Social Networks and Stable Matchings in the Job Market*

Esteban Arcaute[1] and Sergei Vassilvitskii[2]

[1] Institute for Computational and Mathematical Engineering, Stanford University. arcaute@stanford.edu
[2] Yahoo! Research. sergei@yahoo-inc.com

Abstract. In this paper we introduce and study a model that considers the job market as a two-sided matching market, and accounts for the importance of social contacts in finding a new job. We assume that workers learn *only* about positions in firms through social contacts. Given that information structure, we study both static properties of what we call *locally stable matchings*, a solution concept derived from stable matchings, and dynamic properties through a reinterpretation of Gale-Shapley's algorithm as myopic best response dynamics.

We prove that, in general, the set of locally stable matching strictly contains that of stable matchings and it is in fact NP-complete to determine if they are identical. We also show that the lattice structure of stable matchings is in general absent. Finally, we focus on myopic best response dynamics inspired by the Gale-Shapley algorithm. We study the efficiency loss due to the informational constraints, providing both lower and upper bounds.

1 Introduction

When looking for a new job, the most often heard advice is to "ask your friends". While in the modern world almost all of the companies have online job application forms, these are usually overloaded with submissions; and it is no secret that submitting a resume through someone on the inside greatly increases the chances of the application actually being looked at by a qualified person. This is the underlying premise behind the professional social networking site LinkedIn, which now boasts more than 40 million users. And, as pointed out by Jackson [10], has given a new meaning to the word 'networking,' with Merriam Webster's Dictionary's defining it as "the cultivation of productive relationships for employment or business."

Sociologists have long studied this phenomenon, and have time and time again confirmed the role that social ties play in getting a new job. Granovetter's seminal work [7, 8] headlines a long history of research into the importance of social contacts in labor markets. His results are striking, for example, 65 percent

* A full version of this work appears in [2]

of managerial workers found their job through social contacts. Other studies (see, e.g., [15, 13, 9]) all echo the importance of social contacts in securing a new position.

While there are numerous reasons that social ties play such an important role, one may think that the employers themselves would prefer to evaluate *all* candidates for a position before making a hiring decision. This is in fact what happens in some segments of the job market. In the United States, the National Resident Match Program is a significant example of a centralized selection matching mechanism. Such centralized markets have been well studied in *two-sided matching theory*. Indeed, the NRMP is one of the most important practical applications of the celebrated *stable matching problem* in two-sided matching markets [16]. For an overview of two-sided matching markets, see [17].

However, the task of evaluating (and ranking) all possible candidates is often simply not feasible. Especially in today's economy, it is not rare to hear of hundreds of applicants for a position, obviously the vast majority cannot be interviewed, regardless of their qualifications. The recommendation by an employee thus carries extra weight in the decision process, precisely because it separates the specific application from the masses.

Model - In this work, we propose a new model that bridges the rigorous analysis of the two-sided matching theory with the observations made by social network analysis. Specifically, we develop a model of job markets where social contacts play a pivotal role; and then proceed to analyze it through the stable matching lens.

We integrate the usage of social contacts by allowing an applicant to apply *only* to jobs in firms employing her friends. Clearly this limitation depends on the underlying social graph. Intuitively, the equilibrium behavior in well connected social graphs should be closer to that in classical two sided matchings than in badly connected ones. But even in well connected social graphs this limitation leads to behaviors not observed in the traditional model. For example, a firm may lose all of its workers to the competition, and subsequently go out of business.

The model forces us to consider a setting where job applicants have only partial information on job opportunities. The main question we focus on in the paper is: how does the inclusion of such an informational constraint alter the model and predictions of traditional stable matching theory?

Our Contributions [3] - In traditional two-sided matching theory, a matching where no worker-firm pair can find a profitable deviation is called *stable*. Analogously, we call our solution concept a *locally stable matching*, where the locality is qualified by the social network graph. We study structural properties of locally stable matchings by showing that, in general, the set of locally stable matchings does not form a distributive lattice, as is the case for global stable matchings. We also show that, in general, it is NP-complete to determine whether all locally stable matchings are also globally stable. Both of these results exploit a

[3] The proofs of all our results are available online at [2].

characterization of locally stable matchings in the special case of matching one worker per firm, for particular rankings over workers and firms.

We then turn our attention to dynamic analysis. We consider how a particular interpretation of the classic Gale-Shapley algorithm [6] performs under such informational constraints; we refer to our algorithm as the *local Gale-Shapley* algorithm. We first prove that, unlike the standard Gale-Shapley algorithm [17], the existence of informational constraints implies that the output of the algorithm is not independent of the order of proposals. Nevertheless, under weak stochastic conditions, we show that the local Gale-Shapley algorithm converges almost surely, assuming the same particular rankings over workers and firms as before.

Unlike the traditional Gale-Shapley algorithm, the algorithm in the limited information case is highly dependent on the initial conditions. To explore this further we define a minimal notion of efficiency, namely the number of firms still in business in the outcome matching, and quantify the efficiency loss under various initial conditions. Specifically, we show that if an adversary chooses an initial matching, he can ensure that some firms lose all of their workers; conversely there is a distribution on the preference lists used by the firms that guarantees that at least some firms remain in business, regardless of the actions of the adversary.

Related Work

Our work touches on several threads of the literature. Most closely related is the work by Calvó-Armengol and Jackson [3, 4]. They consider how information dissemination through neighbors of workers on potential jobs can affect wage and employment dynamics in the job market. There are several key differences with our model. The most important one is that, in [3, 4], there is no competition for job openings between workers. Unemployment is the result of a random sampling process and not of strategic interactions between workers and firms. Also, all workers learn *directly* about potential job openings with some probability, and *indirectly* through their social contacts, whereas in our model a worker can only learn about potential job openings through her social contacts.

Also related is the work by Lee and Schwarz [11]. The authors consider the stable matching problem in the job market where a costly information acquisition step (interviewing) is necessary for both workers and firms to learn their preferences. Once interviewing is over, the standard Gale-Shapley algorithm is used to calculate the matching of workers to firms. Although the authors use stable matching as their solution concept, and only partial information on jobs and candidates is available, their assumptions imply that the information available to workers and firms is *unchanged* throughout the matching phase. In that sense, their work is related to the equilibrium analysis performed in our model, but is dramatically different when considering the evolution of the job market during the actual matching phase.

Finally, in [1, 12], the authors consider the problem of matching applicants to job positions. A matching is said to be *popular* if the number of happy applicants is as large as possible. This notion is related to the notion of efficiency used in our paper, namely that of maximizing the number of firms in business.

2 Definitions and Notation

Let W be a set of workers, and $G = (W, E)$ be an undirected graph representing the social network among workers. Let F be the set of firms, each with k jobs, for some $k > 0$. We are interested in the case where there are as many workers as positions in all firms, i.e., $|W| = n = k|F|$. Following standard notation, for a worker $w \in W$, let $\Gamma(w)$ be the neighborhood of w in G.

An assignment of workers to firms can be described by a function mapping workers to jobs, or alternatively by a function mapping firms to workers. Following the definition from the two-sided matching literature, we define both functions simultaneously.

We assume that some companies are better to work for than others, and thus each worker w has a strict ranking $\succ_w$ over firms such that, for firms $f \neq f'$, w prefers being employed in f than in f' if and only if $f \succ_w f'$. Note however, that the ranking is blind to the individual positions within a firm: all of the k slots of a given firm are equivalent from the point of view of a worker.

Similarly, each firm f has a strict ranking $\succ_f$ over workers. We assume that all workers strictly prefer being employed, and that all firms strictly prefer having all their positions filled. For any worker w, in a slight abuse of notation, we extend her ranking over firms to account for her being unemployed by setting $f \succ_w w$ for all firms f; in a similar way we extend the rankings of firms over workers.

Definition 1 (Matching).

1. Case 1: $k = 1$. *The function* $\mu : W \cup F \to W \cup F$ *is a* matching *if the following conditions hold: (1) for all* $w \in W$, $\mu(w) \in F \cup \{w\}$*; (2) for all* $f \in F$, $\mu(f) \in W \cup \{f\}$*; and (3)* $\mu(w) = f$ *if and only if* $\mu(f) = w$.
2. Case 2: $k > 1$. *The function* $\mu : W \cup F \to 2^W \cup F$ *is a* matching *if the following conditions hold: (1) for all* $w \in W$, $\mu(w) \in F \cup \{\{w\}\}$*; (2) for all* $f \in F$, $\mu(f) \in 2^W \cup \{f\}$*; (3)* $\mu(w) = f$ *if and only if* $w \in \mu(f)$*; and (4)* $|\mu(f)| \leq k$.

We say that a matching μ *is* complete *if:*

$$\bigcup_{f \in F} \mu(f) = W.$$

Given a matching μ and a firm f, let $\min(\mu(f))$ be the *least preferred* worker employed by firm f (w.r.t. firm f's ranking) if $|\mu(f)| = k$, and $\min(\mu(f)) = f$ otherwise.

To study the notion of stable matchings, we adapt the usual concept of a *blocking pair*. Given the preferences of workers and firms, a matching μ, a firm f and a worker w, we say that (w, f) is a *blocking pair* if and only if $f \succ_w \mu(w)$ and $w \succ_f \min(\mu(f))$. In other words, worker w prefers firm f to her currently matched firm; and firm w prefers worker w to its least preferred current employee.

We now define a generalization of the standard notion of stable matching that accounts for the locality of information. Recall that a (global) matching is said to be *stable* if there are no blocking pairs. However, in our paper we assume that *the workers can only discover possible firms by looking at their friends' places of employment.* This informally captures a significant mechanism of information transfer: although there may exist a firm f that would make (w, f) a blocking pair, if none of w's friends work at f, then it becomes much less likely that w would learn of f on her own. We have the following definition.

Definition 2 (Locally Stable Matching). *Let $G = (W, E)$ be the social network over the set of workers W. We say that a matching μ is a* locally stable matching *with respect to G if, for all $w \in W$ and $f \in F$, (w, f) is a blocking pair if and only if $\Gamma(w) \cap \mu(f) = \emptyset$ (i.e., no workers in w's social neighborhood are employed by firm f).*

Note that for a given worker w, the set of other workers she is competing against depends on both the social network G (i.e., her neighbors), and the current matching.

Example 1 (Indirect Competition). Assume $k = 2$ and G is the path over $W = \{w_1, w_2, w_3, w_4\}$: $w_1 - w_2 - w_3 - w_4$. Consider worker w_4. If $\mu(f_1) = \{w_3, w_4\}$ and $\mu(f_2) = \{w_1, w_2\}$, then w_4 can only see positions in f_1. However, since w_2 is adjacent to w_3, w_2 can see all position in f_1. Hence, if $w_2 \succ_{f_1} w_3 \succ_{f_1} w_4$, w_2 could get w_4's position in f_1, leading to w_4 being replaced by w_2 even though $w_2 \notin \Gamma(w_4)$.

In the remainder of the paper, we characterize static properties of locally stable matchings, and then analyze dynamics similar to the Gale-Shapley algorithm.

3 Static Analysis

For $k = 1$, when the preferences of workers and firms are strict, it is known that the set of global stable matchings is a *distributive lattice*. In general, the distributive lattice structure of the set of global stable matchings is not present in the set of locally stable matchings. We first recall the Lattice Theorem (by Conway), and then show how, in general, it does not hold for locally stable matchings. The exposition of the Lattice Theorem is that found in [17] (Theorem 2.16).

Let μ and μ' be two matchings. Define the operation $\vee_W$ over (μ, μ') as follows: $\mu \vee_W \mu' : W \cup F \to W \cup F$ such that, for all $w \in W$, $\mu \vee_W \mu'(w) = \mu(w)$ if $\mu(w) \succ_w \mu'(w)$, and $\mu \vee_W \mu'(w) = \mu'(w)$ otherwise. For all $f \in F$, $\mu \vee_W \mu'(f) = \mu'(f)$ if $\mu(f) \succ_f \mu'(f)$, and $\mu \vee_W \mu'(f) = \mu(f)$ otherwise. We can similarly define $\wedge_W$ by exchanging the roles of workers and firms.

Theorem 1 (Lattice Theorem (Conway)) *When all preferences are strict, if μ and μ' are stable matchings, then the functions $\lambda = \mu \vee_W \mu'$ and $\nu = \mu \wedge_W \mu'$ are both matchings. Furthermore, they are both stable.*

In general, given strict preferences [4] of workers and firms, Theorem 1 does not hold for the set of locally stable matchings. This is the content of the following example.

Example 2 (Absence of Distributive Lattice). In this example, we assume $k = 1$, $W = \{w_1, w_2, w_3\}$ and $F = \{f_1, f_2, f_3\}$. Further, let the preferences of all workers be $f_1 \succ f_2 \succ f_3$. Similarly, let the preferences of all firms be $w_1 \succ w_2 \succ w_3$. Finally, assume the graph G is the path with w_2 and w_3 at its endpoints.

Let $\mu(w_i) = f_i$ (and $\mu(f_i) = w_i$). It is clear that μ is a 1-locally stable matching. Consider now μ' be such that $\mu'(w_1) = f_1$, $\mu'(w_2) = f_3$ and $\mu'(w_3) = f_2$ (and $\mu'(f_1) = w_1$, $\mu'(f_2) = w_3$ and $\mu'(f_3) = w_2$). The only blocking pair here is (w_2, f_2), but $f_2 = \mu'(w_3)$ and $w_3 \notin \Gamma(w_2)$. Hence μ' is a 1-locally stable matching.

We now construct $\lambda = \mu \vee_W \mu'$. For all i, $\lambda(w_i) = f_i$. Now $\lambda(f_1) = w_1$ but $\lambda(f_2) = \lambda(f_3) = w_3$. Hence λ is not a matching.

Assumption - In the remainder of the paper we focus on a specific family of preferences over workers and firms. Uniqueness of global stable matching is a desirable property in matching markets as it allows for sharp predictions of the outcome at equilibrium. Clark [5] studies thoroughly the question and identifies a set of sufficient conditions on the preferences, called *aligned preferences*, for the global stable matching to be unique. The study of aligned preferences have recently received attention in the economics literature [14, 18].

In this paper we consider a subset of aligned preferences, where all workers share the same ranking over firms, and firms share the same ranking over workers. This assumption is made for technical reasons - we believe our results extend to the case of general aligned preferences.

Assumption 1 *There exist a labeling of the nodes in $W = \{w_1, \ldots, w_n\}$ such that all firms rank workers as follows: $w_i \succ w_j$ if and only if $i < j$. Similarly, we assume there exists a labeling of the firms $F = \{f_1, \ldots, f_{n_f}\}$ such that all workers rank the firms as follows: $f_i \succ f_j$ if and only if $i < j$.*

We first show that, for $k = 1$, the set of locally stable matchings is equivalent to the set of topological orderings over the partial order induced by G and the labeling of the workers.

Theorem 2 (Characterization of Locally Stable Matchings) *Assume $k = 1$, and let $G(W, E)$ be the social network over the set of workers. Let $D(W, E')$ be a directed graph over W such that $(w_i, w_j) \in E'$ if and only if $i < j$ and $(w_i, w_j) \in E$. Let μ be a complete matching of workers to firms. Construct the following ordering ϕ_μ over W induced by μ: the i^{th} node in the ordering is the node w such that $\mu(w) = f_i$, i.e. $\phi_\mu(w) = i$.*

The matching μ is a 1-locally stable matching if and only if ϕ_μ is a topological ordering on D.

[4] The absence of the distributive lattice has been previously observed when the preferences are not strict, see Roth [16].

There are several important corollaries to the characterization from Theorem 2. First, the set complete locally stable matchings can be exponentially large. Thus, by introducing informational constraints, the uniqueness property of global stable matchings under aligned preferences is, in general, lost under locally stable matchings.

Corollary 3 (Number of Locally Stable Matchings) *Assume $k = 1$ and the social network $G(W, E)$ is the star centered at worker w_1. Then there are $(n-1)!$ distinct locally stable matchings.*

It is interesting to ask whether there are specific properties of the social network G that guarantee the existence of a labeling under which there is a unique complete locally stable matching. As shown in the next corollary, it is NP-complete to answer positively such question.

Corollary 4 *Let $(k, G(W, E))$ be given. It is NP-complete to test if there is a labeling $\{w_1, w_2, \ldots, w_n\}$ of the workers such that, if all firms rank the workers according to that labeling, the complete locally stable matching is unique.*

For general $k > 1$, we only have a set of sufficient conditions for complete locally stable matchings to be unique. See [2] for more details.

4 Algorithmic Questions

We are thus interested in decentralized algorithms that can find a locally stable matching. In this section we propose a decentralized version of Gale-Shapley's algorithm. Assumption 1 is again enforced in this section. We first prove that our algorithm converges. Unlike the case without informational constraints, our algorithm does not always select the same locally stable matching.

Recall that the Gale-Shapley algorithm is initialized by an empty matching [17]. Since the empty matching is a locally stable matching, our algorithm requires to be initialized by a non-empty matching. We thus explore our algorithm's performance under adversarial initial complete matchings. We use the number of firms with no employees as a proxy for efficiency[5]. We characterize the potential efficiency loss by providing upper and lower bounds on the number of firms with no employees.

4.1 Local Gale-Shapley Algorithm

One can interpret the Gale-Shapley algorithm from two-sided matching theory as a constrained version of *myopic best response dynamics* in the following way.

The dynamics proceed in rounds, which we index by $q \in \mathbb{N}$. Let $\mu^{(q)}$ be the matching at the beginning of round q. Let $w^{(q)} \in W$ be the *active* worker,

[5] Our bounds naturally translate into unemployment rate, a common indicator of the efficiency of the job market.

where $w^{(q)}$ is sampled uniformly at random from W, and independently from previous rounds. We call such sampling process the *activation process*. Such activation process can be thought of as follows: assume all workers decide to explore employment opportunities according to a random clock with an exponential distribution with a given mean (the same mean for all workers). When the clock of w_i "sets off", w_i becomes active and looks for a better job. It is easy to see that the sequence of active nodes has the same distribution as taking independent uniform samples from W.

In myopic best response dynamics, $w^{(q)}$ would consider its current firm $\mu^{(q)}(w^{(q)})$ and compare it to the best firm f it could be employed by given $\mu^{(q)}$ (i.e. the best firm where the worst employee was worse than w given the matching $\mu^{(q)}$). If its current firm was better, it would pass. Else it would quit its job and get employed by f (leading to a worker being fired, or an empty position being filled).

Gale-Shapley's algorithm is a constrained version of the above dynamics as it requires the active worker to consider the best firm it has not considered before (in other words it requires the active worker to remember what firms he has already failed to get a position at).

We consider a local and decentralized version of the myopic best-response dynamics proposed above. We call it "local Gale-Shapley" algorithm. Instead of restricting the strategy space of the active worker using "memory" as in Gale-Shapley's algorithm, we restrict it using the graph $G(W, E)$ in the following way: $w^{(q)}$ compares its current firm in $\mu^{(q)}$ to the best firm that employs one of its neighbors in G it could be employed by given $\mu^{(q)}$. An alternative way to describe the process is that the active node $w^{(q)}$ applies for a job at all the firms employing its neighbors that she strictly prefers to her current employer, and selects the best offer she gets (that offer might eventually be to stay at her current job).

More formally, the algorithm proceeds in rounds indexed by $q \in \mathbb{N}$. During round $q \geq 0$:

- the active worker $w^{(q)}$ is sampled, independently from previous rounds, uniformly at random from W.
- Next, $w^{(q)}$ applies to all firms she strictly prefers to $\mu^{(q)}(w^{(q)})$, her current employer.
- The active worker receives some offers:
  - if at least one offer is received, $w^{(q)}$ quits her current employer and joins the best firm that sent an offer;
  - if no offers are received, $w^{(q)}$ stays at her current job.

It is important to note that, unlike the Gale-Shapley algorithm, this variant of best-response dynamics can lead to a firm loosing all its employees as demonstrated below.

Example 3 (Firm with no Employees). Let $n = 4$ and $k = 2$. Thus there are four workers and two firms. Assume that $G = K_4$. Consider the following initial matching:

$$\mu^{(0)}(f_1) = \{w_3, w_4\} \text{ and } \mu^{(0)}(f_2) = \{w_1, w_2\}$$

in other words, the best company has the worst workers. Then if we activate workers w_1 and w_2 before activating w_3 or w_4, both w_1 and w_2 would quit f_2 and work for f_1, getting both w_3 and w_4 fired. In that setting, f_2 has no employees, and thus the process ends.

It is also important to understand the need of the activation process. Recall that the matching found by the Gale-Shapley algorithm is independent on the order of activation of the workers [17]. When considering locally stable matchings, this is no longer the case even if the underlying graph is the complete graph. Let us reconsider Example 3.

Example 4. Now consider the resulting matching when the activation sequence is as follows: $\{w_1, w_4, w_2, w_3\}$. First, w_1 leaves f_2 and gets a position at f_1. This makes $w_4 = \min(\mu^{(0)}(f_1))$ unemployed. Next, since we activate w_4, she gets the free position from f_2. Next w_2 leaves f_2 and gets a position at f_1, which results in w_3 loosing her job. Finally, w_3 gets the free position at f_2. Thus the resulting matching is now

$$\mu(f_1) = \{w_1, w_2\}, \text{ and } \mu(f_2) = \{w_3, w_4\}$$

which is a locally stable matching different from that obtained with the activation sequence in Example 3.

An important question is whether this local decentralized version of best response dynamics converges as it is not immediately clear it can't cycle. This is the content of our first result.

Theorem 5 (Convergence of Local Gale-Shapley Algorithm) *Given the social network $G(W, E)$. for any initial matching $\mu^{(0)}$, the local Gale-Shapley algorithm started at $\mu^{(0)}$ converges almost surely to a locally stable matching.*

4.2 Worst Case Efficiency

In this subsection we consider the following question. Given that firms can go out of business when running the local Gale-Shapley algorithm, can we measure the quality of matchings selected by the algorithm. We explore the previous question assuming a given initial complete matching $\mu^{(0)}$.

We consider the following setting. An adversary observes $G(W, E)$ (but not the ranking over workers used by firms) and produces a probability distribution $\mathcal{P}_M$ over initial matchings. The ranking of workers (possibly taken from a distribution) is then revealed, a sample from $\mathcal{P}_M$ is taken to produce $\mu^{(0)}$; and the local Gale-Shapley algorithm run.

To compare the efficiency of different final matchings we simply look at the total number of firms losing all of their employees and subsequently going out of business. One can easily imagine more intricate notions of efficiency, our point here is that even in this austere model, the power of the adversary is non-trivial.

The power of the adversary We first show that even without knowing the relative rankings of the individual workers, the adversary is powerful enough to force some firms to go out of business.

Theorem 6 (Lower Bound on Firms) *Let $G(W,E)$ be given. Let Δ be its maximum degree, and M a maximum matching in G. Then there exist a probability distribution $\mathcal{P}_M$ over complete assignment matchings such that*

$$\mathbb{E}[N_{fob}] \geq \left\lfloor \frac{|M|}{k(2\Delta)} \right\rfloor \frac{1}{2^k k!(2\Delta - 1)^k}$$

where N_{fob} is the number of firms going out of business; and the expectation is taken both over the distribution $\mathcal{P}_M$ and over the activation process.

Further, one can find $\mathcal{P}_M$ in time polynomial in n.

An important observation is that not only does the adversary force some firms to go out of business, but he controls the identities of these firms. Thus, if we measure efficiency by the identity of the firms of the positions filled in a matching, Theorem 6 provides a lower bound on the efficiency loss of the local Gale-Shapley algorithm (under adversarial initial conditions).

The power of the social planner Given the lower bound from Theorem 6 on the expected number of firms going out of business, we can ask the following question: can similar guarantees be proven if a social planner had full control over the ranking used by firms? More precisely, given $G(W,E)$, if the ranking over workers used by firms was a sample from a random variable, can the social planner guarantee, in expectation, a minimal number of firms that will not go out of business *regardless of the power given to the adversary*? The following theorem answers positively that question.

Theorem 7 (Upper Bound on Firms) *Let $G(W,E)$ be given. There exist a probability distribution over the ranking used by firms such that*

$$\mathbb{E}[N_{fob}] \leq n_f - \left\lceil \frac{|I|}{k} \right\rceil$$

where I is a maximum independent set of G (n_f is the number of firms and k the number of positions at each firm)

Note that, just as in Theorem 6 we were able to identify the firms forced out of business (the top firms) but not the unemployed workers, in Theorem 7 we are able to identify the workers that are going to be employed (the top employees), but not which firms will remain in business.

Discussion We have now shown that neither the adversary, nor the social planner have all the power — we can reinterpret the results above as a game between these two players. The game proceeds as follows, the adversary picks the initial assignment matching (possibly random), and the social planner chooses the ordering on the workers (possibly random). Once they both pick an action we run the local Gale-Shapley algorithm.

Theorem 6 then states that, even if the social planner knows the probability distribution selected by the firm adversary, there is a probability distribution over initial assignments that the firm adversary can use such that, in expectation, at least some number of firms go out of business.

Theorem 7 states the converse: Even if the firm adversary knows the probability distribution selected by the social planner, there is a deterministic ordering of the workers such that at least some number of workers will never loose their job.

We note that by looking at the number of firms going out of business, we have used a very minimal notion of efficiency. It is not hard to imagine more complex notions which may take into account the relative rankings of the firms going out of business or workers remaining unemployed. We further note that for dense graphs, where the size of the independent set, and the independent matchings are quite small, our bounds are quite loose. Our main contribution here is not the precise bound on N_{fob}, although that remains an interesting open question, but rather the fact that the adversary has non-trivial power, and the initial matching plays a pivotal role in determining the final outcome.

5 Conclusions

In this work we have introduced a new model for incorporating social network ties into classical stable matching theory. Specifically, we show that restricting the firms willing to consider a worker only to those employing his friends has a profound impact on the system. We defined the notion of locally stable matchings and showed that while a simple variation of the Gale-Shapley mechanism converges to a stable solution, this solution may be far from efficient; and, unlike in traditional Gale-Shapley, the initial matching plays a large role in the final outcome. In fact, if the adversary controls the initial matching, he can force some firms to be left with *no* workers in the final solution.

The model we propose is ripe for extensions and further analysis. To give an example, we have assumed that as employees leave the firm, it may find itself with empty slots that it cannot fill (and go out of business). However, this is precisely the time when it can start looking actively for workers, by advertising online, recruiting through headhunters, etc. This has the effect of it becoming visible to the unemployed workers in the system. Understanding the dynamics and inefficiencies of final matchings under this scenario is one interesting open question.

Acknowledgments

We would first like to thank Ramesh Johari for many discussions on the model studied and its static analysis. We also want to thank David Liben-Nowell for the discussions that lead to the proofs of Theorems 7 and 8. Finally, we would like to thank Ravi Kumar and Matt Jackson for many useful discussions.
This research was supported by the National Science Foundation and by the Defense Advanced Research Projects Agency under the ITMANET program.

References

[1] D. J. Abraham, R. W. Irving, T. Kavitha, and K. Mehlhorn. Popular matchings. *SIAM Journal of Computing*, 37(4):1030–1045, 2007.

[2] E. Arcaute and S. Vassilvitskii. Social networks and stable matchings in the job market. Working paper, arXiv.org, Oct. 2009. Available at http://www.arxiv.org.

[3] A. Calvó-Armengol and M. O. Jackson. The effects of social networks on employment and inequality. *American Economic Review*, 94(3):426–454, 2004.

[4] A. Calvó-Armengol and M. O. Jackson. Networks in labor markets: Wage and employment dynamics and inequality. *Journal of Economic Theory*, 132(1):27–46, 2007.

[5] S. Clark. The uniqueness of stable matchings. *Contributions to Theoretical Economics*, 6(1), 2006.

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

[7] M. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1380, 1973.

[8] M. Granovetter. *Getting a Job: A Study of Contacts and Careers.* University Of Chicago Press, Chicago, Il, USA, 1995.

[9] Y. M. Ioannides and L. D. Loury. Job information networks, neighborhood effects, and inequality. *Journal of Economic Literature*, 42(4):1056–1093, 2004.

[10] M. O. Jackson. *Social and Economic Networks.* Princeton University Press, Princeton, NJ, USA, 2008.

[11] R. Lee and M. Schwarz. Interviewing in two-sided matching markets. In *2008 American Economic Association Annual Meeting*, 2008.

[12] M. Mahdian. Random popular matchings. In *Proceedings 7th ACM Conference on Electronic Commerce (EC-2006), Ann Arbor, Michigan, USA, June 11-15, 2006*, pages 238–242, 2006.

[13] J. Montgomery. Social networks and labor markets outcomes. *American Economic Review*, 81(5):1408–1418, 1991.

[14] M. Niederle and L. Yariv. Decentralized matching with aligned preferences. NBER Working Papers 14840, National Bureau of Economic Research, Inc, Apr. 2009.

[15] A. Rees and G. P. Shultz. *Workers in an Urban Labor Market.* University Of Chicago Press, Chicago, Il, USA, 1970.

[16] A. E. Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, 92(6):991–1016, 1984.

[17] A. E. Roth and M. A. O. Sotomayor. *Two-Sided Matching.* Cambridge University Press, Cambridge, UK, 1990.

[18] M. Sorensen. How smart is smart money? a two-sided matching model of venture capital. *The Journal of Finance*, 62(6):2725–2762, 2007.