that $X^W \cap X^F = \Gamma(X^F)$. Again by stability, we know that $\mu = \Omega(X^W|\mu)$, so:

$$\mu = X^F \cap X^W = \Omega(X^W|\mu) = \Gamma(X^F).$$

Thus we have our result. □

In order to prove that **S2** terminates in a stable outcome, we first need to show the following:

Result 4 (Theorem 2)**.** Suppose that the choice functions satisfy substitutability and the irrelevance of rejected contracts. Then a matching μ is stable if and only if there exists sets of contracts X^W, X^F such that (X^W, X^F, μ) is a fixed point of function f.

Proof. Consider (X^W, X^F, μ) as a fixed point of the function f. The proof of the above will then proceed in three parts.

First, we show that when choice functions satisfy substitutability and the irrelevance of rejected contracts, then the matching μ, which is reached at a fixed point of f, is stable.

Suppose for a contradiction that μ is not stable. This leaves us with three possibilities.

1. If matching μ is not individually rational for some worker i, then we have $u_i(\mu(i)|\mu) < u_i(\emptyset|\mu)$. Since (X^W, X^F, μ) is a fixed point of function f, we know that $\Omega(X^W|\mu) = \mu$ and $X^W \supseteq \mu$. However, substitutability and $u_i(\mu(i)|\mu) < u_i(\emptyset|\mu)$ imply that a contract in $\mu(i) = x \in X^W$ will be rejected in favor of unemployment by worker i. Therefore, $x \notin \Omega(X^W|\mu)$, contradicting the fact that (X^W, X^F, μ) is a fixed point of f.

2. If matching μ is not individually rational for some firm j, then we have $\pi_j(C_j^\mu) < \pi_j(\emptyset)$, implying that there is at least one worker the firm currently employs under μ that it wishes to reject. Again since (X^W, X^F, μ) is a fixed point of function f, we know that $\Gamma(X^F) = \mu$ and $X^F \supseteq \mu$. Analogously to the worker's case, substitutability then implies that the firm j rejects a contract from X^F that was in μ, a contradiction.

3. Finally, if there exists a blocking coalition, there is a group of a firm and some workers with contract set $\bar{X} \subset X$ such that workers and firms weakly prefer their matches under $\bar{X}$ to those under μ, with at least one agent strictly preferring this new match.

Because (X^W, X^F, μ) is a fixed point of function f, we know by **Result 3** (Lemma 3) that $X^W \cup X^F = X$. Then without loss of generality we can assume that $\bar{X} \in X^W$. Suppose, again without loss of generality due to the transferability of utility, that there exists a worker i who strictly prefers this match, with contract $x \in \bar{X}$ such that $u_i(x|\mu) > u_i(\mu(i)|\mu)$, and thus that $x \in \omega_i(\mu \cup \{x\}|\mu) \setminus \mu$. Since again (X^W, X^F, μ) is a fixed point of function f, $\Omega(X^W, \mu) = \mu$ by **Result 3** (Lemma 3), implying that $\omega_i(X^W|\mu) = \mu(i)$. By the irrelevance of rejected contracts, for any set of contracts Y such that $X^W \supseteq Y \supseteq \mu$, we must have $\omega_i(Y|\mu) = \mu$. But then, for $Y = \mu \cup \{x\}$, we have $\omega_i(\mu \cup \{x\}|\mu) = \mu(i)$, which contradicts $x \in \omega_i(\mu \cup \{x\}|\mu) \setminus \mu$.

Therefore, when choice functions satisfy substitutability and the irrelevance of rejected contracts, then the matching μ, which is reached at a fixed point of f, is stable. We now need to show the other direction: that if a matching μ is stable, then there exists sets of contracts X^W, X^F such that (X^W, X^F, μ) is a fixed point of f. We now move to the second part of the proof.

Second, we consider a piece of machinery that will help us show that if a matching μ is stable, then there exists sets of contracts X^W, X^F such that (X^W, X^F, μ) is a fixed point of f. Suppose again that choice functions satisfy substitutability and the irrelevance of rejected contracts. Then define the functions:

$$M^W(\mu) \equiv \max\{\bar{X} \subseteq X | \Omega(\bar{X}|\mu) = \mu\}$$

$$M^F(\mu) \equiv \max\{\bar{X} \subseteq X | \Gamma(\bar{X}) = \mu\}.$$

The maximum here is with respect to set inclusion. These functions describe the maximal set such that the choice function of the side given μ as a reference matching outputs μ. We claim that these functions are well-defined, and that for any contract $x \notin M^W(\mu)$, we have $x \in \Omega(M^W(\mu) \cup x|\mu)$, and likewise for firms. We will prove this to be the case for workers; the proof for firms is directly analogous.

First, we will show that it is well-defined. For this to be the case, it must be true that any single input will always produce the same output; we can demonstrate this by showing that the maximal set is unique. Consider two sets, M' and M'', such that $\Omega(M'|\mu) = \Omega(M''|\mu) = \mu$. By substitutability, we have:

$$\Omega(M' \cup M''|\mu) = (M' \cup M'') \setminus R^W(M' \cup M''|\mu) =$$

$$\left[M' \setminus R^W(M' \cup M''|\mu)\right] \cup \left[M'' \setminus R^W(M' \cup M''|\mu)\right] \subseteq \left[M' \setminus R^W(M'|\mu)\right] \cup \left[M'' \setminus R^W(M''|\mu)\right] = \mu.$$

The set $\Omega(M' \cup M''|\mu)$ cannot be a proper subset of μ, as the irrelevance of rejected contracts would then imply that $\Omega(M'|\mu) = \Omega(M''|\mu) = \Omega(M' \cup M''|\mu)$, a contradiction with our original definitions. Thus, the maximal set must be unique, as we can take the union of any two sets M', M'' with $\Omega(M'|\mu) = \Omega(M''|\mu) = \mu$ to form a weakly larger set with $\Omega(M' \cup M''|\mu) = \mu$.

Next, we show that for any contract $x \notin M^W(\mu)$, we have $x \in \Omega(M^W(\mu) \cup x|\mu)$. Let $x \notin M = M^W(\mu)$. If $x \notin \Omega(M \cup \{x\}|\mu)$, then $\Omega(M \cup \{x\}|\mu) = \Omega(M|\mu)$ by the irrelevance of rejected contracts. Then, however, $\Omega(M \cup \{x\}|\mu) = \mu$, contradicting the maximality of M. Thus, we must have $x \in \Omega(M \cup \{x\}|\mu)$.

Third, we show that if the matching μ is stable and the choice functions satisfy substitutability and the irrelevance of rejected contracts, then there exists sets of contracts X^W and X^F such that (X^W, X^F, μ) is a fixed point of f.

By the second piece above, there exists a maximal set $M^W(\mu) \equiv \max\{\bar{X} \in X|\Omega(\bar{X}|\mu) = \mu\}$. Let $X^W \equiv M^W(\mu)$ and $X^F \equiv X \setminus R^W(X^W|\mu)$. This means that by definition, $X^F = X \setminus R^W(X^W|\mu)$ and $\mu = \Omega(X^W|\mu)$. Therefore, we see that $X^W \cap X^F = X^W \cap (X \setminus R^W(X^W|\mu)) = \Omega(X^W|\mu) = \mu$. Now we need to show that $\mu = \Gamma(X^F)$, and that $X^W = X \setminus R^F(X^F)$.

First, we show that $\mu = \Gamma(X^F)$. Note that $X^F = X \setminus R^W(X^W|\mu) = (X \setminus X^W) \cup \Omega(X^W, \mu) = (X \setminus X^W) \cup \mu$. This implies that $X^F \supseteq \mu$. If $\Gamma(X^F) = Y \neq \mu$, there are two cases, both of which imply that μ is not stable, a contradiction.

1. If $Y \subsetneq \mu$, then the irrelevance of rejected contracts implies that $\Gamma(\mu) = Y$, and thus that some firms would reject some contracts from μ, implying that μ is not individually rational and thus not stable.

2. If $Y \not\subseteq \mu$, there exists some contract $y \in Y \setminus \mu$ and $y \in \Gamma(\mu \cup \{y\})$ by substitutability since $y \in \Gamma(X^F)$ and $X^F \supseteq \mu \cup \{y\}$. Because $y \in (X \setminus X^W) \cup \mu$ and $y \notin \mu$, we know that $y \notin X^W$. However, by the second claim above, this means that $y \in \Omega(X^W \cup \{y\}|\mu)$. Then $\{y\}$ blocks μ, a contradiction.

Thus, the only possible case is that $\mu = \Gamma(X^F)$.

Finally, we show that $X^W = X \setminus R^F(X^F)$. Because $\Gamma(X^F) = \mu$, we must have $X \setminus R^F(X^F) = X \setminus (X^F \setminus \mu) = X \setminus (((X \setminus X^W) \cup \mu) \setminus \mu) = X \setminus (X \setminus X^W) = X^W$. This shows the result, and we thus have the existence of sets of contracts X^W and X^F such that (X^W, X^F, μ) is a fixed point of f. □

We now use **Result** 4 to show that **S2** reaches the desired stable outcome.

Result 5 (Theorem 1)**.** Suppose that the choice functions of workers and firms satisfy substitutability and the irrelevance of rejected contracts. Then **S2** terminates, the outcome is stable for the current set of contracts, and if the algorithm stops at round T, then:

$$\mu(T) = X^W(T) \cap X^F(T)$$

Proof. We need to show that **S2** converges and that the resulting matching is stable. We recall that the preorder $\sqsubseteq$ is defined as :

$$(X^W, X^F, \mu) \sqsubseteq (\tilde{X}^W, \tilde{X}^F, \tilde{\mu}) \iff X^W \subseteq \tilde{X}^W, X^F \supseteq \tilde{X}^F, \mu \preceq^W \tilde{\mu}$$

Let $\hat{\mu}$ be the outcome of **S1**. We see that $f(X, \emptyset, \hat{\mu}) \sqsubseteq (X, \emptyset, \hat{\mu})$, as $\Omega(X|\hat{\mu}) \preceq^W \hat{\mu}$ by construction. By **Result 2**, the function f is monotone increasing, so we can repeatedly apply it to the left side of the inequality above to get $f^t(X, \emptyset, \hat{\mu}) \sqsubseteq f^{t-1}(X, \emptyset, \hat{\mu})$ for every t.

We next consider two cases.

1. Suppose first that the sequence converges, and thus that there exists some t such that:

   $$f^t(X, \emptyset, \hat{\mu}) \sqsubseteq f^{t-1}(X, \emptyset, \hat{\mu})$$

   This means that $f^{t-1}(X, \emptyset, \hat{\mu})$ is a fixed point of the function f. Let $(\bar{X}^W, \bar{X}^F, \bar{\mu}) = f^{t-1}(X, \emptyset, \hat{\mu})$. By **Result 3**, $\bar{X}^W \cap \bar{X}^F = \bar{\mu}$, and $\bar{\mu}$ is a stable matching by **Result 4**.

2. If on the other hand the sequence does not converge, because the number of contracts are finite, there exists a subsequence such that:

   $$f^n(X, \emptyset, \hat{\mu}) \sqsupseteq f^{n+1}(X, \emptyset, \hat{\mu}) \sqsupseteq \ldots \sqsupseteq f^m(X, \emptyset, \hat{\mu}) \sqsupseteq f^{m+1}(X, \emptyset, \hat{\mu}) = f^n(X, \emptyset, \hat{\mu})$$

   By the transitivity of preorder $\sqsubseteq$, we have:

   $$f^n(X, \emptyset, \hat{\mu}) = f^{m+1}(X, \emptyset, \hat{\mu}) \sqsupseteq f^m(X, \emptyset, \hat{\mu}) \sqsupseteq f^n(X, \emptyset, \hat{\mu})$$

   Let $f^n(X, \emptyset, \hat{\mu}) = (X_1^W, X_1^F, \mu_1)$, and let $f^m(X, \emptyset, \hat{\mu}) = (X_2^W, X_2^F, \mu_2)$. By the definition of $\sqsubseteq$, it must be the case that $X_1^W = X_2^W$, $X_1^F = X_2^F$, and $\mu_1 \sim^W \mu_2$. By construction, we have $\Omega(X_2^W, \mu_2) = \mu_1$. By substitutability, we then have that $\Omega(X_2^W, \mu_2) = \Omega(X_1^W, \mu_1) = \mu_1$.

   Furthermore, by substitutability, $X \setminus R^F(X_2^F) = X \setminus R^F(X_1^F)$, and by construction $X \setminus R^F(X_2^F) = X_1^F$, which together imply that $X \setminus R^F(X_1^F) = X_1^F$. Similarly, $X \setminus R^W(X_2^W|\mu_2) = X \setminus R^W(X_1^W|\mu_1)$, and by construction $X \setminus R^W(X_2^W|\mu_2) = X_1^F$, which together imply that $X \setminus R^W(X_1^W|\mu) = X_1^W$. Thus, (X_1^W, X_1^F, μ_1) is actually a fixed point of f, and the sequence converges, returning us to the previous case.

□

We will now show the results used above in Section 7.2, namely that within a fixed set of contracts, workers are made worse off by the closure of a firm, and firms made better off by the retirement of a worker.

First, we repeat the definition of *vacancy chain dynamics*. When an agent leaves the market for labor, their loss may cause cascading vacancies in firms that lead to a shifting in the equilibrium for many more workers; this dynamic recontracting is our vacancy chain dynamic. We will prove the above result for workers; the proof for firms is directly analogous.

Assume that the choice function $\Omega(X|\mu)$ satisfies substitutability and the irrelevance of rejected contracts. We define the choice function $\hat{\Omega}(X|\mu)$ to be the choices made by the workforce after some worker i has retired, a choice function which satisfies substitutability and the irrelevance of rejected contracts under the corresponding preorder $\hat{\succeq}^W$. In math, we can define this choice function using the following. Let worker i's new choice function be $\hat{\omega}_i(X|\mu) = \emptyset \ \forall X, \mu \subseteq X$. Then define $\hat{\Omega}(X|\mu)$:

$$\hat{\Omega}(X|\mu) \equiv \hat{\omega}_i(X|\mu) \cup \bigcup_{w \in M \setminus \{i\}} \omega_w(X|\mu).$$

A consistent preorder for such a choice function would be one that is the same as the preorder defined in Section 3.1, save that worker i is not included.

We can then modify the function f discussed above in order to model this phenomenon. For any $X^W, X^F, \mu \subseteq X$:

$$\hat{f}(X^W, X^F, \mu) \equiv (X \setminus R^F(X^F), X \setminus \hat{R}^W(X^W|\mu), \hat{\Omega}(X^W|\mu))$$

We let $(X^W(0), X^F(0), \mu(0))$ be the initial matching that is stable while worker i is still present in the market. After worker i retires or otherwise leaves the labor force, agents start recontacting dynamically from this starting point, which we call vacancy chain dynamics:

$$(X^W(t), X^F(t), \mu(t)) = \hat{f}(X^W(t-1), X^F(t-1), \mu(t-1)).$$

As $\hat{f}$ fits both substitutability and the irrelevance of rejected contracts, $\hat{f}$ is monotonic.

Before we begin the proof, we present one more definition. A choice function $\hat{\Omega}(X^W|\mu)$ exhibits *weaker externalities* than a choice function $\Omega(X^W|\mu)$ if $\hat{\Omega}(X^W|\mu) \hat{\succeq}^W \Omega(X^W|\mu)$ for any $\mu, X^W \subseteq X$. Note that this is satisfied for the example in which a worker retires, as the choice sets will be the same for all non-i workers in both cases, and the preorder considers only non-i workers, making the two equivalent.

Result 6 (Theorem 6)**.** Suppose that choice function $\hat{\Omega}(X^W|\mu)$ exhibits *weaker externalities* than choice function $\Omega(X^W|\mu)$. Let (X^W, X^F) be sets of stable contracts under Ω, Γ, with stable matching $X^W \cap X^F = \mu$. Then the vacancy chain dynamic starting at (X^W, X^F, μ) converges to $(\bar{X}^W, \bar{X}^F, \bar{\mu})$, where $\bar{\mu}$ is a stable matching under $\hat{\Omega}, \Gamma$, and we have $\bar{\mu} \hat{\succeq}^W \mu$, and $\mu \succeq^F \bar{\mu}$.

Proof. Because X^W, X^F is a stable set of contracts for the original choice functions, by **Result** 4, X^W, X^F is a fixed point of the function f. By **Result** 3, $\mu = \Omega(X^W|\mu) = \Gamma(X^F)$. Because we know that a worker rejects more contracts than before with all other choice functions held constant, we know that $\hat{R}^W(X^W|\mu) \supseteq R^W(X^W|\mu)$, and thus that $X \setminus \hat{R}^W(X^W|\mu) \subseteq X \setminus R^W(X^W|\mu)$. We also know that firm rejection functions are unchanged. By weaker externalities, we also have $\hat{\Omega}(X^W|\mu) \hat{\succeq}^W \Omega(X^W|\mu) = \mu$, which in this case is a relationship of equivalence. Thus, we know that:

$$(X^W, X^F, \mu) = f(X^W, X^F, \mu) \hat{\sqsubseteq} \hat{f}(X^W, X^F, \mu)$$

Because $\hat{f}$ is monotone, we know that $\hat{f}^{t-1}(X^W, X^F, \mu) \hat{\sqsubseteq} \hat{f}^t(X^W, X^F, \mu)$ for all $t \geq 1$. Since the number of contracts is finite, there exists a t such that $\hat{f}^{t-1}((X^W, X^F, \mu)$ is a fixed point of $\hat{f}$, as shown in **Result** 5. By **Result** 3, we know that $\hat{f}^{t-1}((X^W, X^F, \mu) = (\bar{X}^W, \bar{X}^F, \bar{\mu})$. Furthermore, by **Result** 4, $\bar{\mu}$ is a stable matching in the market defined by choice function $\hat{\Omega}$; the market worker i has retired from. Then by the monotonicity of $\hat{f}$, $\bar{\mu} \hat{\succeq}^W \mu$ and $\mu \succeq^F \bar{\mu}$. □

11. Appendix B: Proof of Theorem 5

Here, we show the complete proof of **Theorem** 5 in Section 7 above, the argument of which is directly analogous to the argument of **Theorem** 4.

Proof. Again by **Result** 6, we know that the run of **S2** in which this refusal occurs results in a match $\bar{\mu}$ such that $\mu \succeq^F \bar{\mu}$, and $\bar{\mu} \hat{\succeq}^W \mu$ for remaining workers, the precise reversal of the result used in **Theorem** 4. When a worker retires, the firm may look to fill their spot, and thereby creates a vacancy chain in which workers move only if they are offered a contract with a higher ordinal ranking.

We now consider the full Stealing Process. After we have reached $\bar{\mu}$, the firms may or may not be able to add stealing contracts. If they cannot, as in the firm-closing case, the new equilibrium $\bar{\mu}$

is the final outcome with properties as described above. If they can offer new stealing contracts, we note that any such contract must be higher in the worker's strict ordinal ranking than their match in μ.

We now return to **S1** with a new set of contracts. At the end of **S2**, we have a new equilibrium for this set of contracts, which we will call μ'. We can demonstrate that the worker-offering Stealing Process implies that then $\mu' \hat{\succeq}^W \mu$, and the opposite for firms, by showing that no firm will be the first to reject a contract from its match in μ, and thus that all workers have weakly better matches under μ'.

Because all workers have a fixed ordering over all contracts, and no contract from μ has been rejected yet, if a firm j rejects some worker i with contract $\mu(i)$, every contract in the choice set of firm j, which we will call C, must be weakly ranked higher by the worker associated with that contract to their contract in μ. However, if a firm rejects $\mu(i)$, there must exist a contract x in C such that x is not in C_j^μ, otherwise μ would not be individually rational, as this would imply that removing $\mu(i)$ would raise j's profit. Because $\mu(i)$ is rejected here, we know that if x is replaced with $\mu(i)$, the profit of the firm will fall. Because there are no complementarities in firm production functions, a contract can not be more preferred based on the other contracts the firm takes. If the firm chooses x over $\mu(i)$ in this case, it must still choose x over $\mu(i)$ in any other situation in which it must choose between the two and would be willing to take either. Thus, x is preferred by both firm and worker to $\mu(\mathsf{w}(x))$, and x would block the original matching. Thus, no firm can be the first to reject a μ contract, and no firm will therefore reject such a contract.

This implies that any matching μ' will see all workers matched to weakly higher ranked contracts than under μ. It then follows that all firms will be weakly worse off, or else μ' would contain blocking coalitions for μ, which is disallowed by the stability of μ.

As with the earlier allocation $\bar{\mu}$, workers may enter a new round better off in terms of ordinal ranking and firms worse off in terms of profit. Any new contracts added will again be better than μ in terms of ordinal ranking. Across rounds, the above proof applies, and no firm will be the first to reject a μ contract. Thus, again relying on the finite number of stealing salaries, the algorithm will terminate in final allocation μ^* with the properties that $\mu \succeq^F \mu^*$, and $\mu^* \hat{\succeq}^W \mu$. □