Employment Intermediaries: A model of firm incentives

Debra Hevenstone
University of Michigan
dhevenst@umich.edu

September 11, 2008

This paper introduces a micro-simulation of job-worker matching with intermediaries (i.e. temp agencies). While many suggest that firms hire workers through intermediaries to save money on compensation, this paper finds that in a world of limited information and geographically limited job search, intermediaries' human resources ability could be a strong enough incentive, independent of compensation. The study also has some auxiliary findings showing that traditional fee structures encourage firms to use intermediaries for low-skill hires and that firms are more likely to use intermediaries when there is more worker heterogeneity. In the empirical analysis, it becomes clear that studies' estimates of indirect employment in the United States are inconsistent, partly because individuals are uncertain of their contractual status and their employer.

keywords: outsourcing, contractors, atypical employment, intermediaries

1 Introduction

Temporary or or contingent work has supposedly spread in recent years (Clinton, 1997; Kalleberg et al., 2000b; Kalleberg, 2000; Mangum et al., 1985) raising concerns about their effect on social stratification, as these workers have lower pay, fewer benefits, less on-the-job training, and less job satisfaction (Booth et al., 2002; Kalleberg et al., 2000b; Houseman et al., 2003). "Temp work" or outsourced work means that a firm hires a worker through a second firm. The second firm remains the worker's legal employer, although the worker physically works at the firm that purchased his or her services. (Throughout this paper I refer to these workers as "indirect employees.") Empirically, temp work can be difficult to distinguish from direct employment at firms that provide services to other firms (such as accounting) but where the worker still physically comes to work at the direct employer's. In the US economy, the purchase of labor services, temp work and otherwise, is growing rapidly. In the past decade, firms increased their purchases of services more than they increased direct hires, with the consequence that business services grew at a rate of 5.8% every year from 1988 to 1997, twice the rate of the rest of the economy (Clinton, 1997).[1] The fastest increasing sub-sector within the business services category is the temporary help industry, which grew 11% annually from 1979 to 1995, five times more quickly than all other non-farm employment (Autor, 2000).

This paper presents an agent based model (ABM) of job search and job-worker matching in a labor market with intermediaries.[2] The general term

[1] "Business Services" is a Bureau of Labor Statistics category including: advertising and public relations services; computer system design and related services; employment services; management, scientific, and technical consulting services; and scientific research and development services.

[2] Those unfamiliar with ABM might see Macy and Willer (2002) for an introduction. ABM is a simulation composed of interacting agents that follow micro rules, generating macro system behavior. In this article, the micro agents are firms, jobs, and workers (the basic elements of a labor market) and macro system properties are employment rates, vacancy rates, and other aggregate labor market

"intermediaries" is used here to refer to both "temp agencies" and "contractors" because the functional difference between contracting labor through a firm like a temp agency and simply outsourcing is difficult to distinguish; the difference largely has to do with the implicit length of the contract, the worker's skill, and the level of integration into the primary firm. The model describes how firms might adjust their use of intermediaries in response to (dis)incentives such as intermediaries' better ability to match workers with jobs and tests whether these hypothesized incentives could be one explanation for the patterns of intermediary use in today's US labor market. It further examines how these incentives function differently in different occupational labor markets. This model does not look at wage gaps as an outcome of organizational decisions (only as a determinant) and (like most simulations in the social sciences) does not provide proof that the tested scenarios are necessarily the definitive explanation, but rather, that they are one feasible explanation. Finally, it includes several important assumptions about indirect hiring that will become clear in the explanation of the model.

The model's results suggest that intermediaries can provide a valuable service to firms by increasing the firm's capacity for searching for new workers. The model also finds that incentives to use intermediaries differ depending on the skill variability of workers in the occupational labor market. In addition, organizational ecology is very important in firms' decision to use intermediaries and finally, in an environment where firms pay a percentage of salaries as a fee to intermediaries, lower-skilled jobs are sorted into indirect employment.

indicators.

2 Background

One of the primary concerns about atypical employment is that these workers have lower compensation (Booth et al., 2002; Segal and Sullivan, 1997; OECD, 1999). Figure 1[3] shows the annual wage-gap between direct and indirect workers for four occupations in the US using the March Current Population Survey (CPS).[4] The figure shows that while indirectly employed janitors and clericals are consistently paid lower wages, indirect programmers and accountants earned lower wages only until the mid 1990's, after which they fluctuate around the same level as regular workers. Plotting total income rather than wage income (not depicted here), indirect programmers and accountants consistently earn *more* income than their direct-hire counterparts. This could be because these workers have secondary income sources such as independent contracting. Indirect hires in all occupations are consistently less likely to receive health insurance from their employers as illustrated in figure 1, where the line indicates the difference between the proportion of regular and indirect employees with employer-provided health insurance. Note that some of the workers that are counted as "uninsured," actually have insurance through a secondary source such as a spouse's employer-provided health-insurance scheme. Insurance through a spouse is more likely for high-skill workers who are both more likely to be married and more likely to be married to a partner who has health insurance benefits. To date, evidence suggests that differences between direct hire's and indirect hires' individual characteristics (work effort, education, residential location, age, and gender) do not fully explain the compensation gap. In fact, CPS data suggest that

[3] There is a discontinuity in coding in 1992 and 2002.

[4] Indirect workers are identified by matching industry and occupational codes for workers so that a secretary working for a clerical services firm is assumed to be an indirect employee. This method captures some temp workers, some workers employed through intermediaries, and some workers simply working at service firms.

indirect clerical workers should have higher compensation than their direct-hire counterparts, given that they have more education and reside in more urban areas, two characteristics normally correlated with higher wages. Because of indirect employment's significantly lower wages, many researchers assume that the firm's primary incentive to use intermediaries for firms is to save money on compensation, particularly in low-skill occupations (Houseman et al., 2003; Kalleberg et al., 2000a).

INSERT FIGURE 1 HERE

There are many hypothesized incentives for firms to use intermediaries. In contrast to those arguing that intermediaries are about reducing compensation, some researchers argue that firms *underpay* direct-hires in the high-skill labor market, using indirect employees as a temporary substitute while searching for permanent employees willing to accept lower wages (Houseman et al., 2003).[5] Another potential benefit is that intermediaries could match workers and jobs more efficiently (Katz et al., 1999), decreasing the firm's human resources expenses. Other hypothesized incentives include: maintaining a flexible labor force, testing low-quality or risky workers, hiring specialized workers for short periods, increasing employee-job match quality, and focusing on firms' core competencies (Deavers, 1997; Gramm and Schnell, 2001; Abraham, 1990; Mangum et al., 1985; Mayall and Nelson, 1982; Young and MacNeil, 2000; Davis-Blake and Uzzi, 1993; Osterman, 1999; Pfeffer and Baron, 1988). In addition, in high-skill positions, indirect hires might be more productive working at the intermediary than with their occupational peers (for example a programmer employed at a firm customizing software, rather

[5]While researchers claim that the incentive is saving on compensation, firms themselves disagree. The National Organizations Survey directly asked firms' human resources officers why they hire through intermediaries. HR departments responded that their firm does it primarily because of work fluctuations and because contractors' have specialized skills. Most responded that it does not lower costs.

than the firm using the customized software). On the other hand, there are also many hypothesized disincentives to using intermediaries including the importance of firm-specific skills, intermediaries' fees, large firms' ability to internally smooth labor consumption, and union regulations prohibiting hiring workers through intermediaries.

The most cited incentive in the US context is primarily that intermediaries allow firms to reduce compensation costs, though not primarily through wages but rather by cutting health insurance costs. In the United States, this incentive is embedded in the tax structure: there are federal tax incentives for businesses to provide equal health benefits to all their employees. Firms can qualify for these tax incentives, despite denying part of their workforce health insurance, if they purchase *services* instead of labor, using only direct-hires (who all have health insurance) in the tax calculation.[6] From the employee side, some researchers find that high-skill indirect hires receive the same total compensation in the form of fewer benefits but higher wages; presumably workers with spousal benefits might seek contract work to cash out their benefits (which is illegal in regular employment contracts) (Houseman et al., 2003). While from the worker's perspective, it seems clear that firms save money on compensation, there is actually mixed evidence whether, including the cost of intermediaries, firms actually save money (Benson, 1999;

[6]The US tax code offers businesses tax deductions for health insurance and pension expenditures. However, these tax deductions are only available if these benefits are not provided in a way that favors high skill workers. US Code Title 26, subtitle A, Chapter1, Subchapter D, PartI, Subpart A, Section 401 a(4) states that deductions are granted:

> if the contributions or benefits provided under the plan do not discriminate in favor of highly compensated employees (within the meaning of section414(q)). For the purposes of this paragraph, there shall be excluded from consideration employees described in section 410(b)(3)(A) and (C).

The definition of "highly compensated" is regularly updated and was changed twice while this paper was written. It is defined, generally, as employees earning over some threshold or constituting some top percent of the firm's workforce. The consequent penalty is that firm expenditures on pensions, health insurance, and life insurance are taxed 15%.

Young and MacNeil, 2000; Mayall and Nelson, 1982; Mangum et al., 1985; Deavers, 1997; Davis-Blake and Uzzi, 1993). However, studies do find that those firms with higher wages are more likely to contract out work (Abraham, 1990; Gramm and Schnell, 2001), presumably because indirect employees do not receive those higher wages.

This paper uses a micro-simulation (an agent based model) of the labor market to study the question of which firm incentives could motivate the use of intermediaries. The simulation first matches worker and jobs in a free labor market, looking at the resulting labor market dynamics like unemployment and vacancy rates. The model then introduces intermediaries, tests different incentive theories, and looks at the subsequent overall levels of intermediary use. The mechanism that matches jobs to workers in this artificial labor market is based on the Gale-Shapely marriage matching algorithm (Gale and Shapely, 1962). In this algorithm, men and women rank each other as possible mates. Then, men propose to their highest ranked woman. If they are rejected, they propose to their second choice, and so on. Woman accept proposals if they do not already have a partner or if the new offer is preferable to their current partner. Their prior (jilted) partner must then propose to the next highest ranked woman on his list. Given an equal number of men and women, this algorithm is proven to find a stable solution where everyone is matched and no man and woman would rather be with each other than their current partner (Gale and Shapely, 1962). The solution is optimal for men, leaving them matched to their highest-ranked feasible partner. In the simulation, companies are equivalent to men making offers to workers, instead of to women. The model was implemented with firms playing the male role because I assume that workers generally apply to a broad array of jobs while firms make proposals to individuals chosen from large applicant pools.

There are several labor market models using similar simulation methods. The most similar is Stovel and Fountain (2003), which explores Granovetter's "strength of weak ties" theory (Granovetter, 1973), testing whether workers are more likely to be matched to their jobs through their close friends or through their acquaintances ("weak ties"). Stovel and Fountain test how the shape of a social network limits information in the labor market and affects the quality of worker-job matches. Tesfastion (2001) uses an extension of Gale-Shapely in an agent-based model, testing whether the ratio of jobs to workers or of firms to workers is more important in allocating negotiating power. Tassier and Menczer (2001, 2005) used social networks in a job matching model similar to Stovel and Fountain's, first examining how networks evolve through job matching, and second assessing how employment rates vary between social groups as a function of their network structure. Other models use job matching algorithms to examine frictional unemployment rates (Hosios, 1990), many-to-one matching (Echenique and Yenmez, 2005), or matching in wage posting games (Montgomery, 1991; Peters, 1991; Shi, 1998). Other worker-job matching abms include Fagiolo et al. (2004), Neugart and Storrie (2006), and Richiardi (2003). To date, there is no implementation of these methods examining the role of labor market intermediaries.

3 Model

The model describes the spread of intermediaries, focusing on four labor market scenarios. The model is laid out on a 2-D grid with four types of objects on the grid: firms, jobs, workers, and contractors. Firms and workers stay in fixed locations for the duration of a simulation, while jobs and contractors appear and disappear. Two sets of experiments vary a total of seven parameters, with the parameters related to hypotheses about tradeoffs between incentives and disincentives to use intermediaries. The parameters are listed in the appendix in table 1 and control only four (dis)incentives: 1) intermediaries' better screening capabilities 2) intermediaries' fees 3) compensation differentials and 4) workload variability. The decision to use an intermediary for the next vacancy is based on a utility function that firms regularly update, measuring whether they have had greater utility from past direct or indirect hires. The primary model output is the level of intermediary use under each scenario, though the model also measures the unemployment rate, vacancy rate, firm utility, and job and vacancy duration, most of which are used to qualitatively tune the model using empirical data. The model has several underlying assumptions, which are detailed after a description of the model's algorithm.[7]

The model's algorithm is illustrated in Figure 2. The broad overview is: first workers, firms, and jobs are created; then, workers are matched to jobs using a variation of the Gale-Shapely algorithm; next, workers and jobs suffer separations; and finally, contractor arrangements are updated. Then the model starts the process again starting with the matching step.

INSERT FIGURE 2 HERE

In each run, there are 1000 workers and 138 firms. Jobs are assigned to firms in a skewed distribution with most jobs at a few firms but no firm having more than 10% of the jobs.[8] Workers are created with skill levels

[7]The model was programmed in Java with RePast and the code is available from the author upon request.

[8]The function assigning jobs to firms (C.1 in the appendix) determines the ratio of the number of jobs to the number of firms. Since the model should start with approximately 1,000 jobs (to match the workers), the number of firms were chosen accordingly. Thus there are exactly 138 firms.

sampled from four empirical distributions of occupation (general, minimum wage, programmers, and accountants) and are assigned a skill floor (a random uniform deviation below their skill level, indicating the worst job they would accept). Educational distributions come from BLS and CPS data, and are depicted in figure 3. Jobs are assigned skill levels and floors using the same methods as for workers, but for jobs the skill floor indicates the worst *worker* the job would accept. Workers have a location on the grid, an employment status, an employer and job (when employed), a contractor (if using one), the date they were last employed (if they are unemployed), a random inherent tendency to quit that is time-invariant, and their relative wages when hired through an intermediary. Relative indirect wages is the average percent of a direct hire's wages that indirect hires get, ranging from 90 to 110%. These relative wages are reassigned to a worker each time the worker is hired through a contractor. Firms have locations, jobs (vacant and filled), a contractor (if they are using an intermediary), employees, and a history of their current and past utilities from their direct and indirect jobs. Jobs have skill levels and floors, a firm, an employee (when they are filled), and dates marking the last time they were filled or vacated. Contractors have assigned jobs, workers, vacancies, fee-rates, revenues, and matching rates.[9]

INSERT FIGURE 3 HERE

When the model is initialized, all workers are unemployed, all jobs are vacant, and there are no intermediaries. Workers sort through vacant jobs, calculate the distance to each job, and apply to closer jobs with a higher probability. Workers ignore the match between their skill and the job's skill when sending out applications. This is an unrealistic assumption in those simulations testing the overall labor market but is a realistic assumption in

[9]All equations, variables, and parameters are listed in the appendix.

those testing occupational labor markets. Because the workers apply to so many jobs, this assumption does not make the matching in the overall labor market simulations ineffective; workers still apply for a significant number of jobs they will be considered for. The advantage that intermediaries have is that workers perceive those jobs hiring through intermediaries to be closer, and are thus more likely to apply. This distance effect should not be taken at face value, but rather proxies for intermediaries' better human resources capabilities; they advertise more widely and screen through more applicants. This could also be the equivalent of better screening. High-skill workers also search in a broader radius than low-skill workers. The formulae are specified in detail in the appendix but the general magnitude of the effect is such that all workers apply to adjacent direct-hire jobs with a 100% probability; a graduate-level worker applies to direct-hire jobs at the furthest possible distance across the grid with a 7.7% probability, while the high-school graduate applies to the most distant job with a .09% probability.

Next, firms rank applicants based on the match between the vacant job's skill and the prospective employee's skill, and then offer the job to their top applicant. Workers accept jobs "tentatively," meaning that they accept with the option of taking another offer during the same matching round, just as women in the Gale-Shapely algorithm can dump a suitor. Firms have four chances to make offers in a single matching round. When a round ends, workers must stay with their tentative job for at least one round. The limitation of four offers prevents perfect matching, thus maintaining unemployment and vacancies. If the stock of jobs were not constantly changing, there were no skill floors, the offer process were iterated until matches were stable, and there were equal numbers of workers and jobs, there would be no unemployment or vacancies. While all workers make binding agreements in each round, they have the chance to "think over" an offer, or in other words wait for another offer to come in, before deciding. This is more realistic than a model forcing workers to instantaneously accept a match upon offer.

After workers and jobs are matched, there are quits and fires. The quit function (again, in the appendix) is the sum of three effects. Workers are more likely to quit if they are poorly matched to their jobs, have a high inherent tendency to quit (a time-invariant trait drawn from a normal distribution that could be considered analogous to marital status, age, etc), and if they receive a random shock (a time-variant trait drawn from a random distribution). When a worker "quits," he or she might be matched with the job they just left, since both re-enter the matching pool. As such, "quitting" includes on-the-job search. Direct and indirect hires quit using the same algorithm. Firms fire workers when they suffer random workload shocks (adding or removing jobs) which are proportional to firm size. Shocks are not correlated across firms (economic downturns) but because of the skewed firm size distribution, a negative shock to a big employer strongly influences the overall unemployment rate. When firms fire workers they first remove vacant jobs, then fire indirect hires, and finally fire direct-hires. Firms fire without respect to tenure or match quality.

After matching and separations, the model updates contractor dynamics. Up to two new contractors can be born in a single model step. The first is born if there is a high vacancy rate and the second is born if there is high demand for existing contractors. This represents a continual low level of contractors randomly placed on the grid. The new contractors are allowed to survive for less than 1% of the model duration (presumably on startup capital) before they are forced to meet a revenue threshold. Revenue is calculated as the sum of fee rates times their worker skills divided by the total number of workers they are assigned. If a contractor is earning, on average, less than 10% of the average worker's skill per assigned worker, they a removed from the model. Thus a contractor's health depends on both their ability to match workers and their fee rate. Because contractors are continually born and each has a different number of clients, the number of contractors is not representative of contracting trends but is representative of the service availability because of the way firms find intermediaries. [10]

Firms decide to use intermediaries the first time when they have a persistently vacant job. This is the trigger that introduces intermediaries into the model. Without this element there is no basis for firms to manage future decisions to use intermediaries. However, using *only* this motivation almost no firms use intermediaries in the simulation. This really just serves to introduce the use of intermediaries into the model, so that firms can calculate a preference between indirect and direct hires. The first time a firm looks for an intermediary they look within a local radius, choosing the one that has the best job-worker match rate. Firms ignore the quality of the contractors' matches and their fees when choosing an intermediary, but do consider it in their utility equation, which is later used to determine whether the firm will use an intermediary for another job. Once a firm has experience with an intermediary, they turn over the next open job to an intermediary when their past utility from indirect hires is greater than their past utility from direct hires. Utility is specified two ways, both outlined in detail in the appendix. In both cases the firm assesses its utility using a weighted history function, weighing its more recent hires more heavily than the older. In

[10] Contractors themselves are set up as simply as possible with exogenously set fee rates no realistic profits, and so forth. The model concentrates on the trade of between their arbitrarily set fees (assuming that their fees are not exactly equal to their marginal productivity in the labor market) and their matching efficiency. For a mathematical model with more detailed intermediaries, one might see (Neugart and Storrie, 2006).

the first experiment half of the utility is match quality. Match quality measures the distance between the job's skill and the worker's skill. The optimal match is when the two skill levels are equal. Dissatisfaction is asymmetrical such that the firm would rather have an overqualified worker than an equally under-qualified one. The second half of utility is related to the cost of using an intermediary, or the match fees, which are measured as a percent of the worker's skill. The second experiment uses the same concept where half the utility is related to cost and half is related to match quality. However, in the second experiment costs includes not only fees, but an adjustment for the indirect worker either having a higher or lower salary. Fees are charged on the adjusted salary.

Firms use the same contractors until they either bring their last indirect job in-house or the contractor goes out of business. When the contractor goes out of business, the firm finds a new contractor the same way they found the first one, and continues to use the utility from the prior contractor in historical utility calculations. If the firm cannot find a contractor, they hire directly. Finally, indirect hires that have been at the same firm for more than four periods automatically become direct hires.

The algorithm description included several underlying assumptions in the model's mechanics, some of which are varied in the following experiments. The first assumption is that firms are more likely to fire indirect hires than direct. This setting is premised on the fact that research finds that atypical workers are more likely to transition to unemployment than traditional hires (Corsini and Guerrazzi, 2007; Amuedo-Dorantes et al., 2006; Garcia-Perez and Munoz-Bullon, 2005) and that most OECD countries have stronger limitations on dismissing permanent employees than indirect or short term employees (OECD, 1999, 2003). According to the national organizations survey, firms also report using indirect workers as an adjustable labor force, firing them when demand declines. The second assumption is that when indirect workers are not fired, after some period they must become permanent employees. Temp workers are usually sent to a new assignment, or if they stay on at the same firm, they are hired permanently. For this reason, temp contracts often have clauses specifying a fee the firm will pay the temp agency if they hire the worker. This model is designed to reflect that temp workers cannot remain in the same job as a temp worker indefinitely. They must switch jobs or become permanent hires. The third assumption is that high-skill workers conduct a broader job search than low-skill workers.This assumption is supported by evidence showing that the longer a worker is unemployed (generally less skilled workers), the less likely he/she is to relocate for a job (Herzog et al., 1993) and is also the premise of the spatial mismatch literature that argues that poor inner city workers do not search for or find employment in the suburban ring (Kain, 2004). Further, low-skill workers are generally found to move less for jobs. Fourth, throughout all experiments the utility calculation assumes that firms value match quality, and like to pay less for their workers, both in fees and wages. Finally, in terms of match quality, the model assumes that both workers and firms prefer to find a job/worker that is a perfect match, but that a worker would rather be somewhat under-qualified for his job (this assumes some sort of ambition on the part of the worker), while the firm would prefer a slightly over-qualified worker.[11] The assumptions listed here are not tested, as they are based on a combination of the sociological literature and common sense.

[11] Note that the firms' preference for over rather than under-qualified workers also enters the firm's utility function. Firms' utilities (though not their hiring decisions) are also influenced by wages. Wages are based on the worker's skill (rather than the job's or an average of the two). This means that when the firm hires someone, it just wants a perfect match and prefers the over-skilled to the under-skilled, but when they calculate their future decisions to use intermediaries, they still value match quality the same way, but are also happier with lower skill (cheaper) workers.

4 Experiment One: Fees vs match quality

The first experiment examines a trade-off that firms face between intermediaries' ability to better sort through workers and the fees they charge for their services. The experiment finds that non-wage incentives can be sufficient to encourage the use of intermediaries. There are two parameters controlling these dynamics. One parameter controls intermediaries' ability to screen more workers. This parameter makes intermediaries' jobs appear closer to workers, thus more workers apply to these jobs, and the intermediaries are more likely to make better matches. The parameter ranges from 1.0 to .1 (10 settings), where the worker sees a contract job with the same probability as a regular jobs at 1.0 and where the contract job appears twice as close at .5. To avoid confusion, the reader should be reminded here that this distance effect is a proxy for the breadth of the worker's job search and that contractors' ability to "shorten" the distance measures their ability to sort through more applicants. The second parameter sets contractors' fees, ranging from 5 to 35% of the employee's skill level (7 settings). All combinations of the parameters were run (70 combinations) 20 times each. The consequent transition to using intermediaries under three parameter settings is illustrated in Figure 4 and the final proportion of jobs filled through intermediaries at all parameter settings is illustrated in figure 5. Considering standard error (not depicted in figure 4) the levels are statistically different for the displayed runs after the 100th tick (a model's "time" element). Figure 4, shows that the transition to using intermediaries is rather abrupt; as soon as intermediaries are available, firms rapidly adopt. When intermediaries are most appealing (with lower fee rates and greater search radius enhancement) the transition is quicker.

INSERT FIGURE 4 HERE

The contour plot, figure 5, illustrates the final proportion of jobs matched through intermediaries at all parameter levels. Higher fees discourage firms from using intermediaries while their ability to screen more applicants increases it. Even when there is a low search radius effect and high average fees, firms still use intermediaries over 25% of the time in the first experiment. This comes about for two reasons. First, fees are assigned from a distribution. This means that even if on average fees are high, there are some intermediaries that are cheap. Second, firms base their current decisions on their personal utility histories. Even with no systematic advantage, intermediaries will create better matches 50% of the time. Thus, even when the overall model settings are not advantageous to intermediaries, individual firms can have positive experiences with intermediaries. This is realistic in that firms often make myopic decisions based on their experiences and even though a service might, on average, not be advantageous. Figure 5 suggests that the level of workers hired indirectly at the end of the model decreases steadily as intermediaries' fees go up and increases as intermediaries' search ability increases. As parameter settings move towards the most attractive intermediary scenarios (with low fees and good matching) there are two pockets of higher levels of indirect hires and a small pocket of low levels, though overall the relationship is monotonic with the pockets not deviating more than 1.5 percentage points from the surrounding area.

INSERT FIGURE 5 HERE

Other model outputs behave as anticipated. Firms' average utility increases when they use intermediaries, increasing 37% with medium search enhancement and medium fees or 100% with the lowest fees and highest search enhancement. Utility gains are not monotonically related to parameter settings, with a contour plot (not illustrated here) showing some discontinuities in utilities when radius enhancements are low but fees are in the middle range. Using intermediaries is also associated with lower unemployment, primarily because they give hard-to-match jobs and workers a better chance to find a match. An OLS regression predicting the simulations' unemployment rates suggests that the effect is significant: when there is a nine point increase in the percent of jobs filled through intermediaries, unemployment declines almost a full percentage point. With respect to skill, contractors are more likely to match slightly less skilled workers, although the model design suggests that intermediaries should match workers at either extreme of the distribution, as both will have trouble finding a good match.

There were several important findings from this experiment. First, when intermediaries have better matching capabilities, they improve firms' utilities and reduce frictional unemployment. Second, match quality is a sufficient incentive, in the face of significant intermediary fees, for firms to use intermediaries.

5 Experiment Two: Occupation-specific incentives

While the first experiment focused on firms' decisions to use contractors in the absence of compensation incentives, the second experiment looks at how incentives could vary across occupations. This experiment leaves the contractors' fees constant, and instead allows salaries through contractors to be higher or lower. This form allows the firm to save money through intermediaries.

This model was run with seventy-two permutations of parameter settings, with each setting run 20 times each, for a total of 1,440 simulations. The other parameters such as the grid size and the number of agents are the same as in experiment one. Four cases representing four of the 72 experiments are highlighted in table 2. These four are highlighted, as they are hypothesized to approximate the empirical situation in four occupational labor markets.

The first parameter that was varied is the skill distribution of the workers and the jobs in their labor market. This had four settings, all taken from the Bureau of Labor Statistics and CPS data for programmers, accountants, minimum wage workers, and the overall labor market. The difference between a broad skill distribution like "all labor market" and a narrow skill distribution like "accountants" capture the importance of credentialing. The second parameter indicates the wage premium (or penalty) for an indirect hire and is based on the empirical wage gap between indirectly and directly hired workers. It has three settings: higher, lower, or the same wages. The third parameter that is varied is the variance of the wage premium which also has three settings. Empirically, the standard deviation of wages is highest for high-skill occupations and lowest for low-skill. This is because the labor market is far-right skewed. College graduates can earn average or *extremely* high wages, while high school graduates are concentrated on the narrower lower portion of the wage spectrum. Because of the skewed shape of the wage distribution, we would expect the standard deviation of log wages to be about the same for the two groups, though in CPS data the low skill workers had a slightly higher standard deviation, possibly due to top-coding the highest wages. The last parameter, work variability, was tested at just two levels: high and low. Less skilled occupations have more variable work hours on an individual level (i.e. low skill workers can work 20 or 50 hours a week), while high-skill jobs have consistent individual hours (a little over 40 hours per week) (BLS, 1988-2004).

INSERT TABLE 2 HERE

Hypotheses and experimental results for 4 of the 72 simulations are highlighted in table 2. I hypothesize that workers in the "minimum wage" scenario (with the "minimum wage" skill distribution, the higher compensation savings setting, and more firm shocks) are the most likely to be hired through intermediaries because the model is set such that firms should save more on their compensation, because the wide skill distribution among these workers means that there are more opportunities for intermediaries to improve match quality, and a high variability in firms' demands for workers should prevent the indirect hires from transitioning to direct hires. In contrast, firms hiring accountants might use intermediaries less since there is empirically less of a difference between indirect and direct hires' wages; there is a narrow skills-distribution, so firms are capable of finding a good match without an intermediary; and there might be a relatively constant demand for accountants. Generally, firms should realize greater utility gains from intermediaries in professions with wider skill distributions, larger wage gaps, and more firm work shocks. The last column of the table shows the mean proportion of the labor force that was employed through an intermediary at the end of the 20 runs for each of the four combinations of parameter settings. While the reader might disagree with the hypothetical scenarios displayed in table 2, *every possible scenario* was run 20 times, and the relationships between variables are assessed for *all* runs using multivariate methods, presented below.

INSERT FIGURE 6 HERE

The simulations' predicted levels of indirect hires for these four scenarios, depicted in figure 6, bear out the theoretical expectations. The model finds that firms transition to using intermediaries more quickly for low-skilled occupations. While standard deviations are not depicted in figure 6, there is not a significant difference between the overall labor market scenario and the minimum wage scenario except from approximately tick 100 to 200. There is a significant difference between the 2 scenarios using more intermediaries (all labor market and minimum wage) and the 2 least scenarios using fewer intermediaries (programmers and accountants) throughout the entire run. Programmers and accountants have almost significantly different rates of intermediary use throughout the model (significant at .10) with programmers being hired through intermediaries more, probably because accountants are set with a narrower skills distribution and less workload variability. In addition, it is more expensive to use intermediaries for accountants because they have slightly higher skill-levels so intermediary fees are higher. As in the first experiment, firms rapidly adopt intermediaries and then hover around a final stable level of indirect hires, past the 600 time ticks shown here. Across all scenarios, the proportion of workers hired through intermediaries ranges from 20% to 52%.

INSERT TABLE 3 HERE

A simple multivariate analysis predicting the proportion of jobs filled through intermediaries across all runs of the simulation shows that all the parameters except the compensation variability are significant at the .001 level, while compensation variability is significant at the .01 level. Excluding the predictive power of "time" (ticks) in the model, these parameters of interest explain only 22% of the experimental variance; with "time" they explain 69%. Figure 3 shows the OLS regular and standardized coefficients for a multivariate regression predicting the proportion of indirectly employed workers based on the experiment's parameter settings. Of course model time has the most impact; and is followed by the skill distributions, then workload fluctuations, then cost, and finally cost variance. Most of the results make intuitive sense. The more heterogeneous skill distributions promote intermediary use while more narrow distributions like accountant limit it. When workloads fluctuate there is more turnover both giving firms more opportunities to hire through intermediaries, and preventing indirect hires from moving into full-time positions. The only unexpected result is that higher wages for indirect workers ("relative comp" *increase* the proportion of workers hired indirectly. While this seems odd from the perspective of the firm's utility equation, it makes sense when we consider organizational ecology, or the fact that intermediaries have to make a profit to offer their services.

Organizational ecology matters because when workers hired through intermediaries earn more money, so do the intermediaries. At lower costs firms might want to use intermediaries, but cannot find them. The intermediaries' profits are driven by an interaction between the worker's skill and the compensation differential, with them earning the most from a high-skill worker with a contract wage premium and the least from a low-skill worker with a contract wage penalty. If we recall, intermediaries are forced out of business if their revenues per assigned job are less than 10% of the average skill level in the model. In a run of minimum wage workers, where many workers have a high school education, the average contractor who matches all of his clients under a wage penalty scenario, can still expect to earn about 17% of the average skill level in that occupation, well above the minimum requirement. But if the intermediary sets a slightly lower fee rate (remember fee rates are set by a random distribution around an average level), or their workers draw slightly worse penalties, or if they fail to match some jobs, they go out of business. When they go out of business, the firms working with them have to look for another local contractor and if none is available, they hire workers directly. Thus, the firms don't use intermediaries simply because they are cheaper, but also when they are expensive enough to sustain a healthy organizational ecology.

This model excludes many real world dynamics that would also influence the predicted level of indirect hires. For example, if we were to include firm specific-skills or employer loyalty, there would be less turnover, longer tenure, and thus there would be fewer indirect hires since indirect hires would have a better chance to move into direct employment. It would also increase the incentive to move to direct employment, since indirect hires would be less productive and paid less. This exclusion might be one reason that the model's predicted level of indirect employment exceeds empirical levels.

As in the first experiment, the model also produces other experimental labor market statistics including: unemployment rates, vacancy rates, job duration, and vacancy duration. These measures are used to verify that the model is a reasonable approximation of the labor market. The model has unemployment and vacancies fluctuating around 5% regardless of parameter settings or the time in the model run. Since most jobs are at a few firms, this means that the overall unemployment and vacancy rates are autocorrelated, and generally resemble a real-world labor market having business cycles of high unemployment and low unemployment rather than random noise in the unemployment rate.

In terms of skill level, the experiments find that unemployed workers are consistently the least educated, then indirect hires, and then direct hires are the best educated. In the model, one might expect the *mean* skill level of indirect hires would not differ from that of direct hires since firms should use intermediaries at both extremes of the skill distribution since the least-skilled workers should receive fewer job offers and the most educated should be the least likely to receive an offer filling their minimum requirements. Thus, both the most and least educated should be the hardest to match. On average, these characteristics should balance out leaving no difference between direct hires and indirect hires. While the differences in figure 7 look small; they are statistically significant, with indirect hires having significantly less skills, and direct hires having more. This is the second unexpected and interesting finding: intermediaries encourage firms to sort workers, keeping the best inside the firm.

INSERT FIGURE 7 HERE

A multivariate regression comparing the different parameters' effects on the skills gap between indirect and direct hires suggests that when indirect employees are paid more relative to direct hires, the skill gap increases. Second, the regression finds that when firms have more workload shocks there is a lower skill gap (a low rate of shocks increases the skills gap by about 1 degree level or the difference between an MA and a BA). This occurs because the low-skill workers who would stay on and transition into direct hires (thus diminishing the skill difference) are the first to be fired in a volatile market. Third, using the all labor market distribution results in the smallest skill gap, while accountants have the highest skill gap. This is unexpected because, of course, accountants have the narrowest distribution of skills among workers.

Firm utility is one of the most important measures, as it motivates all the model's dynamics. Utility ranges from 0 to 1, with the experiments including contractors increasing utility on average .05 points compared to a baseline model with no contractors. A multivariate regression predicting utility levels based on parameter settings suggests that high contract premiums and more workload shocks increase firms' utilities as does using the programmer or accountant skill distributions (because they are narrower). The premium's role is attributable to the aforementioned organizational ecology effect. Workload shocks increase utility because they increase turnover, which in turn increases the chance the firm makes a better match and better matches are more likely to survive since the worker is less likely to quit. Thus, turnover is good for a firm in the simulation since the model ignores the importance of on the job training and firm-specific skills. It is surprising that the narrower skill distribution increases utility since utility should improve through intermediaries more in the broader skill distributions. Other model output includes the unemployment rate, vacancy rate, average vacancy duration, and average unemployment duration. All roughly matched the labor market, with unemployment hovering around 5% and most unemployment spells being short, but with a few chronically unemployed.

The model's primary limitation is that as a model, it excludes many incentives and makes many assumptions. This model focuses on the trade-off between the cost of using intermediaries and their better matching ability, omitting dynamics like firm-specific skills and worker substitutability. The model showed that there are strong incentives for firms to use intermediaries even in the absence of wage premiums; that wage premiums and fees can have unexpected effects because of organizational ecology; that firms are more likely to use intermediaries to fill their least-skilled jobs (stratifying workers by contact type); and that even when, in expectation, intermediaries are not advantageous, some firms will persist in using them, misinterpreting natural variability as a systematic advantage.

6 Empirical trends and model validation

A significant amount of research has measured the number of outsourced workers (presumably indirect-hires), the types of firms most likely to out-source, and the number and size of firms offering job-matching services. Generally, studies using employer data estimate higher growth rates for atypical employment than those using employee data. I examine three measurements of the trend of using indirect hires in the US economy, and use one to validate the model.

The Economic Census is a survey of businesses conducted by the US Census Bureau every five years. It collects information on firms' industries and employment. Figure 8's panel A shows that between 1997 and 2002, employment at companies providing contract services grew more rapidly than the rest of the economy. (This measure includes direct hires at the service companies, but assuming administrative costs are a constant proportion of staff, this does not effect the data.) It is somewhat ambiguous whether the workers at these firms are actually indirect hires. At a janitorial services firm it is very likely that the workers perform their duties at clients' sites while in accounting firms it is less likely. This data also does not provide information about the total number of workers in an occupation and only has information for two periods. Thus, we can only estimate rate of employment growth for the workers at the service firms, not the growth rate for the proportion of an occupation that is hired through intermediaries. As such, the data is relevant to the model, but not directly comparable.

INSERT FIGURE 8 HERE

There are two measurements approaching the question from employee-side data that could be used for validation. The first technique uses the March Current Population Survey (CPS) to construct a time series of the proportion of an occupation that is indirect hires. Indirect workers are identified by matching occupational and industry codes, positively identifying any worker in an occupation working for a firm specializing in providing that same occupational service. Figure 8's panel B shows the proportion of workers in a given occupation who are indirect employees using this method.[12] In contrast to the employer side data, this data suggests a much slower growth pattern, and even a slight decline in the clerical sector in recent years. This method is superior because it incorporates the general growth rate of an occupation and can estimate the proportion of the occupation hired indirectly. The method is limited in that it can only identify indirect workers when there are matching codes for occupations and industries (i.e. clerical workers and clerical services) and that it misidentifies direct-hires working at an intermediary in the same occupation that they rent out labor in (i.e. an accountant working in an accounting firm). As with the employer-side data, the method is likely to misidentify high-skill workers who actually work at the contractor's site.

The third technique uses the CPS Contingent Worker Supplement (February 1995, 97, 99, 2001, and 2005) which directly asks workers about their employment status. This method counts temporary workers, on-call, casual laborers, day laborers, or any worker reporting that their employer leases out their services. Figure 8's panel C shows that for programmers, accountants, janitors, and clericals, this method suggests the opposite trend as the two prior techniques, suggesting employment through intermediaries has actually declined since 1997. In panel C, the level of intermediary use for all occupations is also shown, although this is a very small percent of the total workforce and one should note the small scale on the y axis.

Why do the estimates vary? Theoretically, the method using the Contingent Work Supplement should be the most accurate since there is no proxy measurement; the survey directly asks the workers about their employment

[12]CPS occupational and industrial codes changed in 1992 and 2002, leaving a slight discontinuity.

status rather than inferring it. However, individuals often misreport the firm they commute to as their employer rather than the intermediary, who is actually their legal employer (Bjelland et al., 2006). This bias is confirmed by examining a single question from the same survey. Early in the survey, the worker is asked to report his or her employer. Later, the respondent is asked whether they were paid by their employer or a temporary help agency. If they were paid by a temporary help agency, the interviewer then asks them whether their reported employer was the agency or the agency's client. Surprisingly, the majority of respondents report the client as their employer. Figure 9 shows that this bias has gotten worse over time. This distorts measurements using both the inferred and direct CPS estimates since the indirect method relies on workers accurately reporting their employer's industry.

INSERT FIGURE 9 HERE

The empirical evidence on intermediary use is of mixed quality (figure 8). Nevertheless, the data from the CPS was compared to the experimental data. The first 200 ticks (or time periods in the model) were discarded as "burn in."[13] Across all the models the average level of workers hired through intermediaries, as of the end of the model, was 36%. This far exceeds the level found in the empirical data, indicating that incentives are set too high, or important disincentives are omitted. The difference in occupations found in the model is different than the empirical data. In the model the low skilled workers in wider skill distributions seems to be more likely to work in indirect

[13] Originally, the diffusion of intermediary use in the model from tick 200 on was fit to the time series from the CPS data from 1983 to 2005, with the assumption that hiring through intermediaries really began in the 1980's. The fits are not presented here primarily because the fitting of the time axis (linking tick 200 to 1983 and the last tick to 2005) was arbitrary. The experimental and empirical data matched for some experiments such as the prediction of programmers' indirect hire rates using those experiments with the programmers' skill distribution, no indirect hire wage premium, a low variance of the wage premium, and few firm shocks. Nevertheless, the same experiment fits the growth in indirect employment for clerical workers and janitors well too.

employment while in the empirical data programmers and accountants seem to. I would speculate that this mismatch is largely driven by miscalculation in the empirical data. The empirical data calculation includes, for high skill jobs, a large portion of workers who have regular jobs at service firms (i.e. accountants at accounting firms). In sum, the overestimation in the model is probably due to the omission of more indirect hiring disincentives than incentives while the mis-ordering of occupations is likely due to the empirical measurement problems.

While the predicted levels of indirect hiring do not exactly predict reality, it is first, unlikely to do so in a simple model and second, could easily be remedied by adjusting the magnitude of the effects of tested incentives and disincentives on firms' utilities. It is perhaps more important that the model generally functions like a real labor market. One way to test that is using the unemployment and vacancy rates in the model. In the model, across simulations the mean unemployment rate is 5.3% while vacancies are 5.2%. These are relatively realistic values; unemployment in the US was 5.5% as of May, 2008.

Another way to compare this simulated market to the real market is using the relationship between unemployment and vacancies (Fagiolo et al., 2004). Unemployment rates and vacancy rates are hypothesized to follow the "Beveridge Curve," named after the English Economist William Beveridge. The Beveridge curve is a curve relating unemployment rates (x-axis) and vacancy rates (y axis) that has a negative first derivative and a positive second, so that it is concave, with vacancies approaching zero as unemployment increases and unemployment approaching zero as vacancies increase. Empirically we never know the exact shape of the curve, as it slowly migrates in towards the origin and away from the origin as worker-job matching becomes more and less efficient, respectively. Nevertheless, the curve is shown to exist as employment and vacancy rates move in a counterclockwise fashion following the theorized concave curve. During the 1990s and 2000's the Beveridge curve wandered inward towards the origin, in the general range of unemployment around 4 to 8 % and vacancies around 2 to 4% (Valletta, 2005).

Unlike using real data, the ABM can generate full Beveridge curves because multiple experiments are run under each regime. Using data from all experiments, at each combination of parameter settings, at multiple time points (ticks), for those observations beyond the 200th tick, we can generate Beveridge curves from the ABM. In figure 10 we see in the lower left hand corner, the plot of the natural log of unemployment versus vacancy from the experiments, suggesting that the natural log of unemployment might be a good fit to predict vacancies, a functional form that naturally fits the Beveridge curve. Next, a model of $v = c + \alpha ln(u) + \beta X$ (where u is unemployment, v is the vacancy rate, X is the vector of experimental conditions, and c is a constant) was fit, finding the significant coefficient of -.0359 for ln(unemployment) with an R-square of .82. (The R-square just using ln(unemployment) is .20) Among the experimental conditions, the skill distribution and the firm shocks also change the prediction. The upper panel of figure 10 shows the ABM's combinations of vacancy and unemployment rates and the predicted values. Even though ID's were omitted from this diagram for clarity, it is clear that the experimental and predicted values mirror each other. The bottom right hand plot shows this by plotting residual against the predicted vacancy rates- showing a good fit. The most efficient markets (with a Beveridge curve closest to the origin) in the simulation are those simulations with the accountants' skill distribution and more firm shocks. This makes sense because it is easier to match workers and jobs in a market.

It also suggests that the market is more efficient in a more fluid labor market where firms are often hiring and firing people. The least efficient simulations are those using the all labor market distribution and fewer firm shocks- the polar opposite scenario.

What this exercise shows, is that the simulated labor market is more or less functioning as a normal labor market does. Unemployment and vacancies are not related linearly, but in a concave curve in the shape of the Beveridge curve. In sum, both the absolute levels of unemployment and vacancies, and their relationship, in the simulation, mirror that which we see in the real world, suggesting that the model does a reasonable job of copying the real labor market.

INSERT FIGURE 10 HERE

While the estimates of intermediary use in the model do not match the empirical data, more importantly, its other characteristics such as employment and vacancies resemble a real labor market.

7 Conclusion

The agent based model had several important findings. The first finding is that independent of compensation differentials, intermediaries' better ability to match workers and jobs is a sufficient incentive for firms to use intermediaries. The second finding is that organizational ecology matters, and consequently higher intermediary fees can increase firms'propensity to use intermediaries by sustaining an organizational ecology of contractors. Another finding is that a percentage fee structure turns indirect employment into a sorting mechanism, where firms hire their less skilled workers through intermediaries. Finally, the model found that in occupations with more heterogeneity among workers, firms are more likely to use intermediaries.

The empirical data suffers many flaws, but did suggest some conclusions and directions for further research. First, using intermediaries is not necessarily an exploding trend; this could be a misperception based on using employer-side data. Second, the data suggests that the major transition has most likely already happened and that intermediary use seems to have stabilized in the 1990s. Finally, it is difficult to make empirical conclusions because the definition of an indirect hire is unclear (particularly for high-skilled jobs) and because workers in these arrangements misreport their employers.

This paper was limited in that it did not explore the role of other incentives to hire through intermediaries, like firm-specific skills and tested a context with legal limitations on indirect employment duration. Finally, the model used compensation gaps calculated from a simple method of taking the difference between indirect workers' wages and direct workers wages rather than controlling for selection on basic worker characteristics like age, race, education, and location. As such, the simulation may benefit from a better measure of the compensation gap.

Atypical employment is an important area of research. If there are trends towards using intermediaries with strong repercussions for compensation gaps, and these trends sort workers by ability, we are moving towards a two-tier system of employment with one group of workers enjoying the benefits of direct employment and another group suffering the penalties of indirect employment. Many European countries have moved towards guaranteeing these types of workers equal compensation and union negotiating power in an attempt to combat this labor market segmentation. This same question is equally if not more important in the United States, where employer-provided health benefits are at stake.

References

Abraham, Katherine. 1990. *New Developments in the Labor Market: Toward a New Institutional Paradigm*, chapter Restructuring the Employment Relationship: The growth of market-mediated work arrangements, pp. 85–119. MIT Press.

Amuedo-Dorantes, Catalina, Miguel A. Malo, and Fernando Munoz-Bullon. 2006. "The Role of Temporary Help Agencies in Facilitating Temp-to-Perm Transitions." Technical report, IZA discussion paper.

Autor, David. 2000. "Outsourcing at Will: Unjust Dismissal Doctrine and the Growth of Temporary Help Employment." Technical Report 7557, NBER.

Benson, John. 1999. "Outsourcing Organisational Performance and Employee Commitment." *Economic and Labor Relations Review* 10:1–21.

Bjelland, Melissa, John Haltiwanger, Kristin Sandusky, and James Spletzer. 2006. "Reconciling Household and Administrative Measures of Self-Employment and Entrepreneurship." Technical report, CES census.

BLS. 1988-2004. "Industry Productivity and Costs." Technical report, Bureau of Labor Statistics.

Booth, A.L., M. Francesconi, and J Frank. 2002. "Temporary Jobs: Stepping stones or dead ends." *Economic Journal* 112:F189–F213.

Clinton, Angela. 1997. "Flexible labor: Restructuring the American Work Force." *Monthly Labor Review* 120:3–27.

Corsini, Lorenzo and Marco Guerrazzi. 2007. "The Transition from Temporary to Permanent Employment: Evidence from Tuscany." *Labour* 21:303–332.

Davis-Blake, Alison and Brian Uzzi. 1993. "Determinants of Employment Externalization: A Study of Temporary Workers and Independent Contractors." *Administrative Science Quarterly* 38:195–223.

Deavers, Kenneth. 1997. "Outsourcing: A Corporate Competitiveness Strategy, Not a Search for Low Wages." *Journal of Labor Research* 18:503–19.

Echenique, Federico and Mehmet B. Yenmez. 2005. "A Solution to Matching with Preferences Over Colleagues." Technical Report 120, Fondazione Eni Enrico Mattei.

Fagiolo, Giorgio, Giovanni Dosi, and Roberto Gabriele. 2004. "Toward an Evolutionary Interpretation of Aggregate Labor Markets Irregularities." Technical report, Laboratory of Economics and Management Sant' Anna School of Advanced Studies.

Gale, David and Lloyd S Shapely. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Quarterly* 69:9–15.

Garcia-Perez, J. Ignacio and Fernando Munoz-Bullon. 2005. "Temporary Help Agencies and Occupational Mobility." *Oxford Bulletin of Economics and Statistics* 67:163–180.

Gramm, Cynthia and John F Schnell. 2001. "The Use of Flexible Staffing Arrangements in Core Production Jobs." *Industrial and Labor Relations Review* 54:245–258.

Granovetter, Mark S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78:1360–1380.

Herzog, Henry W, Alan M Schlottmann, and Thomas P Boehm. 1993. "Migration as Spatial Job-Search." *Regional Studies* 27:327–340.

Hosios, Arthur J. 1990. "On the Efficiency of Matching and Related Models of Search and Unemployment." *Review of Economic Studies* 57:279–298.

Houseman, Susan, George A. Erickcek, and Anne L. Kalleberg. 2003. *Low Wage Amerca: How Employers are Reshaping Opportunity in the Workplace*, chapter The Effects of Temporary Services and Contracting Out on Low-Skilled Workers: Evidence from Auto Suppliers, Hospitals, and Public Schools. Russel Sage Foundation.

Kain, John. 2004. "A pioneer's perspective on the spatial mismatch literature." *Urban Studies* 41:7–32.

Kalleberg, Arne L. 2000. "Flexible Firms and Labor Market Segmentation." *Work and Occupations* 30:154–175.

Kalleberg, Arne L, Barbara F. Reskin, and Ken Hudson. 2000a. "Bad Jobs in American: Standard and nonstandard employment relations and job quality in the United States." *American Sociological Review* 65:256–278.

Kalleberg, Arne L, Barbara F. Reskin, and Ken Hudson. 2000b. "Nonstandard Employment Relations: Part-time, temporary, and contract work." *Annual Review of Sociology* 26:341–365.

Katz, Lawrence F, Alan B Krueger, Gary Burtless, and William T Dickens. 1999. "The High–Pressure U.S. Labor Market of the 1990s." *Brookings Papers on Economic Activity* pp. 1–87.

King, Miriam, Steven Ruggles, Trent Alexander, Donna Leicach, and Matthew Sobek. 1983-2006. "Integrated Public Use Microdata Series, Current Population Survey: Version 2.0." Technical report, Minnesota Population Center.

King, Miriam, Steven Ruggles, Trent Alexander, Donna Leicach, and Matthew Sobek. 1995-2005. "Integrated Public Use Microdata Series, Current Population Survey: Version 2.0." Technical report, Minnesota Population Center.

Macy, Michael W. and Robert Willer. 2002. "From Factors to Actors; Computational Sociology and Agent-Based Modeling." *Annual Review of Sociology* 28:143–66.

Mangum, Garth, Donald Mayall, and Kristin Nelson. 1985. "The Temporary Help Industry: A response to the dual internal labor market." *Industrial and Labor Relations Review* 38:599–611.

Mayall, Donald and Kristin Nelson. 1982. "The Temporary Help Supply Service and the Temporary Labor Market." Technical report, Office of Research and Development Employment and Training Administration US Department of Labor.

Montgomery, James D. 1991. "Equillibrium Wage Dispersion and InterIndustry Wage Differentials." *Quarterly Journal of Economics* 106:163–179.

Neugart, Michael and Donald Storrie. 2006. "The emergence of temporary work agencies." *Oxford Economic Papers* 58:137–156.

OECD. 1999. *Employment Outlook*. OECD.

OECD. 2003. *Employment Outlook*. OECD.

Osterman, Paul. 1999. *Securing Prosperity*. Princeton University Press.

Peters, Michale. 1991. "Ex Ante Price Offers in Matching Games: Non-Steady States." *Econometrica* 59:1425–1454.

Pfeffer, J and J. N. Baron. 1988. "Taking the Workers Back Out – Recent Trends in the Structuring of Employment." *Research in Organizational Behavior* 10:257–303.

Richiardi, Matteo. 2003. "A Search Model of Unemployment and Firm Dynamics." Technical report, LABORatorio r. Revelli Centre for Employment Studies.

Segal, Lewis M and Daniel G Sullivan. 1997. "The Growth of Temporary Services Work." *Journal of Economic Perspectives* 11:117–136.

Shi, Shouyong. 1998. "Unskilled Workers in an Economy with Skill Biased Technology." Technical Report 73, CREFE.

Stovel, Katherine and Christine Fountain. 2003. "Hearing About a Job: A Model of Differential Information Flow and Job Matching." Technical Report 30, Center for Statistics and the Social Sciences, University of Washington.

Tassier, Troy and Filippo Menczer. 2001. "Emerging Small World Referral Networks in Evolutionary Labor Markets." *IEEE Transactions On Evolutionary Computation* 5:482–492.

Tassier, Troy and Filippo Menczer. 2005. "Social Network Structure, Segregation, & Equality in a Labor Market with Referral Hiring." Technical report, Fordham University.

Tesfastion, Leigh. 2001. "Structure, Behavior, and Market Power in an Evolutionary Labor Market with Adaptive Search." *Journal of Economic Dynamics and Control* pp. 419–457.

Valletta, Robert G. 2005. "Why has the US Beveridfe Curve Shifted Back?" Technical report, Federal Reserve Bank of San Fransisco.

Young, Suzanne and Johanna MacNeil. 2000. "When Performance Fails to Meet Expectations: Managers' Objectives for Outsourcing." *Economic and Labor Relations Review* 11:136–168.

8 Technical Appendix

Parameter list

The first part of the list includes the parameters which were not tested in their entire parameter space. The second part of the list includes the 5 parameters that were tested. Many in the first list (like model length or grid size) are arbitrary and do not influence the model outcomes (they were not tested rigorously, but were varied in a few trial runs). Some of these parameters, like the distribution of jobs across firms, are based on specific empirical US data, while others, like the search radius of a worker, are more loosely based on empirical research (i.e. studies find that skilled workers look for jobs in a broader radius.) Skill floors, the continual generation of contractors, and the contractor's startup grace period were tested and found to have no effect on the model's findings, so the various tested parameter settings are not shown here.

Parameter	Definition	Default Value
stopTicks	length of run	600 or 1000
numWorkers	number of worker agents	1000
numFirms	number of firm agents	142
sizeX sizeY	grid size	100
feeRVar	variance for contractor fee rates	.05
maxCDistance	firms' search radius for contractors	.2
ξ	exponent distributing jobs across firms	2.1
γ	history weighting	.75
tPerm	contract worker's transition to direct hire	5
rDeath	revenue a contractor must maintain	.1
vSContractors	vacancy rate generating contractors	.04
oSContractors	outsourcing rates generating contractors	.02
cSTime	contractors' startup period to generate revenue	3
ceiling	a ceiling on unemployment and vacancy	.15
sSearchRWorker	worker's skill effect on search radius	5
maxWSTolerance	maximum deviation for job floor	3
minWSTolerance	minimum deviation for job floor	1
maxJSTolerance	maximum deviation for worker floor	3
minJSTolerance	minimum deviation for worker floor	1
hWeighting	weights firm's utility histories	.75
fRFloor	a floor on contractors' fee rates	.025
vDisutility	disutility for firms for vacancies	$-.1$
settings for tested parameters in those experiments in which they were not varied		
fRMean	contractors' mean fee rates	.2 (exp 2)
		.05 to .35 (exp 1)
cRWorker	contractors' mean search radius	.5 (exp 2)
		.1 to 1.0 (exp 1)
wSDist	workers' edu distribution	all labor (exp 1)
	all, accountant, programmer, min wage	(exp 2)
jSDist	jobs' edu distribution	all labor (exp 1)
	same as wsDist	(exp 2)
cAlphaMean	contracting's effect on compensation	NA (exp 1)
		.9 to 1.1 (exp 2)
cAlphaVar	variance of above	NA (exp 1)
		.05 to .15 (exp 2)
wVar	firms' workload fluctuations	.05 (exp 1)
		.05 to .1 (exp 2)

In the first experiment feeRateMean and contractorRadiusWorker were varied
In the second worker experiment SkillDist, jobSkillDist, contractedAlphaMean, contractedAlphaVar, and workVar were varied .

Table 1: Simulation parameters

Classes and their instance variables

- Firms have:
  - X and Y locations
  - a list of their jobs

  - a list of their vacant jobs
  - a list of their employees
  - a change in workload (updated each round)
  - a pointer to their contractor
  - a utility (from contracted and direct hires as well as vacancies)

- Jobs have:
  - a pointer to their firm
  - a pointer to the contractor
  - a skill level
  - a skill floor for the least qualified worker they will accept
  - a pointer to their worker
  - the tick the job was last filled
  - the tick the job was last vacated
  - a comparator used to sort workers by how well they match the job
  - a list of unemployed workers, sorted by how well they match the job
- Workers have:
  - x and y locations
  - skill levels
  - a skill floor for the lowest job that they would accept
  - a quit propensity
  - the date they were last employed if currently unemployed
  - the date they were last hired
  - a list of vacant, visible jobs

  - their employer
  - their job
  - an effect on their salaries for a contractor match

- Intermediaries have:
  - x and y locations
  - a list of the firms employing them
  - a list of their assigned jobs
  - a fee rate (a percent of the worker's skill level)
  - the percent of assigned jobs they matched in the last round
  - revenue (based on their fee rate and their employees' skills)

Equations In the notation below normal(x,y) means a draw from a normal distribution with mean x and standard deviation y. Similarly, uniform(x,y) is a draw from a uniform distribution ranging between x and y. The notation uses the following: i indicates indirect hires, d indicates direct hires, w indexes workers, f indexes firms, j indexes jobs, z indexes intermediaries, and t indexes time. Equations refer back to the parameter table 1 as necessary.

- Initial job creation

  For each firm f, draw a number of jobs at the firm. If the number of jobs exceeds 10% if the workforce, redraw. In table 1, ξ is the parameter that distributes jobs across firms.

$$nJobs_f = [1 - uniform(0,1)]^{\frac{-1}{\xi-1}} \tag{1}$$

- Probability of worker w quitting in time t

  Note that variables with no subscript t are drawn just once, during the model setup.

  - Experiment 1

$$pQuit_{w,t} = .333(\rho_{w,t} + \tau_{w,t} + \sigma_{w,j,t}) \tag{2}$$

  - Experiment 2

$$pQuit = .5 * (\rho_{w,t} + \sigma_{w,j,t}) \tag{3}$$

  - For both:

iff $pQuit >$ uniform(0,1), quit

iff $pQuit <$ uniform(0,1), stay

where,

$\rho_{w,t}$	random quits	normal($qP_{w,i}$, .05)
qP_w	quit propensity	uniform(0, .3)
$\tau_{w,t}$	tenure effect	normal(1-$\frac{current job sticks}{total\ life\ ticks}$,.05)
$\sigma_{w,j,t}$	match quality	normal(ψ,.05)
ψ	if $ws_w > js_j$	$\frac{ws_w - js_j}{ws_w}$
ψ	if $ws_w < js_j$	$\frac{(js_j - ws_w)^2}{(js_j)^2}$
ws_w	worker skill	
js_j	job skill	

There are four parameters here not listed in the initial table including the variance of random quits (.05), the variance of the quit propensity (.3), the variance in match quality (.05), and the variance in tenure effect (.05). All are the variance of another main parameter and do not strongly effect the model.

- Fluctuation in the number of jobs at firm f in time t

$$\psi_{f,t+1} = \psi_{f,t} + \Delta * \psi_{f,t} \tag{4}$$

$$
\begin{array}{lcl}
\psi_{ft} & = & \text{firm f's number of jobs in time t} \\
\Delta & = & \text{normal(0,wVar)} \\
\text{if unemployment} > 15\% & & \Delta = |\Delta| \\
\text{if vacancy} > 15\% & & \Delta = -1 * |\Delta|
\end{array}
$$

Note that wVar was a swept parameter. It is listed in table 1. Also note that the rules limiting unemployment and vacancies are a simple proxy for economic dynamics in the real world that hold unemployment and vacancies in a tolerable range, a fact that is empirically observable.

- Worker's and job's skills

  Skill distributions are set based on empirical educational distributions for workers in different occupations. Skill floors are assigned to workers or firms in the beginning of the model and remain constant. The skill floor is a uniform deviation from -1 to -3 plus the worker's or job's skill (the education scale ranges from 1 (less than fifth grade) to 11 (PhD)) to a minimum of 1.

- How firms search for intermediaries

  Firms find the intermediary within a static search radius and pick the one who had the best match rate last round.

- Workers apply to all jobs they "see"

$$P_{wjt} = e^{\frac{-\delta * d_{wjt}}{ws_w * \nu}} \tag{5}$$

P_{wjt} = probability of worker w seeing job j in time t

ws_w = worker skill level for worker w

d_{wjt} = distance between worker w and job j in time t

ν = sSearchRWorker (skill's effect on search radius)

δ = if $\neq$ indirect job = 1

if = indirect job = cRWorker

- The cost of contracting

Firm f's cost of employing worker w in job j though intermediary z in time t is

$$TCC_{w,j,z} = ws_w^{\alpha} * (1 + fRate_z) \tag{6}$$

$WCI_{w,f,z}$ = worker w's cost to firm f using intermediary z

ws_w = worker w's skill level

$fRate_z$ = normal(fRMean, .05) (fee rate from firm z)

$\alpha_{w,z}$ = normal(cAlphaMean, cAlphaVar), for an indirect hire

$\alpha_{w,z}$ = 1 for a direct hire

$\alpha_{w,z}$ = 1 in experiment 1

Intermediary z's fee rate is held constant throughout its life. In experiment 1, the fee rate is varied, but α is held at one. In experiment 2, the fee rate is held at .2. value $\alpha_{w,z}$ measures how much relatively more or less the worker gets paid through an intermediary. This α is redrawn every time a worker is rematched through an intermediary, although the settings for alpha's distribution are held constant through every experiment.

- The decision to use an intermediary:

$$ind = \frac{\beta_{ipast}}{\beta_{ipast} + \beta_{dpast}} \tag{7}$$

iff $ind >$ normal(.5,.2), use intermediary
iff $ind \leq$ normal(.5,.2), hire directly

In experiment 1 a standard deviation of .1 was tested while experiment 2 tested a standard deviation of .2.

$$\begin{aligned}
\beta past_{i,t} &= \gamma * \beta past_{i,t-1} + (1-\gamma) * \beta_{i,t-1} \\
\beta past_{d,t} &= \gamma * \beta past_{d,t-1} + (1-\gamma) * \beta_{d,t-1} \\
\gamma &= \text{history weighting} \\
\beta_i &= \overline{MQ}_i \text{ - averagefeecost}_i \\
\beta_d &= \overline{MQ}_d
\end{aligned}$$

The average fee cost$_i$ is the average fee payment over all workers hired through intermediaries (n_i is the number of indirect workers) $= \frac{\sum feerate_{i,z,j} * ws}{n_i}$

The match quality from the firm's perspective for matching worker w with job j, $MQ_{w,j}$, where js_j is job skill for job j and ws_w is worker skill for worker w, is:

if the worker is underskilled: $1 - \frac{js_j - ws_w}{js_j}$
if the worker is overskilled: $1 - \frac{(ws_w - js_j)^2}{ws_w^2}$

For experiment two the calculations are the same, except the TCC cost function is used instead of the simple fee rate times skill. The means here:

$\beta_{i,t} = \overline{MQ}_i - .5\overline{TCC}_i$ is used instead of $\overline{MQ}_i - averagefeecost_i$

- Utility

Firm utility was calculated as a model output. This is basically the same calculation as used in the decision of whether or not to use an intermediary, except a small negative amount is added for vacancies.

$$U_f = N_{i,f}(\overline{MQ}_i - .5\overline{TCC}_i) + N_{d,f}(\overline{MQ}_d - .5\overline{TCC}_d) + (-.1N_{v,f}) \quad (8)$$

U	utility
$N_{i,f}$	number indirect hires at firm f
$N_{d,f}$	number direct hires at firm f
$N_{v,f}$	number vacancies at firm f
$\overline{MQ}_{i,f}$	average match quality for indirect hires at firm f
$\overline{MQ}_{d,,f}$	average match quality for direct hires at firm f
$\overline{TCC}_{i,f}$	average total compensation cost indirect hires at firm f
$\overline{TCC}_{d,f}$	average total compensation cost direct hires at firm f

- Intermediary death

Intermediaries' economic health is measured by dividing their total revenue by the number of jobs they have been assigned. If this revenue is less than 10% of the average worker's skill (remember that depending on the experiment being run, fee rates average around 20% of a worker's skill), the contractor dies. Thus the contractor's health depends on their ability to match workers with jobs and their fee rates.

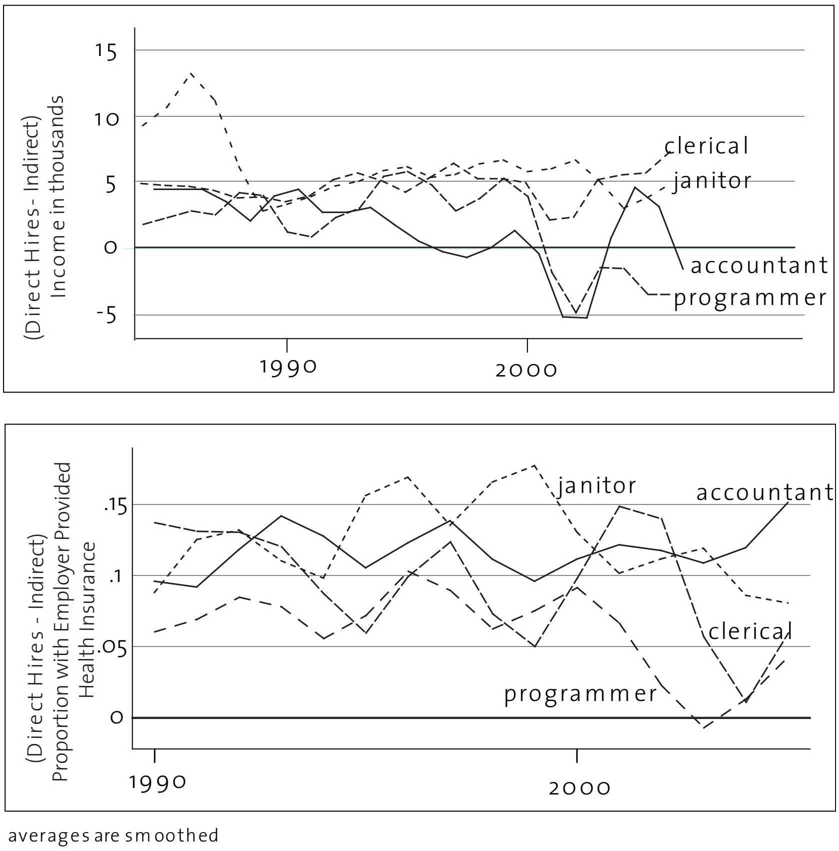

Figure 1: Indirect hires' relative compensation (King et al., 1983-2006)

parameters				*outcome*	
				hypothesized	*experimental*
skill distribution	intermediary premium (3)	variance of premium (3)	firm work shocks (2)	level of indirect hires	percent indirect
minimum wage	negative	low	normal	+ + ++	39.6%
all labor	none	high	normal	+ + +	41.5%
programmers	positive	medium	normal	++	32.7%
accountants	none	high	low	+	28.0%

Table 2: Model predictions and outcomes

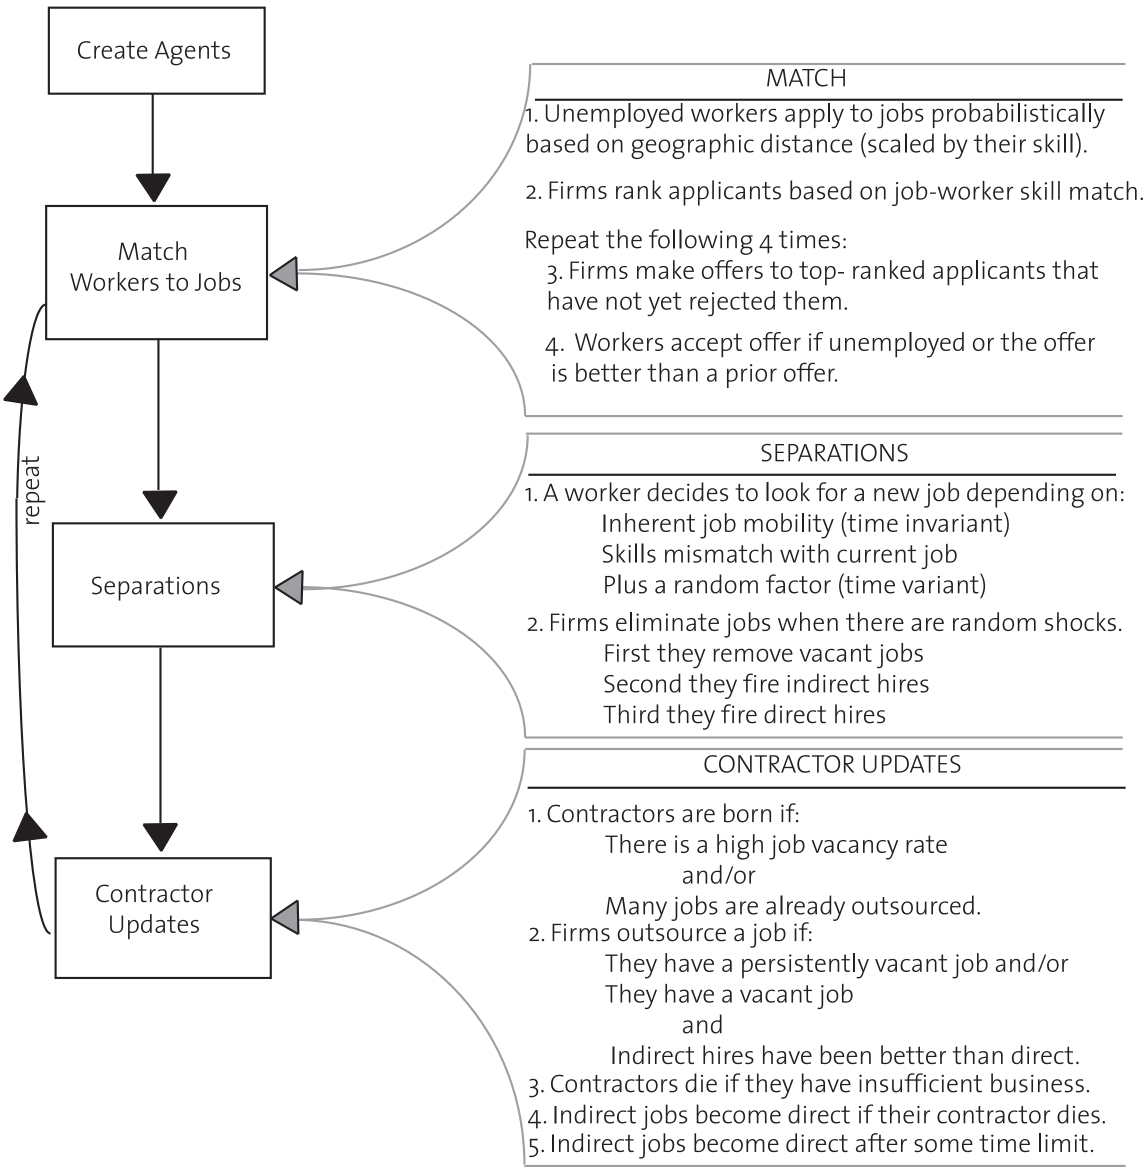

Figure 2: Program structure

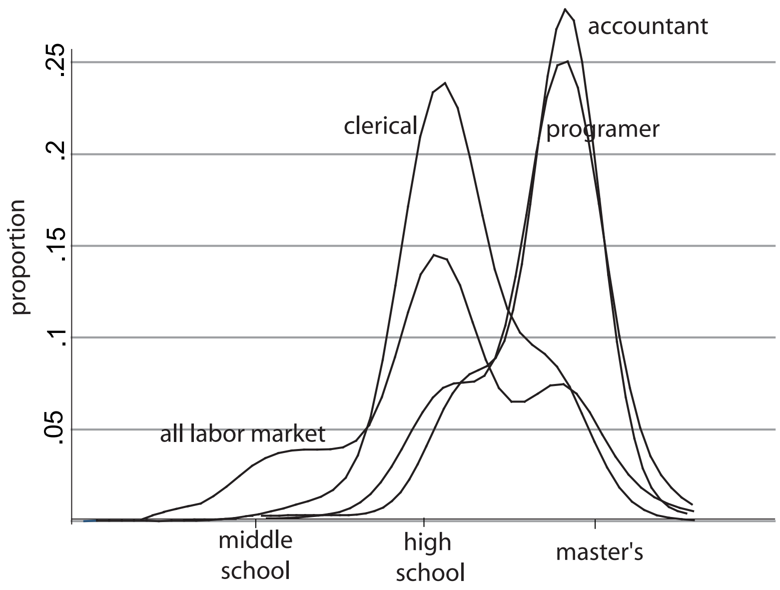

Figure 3: Workers' educational distributions

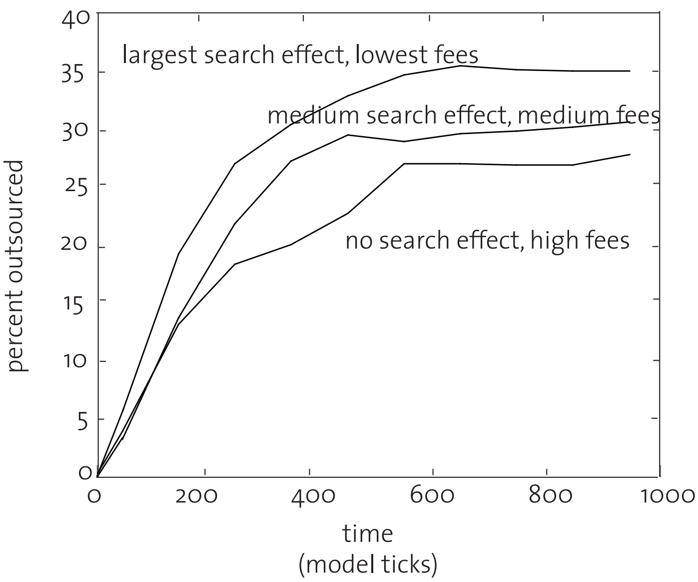

Figure 4: Transition to using intermediaries, mean of 20 runs per scenario

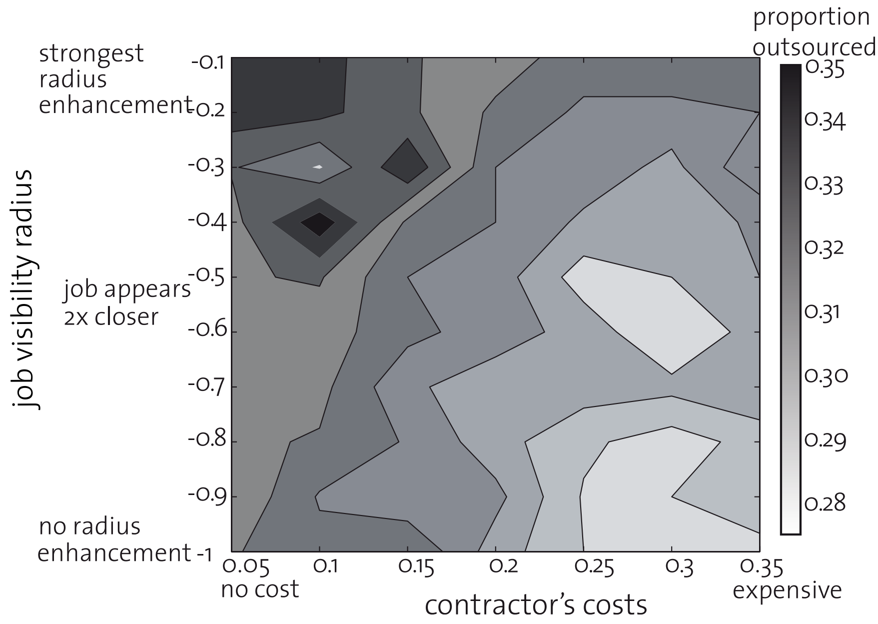

Figure 5: End of run proportion of jobs matched through intermediaries

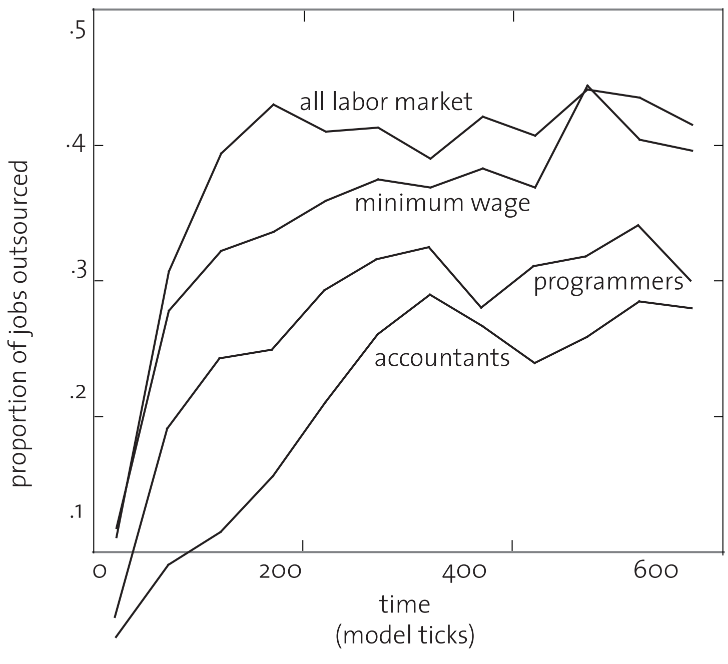

Figure 6: The experimental spread of indirect hiring under four scenarios

parameter	β	standardized β	min	max
intercept	-.027	-	1	1
relative comp	.1701***	.1239	.9	1.1
var of above	.1420**	.0517	.05	.15
work shocks	.925***	.2063	.05	.1
ticks	.0004***	.6789	25	575
skill distribution				
min wage	-.0298***	-.1153	0	1
programmer	-.0916***	-.3541	0	1
accountant	-.1147***	-.4432	0	1
all labor	baseline			

OLS regression, observations are model runs
864 observations (72 settings, 12 intervals)
R-square .69

Table 3: How do ABM parameters influence indirect hiring?

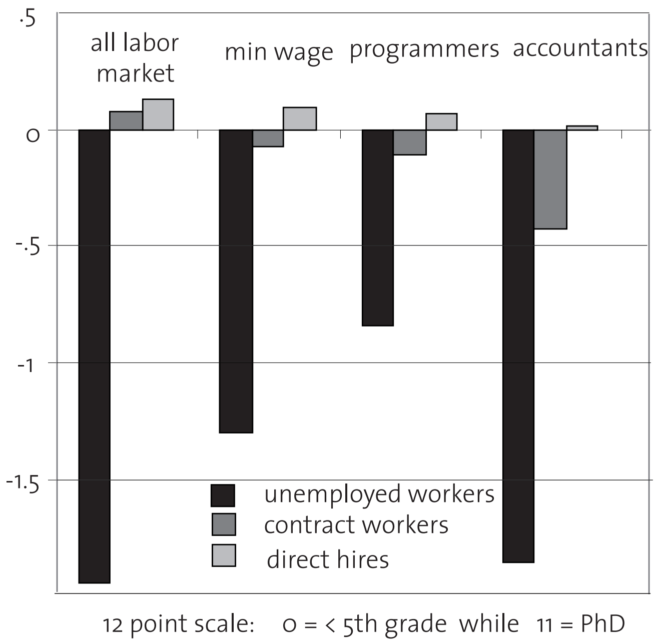

Figure 7: Deviations from the average worker's skill

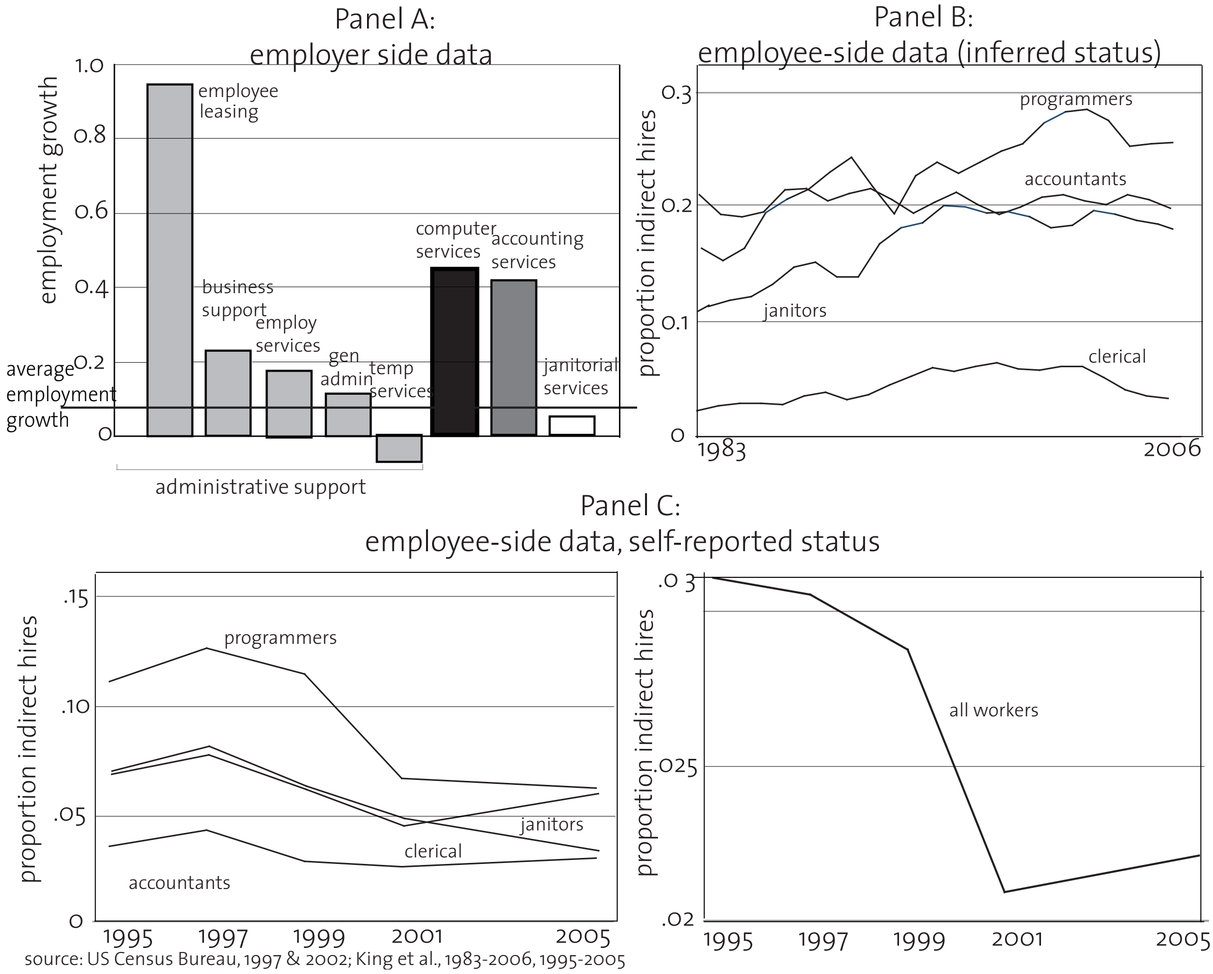

Figure 8: Empirical measures of indirect hiring.

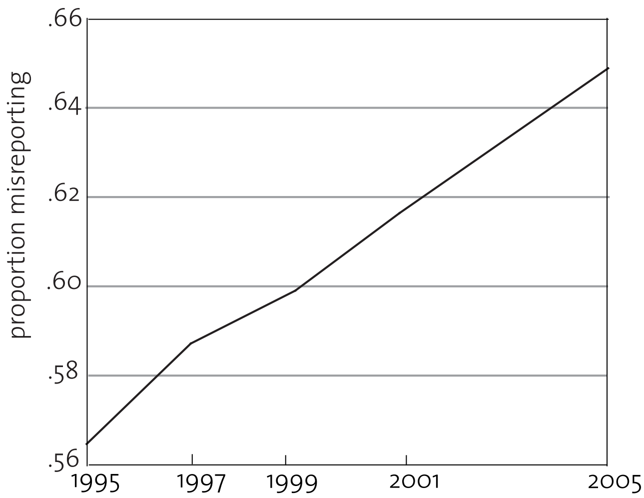

Figure 9: Proportion of indirect hires misreporting their indirect employer as their primary employers (King et al., 1995-2005).

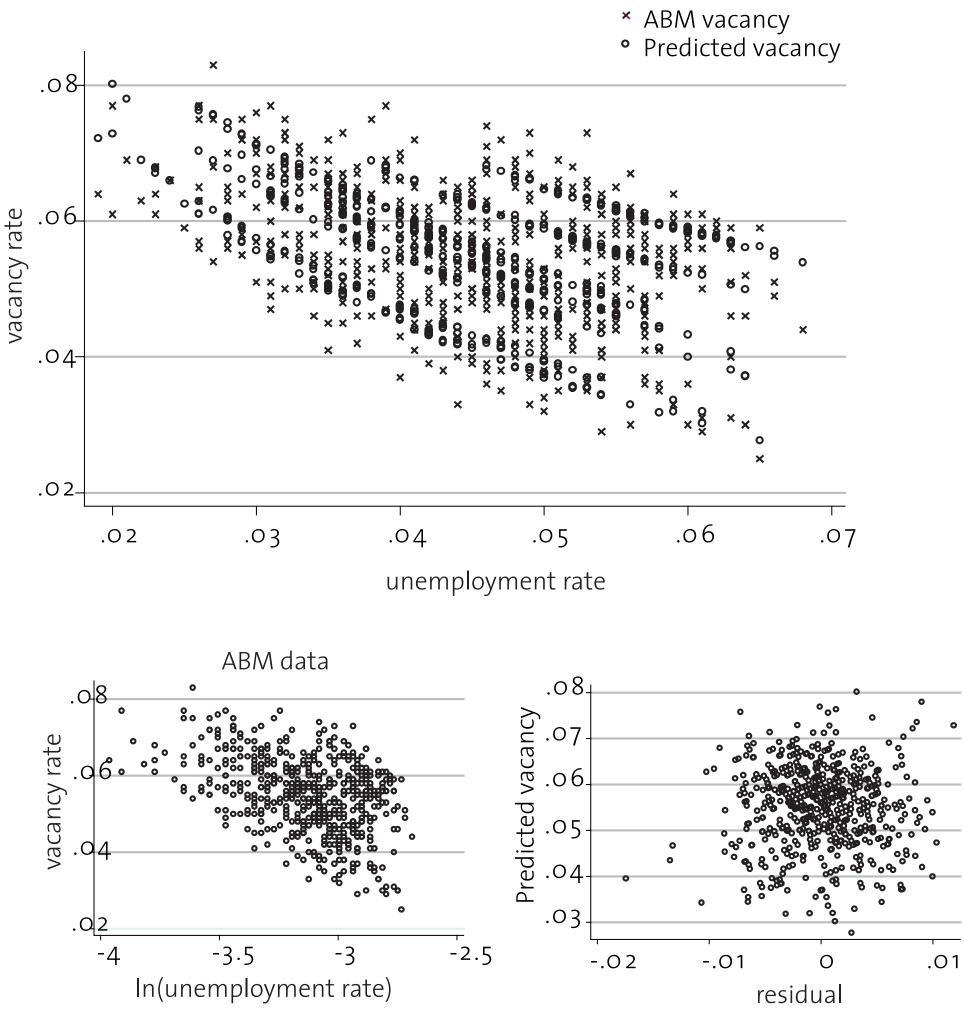

Figure 10: ABM Beveridge curves