Hearing about a Job:
A Model of Differential Information Flow and Job Matching

Katherine Stovel*
and
Christine Fountain*

Working Paper no. 30
Center for Statistics and the Social Sciences
University of Washington
Seattle, WA 98195

April 17, 2003

*Department of Sociology, University of Washington. Email: **stovel@u.washington.edu**. This work is supported by a Seed Grant from the Center for Statistics and the Social Sciences at the University of Washington, funded by the University Initiatives Fund. We gratefully acknowledge useful comments and suggestions from members of the CSSS Seminar, and programming assistance from Peter Hoff.

Abstract

Which people end up in which jobs is not merely a matter of the individual and human capital characteristics of workers and the requirements and rewards of jobs, but is also a function of the process by which persons and jobs are matched with one another (Granovetter 1981; Sørensen and Kalleberg 1981). A poorly understood component of the matching process is how workers and employers find information about each other. We propose a framework for analyzing the dynamics of labor market behavior that emphasizes both the two-sided nature of the matching problem and the potential effects of different information structures. The basic framework for describing labor markets is flexible, and can be used to represent common types of labor markets including spot markets for labor and systems characterized by competition over vacancies. Through simulation, we use this framework to explicitly vary the structure of information flow between workers and employers, thereby allowing us to explore how different recruitment strategies influence labor market performance in different institutional environments. Preliminary results suggest that full information is the most efficient and egalitarian information regime. However, in key respects, information that flows through socially structured networks produces macro-level and individual-level outcomes quite similar to those observed in the full information regime, although with lower mobility rates. In contrast, unstructured limited information produces more egalitarian outcomes than either full information or social networks, by giving workers equally bad opportunities.

Introduction

Which people end up in which jobs is not merely a matter of the individual and human capital characteristics of workers and the requirements and rewards of jobs, but also the process through which people and jobs are matched (Granovetter, 1981; Sørensen and Kalleberg, 1981). Thus understanding labor market outcomes requires not only an understanding of supply and demand, but also making sense of this matching process. While difficult to study empirically, the first step in the matching process is the flow of information between employers and workers about potential vacancies. In most circumstances, those with jobs and those looking for jobs do not have perfect information about one another's existence. To the extent that access to information is limited, workers and employers are forced to choose from a set of possibilities that are incomplete and potentially biased. That is, if a worker does not hear about a job for which she would be an attractive candidate, that job will not be member of the set of offers she chooses from. And of course, the same applies to employers searching for qualified workers. Given the consequences of imperfect information, it is important to recognize that the structure of information flow is probably neither random nor uniform; rather, access to information about employment opportunities is typically structured by a variety of factors such as geography, social and familial relationships, professional and occupational relationships, and access to media and information technology.

This paper describes a labor market model that emphasizes both the two-sidedness of the matching problem and the potential effects of different information structures on both individual and macro-level employment patterns. Our focus on information availability highlights the importance of strategies that employers use for advertising vacant positions to potential workers, and how these strategies influence the matching process in different institutional environments.

Background

Recruiting for workers

Employers use a variety of different strategies to recruit applicants for vacant positions. Some use classified advertisements, some recruit through their social and business contacts, some post vacancies in school or professional registers, others use job fairs or employment agencies, and an increasing number take advantage of the internet bulletin boards and job search services to recruit potential workers. Yet, little is known about the relative efficacy of these strategies and what their relative impact is on the efficiency, quality, and rewards of the matching process. Some research suggests that recruitment through personal networks results in better paying and more satisfying jobs for those workers who are lucky enough to get them that way (Granovetter 1995) but there is also reason to believe that the advantage of social networks is not uniformly distributed (Lin, Ensel & Vaughn 1981).

The simplest versions of neo-classical theory suggest that broadcasting information about a vacant position as widely as possibly results in the best job-match, because this brings in the largest number of applicants from which an employer could choose. However, since employers must determine which of these applicants are qualified, and then choose among qualified candidates (both costly endeavors), employers might in fact prefer to choose from among a small group of qualified employees rather than a large group with more variance in qualification level. This refined pool might be better obtained through personal networks than through broadcast means. Fernandez, Castilla & Moore (2000) show that in a large bank, recruiting through existing employees' contact networks resulted in a pool of applicants who were better qualified than those who come in without a referral.

However, recruiting through personal referrals has implications for the diversity of organizations, because homophily in contact networks tends to result in the reproduction of the current socioeconomic make-up of organizations (Doeringer & Piore 1971, Fernandez, Castilla & Moore 2000). Because of this, many large firms and public organizations are required to widely disseminate information about all job vacancies, so that applicants unconnected to the current employees have the chance to apply as well. Of course, in spite of bureaucratic rules designed to keep the "old boys'"

network from excessively influencing hiring decisions, those applicants with personal referrals may still have an advantage, since employers have more trustworthy information about these workers and may have lower monitoring costs once referrals are hired.

In certain segments of the labor market, however, broadcast recruitment may be an effective recruitment strategy. When an employer has many jobs with low or variable skill levels, it may be cheaper to place classified ads or post openings on the web since there is little pay-off to worrying about finding the 'best' worker. Widespread recruitment strategies are common among large firms seeking entry level workers in many divisions, temporary help or business services/contractor firms that require a constant stream of new recruits at all skill levels, and in industries-- such as the fast food industry—where jobs have very high turnover and low skill requirement.

This suggests that different recruiting strategies employed by firms might result in a different set of applicants for vacant positions, as well as a different set of job opportunities for workers. Even under the assumption that all workers are looking for the most rewarding jobs and employers are looking for the most qualified workers, if workers and jobs do not know about each other then there is no possibility for them to be matched.

Workers' opportunities

On the worker side of the matching process, types of job-searching strategies as well as the structure of personal networks can influence the size and quality of a worker's opportunity set. Granovetter (1973, 1995) showed that individuals with many "weak" ties might have an advantage in the labor market, since these acquaintances give them access to more diverse information than if they had to rely on close friends and family. Building from Granovetter, Montgomery (1992) argued that the proper focus should be the overall structure of a worker's network--not just the actual source of the information that led to a successful job match. Using an analytic model, he also showed that the level of weak versus

strong ties in a population's social network has important implications for both employment level and wage inequality.

Labor Market Structures

Complicating this situation even further is White's insight that when a worker finds another job, her former job becomes vacant, which creates a specific opportunity for another worker (White, 1970). In this way, worker mobility creates a vacancy chain that moves in the opposite direction the workers' movement. As long as jobs are somewhat durable even when vacant, a worker's opportunity is constrained by the existence of opportunities in the form of vacancies, which are created by the movement of other workers. Sørensen and Kalleberg (1981) and Sørensen (1977) described conditions under which we might observe opportunity structures where this type of movement is especially important, which they call "vacancy competition."

Of central importance in such "closed" systems is the fact that workers can only fill jobs that are not currently held by another worker: thus workers are competing for vacant jobs. Jobs are more or less attractive because they are associated with certain wages and other benefits, so that it is the particular job a worker holds that determines wages, not the productivity of the worker per se. Of course, the worker's productivity helps to determine which jobs she might fill, and especially the highest potential attainment, but is not the only factor. Thus the opportunities of a worker depend not only on her own productivity, but also on the existence of opportunities, which are themselves a product of the opportunities faced by other workers.

Vacancy competition can be contrasted with neoclassical free market for labor, where workers compete amongst themselves for the highest wage level at a certain skill level (thus keeping the price of labor at the point of marginal productivity). Free markets for labor (often termed spot-markets) are "open" labor markets, since there is no expectation by either employer or employee that the worker should continue to hold a job. When the worker cannot restrict access to their job, wage competition will

result where employers hire workers on a daily basis at the market rate. Although these markets might be empirically rare (Kalleberg & Sørensen 1979), they are theoretically important for economic analysis of labor markets and useful as a point of reference for more structured markets.

These two models of labor market structures can easily be integrated into to the literature on segmented or dual labor markets. In this line of research, labor markets are considered to be divided into two or three distinct sectors, with mobility between sectors relatively rare. In general, there is a primary market for good jobs which are characterized by higher wages, benefits, and opportunities for advancement, and a secondary market for bad jobs which are low-paying, unfulfilling, unstable, and dead-end. The primary market can be considered roughly equivalent to a regime of vacancy competition, where workers have some expectation of stability in their job until they choose to leave for a better opportunity, and that incumbency at a certain level gives one access to opportunities at the next level. In contrast, the secondary labor market is more similar to a wage competition regime (with some important differences), where low wages, involuntary separations, and a notable lack opportunities for upward mobility (given stable human capital levels) constrain attainment. Most researchers distinguish between these labor markets according to industrial sector, however others focus on the characteristics of the workers, skill requirements, and employer-specific training.

Another useful analog for this framework is in comparing firm internal labor markets to a neoclassical unstructured labor market. In an internal labor market there is typically a "ladder" of jobs within a firm, with a clear entry level (although workers can enter at any level) and a hierarchy of promotions associated with higher pay and greater responsibility. These promotions can be used as a reward for good performance, so that although wages do not necessarily reflect the worker's intrinsic marginal productivity, they are not necessarily inefficient either. In a labor market structured in this way, opportunities for advancement depend not only on the worker's performance, but also on the availability of a position on the next rung of the ladder.

If access to jobs is limited by the opportunity structure of the labor market, it is further constrained by the availability of information. Even if a vacancy exists for which a worker is qualified, she cannot be hired if she does not know about the vacancy. Conversely, if a worker knows about a particular job that she prefers to her own, she cannot be hired if the incumbent cannot be easily fired. In a regime where vacancies, and thus opportunities for advancement, are scarce, timely and complete information becomes even more important. A missed opportunity might mean an indeterminate wait for another appropriate vacancy, thus limiting the lifetime attainment of a particular worker. The consequences of a missed movement filter through the system and impact the opportunities of other workers who might hope to fill vacancies in the chain created by a potential job change.

Modeling Framework

Our dynamic modeling framework takes as fundamental that the problem of matching in a labor market has two sides: a set of workers choosing from a set of jobs, with the same set of jobs choosing from the set of workers. The actual matching of workers-to-jobs is conceptualized as an iterative process in which workers and jobs repeatedly evaluate their utility for one another. Characteristics of the labor market structure and the preference structure are incorporated into the utility calculation; the utilities of workers for jobs and jobs for workers, combined with the access structure, drive the matching process.

Different institutional arrangements in the labor market –ranging from pure spot-markets to pure vacancy competition—control the "stickiness" of the match between particular workers and particular jobs. Thus in an open labor market system, qualified workers can potentially be hired for any job, vacant or filled, because the employment relationship is constantly re-negotiated, while in a closed system a qualified worker can only move to a position if it is vacant. Beyond incorporating different labor market structures, however, the main advantage of our modeling framework is that employers and workers may not have access to information about one another; rather, access to information about

workers and jobs is explicitly varied in controlled ways. In addition to a baseline of complete information, access to information about employment opportunities can flow randomly, through job- or occupation-based networks, or through social networks of different configurations. This enables us to explore how various recruitment strategies influence labor market performance in different institutional environments. We discuss details of the model's parameter structure below; first, however, we outline the model's basic flow control.

Flow Control

We implement the model as a dynamic simulation, as shown in Figure 1. The initialization phase has two components: here we set basic parameters describing the labor market structure (including size, number of firms, number and distribution of occupational classifications, and initial employment level), the access regime(s), and the number of iterations; and impose the initial match of workers to jobs.

Figure 1: Flow Control for Model

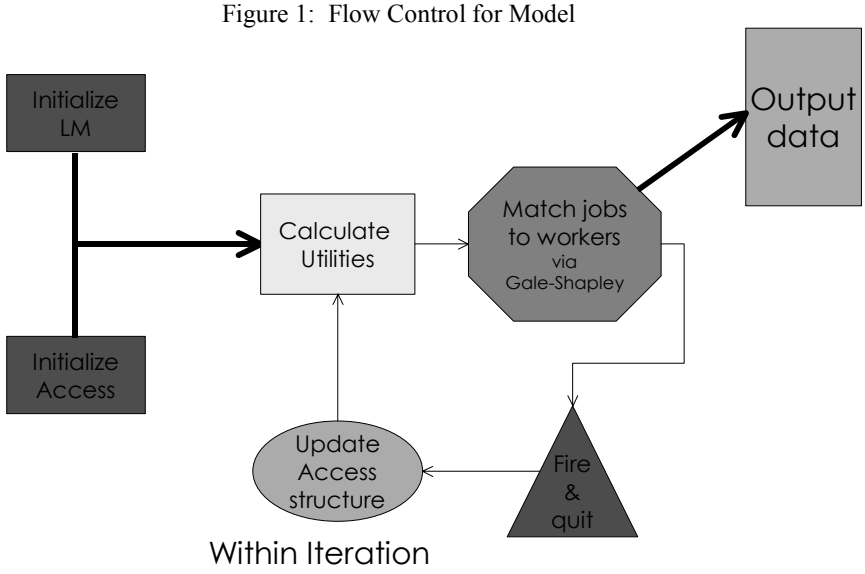

After initialization, the simulation proceeds to the first iteration. We begin by calculating the utilities of workers and jobs for one another. Utilities must be recalculated for each iteration because individual and job-related characteristics vary over time, and because the flow of information may

change. After utilities are calculated, workers and jobs are matched to one another. Some workers or jobs may remain unmatched if their utility for being unemployed/vacant is higher than their utility for the best offer or potential workers. The labor market created by this matching algorithm is a stable, optimal match given the utilities created at the beginning of that iteration. We then begin initialization for the next iteration. We first run a routine that fires or resigns a small proportion of workers. Firing and resignation are stochastic functions of tenure and occupational classification (the shape of these functions depends on the institutional regime of the labor market). These processes create additional vacancies and unemployed workers for the next iteration. If appropriate, the access structure is updated at this point; then the next iteration begins. The simulation proceeds over the specified number of iterations, detailed data describing the state of the labor market at each iteration is saved to a file.

Outcomes

Our substantive goal is to examine how the dynamics of labor markets are altered by various recruitment strategies. Of particular interest is the extent to which future reductions in the cost of widely disseminating information about vacancies -- such as might be realized with web-based job-listings – may alter the quality, speed, and rewards of the matching process for various types of jobs and workers. To this end, a variety of individual and macro-level outcomes and statistics are saved for graphical and numerical analysis. Some outcomes of interest include employment level and stability, transitions into and out of employment and unemployment, quality of matches, efficiency of matching process, and employment segregation. We are continuing to develop strategies for assessing the relationship between model parameters and these and other outcomes; later in the paper, we describe patterns that have emerged in our early experimentation with the model.

Model Details

Three broad families of parameters govern the model: labor market structure, preference

structure, and access structure. These are linked via the matching algorithm, which allocates workers to jobs.

Labor Market Parameters

The model begins with a population of workers and jobs, of sizes *w* and *j*, respectively. The populations of workers and jobs are structured across several dimensions. Both workers and jobs are allocated into *k* classes, which are roughly analogous to occupations or type of work (but can alternatively be thought of as skill levels). The number of classes can be altered to represent certain labor market characteristics, such as markets for workers with and without college degrees, or high, medium and low-skill workers. These classes are qualitatively different from one another, and can be treated as ordered or unordered. In the extreme cases, when $k = 1$ all jobs and workers are indistinguishable, while as $k \to n$ each worker or job is unique. Workers and jobs are assigned to classes by an ordered vector of probabilities.[1] Since they are assigned probabilistically, the classes include roughly, although not exactly, the same numbers of workers and jobs.

In addition, each worker is also assigned a restlessness score (drawn from a normal distribution, $\mu = 0$) to reflect individual heterogeneity in propensity to change jobs. When restlessness is positive, the worker will be inclined to change jobs even when their new job is not substantially better than the old; if restlessness is negative, she may stay in a job even when "better" opportunities exist. An analogous score is created for jobs, which reflects differing returns to tenure in jobs.

Finally, each job is assigned to a firm. The number of firms in the labor market is determined by the parameter *nf*; the number of jobs in each firm is variable, though firms can contain an assortment of job classes.

Preference Structure and Utility Calculation

[1] Proportions of workers and jobs in each class can be drawn from using data such as the Current Population Survey, so as to be representative of a national, regional, occupational, or industrial labor market.

Our model assumes rational workers and jobs who seek to maximize the utility of their match. In each iteration of the dynamic model, utilities are calculated for each job-worker and worker-job pair. Utilities of workers for specific jobs are a linear function of worker characteristics (such as employment status, utility for current job, worker inertia), job characteristics (class, vacancy), and worker-job pair characteristics (tenure, mismatch). Utilities of jobs for specific workers are also function of job characteristics, worker characteristics, and job-worker pair characteristics. In the most general form, each worker *i*'s utility for job *j* can be described as:

$$U_{ij} = \alpha x_{ij} + \varepsilon_{ij}$$

where x_{ij} includes a variety of worker and job characteristics and ε_{ij} is a random term for individual variation in preferences. U_{ij} for unemployment will be nonzero for many workers, which implies that some workers may choose unemployment over any available options.[2]

Similarly, each job *j* has a utility for each worker *i*, given by:

$$V_{ji} = \beta x_{ji} + \varepsilon_{ji}$$

where x_{ji} includes job and worker characteristics and ε_{ji} is a random term for job-specific variation in preferences. V_{ji} can equal zero if the employer would prefer to have the job remain empty rather than fill it with *i*.

The sets of parameter values α and β in the utility equations can be thought of as a preference structure; in the trials we describe in this paper, the preference structure is fixed (though the variable values vary across iterations). Once iteration-specific utilities are calculated, values are transformed into a rank-ordered sets of preferred options for each actor. Sets can be empty if no possible matches are preferred to remaining unmatched, or can include the entire population of potential matches.

Information-access related structures

[2] In this sense, U_{ij} for unemployment is analogous to a reservation wage, where the value of welfare, unemployment insurance, leisure time, or not having to pay for child care is greater than the value of all jobs from which a worker can choose.

Access to information about specific jobs or workers is structured by various *information regimes*. Each information regime determines how likely it is that a particular worker will hear of a particular job or vacancy. Conceptually we represent an information regime with a set of n_W x n_W matrices C_T, where n_W is the number of workers in the system; and T denotes particular a type of tie between workers that may be used for recruitment (e.g., social ties, shared occupation, shared firm, subject to same broadcast media). Hence $C_{T:i,j} = 1$ if worker i and j share a tie of type T. Information contained in the set of connectivity matrices C_T can be used to generate an n_W x n_J matrix of access A such that $A_{i,j} = 1$ if worker i and job j have access to each other. By convention, all access ties are symmetric, such that if worker i is aware of job j, then job j is aware of worker i. The matrix A is then used to adjust the utilities $U_{i,j}$ and $V_{j,i}$ such that if $A_{i,j} = 0$ then $U_{i,j} = 0$ and if $A_{j,i} = 0$ then $V_{j,i} = 0$.

In our preliminary use of the model, we have developed seven theoretically distinct information regimes. Our baseline information regime is the *full information regime*, in which every worker and every job know about each other, and $A_{:i,j} = 1$ for all workers i and jobs j. All the remaining information regimes introduce an information constraint; our theoretical interest is in the consequences of variation in the structure of information flow across nodes. While the total amount of information flowing in the labor market under a given information regime is a tunable parameter *rp*, we hold this constant in our comparison of regimes.

The least structured of our incomplete information regimes are the *random incomplete regimes***.** For a given value of *rp*, access information is distributed randomly across workers and jobs independent of their characteristics, with the caveat that all workers know about the job they currently hold. In the *fixed random* regime, the access structure is constant across all iterations; in the *updated random* information flows are re-assigned at the start of each iteration.

In the *firm* information regime, workers know only about jobs in their own firm. If a worker quits or is fired, she no longer has access to information about any new jobs, and therefore stay

unemployed for the remainder of the simulation. Since there is no mechanism for movement between firms in this regime, the unemployment/vacancy level tends to rise with each iteration; absent market-level population dynamics, the labor market will eventually fail. By itself, this regime is a poor descriptor of real labor markets, but because it captures key features of strict firm-internal labor markets, it is interesting in conjunction with other information regimes.

The *occupational information* regime structures information by worker-job classification. In its pure form, this regime prohibits movement between classes, essentially forming k distinct labor markets in which all job turnover is within a particular worker-job sector. Theoretically. this regime approximates a strict caste system or a guild, where assignments to jobs are made through an occupational organization. As with the firm-based regime, the occupational regime is less interesting in itself, but can be combined with other information dissemination structures.

Finally, we use a class of information regimes that are structured by workers' *social networks*. Following Watts, we adopt a network structure characterized by both a high degree of local clustering and high overall connectivity (1999a, 1999b). This "small-world" structure is common among human actors, and has been shown to have high potential for diffusion of information. In these networks, actors tend to have ties to actors with whom they share alters, but a tunable proportion (controlled by the parameter α) of ties are random. We use two variants of a small-world structure: one in which clustered ties are independent of worker characteristics (*social network no ring*), and one in which workers tend to know other workers of their type (*social network with ring*).

Since social networks are ties between workers, we use firms to connect workers to jobs. Under the social network regimes, a worker knows about a job if he or she is socially tied to someone who works in the same firm as that job. This means that access to information about jobs is limited by actors' social networks so that workers who know many people, or people in large firms, have more job opportunities.

It also means that if workers tend to know others in their own class, inter-class mobility will be even more limited that would be expected based purely on the preferences of workers and jobs.

Currently, both of the network structures we use are static, though we are preparing to implement a flexible procedure that makes the networks themselves dynamic over iterations of the simulation. In this procedure, the propensity of ties to form and dissolve is a function of parameters governing clustering, connectivity, and labor market experiences (including firm tenure). We believe that allowing networks to change dynamically will more realistically simulate the way that social networks change as persons move in and out of jobs, make new friends, and lose track of old ones.

Matching algorithm

The matching algorithm we use is a variation on the Gale-Shapley algorithm for two-sided markets. The Gale-Shapley procedure has many advantages, not the least being that it has been frequently used in two-sided matching models, and its properties are fairly well understood. The basic logic of Gale-Shapley is a deferred acceptance model. Given a set of rank orders for all potential actors in the system, a job is randomly selected to propose a match to its most preferred worker. The worker provisionally accepts the job if it is in the worker's ranked set and is more preferred than the offer they are currently holding. If a worker has already provisionally accepted a job but receives an offer from a more preferred job, the second offer is provisionally accepted and the first job is returned to the pool of vacant jobs. The matching proceeds in this fashion until all jobs have made offers to all the workers they prefer to remaining vacant. All workers need not have received offers, and some jobs may remain vacant. It has been shown that given a fixed set of preferences, the conventional G-S matching algorithm produced a stable, optimal match (see Roth and Vande Vate 1990).

We make two modifications to the Gale-Shapley algorithm.[3] First, and most importantly, we limit the set of actors who can be matched in any given iteration to those pairs of jobs and workers who have access to information about one another. Technically, this means that if $A_{ij} = 0$ then $U_{ij}=V_{ij}=0$. Second, the utility inputs have a stochastic component, which means that neither utilities nor the actors available to be matched are necessarily stable between iterations. This produces the dynamic nature of the system that is crucial for simulating a labor market. A worker who is matched to her most preferred job at time 1 might find that an even better job has opened up in time 2. Thus, mobility can be more explicitly vacancy-driven than is possible in the conventional Gale-Shapley where the matching proceeds among all actors in the system.

Preliminary Trials

Preliminary experiments with this simulation model have involved testing the model and outcomes for face validity by running trials over various ranges of parameter values. The simulated data discussed here all come from trials of a small, symmetric labor market (number of workers = number of jobs = 100) with an initial employment level of 0.9.[4] As summarized in the table below, we ran the simulation across the seven "pure" information access regimes, varying the number of firms. Each configuration of parameter values was run fifteen times, with 100 iterations per trial.[5]

---

[3] Details about the simulation, including both a flow chart describing the algorithms steps and computer code in R, are available from the authors.

[4] Early experiments with this parameter showed that the initial employment level had no impact on the simulation after the first few iterations. The employment level resulting from the simulation depends much more on other factors, including information regime and the preference for remaining unmatched.

[5] The specific utility equations we use in our preliminary trials are:

$UW_{ij}=\pi + \psi_1$(current unemployment rate)$+ \psi_2$(job j is vacant)$+ \psi 3$ (worker and job in same class) $+ \psi_4$ (job is in class 1)$+$ turnover $_j$* (tenure$_{ij}$) + restlessness of worker $i + \mu_{ij}$

and

$UJ_{ij}=\Omega + \lambda_1$ (current unemployment rate)$+ \lambda_2$ (job j is vacant)$+ \lambda_3$ (worker and job in same class)$+$ turnover $_j$* (tenure$_{ij}$) + tenure reward for job $j + \mu_{ji}$

Where

π = . 3 (base-line preference for employment)

Number of Firms	Information Access Regime						
	Full Info	Random (fixed)	Random (updated)	Firm	Occupation	Network (connected)	Network (unconnected)
5	*15*	*15*	*15*	*15*	*15*	*15*	*15*
10	*15*	*15*	*15*	*15*	*15*	*15*	*15*
15	*15*	*15*	*15*	*15*	*15*	*15*	*15*
20	*15*	*15*	*15*	*15*	*15*	*15*	*15*

After running these trials, we examined patterns in both individual- and macro-level labor market outcomes. Our simulated data suggests that different information transmission regimes are associated with markedly different labor market performance.[6] Thus we are confident that the model offers a useful tool for investigating important and yet empirically difficult research questions.

Macro-level Outcomes

The mean proportion of iterations spent unemployed is a system-level measure of the efficiency of the labor market at keeping workers and jobs matched over the course of the trial. Figure 1 displays boxplots of the mean unemployment level for each of seven access regimes. The full information regime has an unemployment rate similar to that of the random updated regime, however the random updated regime appears to have slightly less variation. Although full information is an important theoretical

---

ψ_1 = .1 (preference for all jobs if unemployed)
ψ_2 = .2 (preference for vacant jobs)
ψ_3 = .05 (preference for taking job in appropriate occupational category)
ψ_4 = .1 (preference for type 1 jobs)

and

Ω = .3 (base-line preference for hiring)
λ_1 = -.1 (preference for unemployed workers)
λ_2 = .2 (preference for hiring if job is vacant)
λ_3 = .2 (preference for hiring worker in appropriate occupational category)

In addition, we added some random noise (μ_{ij} and μ_{ji}) to each utility value to reflect variations (and perhaps the bounded rationality) in worker/job preferences between iterations. The noise is small enough to not overwhelm the preferences for certain characteristics, but should be large enough to discriminate between choices and to introduce some instability in preference ordering between iterations.

[6] Number of firms produced few interesting patterns, and therefore we have suppressed variation in this parameter for ease of presentation.

construct because it corresponds to a crucial neoclassical assumption for a perfectly competitive market, this suggests that it might not be necessary to maintain stable and efficient matching, at least as measured by mean unemployment. In fact, the more limited information regime seems to create more stable conditions, possibly as a by-product of reducing the number of job changes.

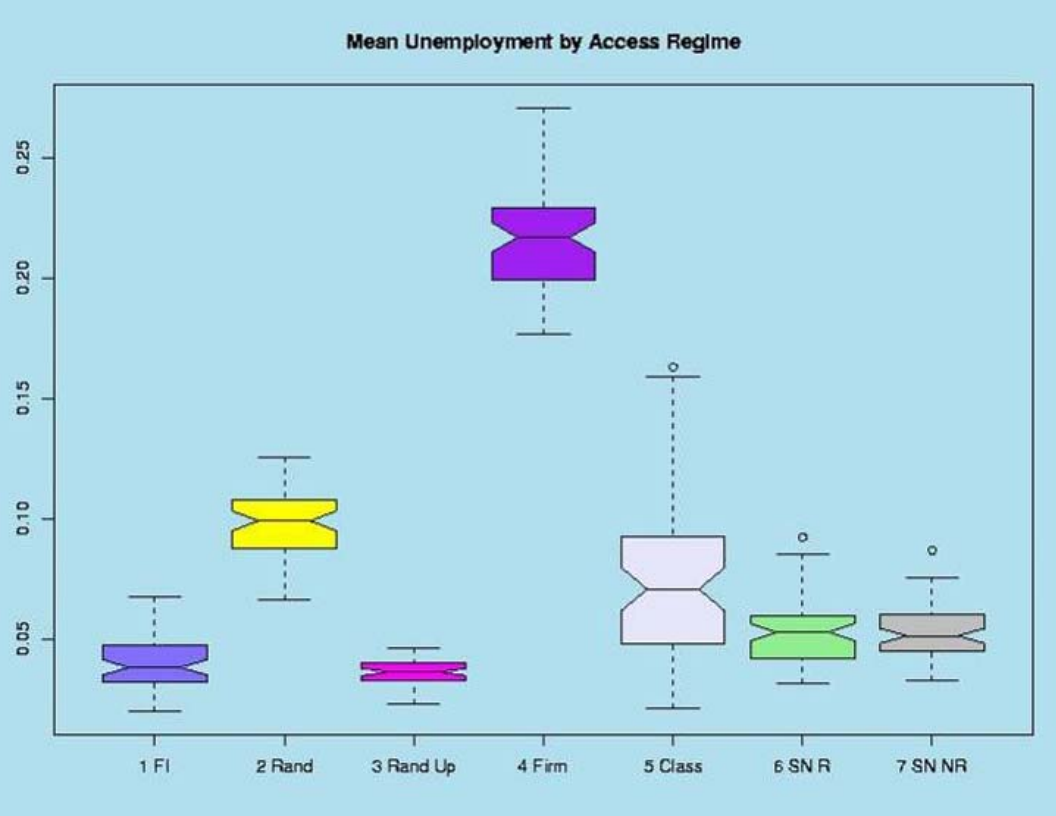

The firm-based information regime results in very high unemployment because since information is only transmitted through the firm one already works in, the unemployed are never able to find out about jobs and so the unemployment rate rises in each iteration. The class-based regime also has fairly high unemployment, with a substantial amount of variation about the mean. The variance is likely due to the fact that mismatch between the size of worker and job populations can result in the absence of appropriate jobs that persists throughout the simulation. Finally, the social network regimes (which provide the same amount of total information as the random regimes) result in unemployment rates similar to those found in the full information and randomly limited regimes. Mean unemployment is

only slightly higher, and has a similar amount of variance as the full-information regime. Thus in terms of mean unemployment, the simulated data suggests that there is little difference between full information and randomly or socially constrained information structures. Mean unemployment tells us nothing about the distribution of unemployment spells across individual workers, however.

Stability in worker-job matches, measured by the *frequency of job transfers*, is another outcome of interest. In this context, stability in matching suggests minimal worker-job mismatch. In addition, to the extent that frequent job changes mean additional search and screening costs for both workers and jobs, as well as training costs for jobs and potential future income loss for workers, information regimes resulting in lower rates of job transfer represent a more efficient and effective matching process. The figure below reports the mean rate of job transfer across trials for each information regime.

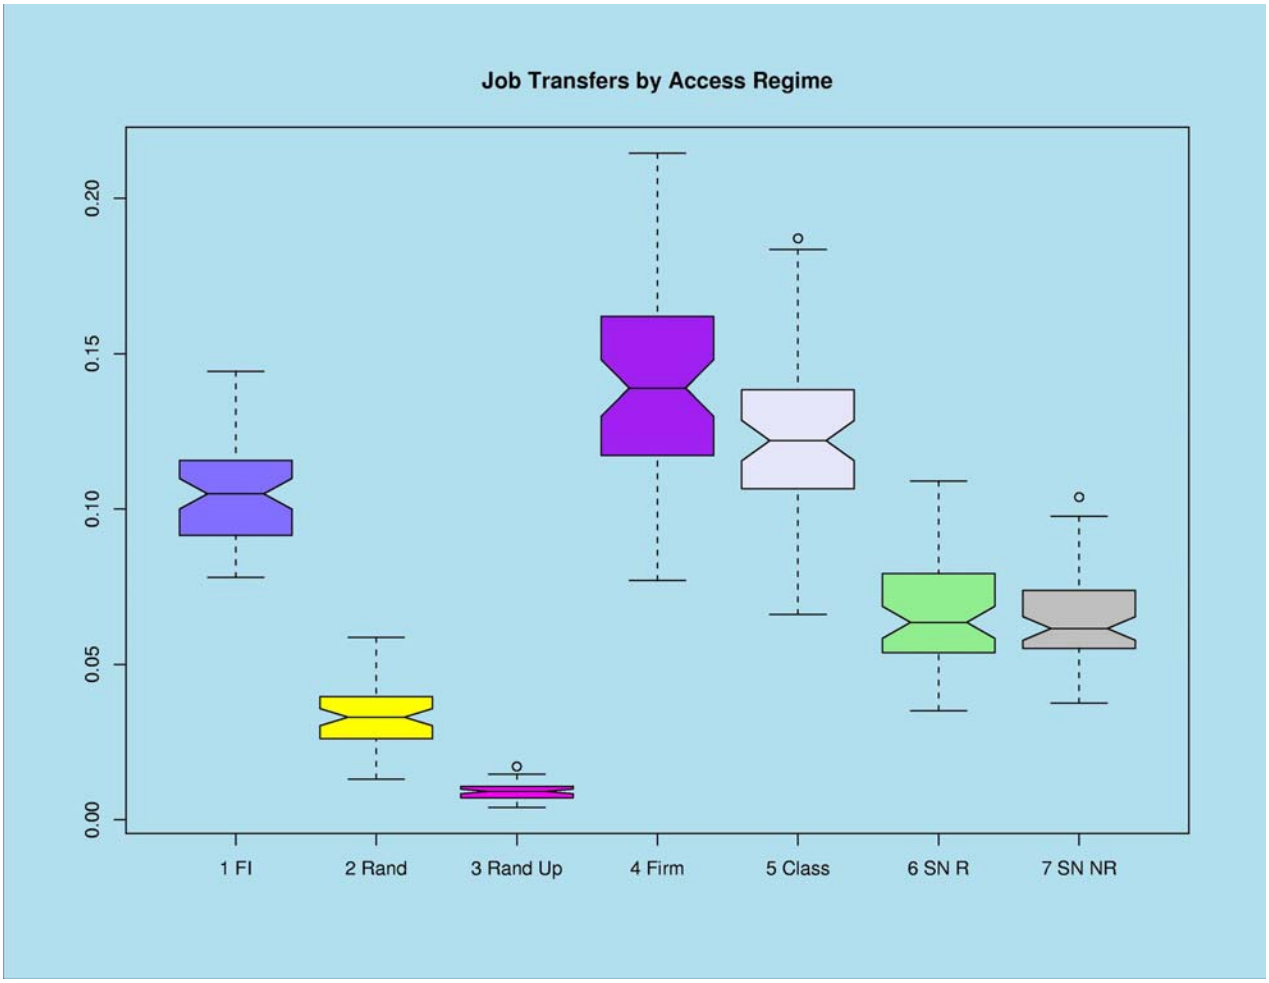

The firm and class-based regimes result in the most instability in job-matching, though there is quite a bit of variability within these regimes. This is interesting because access to information is severely structured in both these regimes, yet frequent job changes can still occur, suggesting a large amount of mismatch. Full information also results in fairly high mobility rates, which is not surprising since all actors know about all potential matches and have slight instability in preferences. The updated random regime has very low mobility, also not surprising given the limitations each actor has on their queue of potential partners. Although very high rates of mobility might not be optimal in a labor market, very low rates might also suggest a lack of opportunities facing the actors in this simulation. Interestingly, the social network regimes, with the same absolute amount of information as the random regimes, produce

substantially higher mobility that random, but lower than full information. Since the limited information in the network regimes cross-cuts other relevant dimensions, these regimes can provide workers and jobs with more expansive pools of information about attractive job opportunities, while the random regime is less likely to give workers information about attractive jobs for which they are qualified. On the other hand, since information is still limited, the social network structure produces more stability than full information. Again, this macro-level statistic does not tell us about the distribution of job transfer among individual actors; these regimes might produce radically different amounts of inequality in career stability if job transfers are concentrated among a segregated group of workers and/or jobs.

Under some information regimes, it might be more difficult to find a job once unemployed than under others. The next figure reports the *mean number of workers in each iteration who moved from unemployment to being matched to a job*. Again the firm-based regime is an outlier: here is zero of these movements, since once one becomes unemployed in this system, it is impossible to find a job. As we would expect, full information provides the most opportunity to move out of unemployment. Workers in the random regime are less able to move out of unemployment, which is not surprising since there is little reason to expect that many of the small number of jobs these workers know about will be desirable. The class-based and social network regimes produce approximately the same rates of movement out of unemployment, although the class distribution is broader. Social networks seem to result in less opportunity to move out of unemployment than the full information regime, but more than the random regimes. However, differences in all of these statistics are quite small.

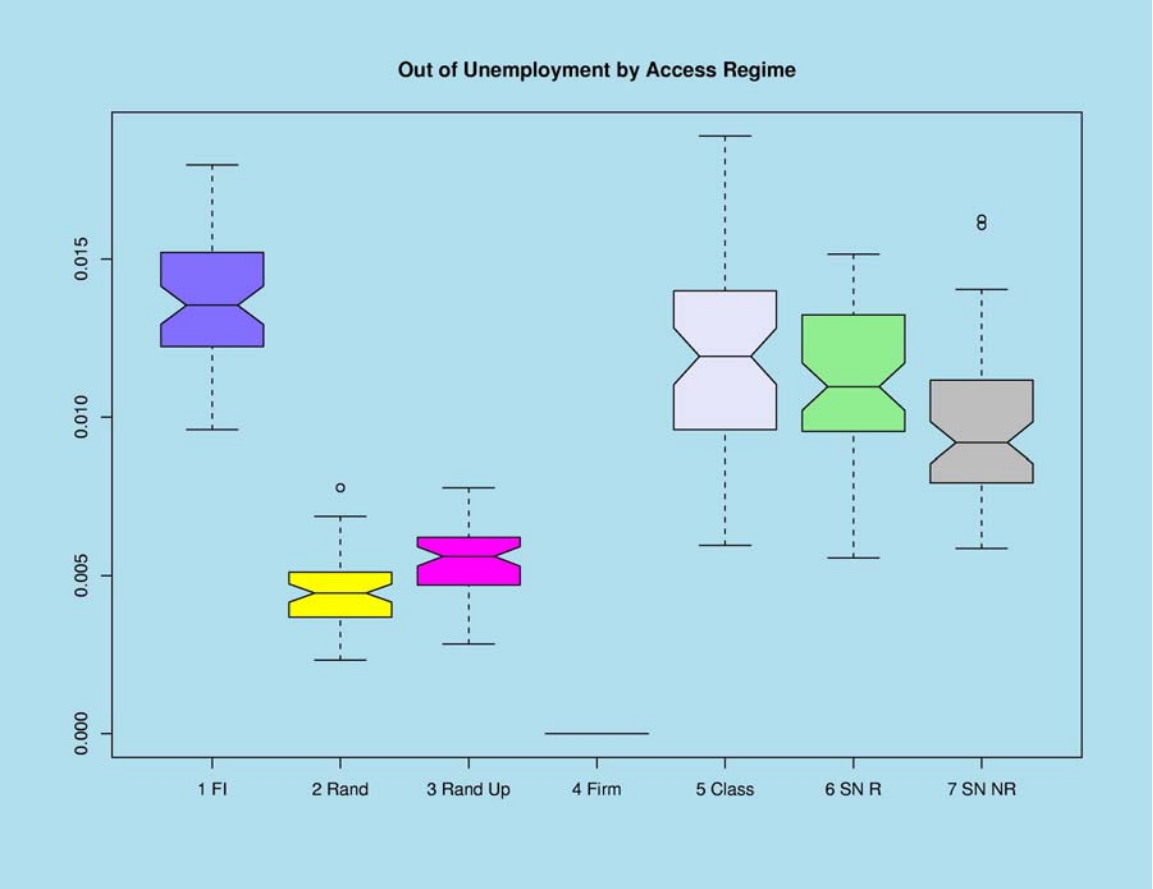

Individual Inequality in Outcomes

The next figure reports the number of different jobs a worker might hold during her 100 iteration "career," under each of the information regimes. The idea is to learn something about how instability and turnover are distributed among actors. We see that both full information and the two social network regimes result in similarly wide, flat distributions. There is a definite mode around four jobs, yet there is a great deal of inequality represented by the long right tail. A small number of workers are experiencing a great deal of instability under each of these three regimes. In contrast, the two random regimes produced a much more egalitarian distribution, with overall instability quite low (as we have seen above). Again, workers are more likely to stay in current jobs when faced with limited opportunities, and are unlikely to be stuck in careers characterized by frequent job changes.

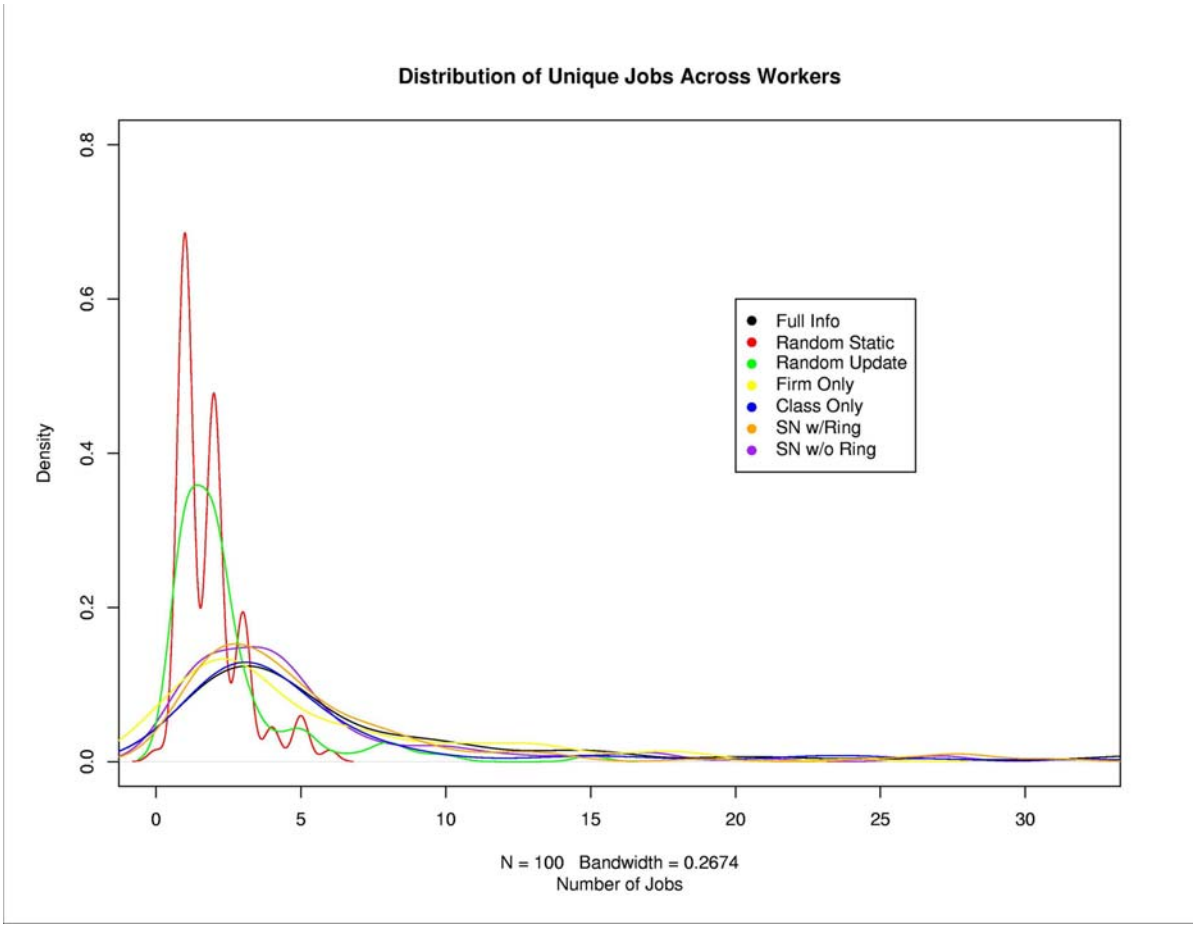

Discussion

Overall, we find that when information is limited in any way, mobility and turnover decrease. This suggests that if there is any penalty for frequent job changes, either as search costs or rewards to tenure, full information may not produce optimal outcomes. However, this might also signify a limited amount of opportunity for actors in these labor markets. The pattern evident in our early simulation trials is that the distribution of number of jobs held during a career does not differ greatly between social network and full information regimes, while the limited unstructured information regime reduces both the mean and variance substantially. Thus in key respects, social networks produce macro-level and individual-level outcomes quite similar to those observed in the full information regime, although with

lower mobility rates. In contrast, unstructured limited information produces more egalitarian outcomes than either full information or social networks, by giving workers equally bad opportunities.

Our experience with this model suggests that it holds the potential to investigate a wide variety of different research questions. In one sense, the algorithm that stochastically matches a population of workers to a population of jobs assumes a standard economic model of utility maximization for both sides. However, in classic sociological fashion, we have relaxed the unreasonable assumption of perfect information, instead structuring the flow of information between people and jobs in more realistic ways. By allowing information dissemination strategies to vary, we can examine how these factors influence the process of matching individuals to jobs, thereby learning a great deal about occupational outcomes, social mobility, inequality, and the potential implications of recruitment strategies and new broadcast information technology.

There are many directions to go with this model in the near future. We are currently working on refining the parameterization of the labor market itself (especially our representation of labor market institutions), the preference structure of workers and jobs, and the various information regimes (including introducing dynamic networks). In addition, we plan to introduce more structure into the population of workers, in order to represent other types heterogeneity (i.e., sex, race or ethnicity, or some other worker attribute like possession of a particular skill) that are important for labor market outcomes.

Beyond these obvious refinements, however, our next major step is to incorporate population dynamics into the model. In order to make the model fully dynamic and more realistic, we feel it is important to model the process by which workers enter and exit the labor market, and the process by which jobs are created and destroyed. This requires both incorporating substantive information about the processes governing population dynamics in labor markets, and an adjustment of the existing simulation architecture to allow for variation in population size across iterations. Incorporating population dynamics will have both substantive and technical implications for the structure of

information flows as well; for instance, it will require that the composition of social networks are revised in each iteration. The payoff of this development is that we will be able to examine additional interesting research questions, such as what happens when there is a change in the supply or demand for a certain type of labor, or what happens when the populations of workers and jobs are of significantly different sizes.

References

Doeringer, Peter B., and Michael J. Piore. 1971. *Internal Labor Markets and Manpower Analysis.* Lexington, Mass.: D.C. Health.

Fernandez RM & Weinberg N. 1997. Sifting and sorting: Personal contacts and hiring in a retail bank. *American Sociological Review* 62: 883-902.

Fernandez RM, Castilla E & Moore P. 2000. Social capital at work: Networks and employment at a phone center. *American Journal of Sociology* 105(5): 1288-1356.

Gale, D. and Shapley, L.S. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Quarterly* 69(1):9-15.

Granovetter MS. 1973. The strength of weak ties. *American Journal of Sociology* 78(6): 1360-1380..

Granovetter MS. 1981. Toward a sociological theory of income differences. In: I Berg, ed. *Sociological Perspectives on Labor Markets*. New York: Academic Press.

Granovetter MS. 1995. *Getting a Job: A study of contacts and careers*, 2nd ed. Chicago: University of Chicago Press.

Kalleberg AL, Sørensen AB. 1979. The Sociology of Labor Markets. *Annual Review of Sociology* 5: 351-379.

Lin N, Ensel WM, Vaughn JC. 1981. Social resources and strength of ties: Structural factors in occupational status attainment. *American Sociological Review* 46: 393-405.

Marsden PV, Hurlbert JS. 1988. Social resources and social mobility: A replication and extension. *Social Forces* 66: 1038-1059.

Montgomery JD. 1992. Job search and network composition: Implications of the strength-of-weak-ties hypothesis. *American Sociological Review* 57(5): 586-596.

Montgomery JD. 1994. Weak ties, employment, and inequality: An equilibrium analysis. *American Journal of Sociology* 99(5): 1212-1236.

Roth, Alvin E. and Vande Vate, John. 1990. "Random Paths to Stability in Two-Sided Matching." *Econometrica* 58(6):1475-1480.

Sørensen AB, Kalleberg AL. 1981. An outline of a theory of the matching of persons to jobs. In: I Berg, ed. *Sociological Perspectives on Labor Markets*. New York: Academic Press.

Sørensen AB. 1977. The Structure of Inequality and the Process of Attainment. *American Sociological Review* 42: 965-978.

Watts DJ. 1999a. *Small Worlds: The Dynamics of Networks Between Order and Randomness*. Princeton, NJ: Princeton University Press.

Watts DJ. 1999b. Networks, Dynamics, and the Small World Phenomenon. *American Journal of Sociology* 105: 493-527.