

The Pricing of Job Characteristics When Markets Do Not Clear: Theory and Policy Implications¹

Kevin Lang

Department of Economics, Boston University

270 Bay State Road, Boston, MA 02215

and

National Bureau of Economic Research

Sumon Majumdar

Department of Economics, Queen's University

Kingston, Ontario. Canada K7L 3N6.

This version: August 2003.

¹We are grateful to participants in seminars at ITAM, the NBER, Queen's University, University of Rochester, Syracuse University, Virginia Tech and the 1999 European Econometric Society meeting and to Eli Berman, Denise Doiron, Hsueh-Ling Huynh, Shulamit Kahn, Peter Kuhn, Michael Manove, Dilip Mookherjee, Bob Rosenthal and an anonymous referee for helpful comments and discussions. The usual caveat applies. Financial support from National Science Foundation grant SBR-9515052 is gratefully acknowledged. E-mail: lang@bu.edu and sumon@qed.econ.queensu.ca

Abstract

This article examines nonsequential search when jobs vary with respect to nonpecuniary characteristics. In the presence of frictions in the labor market, the equilibrium job distribution need not show evidence of compensating wage differentials. The model also generates several pervasive features of labor markets: unemployment and vacancies, apparent discrimination, and market segmentation. When workers are homogeneous, restrictions on the range of job offers decrease welfare and cannot reduce unemployment. However, when workers have heterogeneous preferences, such restrictions may lower unemployment and can even lead to a Pareto-improvement in welfare. We consider the impact of policies banning discrimination and regulating working-conditions.

1 Introduction

Since the publication of his seminal work, Rosen's (1974) model has provided the basis for labor economists' understanding of how worker and job characteristics are priced by the labor market. Rosen established that with many firms and workers, an implicit market arises for each characteristic, with prices adjusting until all markets clear. All workers get matched to their desired jobs; the market is efficient, and workers are compensated for taking up more onerous jobs. However, there is limited empirical support for either perfect matching or compensating wage differentials. Both casual empiricism and research show that matching is imperfect. Typically, 35-40% of workers report dissatisfaction with their current hours, suggesting imperfect matching (see Lang and Kahn, 2001). Altonji and Paxson (1988) find that workers who are dissatisfied with their working hours tend to move to jobs more in line with their stated preferences, and receive a smaller wage increase when they move. Moreover, empirical work has often failed to find compensating differentials for bad job characteristics (e.g. Brown, 1980).

Much recent theoretical work recognizes frictions and imperfect information in labor markets so that firms must search for workers and workers must search for jobs. In this paper, we develop a model that incorporates frictions into a labor market in which jobs have a non-pecuniary element (e.g. hours of work, flexibility, safety or physical effort). We extend the Butters (1977) framework of nonsequential search to model market imperfections when compensation is multidimensional and workers search for employment taking into account both the wage and the non-wage dimension(s) of jobs. We begin with a model in which all workers share the same taste for the job-characteristic. In equilibrium there is a dispersion of jobs even among identical workers, and wage differentials between disparate jobs need not be "compensating". In fact, we show that under reasonable restrictions, salary compensation will be positively related to the goodness of the job characteristic offered. For example, although firms make tied salary/hours offers, salary will generally be decreasing in hours worked. The model thus provides an explanation for our difficulty

in finding empirical evidence of compensating differentials. Hwang, Mortensen and Reed (1998) also show that in a dynamic sequential search model with heterogeneous firms (and homogeneous workers), salary need not increase as the level of disamenity rises. We use a static model of job-search to derive this result more directly and thus obtain a set of sufficient conditions, and even for the case where both workers and firms are homogeneous.¹

In section 3, we allow for workers with heterogeneous tastes and therefore imperfect matching. We consider the case of a dichotomous job characteristic which we term “clean” and “dirty” (e.g. jobs with and without flexible hours, or jobs involving physical exertion or not), and two groups of observably different but equally productive workers (say, men and women) with different degrees of aversion to dirty work. In equilibrium, for each type of job, the range of accepted offers differs across the two groups of workers. Under fairly general conditions, there will appear to be discrimination against the group that is more averse to dirty work. If, say, women are more averse to “dirty” work than men, then women will, on average, earn less than men in both clean and dirty jobs even though men and women are equally productive. Like the neoclassical model, our model ascribes the over-representation of women in certain jobs and their lower average wage to their preferences. However, in contrast with the neoclassical model, conditioning on working conditions does not eliminate the wage differential between men and women. We thus obtain a novel explanation for apparent discrimination as would be measured by comparing average wages across different groups of workers even in the same occupation.

Since our model generates several pervasive features of labor markets such as unemployment, vacancies and imperfect matching, it is also possible to discuss the effect of policy interventions involving regulation of job-characteristics, such as caps on the length of the working week or workplace-safety regulations. In Rosen’s model, the market is efficient, making such interventions

¹Because Hwang et. al.’s model constrains firms to make a single wage offer for the duration of the employment relation, the difference between their dynamic model and our static formulation is relatively minor. The major advantage of the one-shot formulation is its simplicity which, in turn, allows us to obtain this result analytically.

inefficient. In nonsequential search models such as ours, the market is not fully efficient (Lang and Rosenthal, 1991). Nevertheless, we show that in the homogeneous worker environment, it will generally not be socially desirable to restrict the range of offers firms can make. Simple policies such as maximum hours or minimum wages will reduce workers' expected utility.

However, the results are quite different when workers have heterogeneous tastes. In section 4, we derive a number of interesting possibility results. In particular, requiring that firms make offers without knowing a worker's type (i.e. outlawing discrimination) *can* make both types of worker better off. Each type may exert a positive externality on the other and thereby generate a more attractive offer distribution for *both*. Further, the labor market may show segmentation in the sense that there is a discrete difference in the level of utility (for both types of workers) between the worst job in the (endogenously obtained) "good" segment and the best one in the "bad" segment. This discontinuity in the distribution of utilities offered occurs despite the fact that the technology is continuous and workers of all types are found in both sectors. With heterogeneity in the workforce, other policies such as minimum wage laws and laws to regulate working conditions can also reduce unemployment, help one type of worker at the expense of the other type or even increase expected utility for all types of workers.

We begin with the simplest case i.e. where workers are homogeneous.

2 Homogeneous Workers

We consider a labor market variant of the Butters equilibrium search model incorporating non-wage dimensions of jobs. Each firm decides simultaneously whether or not to make a job-offer to a worker and, if so, what offer to make. Making an offer entails paying a fixed cost. This may be construed either as the administrative cost of making an offer or the cost of setting up the job characteristic or renting a machine prior to searching for a worker. A potential worker collects all the job offers he receives, and then chooses the firm making him the best offer. Each firm

recognizes that making a more desirable offer raises its probability of hiring a worker, but lowers profits conditional on getting the worker. Because some workers randomly fail to receive any offers, there is unemployment. Similarly, some firms' offers are turned down, resulting in unfilled vacancies.²

A firm's job-offer consists of salary S , and a non-wage dimension h . For concreteness, we will typically treat this characteristic as hours of work. However, the model can apply to many other job-characteristics such as job safety, flexibility, etc. We assume that output is a nondecreasing function $V(h)$, of hours worked. This implies that production is constant returns to scale since the productivity of a worker does not depend on the number of other workers in the firm. We can therefore treat separately the problem for each worker. Allowing for increasing or decreasing returns to scale greatly complicates the problem (for an example in a different context, see Lang and Rosenthal, 1991).

We assume that the number of firms with vacancies is given by Z , which is a random variable with a Poisson distribution with mean $z \equiv E(Z)$. This is a natural assumption in that a Poisson distribution would result if firms from a large population made independent and equally probable decisions to enter the job market. Although having a fixed number of firms does not change the qualitative features of the model, use of this random variable simplifies exposition.³ While firms know z , they do not know the realization of Z when deciding on offers.

²For those not familiar with the Butters model, it may be helpful to explain the intuition in the standard formulation (Hosios, 1986; Lang, 1991). Suppose there is one worker, who produces output valued at 10. It costs 1 to make an offer. If firms knew that their offers would be accepted, they would each be willing to pay up to 9 and the wage would be bid up to this level. However, if all firms offer 9, only one of them will get the worker; due to the fixed cost of making a job offer, this implies all firms would be making an expected loss. If there were a mass of offers at any other wage, say 5, a firm could offer just a little more than the other firms, hire the worker for sure and make a profit. Thus the equilibrium will involve a continuous distribution of offers. Note that the lowest offer any firm makes must have a positive probability of acceptance or it will make a loss. If there are no mass points, the only way the lowest offer can have a positive probability of acceptance is if there is some probability that the worker receives no other offers. Thus, there must be a positive probability of unemployment.

³As will become evident, a fixed number of firms in the job market would yield a random number of job offers to individual workers, but these would have a binomial rather than a Poisson distribution. However, the Poisson distribution closely approximates the binomial distribution for large economies. Setting up the model so that the Poisson distribution is exact rather than approximate greatly simplifies the proofs.

Consider the market for a single potential worker. All firms simultaneously decide whether or not to make the worker an offer. If a firm does not make an offer, it incurs no costs and makes zero profit. Otherwise it makes a tied salary/hours offer (S, h) , and incurs a fixed-cost c of making the offer. This cost c is normalized to be 1 in this section. Firms seek to maximize their expected profit.

The worker's utility function on salary S and hours of work h is given by $u(S, h)$, where u is assumed to be continuous and differentiable both in S and h , with $u_s > 0$ and $u_h < 0$. We further assume that $u(0, h^*) < u_r$ where h^* is the worker's endowment of leisure and u_r is his reservation utility. The worker collects all the offers he receives and chooses the firm offering the highest utility level, provided this exceeds his reservation utility. In the event of a tie, he chooses randomly among the best offers. If he receives no offer exceeding his reservation utility, he remains unemployed. If the worker accepts employment, he works the specified number of hours h and receives the specified salary payment S .

Since workers evaluate offers based on their utility level, the probability of attracting the worker by a particular salary/hours offer (S, h) depends on how this offer ranks in the distribution of utilities from the offers that the worker has received. Hence, in calculating its expected profit, a firm only needs to consider the distribution of *utility levels* resulting from the salary/hour offer distributions of all the other firms.

Here, a firm's pure strategy is either a salary/hours pair or the decision not to make any offer, while a mixed strategy is a probability distribution over all possible salary/hours pairs including the possibility of no offer. The firm's expected profit from making an offer (S, h) is

$$(1) \quad E(\pi) = P(u(S, h))(V(h) - S) - 1$$

where $P(\cdot)$ is the probability of getting the worker. This probability depends on the utility level

$u(S, h)$ from the firm's offer of (S, h) , and on what other firms exist and the offers (if any) that they make.

Thus the firm's problem can be separated into the decision of what level of utility to offer (if at all), and how best to offer it. The solution to the second problem is simple: for a given level of utility u , the firm will choose S and h to maximize its (gross) profit $V(h) - S$. We make the following assumption to ensure that the profit function of the firm subject to a utility constraint is strictly concave. We will maintain this assumption throughout this section.

Assumption 2.1: $u_S V''(h) + u_{hh} + 2u_{Sh} V'(h) + u_{SS} (V'(h))^2 < 0$

Although this assumption has no direct interpretation, it is satisfied for example, if the utility function $u(S, h)$ and the production function $V(h)$ are both concave. Under assumption 2.1, any offer that the firm makes will satisfy the following condition, obtained from maximizing $V(h) - S$ subject to $u(S, h) = u^*$:

$$(2) \quad V'(h) = -\frac{u_h}{u_S}$$

This is simply the standard result that offers set the value of marginal product equal to the marginal value of leisure. It is also an implicit function linking S and h . Given assumption 2.1, the second order conditions for the profit-maximization problem are satisfied; thus, given any utility level u^* , there exists a unique profit-maximization pair (S^*, h^*) associated with u^* . Therefore, in equilibrium, the distribution on salary/hour pairs can be obtained from that on the utility levels.

Define $\tilde{\pi}(a) = \max_{\{S, h : u(S, h) = a\}} V(h) - S$, and the utility level u_{\max} by $\tilde{\pi}(u_{\max}) = 1$. Clearly, $\tilde{\pi}(u)$ is decreasing in u , and u_{\max} is the maximum level of utility that the worker can hope to achieve.

The following two lemmas are similar to results in other nonsequential search models (see, for example, Lang 1991) and are stated without proof.

Lemma 1 *In any equilibrium, the offer distribution must be continuous, have support $[u_r, u_{\max}]$, and have no mass points.*

Lemma 2 *If $e^{-z}\tilde{\pi}(u_r) < 1$, all offers in the support of any equilibrium distribution must make zero expected profit. All offers that are outside the support must make nonpositive expected profit.*

Lemma 1 follows from the fact that if there were gaps in the distribution, a firm offering a package at the top of the gap could lower its utility offer slightly without affecting its probability of acceptance (thereby increasing its expected profit). Similarly, if the distribution contained mass points, a firm could increase its probability of acceptance discretely and thus its expected profit, by an infinitesimal increase in the utility offer. To understand lemma 2, note that u_r is the lowest utility offer in any equilibrium distribution. Given that the offer distribution must be continuous, the offer u_r will only be accepted either if there are no other firms *or* other potential firms choose not to make an offer. Recall that z is the expected number of firms with the potential to make offers; therefore, from a particular firm's point of view, e^{-z} is the probability that there are no other such firms. If $e^{-z}\tilde{\pi}(u_r) < 1$, then the only way that an offer of u_r can make positive expected profit is if there is a positive probability that firms choose not to make an offer. But for firms to be indifferent between making an offer or not it must be that the expected profit from making an offer must be zero. We assume that z is large enough so that $e^{-z}\tilde{\pi}(u_r) < 1$.

It follows that the equilibrium will involve mixed strategies. We consider the case where all firms pursue the same mixed strategies.

Proposition 1 *Assume that there exists some (S', h') such that $V(h') - S' > 1$ and $u(S', h') > u_r$; then there is a unique symmetric equilibrium.*

Proof. From (2) we know that there is a unique pair (S, h) for each u . If $P(u)$ is the probability of getting the worker by offering utility u , then from Lemmas 1 and 2, in any symmetric equilibrium we must have: $E(\pi) = P(u)\tilde{\pi}(u) - 1 = 0$ for $u_r \leq u \leq u_{\max}$.




Figure 1: The equilibrium salary/hours distribution

Therefore $P(u) = \frac{1}{\tilde{\pi}(u)}$. Since $\tilde{\pi}(u)$ is decreasing in u , $P(u)$ is a valid distribution function and is uniquely defined over the range $u \in [u_r, u_{\max}]$. Outside of this range all offers make negative profit and therefore are not made. ■

Figure 1 shows the equilibrium graphically (for a concave production function). The points of tangency between the worker's indifference curves and the firm's isoprofit curves map out the set of efficient offers, the contract curve. The set of offers is limited to those to the right of the $\pi = 1$ isoprofit curve and to the left of the worker's reservation indifference curve.⁴

In the above proposition, $P(u)$ is the probability that the worker accepts an offer of utility level less than or equal to u . Since the equilibrium is symmetric, $P(u)$ is also the distribution of the highest order statistic resulting from the offer distributions of all the firms (although from a particular firm's point of view, $P(u)$ is the probability that the best offer made by the *other* firms is less than or equal to u , the assumption that the number of firms in the market is Poisson rather than fixed ensures that $P(u)$ is also the distribution of the highest wage offer made by *all* firms rather than that for all firms minus one); thus, $P(u)$ is the equilibrium distribution of

⁴Although we focus on symmetric equilibria all throughout the paper, there can be asymmetric equilibria as well. For example, by having some subset z' of the firms make no offers (even when they have vacancies) and the rest follow the symmetric equilibrium of the game with mean number of firms equal to $z - z'$, an asymmetric equilibrium is trivially generated.

utility levels for the worker. Using the zero-profit condition $P(u)\tilde{\pi}(u) - 1 = 0$, the actual offer distribution followed by any firm is as follows:

- (i) with probability $1 - p$, make no offer, where p is given by the condition: $e^{-zp}\tilde{\pi}(u_r) - 1 = 0$,
- (ii) with probability p , make an offer with utility $u \in [u_r, u_{\max}]$ according to the distribution $G(u)$ given by⁵: $e^{-zp(1-G(u))}\tilde{\pi}(u) - 1 = 0$.

We have assumed that firms are free to randomize over both wages and the job-characteristic. Perhaps more realistically, firms would first choose their technology (and thus a job-characteristic) and then choose what wage to offer. So long as firms do not learn other firms' technology choices before making their own wage decisions, the sequential decision is isomorphic to the simultaneous decision. Moreover, even if each firm were constrained to choose the same technology for all of its workers, with constant returns to scale, firms cannot take advantage of the resulting correlation in the offers.

Assumption 2.1 places few restrictions on the production and utility functions, and therefore on the equilibrium salary/hours locus. Salary may be increasing or decreasing in hours; indeed there can be more than one salary associated with a given hours offer. However, in a sense made precise in the following proposition, a negative salary/hours relation is the most plausible.

Proposition 2 *If $V'' \leq 0$, and leisure and consumption are normal goods, then equilibrium salary is declining in hours worked.*

Proof. The equilibrium salary/hours pairs are obtained from equation (2). Differentiating it completely, we obtain: $\frac{dS}{dh} = \frac{-(u_{hh} + u_{Sh}V') - u_S V''}{u_{Sh} + u_{SS}V'}$. Since leisure is assumed a normal good, then for any wage w , we have: $u_{Sh} + wu_{SS} < 0$. Similarly, since consumption is assumed to be a normal good, we must have: $u_{hh} + wu_{Sh} < 0$. Now, using these two inequalities and the assumption that $V'' < 0$ proves that $\frac{dS}{dh} < 0$ i.e. that the equilibrium salary is declining in hours worked. ■

⁵ $P(u) = \text{Prob.}(\text{best offer by any other firm} \leq u) = \sum_{n=0}^{\infty} \frac{e^{-z} z^n}{n!} [1 - p + pG(u)]^n = e^{-zp(1-G(u))}$

To understand the intuition for this result, consider the dual of the profit-maximization problem with a linear production function $V(h) = vh$: $\max_{S,h} u(S, h) \text{ s.t. } vh - S = \pi^*$. This is just the standard consumer demand problem with v playing the role of the wage (or price of h) and $-\pi^*$ being the nonlabor income. Therefore if both consumption and leisure are normal, $\frac{dS}{d\pi^*} < 0$ and $\frac{dh}{d\pi^*} > 0$, so that along the contract curve $\frac{dS}{dh} < 0$.

Note that from the perspective of the market-clearing model, a downward-sloping salary/hours locus is anomalous. To “explain” it, one would need to argue that labor supply is, in fact, negatively sloped and that the lower wages among long hour workers arose not because of different choices by equally productive workers but because less productive workers were choosing longer hours. In our model, even identical workers get jobs with differing utility levels. Assuming that leisure and consumption are normal implies that higher utility jobs have higher salaries *and* lower hours; hence, in equilibrium, there is a negative relation between the two.

If the conditions of the above proposition do not hold, one can construct examples with a positively sloped salary/hours locus. For example, with $V(h) = vh^a$, $a > 1$, and $u(S, h) = \ln S + b(T - h)$, the efficient-offer condition (2) gives $S = \frac{avh^{a-1}}{b}$, which is upward-sloping.⁶ If $a = 1$, all jobs offer the same salary and vary only with respect to the hours required. Clearly jobs offering lower hours are preferred to those with higher hours.

It is also possible to mimic a perfectly inelastic labor supply curve by considering a utility function that is linear in S . For example, with $u(S, h) = S + b \ln(T - h)$, $V(h) = vh$, the efficient-offer condition (2) gives $v = \frac{b}{T-h}$. All jobs now specify the same number of hours, but differ in their salaries. In this case, the contract curve is vertical because leisure is not a normal good. We can again obtain an upward sloping salary/hours profile if the income elasticity of the demand for leisure is small but positive and $V'' > 0$.

The result that with labor market frictions, salaries need not be compensating does not seem

⁶In this case, assumption 2.1 is satisfied in the range where $vh^a > S$ i.e. over the range where profits are positive.

to rely on the Butters mechanism and is likely to generalize to other search models which generate jobs differing in utility levels. A referee has pointed out that even in models where the compensation is determined after a match is formed, either due to the realization of some match-specific productivity and/or a stochastic bargaining parameter, lucky workers will bargain to receive part of their higher compensation in the form of higher wages and part in the form of more pleasant working conditions, yielding results analogous to our model.

Since the equilibrium produces a range of utilities, it is natural to ask whether we can increase worker welfare by restricting the set of offers that firms can make. The following proposition shows that such restrictions must be welfare reducing if the constrained equilibrium is also symmetric and has a continuous utility distribution with no mass points.

Proposition 3 *Consider a symmetric equilibrium when firms are restricted on the (S, h) pairs that they can offer. If the equilibrium utility distribution under the restriction does not contain any mass points, then it cannot have higher welfare or lower unemployment than the symmetric equilibrium in the unrestricted case.*

Proof. Let $\pi^R(u^*)$ be the profit at utility level u^* under the set of restrictions. Then if this utility level is offered both in the presence and absence of restrictions, the probability that the offer is accepted is

$$P(u^*) = \frac{1}{\tilde{\pi}(u^*)} \leq \frac{1}{\pi^R(u^*)} = P^R(u^*)$$

The inequality follows from the fact that $\tilde{\pi}(u)$ is the highest level of profit that can be attained when the worker receives utility level u . Since the utility distribution (both with and without restrictions) contains no mass points and the equilibrium is symmetric, this implies that the probability that the worker receives an offer of utility exceeding u^* is no higher under the restriction. Further, in the unrestricted equilibrium, all utilities between u_r and u_{\max} are in the support; therefore, there cannot be a utility level that is offered in the restricted equilibrium but not in

the unrestricted equilibrium. Thus, the utility distribution in the unrestricted case first-order stochastically dominates that under restrictions.

Finally, suppose a utility level u' is offered in the unrestricted equilibrium but not in the restricted equilibrium. Let u'' be the lowest utility level above u' offered in both equilibria. Since we know that the probability of u'' or better being offered is no higher in the restricted equilibrium, the probability of u' or better being offered must be strictly lower. Applying these arguments to $u^* = u_r$ establishes the result for unemployment. ■

Proposition 3 falls short of establishing that the unrestricted equilibrium is efficient for two reasons. First, the restricted equilibrium need not be free of mass points. Restrictions that create mass points may be welfare improving even if compensation is unidimensional (see the working paper version for an example). In our model, for mass points to exist in equilibrium, firms must be unable to increase utility infinitesimally and thereby obtain a discrete increase in the probability of their offer being accepted. This requires, in turn, that policy must either constrain the utilities associated with offers directly or constrain both hours *and* salary by establishing minimum hours and maximum salary. Such policies are not particularly relevant to the policy debate. Most policy proposals involve setting a *minimum wage* but leaving hours unconstrained, or setting *maximum hours* but leaving salary unconstrained.

Second, as discussed in a related model by Lang and Rosenthal (1991), if the government is able to reduce the number of firms participating in the labor market, it can increase welfare. In the context of our model, this can be seen as follows: suppose the number of firms in the labor market is N (thus it is a departure from our assumption of a stochastic number of firms); then the equilibrium cumulative distribution function for utility levels is given by: $P(u) = [\frac{1}{\pi(u)}]^{\frac{N}{N-1}}$. Now, as N decreases, the utility distribution becomes better in the sense of first-order stochastic dominance. Thus regulations that force some firms to be inactive could lead to welfare improvements.

Thus proposition 3 is quite restrictive. It rules out policies that create mass points in the

offer distribution or those that restrict the number of competitors. Nevertheless, the proposition creates a presumption in favor of laissez-faire. Many policies designed to change firms' offers will be welfare-reducing. As we will see, even this restricted result does not hold when workers have heterogeneous tastes.

3 Discrimination

We now consider the implications of our model when workers can differ in their taste for the job characteristic. For simplicity we assume that jobs can be of two types $i \in \{0, 1\}$. Following Rosen (1986), we term them “dirty” ($i = 1$) and “clean” ($i = 0$). Instead of a literal reference to the job-environment, “dirty” and “clean” may also refer to long- and short-hour jobs respectively, or to jobs with and without physical exertion, for example.

There are two types of workers, f and m , with preferences given by $\Phi(S) - i\Psi_j$, $j = f, m$; $\Phi' > 0$, $\Phi'' < 0$. With a slight abuse of notation, we take $\Psi_j = j$, with $f > m$. Although not completely general, this utility function captures two essential features of the model: one group dislikes dirty jobs more than the other, but for a sufficient compensating differential each would be willing to take a dirty job. All workers are equally productive, with productivity v at clean jobs and $v + \delta$ at dirty jobs. The fraction of f and m types in the population is assumed known to all firms. In this section, we study the case when firms can identify the type of a worker and make type-contingent offers.

As a benchmark, consider first the equilibrium in Rosen's implicit market model. There will be separate wages for clean and dirty jobs with a single equilibrium wage in each market. Depending on the parameters v , $v + \delta$ (which are the possible equilibrium wages in the two markets) and Ψ_f, Ψ_m , either each group will be in a different type of job or all workers from both groups will have the same job characteristic. In the interesting case where type f workers are employed in clean jobs and type m workers in dirty jobs, the wage differential between dirty and clean jobs

will be δ . If two demographic groups have different proportions of type m 's and type f 's, the one with a higher proportion of m 's will have a higher wage. However, conditional on the type of job in which they are employed, there is no wage difference between the two groups. Additionally, the equilibrium is efficient.

3.1 Search Equilibrium

When firms can identify the worker's type and make different offers to the two types, the equilibrium for each worker group is a simple extension of the case in section 2, now allowing for a discrete job characteristic. As before, the equilibrium consists of a continuous distribution of utilities, ranging from the reservation utility level up to a maximum. We state without proof the nature of equilibrium in this case.

Claim 1 *Let u_j^r denote the reservation utility level for a type j worker. If dirty jobs are more profitable when the lowest possible wages are paid i.e.*

$$(3) \quad v + \delta - \Phi^{-1}(u_j^r + \Psi_j) > v - \Phi^{-1}(u_j^r)$$

then the lowest utility offer for group j specifies a dirty job.

If workers prefer clean jobs when the highest possible wages are paid i.e.

$$(4) \quad \Phi(v - c) > \Phi(v + \delta - c) - \Psi_j,$$

then the highest utility offer for group j specifies a clean job.

If both (3) and (4) hold, then all offers above some cut-off utility level u_j^ specify that $i = 0$ and all offers below u_j^* specify that $i = 1$.*

u_j^ is given by the equation: $v + \delta - \Phi^{-1}(u_j^* + \Psi_j) = v - \Phi^{-1}(u_j^*)$.*

u_j^ is nondecreasing with respect to δ , and $u_m^* > u_f^*$.*




Figure 2: Equilibrium distribution for (a) m types, (b) f types.

If workers are too averse to dirty work, there will be no dirty jobs offered in equilibrium while if they are insufficiently averse, there will be no clean jobs. Conditions (3) and (4) ensure that both types of jobs exist in equilibrium. Since Φ is concave, the best (i.e. highest utility) offers will involve a clean job and pay a relatively high salary. Marginally less attractive jobs will pay a lower salary but still be clean. As we go down the utility distribution, at some point u_j^* , it becomes efficient to switch to offering a dirty job. Since utility is continuous, this jump to the discretely worse job characteristic must be accompanied by an upward jump in the salary offer (this difference depends on the value of δ and Ψ_j). As the utility offered by the firm falls beyond u_j^* , it continues to offer a dirty job and lowers the salary level. Depending on the value of δ and Ψ_j , dirty jobs may pay, on average, more or less than clean jobs. Thus there need not be a compensating differential for clean jobs, but in contrast with the continuous case, a compensating differential is not ruled out.

Figure 2(a) depicts the equilibrium for the m -type workers. Note that a job-offer with salary S and characteristic i will give an m -type worker utility $u_m = \Phi(S) - im$ and an f -type worker

utility $u_f = \Phi(S) - if = u_m - i(f - m)$. Hence, we can equivalently characterize an offer in terms of the utility it offers to the two types. Since both groups suffer no disutility in clean jobs, the line $u_m = u_f$ in figure 2 represents the clean jobs, with higher levels of utility corresponding to higher salary jobs. In dirty jobs, the f -type workers suffer a greater disutility than the m -types. This difference is given by $f - m$; hence, the dirty jobs are characterized in the figure by the line $u_f = u_m - (f - m)$. As claim 1 shows, in equilibrium, for utility below u_j^* only dirty jobs will be offered, while those with utility above u_j^* will be clean.

3.2 Apparent Discrimination

Under conditions (3) and (4), both groups of workers are found in both types of jobs. The following proposition compares the conditional wage distribution for the two groups. A formal proof of the proposition can be found in the working paper; however, the intuition for it is relatively straightforward, and follows the proposition.

Proposition 4 *Suppose (3) and (4) hold. Then the expected salary for a type f worker in an $i = 0$ job is lower than that for a type m worker. If both types have the same reservation salary in $i = 1$ jobs, then the expected salary for a type f worker in a $i = 1$ job is also lower than that for a type m worker.*

It is not surprising that type f workers, who will be found disproportionately in the clean jobs, have lower average wages overall. After all, they are more averse to dirty work and pay a penalty for being less likely to accept such jobs. Faced with knowledge of this wage differential, the empirical researcher could control for working conditions. The proposition states that even with controls for working conditions, type f workers have lower salaries on average. This result follows from the fact that for f types, the cut-off utility below which offers specify that the job is dirty, will be lower (i.e. $u_f^* < u_m^*$). Therefore, as can be seen from figure 2, the wage distribution for clean jobs contains a lower tail for group f that is not present for group m . Since the probability

of acceptance must be the same for any offer that is made to both groups, it implies that the average salary for a type f worker in a clean job is lower than that for a type m in the same kind of job. Intuitively, on the margin, there is a set of clean jobs with the lowest salaries in which only type f workers are employed; for the corresponding utility levels, it is more profitable to employ a type m in a dirty job rather than in a clean one, as their relative aversion to the dirty job is lower. Hence, even in the absence of any taste for discrimination on the part of employers, the average wage of type f workers in clean jobs is lower.

Similarly the wage distribution for dirty jobs contains an upper tail not present for group f , resulting in a higher average salary for a type m worker. However, in the comparison for dirty jobs, the assumption that both types have identical reservation salary is important; if instead, they had the same reservation utilities, then the distribution for the m -types would also contain a lower tail that would not be present for group f .

This result on apparent discrimination derives chiefly from the difference in the cut-off utility levels for the two types of workers. It seems likely to also hold in other search models with an endogenous distribution of jobs, provided workers cannot ex-ante distinguish between firms with “clean” and “dirty” jobs and accordingly direct their search. For example, in models with match-specific productivity, a match between a m -type worker and a clean job with a low match-specific productivity is more likely to be terminated in favor of a dirty job, than a similar match between a f -type and a clean job; the reasoning here is analogous to our model, thus again leading to a difference in average wages between the two groups of workers in the same type of job.

4 Policy Considerations

4.1 Anti-Discrimination Policy

Faced with the evidence of apparent discrimination as seen in the previous section, government might be tempted to pass legislation prohibiting conditioning offers on the basis of group membership, perhaps by making it illegal to require information about sex or marital status on job applications. Given the efficiency result for the case of homogeneous workers and the absence of discriminatory tastes, it may be tempting to conclude that such an anti-discrimination policy must lower welfare. In fact, we show that the policy can even be Pareto-improving.

Of course, in some cases such a policy will be Pareto-deteriorating. If, for example, type f 's are so averse to dirty jobs that they will never accept one and type m 's are so averse to clean jobs that they will never accept one, forcing firms to offer dirty jobs to f 's and clean jobs to m 's is not only pointless but costly since firms must bear the cost of offers they know will not be accepted. It is also not surprising that there are cases where banning discrimination helps one type and hurts the other. What is more interesting is that there are cases where eliminating discrimination can make both types better off.

In general, the precise nature of the equilibrium when firms do not know the type of the worker, is heavily dependent on parametric specifications.⁷ In our working paper version, we establish a set of sufficient conditions under which eliminating discrimination lowers the unemployment rate for both types of workers. Here we limit ourselves to an example. The calculations for this example are also discussed in the working paper.

Example 1 *Consider the model in section 3 with $v = 10$, $\delta = 6$, $c = 1$. Suppose that utility for*

⁷The problem now is that the same distribution of job offers is ranked differently by the two groups. Arguments which guarantee continuity in the one-type case do not hold here since a gap in the utility distribution of one type does not necessarily imply a similar gap in the utility distribution of the other. Lowering the utility offer over the gap for one type may now reduce the probability of attracting the other type (and could thus even lead to lower profits). Thus the equilibrium distribution may not even be continuous.

m-types is given by $u_m = \ln S$ (so they have no disutility from dirty jobs), while that for *f*-types is given by $u_f = \ln S - i \ln 4$, where $i = 1$ for dirty jobs and $i = 0$ for clean jobs. Both type of workers are in equal proportion, and we assume that both have the same reservation utility $u^r = -25$.

If firms could condition their offers on the type of the worker, the *m*-types would be offered only dirty jobs with salaries ranging from a little above 0 to 15, the unemployment rate among them would be 6.25%, and they would enjoy an expected utility of 0.803. The *f*-types would be offered $i = 1$ jobs with $u \leq u_f^* = \ln(2)$ and $i = 0$ jobs above it; salaries in the clean jobs would range from 2 to 9, and in dirty jobs would range from a little above 0 to 8; their unemployment rate would be 6.25% and their expected utility would be 0.1897. We note that both unconditionally as well as conditional on being in a dirty job, *f*-types would have lower wages than *m*-types.

When offers cannot be conditioned on type, the equilibrium involves firms offering some clean jobs with salaries ranging from 3.6028 to 8.3436 and some dirty jobs with salaries ranging from 8.3436 to 14.4112. The probability that a worker receives an offer of a clean job is 0.74, while the probability of an offer of a dirty job is 0.79. The lowest utility that an employed *m*-type gets is 1.28; similarly, the lowest utility enjoyed by an employed *f*-type is 0.735, while the unemployed of both types get an utility of -25 . The unemployment rate is 5.37%, which is lower than that under discrimination. Further, expected utility is now 0.97 for *m*-types and 0.32 for *f*-types — both types are better off than under discrimination.

Why does eliminating type-contingent offers lower unemployment and make workers better off? Intuitively, the equilibrium now has two sets of “good” jobs — the jobs of type 1 are primarily for the *m*-types, while those of type 0 have a higher probability of attracting the *f*-types. Now the two types of workers exert an externality on each other — the presence of *m*-types makes it possible for unlucky *f*-types to get jobs aimed at *m*-types but which are nevertheless more attractive than the jobs designed to capture only the very unlucky *f*-types. The presence of the *f*-types has a similar beneficial effect for the *m*-types.

Thus, in contrast with the homogeneous worker case, the distribution of utilities (for both types of workers) after restricting firms to type-independent strategies is not strictly first-order stochastically dominated by that under no restriction. Outlawing discrimination makes the least desirable outcomes (unemployment) less likely but it also eliminates the most desirable outcomes for each group. In the above example, the highest wage in clean jobs drops from 9 to 8.34 and that in dirty jobs drops from 15 to 14.4.

It is not possible to establish a general set of conditions sufficient to ensure that outlawing discrimination raises the expected utility of both types. The reason is that while the equilibrium is invariant to any monotonic transformations of the utility functions (provided of course that the reservation utilities are also adjusted appropriately), the expected utility change depends on the cardinality of the utility function. Thus if we use a social welfare function that depends on the expected utility of the participants, the welfare analysis of the policy change cannot be resolved only by examining the income distribution and unemployment rate. We must actually know each individual's utility function. This result will also arise in our other policy discussions. We believe that it casts light on the Fuchs et. al. (1998) finding that policy differences among economists reflect differences in values rather than differences in scientific judgment, a point to which we return in the conclusion.

4.2 Labor Market Segmentation

Another interesting feature of the example above is that in both clean and dirty jobs, all job-offers involved a strictly higher utility than the reservation utility level. This suggests that there could be a second segment of the market (a secondary labor market) with utility substantially below that received by workers in the first segment (a primary labor market). In the above example, there is no secondary labor market (and all employed workers are discretely better off than the unemployed workers) because the externality imposed by the presence of each type on the other

makes the resulting unemployment rate sufficiently low so as to make *all* lower utility offers unprofitable. However with a slight modification, a secondary labor market may arise.

Suppose all parameters are identical to example 1 except that $\delta = 2$. Equilibrium salaries in the primary sector now range from 2.5757 to 8.6761 for clean jobs, and from 8.6761 to 10.303 in dirty jobs. There is however also a secondary sector of dirty jobs with low salaries ranging from a little above 0 to 1.0181. Both types of workers work in both types of jobs.

4.3 Regulating working conditions

The prototypical argument between an economist and an advocate for regulating working conditions might go as follows. The advocate argues that no worker should be subject to such bad working conditions. The economist argues that workers in bad working conditions have “chosen these jobs”. The noneconomist counters that the workers “had no choice”. The economist responds that eliminating “bad jobs” will not increase the workers’ choice. At best, the regulation will reduce the number of jobs; at worst it will eliminate any choice for some types of workers.

The example below should give both parties cause for hesitation. When firms cannot differentiate between worker types, a total ban on one class of jobs can also have the effect of lowering the overall unemployment rate.

Example 2 *We consider the same parameters as in example 1, but now the reservation utility for the f -types is taken to be $u_f^r = \ln(7)$, which is greater than $u_f^* = \ln(2)$. If firms are not able to differentiate between the two types of workers, the equilibrium will now consist of some clean jobs with utility levels ranging from u_f^r to $\ln(8.4)$, and some dirty jobs with $u_m \in [\ln(8.4), \ln(14)]$ and $u_m \in [u_m^r, \ln(1.6)]$. The dirty jobs will attract only the m -types, while the clean jobs attract both types (but have a higher probability of attracting a type f , rather than a type m worker). In this case, the unemployment rate among f -types will be 0.5275, while that among m -types is 0.125.*

Now if dirty jobs are completely banned, then only clean jobs will be offered. Those clean

jobs which are above the reservation utility level of the f -types now have the same probability of attracting both the types; hence the unemployment rate for the f -types falls to 0.33. For the m -types however, the lowest utility jobs are now of the clean type. Therefore, their unemployment rate rises to 0.2. Overall, the unemployment rate drops from 0.3265 to 0.265.

As the economist argued, eliminating dirty jobs hurts workers who would have accepted one. However, by increasing the likelihood that a low-wage clean job will be accepted, the ban on dirty jobs increases the availability of clean jobs as the advocate for regulation argued, making the type f workers better off. The outcome can be a reduction in overall unemployment. In terms of utility, the two situations are not Pareto-rankable. The expected utility for the f -types increases due to the regulation, while that for the m -types falls.

Whether the intuition in the above two examples generalizes to other matching models in which wages are determined through bargaining and/or match-specific productivity, we leave for future work. Speculatively, within a multi-sector model with wages determined by bargaining, restrictions on bargaining might reduce the surplus in one sector and encourage increased job search and thus the opening of vacancies in another, thus aiding one type of worker and harming the other type.

4.4 Minimum Wages

Sometimes regulation can have the effect of raising the expected utility of all workers. In an example with a continuous job characteristic (namely, hours of work) and two groups of indistinguishable workers with differing preferences towards leisure, we show (see the working paper version for details) that the hourly wage profile is U-shaped in hours. The most preferred hours for the two groups lie at the two ends of the hours distribution.

Now, setting a minimum wage above the lowest wage in equilibrium makes job offers with intermediate hours (which had the lowest wages) less prevalent and increase offers towards the

extremes. Therefore the minimum wage increases the likelihood that a worker will get a job near her preferred hours, lowers the probability for somewhat less preferred jobs and raises it again for her least preferred jobs (and also increases the probability of unemployment). As a consequence, a minimum wage law can raise or lower expected utility for each type of worker. As discussed in section 4.1, since neither distribution first-order stochastically dominates the other, there will be monotonic transformations of each utility function giving rise to either conclusion.

5 Conclusion

The model presented in this paper is a relatively modest departure from existing models in the literature, notably Butters' original nonsequential search model but also the closely related work by Burdett and Mortensen (1998). We have simply modified the Butters model to allow for multidimensional compensation and heterogeneity in worker tastes. Yet these seemingly minor innovations generate some surprising results. First, even among homogeneous workers, wage differentials between disparate jobs need not be “compensating”. Indeed, a positive relation between job quality and earnings is the more natural outcome. Secondly, when workers differ in terms of their taste for the job-characteristic, there will appear to be discrimination against the group that is more averse to “dirty” work. Even controlling for type of job, this group (who are equally productive) will earn lower wages on average. Now if firms are unable to condition their offer on the worker's type, there can be a discontinuity in the support of the utility distribution; the labor market may exhibit segmentation in the sense that all workers prefer the jobs in one segment to those in the other. This segmentation is not merely a separating equilibrium — both types of workers are found in each segment. Further, in some cases, this segmentation may be extreme enough so that **all** employed workers are strictly better off than the unemployed.

From a policy perspective, the model shows a limited role for policy interventions in the case of homogeneous workers. However, with a heterogeneous workforce, there may be a positive role

for policy. Policies designed to limit the set of possible offers either by preventing firms from differentiating between different types of workers or by limiting the range of compensation and working conditions may be welfare improving. In some cases the welfare judgment depends on the relative weighting of different types of workers. In others, restrictions may lead to a Pareto improvement. The model therefore generates a potential justification for such diverse policies as occupational safety legislation, minimum wage laws and maximum hours laws.

We wish to emphasize that the justification is only potential. Although we have developed examples in which such laws improve welfare, we have not established that they would be so with realistic parameter values. Indeed, a more realistic model would require considerably more dimensions of job quality and more worker heterogeneity with respect to both tastes and productivity.

Perhaps more significantly, the model suggests that the relative merits of policies often cannot be resolved on the basis of observable effects. In our minimum wage example, it is possible for everyone to agree on the effect of the policy on unemployment and the wage distribution. However, whether or not the policy is desirable depends on how we assess the trade-off between having more “good” jobs but also more unemployment. In principle, if our objective is to maximize expected utility, we could resolve this with knowledge of individuals’ utility functions. In practice, our judgment of the policy will be driven by our values. Thus the model is consistent with the Fuchs et. al. (1998) report that policy differences among economists are primarily driven by values, not scientific disagreements.

Still, the model emphasizes an important element of policy — many policy decisions influence the distribution of jobs offered. By increasing the availability of jobs with certain characteristics, policies favor workers who are most drawn to such jobs. In some cases, whether or not a policy will be desirable will therefore depend on the social welfare function of the individual evaluating the policy. In others, the shift in job characteristics may be more clearly beneficial.

References

- Altonji, J. G. and C. H. Paxson. "Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-Offs." *Journal of Labor Economics* 6(2) (1988), 254-76.
- Brown, C. "Equalizing Differences in the Labor Market." *Quarterly Journal of Economics* 94 (1980), 113-34.
- Burdett, K. and D. T. Mortensen. "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39 (1998), 257-73.
- Butters, G. R. "Equilibrium Distributions of Sales and Advertising Prices." *Review of Economic Studies* 44 (1977), 465-91.
- Fuchs, V. R., A. B. Krueger and J. M. Poterba. "Economists' Views about Parameters, Values, and Policies: Survey Results in Labor and Public Economics," *Journal of Economic Literature*, 36 (1998), 1387-1425.
- Hosios, A. J. "Layoffs, Recruitment and Interfirm Mobility." *Journal of Labor Economics* 4 (1986), 473-502.
- Hwang, H., D. T. Mortensen and W. R. Reed. "Hedonic Wages and Labor Market Search." *Journal of Labor Economics* 16 (1998), 815-47.
- Lang, K. "Persistent Wage Dispersion and Involuntary Unemployment." *Quarterly Journal of Economics* 106 (1991), 181-202.
- Lang, K. and S. Kahn. "Hours Constraints," in G. Picot and G. Wong, eds., *Working Time in Comparative Perspective, Volume I* (Kalamazoo, MI: Upjohn Institute, 2001).
- Lang, K. and S. Majumdar. "The Pricing of Job Characteristics When Markets Do Not Clear: Theory and Policy Implications," *NBER Working Paper No. 9911*, August 2003.
- Lang, K. and R. W. Rosenthal. "The Contractor's Game." *RAND Journal of Economics* 22(3) (1991), 329-38.
- Rosen, S. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition."

Journal of Political Economy 82 (1974), 34-55.