BRUSSELS ECONOMIC REVIEW - CAHIERS ECONOMIQUES DE BRUXELLES
VOL. 48 - N°1/2 SPRING - SUMMER 2005

MONOPSONY AND LABOUR DEMAND+

ALAN MANNING*

ABSTRACT:

'Monopsony and Labour Demand' might strike many as a contradiction in terms as monopsony is often thought to mean an outcome on the labour supply and not the labour demand curve. This paper argues that, despite initial appearances to the contrary, there is no inevitable contradiction between the view that employers have some monopsony power over their workers and the enormous literature on the economics of labour demand. Indeed, popular models of 'labour demand' with convex costs of adjustment are really better thought of as models of monopsonistic firms. But the 'labour demand' literature could do with some re-focusing as the study of the wage and employment decisions of employers.

JEL CLASSIFICATION: J23, J42.

KEYWORDS: Monopsony, labour Demand.

+ This paper is based on a lecture given at the Applied Econometrics Association Conference on the 'The Econometrics of Labour Demand' held at the University of Mons-Hainaut on October 7-8 2004.

* Department of Economics and Centre for Economic Performance, LSE, Houghton Street, London WC2A 2AE.

INTRODUCTION

The title 'Monopsony and Labour Demand' might strike many as a contradiction in terms. For example, in the textbook model of labour demand, the imposition of a just-binding minimum wage on the free market equilibrium results in a fall in employment as employers move along a downward-sloping labour demand curve. In contrast, in the textbook model of monopsony, the impact of a just-binding minimum wage is to raise employment as the firm moves along an upward-sloping labour supply curve. From this perspective 'Monopsony Or Labour Demand' might appear a more sensible title. And faced with a choice between 'labour demand' and 'monopsony' it might appear to be a very obtuse soul who would choose monopsony in the face of the vast amount of research that appears to successfully estimate labour demand curves.

But, this paper argues that it is sensible to think in terms of 'monopsony and labour demand', that it is a false dichotomy to insist on one or the other. The first part of the paper argues that much of the literature on labour demand curves is perfectly consistent with the view that labour markets are pervasively monopsonistic as argued by Manning (2003). The second section then argues that many models that are generally thought of as being models of labour demand (with the comparative statics that those words conjure up) are, in fact, better thought of as models of monopsonistic firms. It is shown how models of employment adjustment with convex costs of adjustment are really best thought of as models of monopsony. The third section then compares the monopsony view of labour markets not with the perfectly competitive model but with a model of the labour market in which wages are set by employers at such a level that there is a chronic excess supply of labour to the firm – what we call, rather loosely, 'efficiency wage' models.

The conclusion is that there is little in the existing literature that is inconsistent with the view that labour markets are pervasively monopsonistic. But, in part, this may be because we do not have much in the way of evidence on the crucial issues that would provide more definitive tests of the relevance of monopsony – we do not have good estimates of the elasticity of the labour supply curve facing a firm, we do not have good estimates of the form of hiring costs, and we do not have good information on the recruitment process. The paper concludes by laying out a research agenda that might help to answer these questions.

1. Reconciling Labour Demand and Monopsony

A very basic first problem for anyone who wants to argue that employers possess a substantial amount of monopsony power over their workers is how to deal with the enormous number of papers that seek to estimate of "labour demand curves" and claim some success in this quest (see Hamermesh, 1993, for a survey)[1]. These labour demand curves can all be thought of as an attempt to estimate the first-order condition for the profit-maximising level of employment when faced with a parametric wage. A canonical way of writing this is:

$$pF'(N)=w \tag{1}$$

though dynamic models are common and (1) does not reflect this aspect of the problem. But one can summarize this literature by saying it seeks to estimate an outcome on the MPL=w curve.

On the face of it such an outcome is not consistent with monopsony. Consider the simple textbook model of monopsony as presented in Figure 1. The outcome is on the labour supply curve at the point where the marginal product of labour is equal to the marginal cost of labour and not on the MPL curve. But, it is simple to reconcile the apparent existence of labour demand curves with monopsony. For employers with some market power, one can write the first-order condition for employment as:

$$pF'(N) = w(1+\varepsilon) \tag{2}$$

where ε is the inverse of the elasticity of the labour supply curve facing the firm. (2) makes it apparent that anything that changes the wage but leaves ε unchanged (i.e. shifts the labour supply curve facing the firm but leaves its elasticity unaltered) will lower employment leading to an apparent labour demand curve. But, these labour demand curves will not be a good guide to the impact of policies like the minimum wage and equal pay legislation that are likely to change ε as well as the wage. So, there is no real problem in reconciling the existing labour demand literature with the monopsony perspective – the correlation between wages and employment depends on the source of shocks to wages.

---

[1] One should not conclude form this that the empirical correlation between employment and wages is always negative – indeed, a simple correlation of the two is almost always positive, the employer size-wage effect.

2. Are Models of Dynamic Labour Demand Really Models of Monopsony ?

The argument of the previous section is rather defensive, arguing that the empirical literature on labour demand curves is not inconsistent with the view that labour markets are monopsonistic. That may be true but it is a little feeble. In contrast, this section is rather more aggressive arguing that many models that economists generally think of as models of the demand for labour by a competitive firm are best thought of as models of a monopsonistic firm.

Suppose one is interested in modelling the demand for labour of an individual firm, i.e. interested in studying the level of employment in that firm. One might start by specifying a revenue function pF(N) where p is the price of the product and N the level of employment (other inputs being suppressed in the interest of simplicity). We might use this production function to estimate a labour demand curve along the lines of (1). But, as a practical matter one would quickly find that this is not a very adequate empirical model as it would imply a responsiveness of employment to shocks that is far more rapid than seen in reality. So, one needs to introduce some factors that make the employment response to shocks sluggish – what we might loosely call employment adjustment costs. These employment adjustment costs represent the costs of recruiting and training new workers[2]. In the early literature on employment adjustment costs, it was assumed that they were quadratic so let us follow this here and assume that the labour adjustment costs are given by:

$$C = \frac{1}{2} c H_t^2 \tag{3}$$

The number of hires needed depends not just on the level of employment today compared with yesterday but also on the number of quits so that we have:

$$H_t = N_t - (1-q) N_{t-1} \tag{4}$$

Given (4), one needs to think about the determinants of quits. In the context of a perfectly competitive model of the labour market in which the employer has no discretion over the wage paid this is a little hard to do as there is not really an adequate theory of job mobility. But a sensible, pragmatic approach is to assume that the quit rate depends negatively on the wage paid among other factors. This is in line not just with common sense but also with a very large number of studies that find a negative relationship between quit rates and wages (see Manning, 2003, ch4, for a review of these studies). The alert reader will be aware that a monopsonistic element has crept into the analysis here for as soon as one assumes that the quit rate depends in a continuous way on the wage, the employer is inevitably seen as having some discretion over the wage paid,

[2] In many countries there also sizeable costs of firing workers but we will not consider these here as these are the product of policy and not inevitable in the way that hiring costs would seem to be.

something that is different from the wage-taker behaviour in a strictly perfectly competitive labour market. But there is nothing artificial about the model set up – one has to think about hiring costs, one has to think about quit rates and one ends up with the model described. The model with quadratic adjustment costs and quit rates depending on the wage is the 'dynamic monopsony' model of Mortensen (1970).

Let us consider the optimal wage and employment decisions of the employer. The state variable is last period's level of employment so the value function can be written as:

$$\Pi(N_{t-1}) = pF(N_t) - w_t N_t - \frac{1}{2} c\left[N_t - (1-q(w_t))N_{t-1}\right]^2 + \delta\Pi(N_t) \tag{5}$$

From this we have the first-order condition for employment that can be written as:

$$pF'(N_t) - w_t - c\left[N_t - (1-q(w_t))\,N_{t-1}\right] + \delta\Pi'(N_t) = 0 \tag{6}$$

a first-order condition for wages that can be written as:

$$-w_t - q'(w_t)N_{t-1}c\left[N_t - (1-q(w_t))\,N_{t-1}\right] = 0 \tag{7}$$

and an envelope condition:

$$\Pi'(N_{t-1}) = (1-q(w_t))\,c\left[N_t - (1-q(w_t))\;N_{t-1}\right] \tag{8}$$

For current purposes we are only interested in the steady-state levels of wages and employment, (w,N). Using (6)-(8) the first-order condition for steady-state employment can be written as:

$$pF'(N) = w + cq(w)N\left[1-\delta(1-q(w))\right] \tag{9}$$

and the first-order condition for wages can be written as:

$$cq(w)Nq'(w) = 1 \tag{10}$$

My claim is that this is best thought of as a model of monopsony and not as a model of a perfectly competitive labour demand curve. The reason for this claim is that the comparative statics of this model are in line with the textbook model of monopsony and not the textbook model of labour demand. I will give two examples of this.

First consider the imposition of a just-binding minimum wage. To work out the impact of this change we can simply use the first-order condition for steady-state employment and differentiate to give the response of employment to changes in the wage. This gives the following:

$$\left[pF''(N) - cq(w)\left[1-\delta(1-q(w))\right]\right]\frac{\partial N}{\partial w} = 1 + cq'(w)N\left[1-\delta+2\delta q(w)\right] \tag{11}$$

The sign of the term on the left-hand side is negative so we have that:

$$\text{sgn}\left[\frac{\partial N}{\partial w}\right]=-\text{sgn}\left[1+cq'(w)N\left(1-\delta+2\delta q(w)\right)\right] \tag{12}$$

To work out the sign of the employment response to a just-binding minimum wage we need to use the first-order condition for the wage in (10). Using (10) to eliminate $cq'(w)N$ from (12) we can re-write it as:

$$\text{sgn}\left[\frac{\partial N}{\partial w}\right]=\text{sgn}\left[\left(1-q(w)\right)\left(1-\delta\right)+\delta q(w)\right]>0 \tag{13}$$

i.e. the impact of a minimum wage that just binds is to **raise** steady-state employment. This, of course, is in line with the prediction that would be made using a model of the employer as a monopsonists and the opposite of the prediction that would be made using a model of the employer as a competitive firm operating on their labour demand curve. Two notes of caution are needed here: first the exercise we have done is the application of a minimum wage to a single firm. In reality minimum wages apply across all or many firms in a market and one needs to model the interactions between firms in order to adequately investigate the employment impact of the minimum wage (e.g. the quit rate of this firm probably depends negatively on the wage paid by other firms). And secondly, the result only applies to minimum wages that just bind – push the minimum wage up and up and eventually one will enter a region where employment starts to fall.

As a second example consider the wage response to an increase in the price the firm gets for its output (or any revenue shock that raises the MPL within the firm). For a competitive firm this has no effect on the wage paid because the labour supply curve is infinitely elastic at the market wage so the firm can expand without paying higher wages. In contrast, a monopsony model predicts a rise in the wage paid as the firm wants to hire more workers. The model of quadratic adjustment costs predicts that the wage rises so is more like the monopsony model.

That the static textbook model of monopsony gives the right answers is no coincidence – there is a good reason for it. The following discussion is based on Manning (2006). There a 'generalised model of monopsony' is presented. That model introduces the idea of a labour cost function C(w,N) which represents the per-worker recruitment and training costs of maintaining employment at N when the employer pays wage w. Some examples of how this could be derived from an underlying model of recruitment costs and labour turnover might make the idea clearer. If recruiting and training a worker costs H(w) and the separation rate is s(w), a flow of sN recruits is needed to maintain employment at N so that C(w,N)=H(w)s(w). In this case, the labour cost function is independent of N. But, if it becomes increasingly hard to recruit and train workers so that the costs of training and recruitment are H(R,w) where R is the flow of recruits then the labour cost function will be of the form C(w,N)=H(s(w)N,w)s(w) in which case it will depend on employment.

What Manning (2006) shows is that if there are constant marginal costs of recruitment i.e. that C(w,N) does not depend on N then the model behaves like a competitive firm even though wages are a choice variable for the firm. But, if there are increasing marginal costs of recruitment then the model behaves like one of monopsony.

One can understand the earlier result about the model of dynamic employment adjustment with quadratic adjustment costs as a special case of this general result. To see this, let us derive the labour cost function for this model. If the firm pays wage w and wants a steady-state employment level of N then the total adjustment costs payable each period will be:

$$\frac{1}{2}c\left[\left(1-q(w)\right)N\right]^2 \tag{14}$$

which implies that the labour cost function (which is defined on a per-worker basis) is:

$$C(w,N)=\frac{1}{2}c\left[1-q(w)\right]^2 N \tag{15}$$

which is increasing in N. The result in Manning (2003) then says this is a 'monopsony' model as was argued earlier.

This section has argued that some labour market models that would commonly be thought of as 'competitive' models are really better thought of as models of monopsony.

3. Vacancies

So far I have contrasted the monopsony view of labour markets with the traditional perfectly competitive view of labour markets. It is helpful to think about what these two paradigms imply about the way the labour market looks from the perspective of agents on each side of the market, workers and employers. In the perfectly competitive model there is a going market wage for each type of labour and workers can always find a job paying that wage. Employers can hire any amount of workers at that wage but any attempt to cut wages would result in all workers immediately leaving the firm and any choice to pay higher wages only reduces profits. This is summarized in Table 1. In contrast the paradigm of monopsony implies that, from the perspective of workers, jobs are hard to find. As a conceptual framework, a job search model is the right way to think about the labour market from this point of view. Employers have choice over the wage paid.

The model of perfect competition has a strong hold over the minds of economists in general and labour economists in particular. But it simply does not pose much a contest. Workers clearly do think of getting and losing jobs as a big deal and employers do think they have discretion over the wages they pay as a look at any textbook n human resource management will confirm.

But there is another labour market paradigm that does pose more of a threat to the monopsonistic perspective. I shall loosely call this the efficiency wage model. In this view of the labour market (see Akerlof and Yellen, 1984; for a summary) employers choose to set wages substantially above market-clearing levels so that there is a chronic excess supply of labour to the firm or, to put it another way, involuntary unemployment. From the perspective of workers jobs are hard to find in this labour market so a job search model is the right way to think about things as in the monopsonistic framework. And employers have discretion over the level of wages they choose so this is also similar to the monopsony perspective. But, in contrast to monopsony, employment will be chosen so that the marginal product of labour is equal to the wage so the firm will be on its 'labour demand curve'. This is also summarized in Table 1.

How can one distinguish between monopsony and efficiency wage paradigms? Table 1 makes clear that they differ in whether employment is equal to the supply of labour to the firm or there is an excess supply of labour. Things are not quite that simple – the model of the previous section did not impose the constraint $H<R(w)$ where $R(w)$ is the flow of recruits to the firm so implicitly assumed an excess supply of labour to the firm yet had 'monopsonistic' results (see Manning, 1995; for a similar discussion using an efficiency wage model as its foundation).

But, the issue of whether there is a chronic excess supply of labour to a firm is a question of some interest. What information could be used to shed some light on this question? If there is an excess supply one might expect that firms find it very easy to fill vacancies so that vacancies are rare and vacancy durations short. Most data suggests that this is the case. For example, Table 2 presents some information on the duration of vacancy durations for vacancies placed with the UK employment service. They are much shorter than the corresponding duration for the unemployed being a handful of weeks and not months. Part of this is probably the result of a difference in the definition of 'duration' – a vacancy is recorded as filled when someone has accepted a job but probably before they have started whereas an unemployment spell only ends when a worker actually starts work (see van Ours and Ridder, 1992; for evidence that much of recorded vacancy durations are administrative delays). But, even after taking account of these problems, the duration of an average vacancy still seems much shorter than the duration of an unemployment spell suggesting a chronic excess supply of labour to vacancies.

However Manning (2003, chapter 9) argues that vacancy rates and durations contain little useful information about whether there is a chronic excess supply of labour. It sets up a model in which employers must invest money to create jobs ex ante but they cannot perfectly predict when workers will leave these jobs and how many workers will be in the pool of potential applicants at such time. In this model vacancies are simply 'accidents'. Manning shows that vacancy rates and durations can be very low if labour supply is close to predictable and job creation is costly. One way to understand this result is to think what happens when there is no uncertainty in the supply of labour to the firm.

In this case the number of jobs created will be equal to the optimal level of employment, the vacancy rate will be zero and 'miraculously' one applicant appears at the same time as a worker quits so that the vacancy duration is zero. But this is a monopsony model with all that implies so this means that vacancy rates and durations can tell us nothing about whether a firm is on its labour supply curve or not.

But there is one piece of information that is informative and that is the number of applicants for vacancies. In the example just given each vacancy that arrives has exactly one applicant i.e. a low number. But if there was a chronic excess supply of labour to the firm then we would expect to see large numbers of applicants. So the number of applicants is informative about the state of the labour market.

This does depend on a particular view of the process that matches vacancies and job seekers, namely stock-flow matching. If this is the right way to think about the labour market then the stock of job seekers see a new vacancy, a certain proportion apply so that a group of applicants arrive en mass. I think this is the right way to think about the matching process but the most common way in which it is modelled is the Diamond-Pissarides model in which job applicants arrive at the firm one-by-one. In this case the number of applicants is simply the number required to meet an acceptable one. In models with homogeneous workers (e.g. Pissarides, 1985) this will always be one. But, although this model has prominence in the literature it is simply not a good way to think about the matching process – see Coles and Smith (1998), Gregg and Petronglo (1997) and Coles and Petrongolo (2002) for evidence of the superiority of stock-flow matching.

The few studies of the number of applicants tend to find rather small numbers – less than 10 (see Holzer, Katz and Krueger, 1991; and Holzer, 1994; for US evidence and Brown et al, 2001; and Manning, 2000; for UK evidence). It is also not just the number of vacancies but the sensitivity to the wage that is of some interest. If there is a chronic excess supply of labour because the wage is 'so high' we would not expect to see much sensitivity in the number of applicants to the wage because all workers will be interested in all jobs.

But the evidence we have on the number of applicants is rather sparse so we supplement it here using two different datasets. To consider the number of job applicants and the sensitivity to the wage, we will use data from the UK Survey of Employer Recruitment Practices conducted in 1991/2[3]. This, as its name suggests, was a survey designed to provide information on the process by which workers are recruited. It asked some general questions about the way in which the employer recruited workers and then some questions about up to 5 specific engagements made in the previous 12 months. The method by which engagements were selected was a little complicated:

[3] I am very grateful to Michele Pellizari for drawing my attention to this data set.

occupational groups were ordered by the number of engagements made in that group and the employer was asked to select the most recent engagement form each of these groups, followed by the second most recent if there had been recruitment from less than five occupational groups.

Table 3 reports the information on the median and mean number of applicants by broad occupational group. This information is only collected for the first selected recruit so they may not be representative of recruits as a whole[4]. In this data set the average number of applicants for vacancies is larger than in some others with the mean being 40 and the median being 15.The big gap between the mean and the median is the result of approximately 10% of vacancies reporting more than 100 applicants – the mean among the other 90% is 20. There seems little systematic relationship between the occupational level of the vacancy and the number of applicants. It should be remembered that this data was collected in the depth of a severe recession in which the overall unemployment rate peaked at 10.9% at the end of 1992.

Table 4 then investigates the sensitivity of the number of applicants to the offered wage. As the dependent variable is a count variable we estimate a negative binomial model. In the first row we simply regress the number of applicants on the log wage. One can see that the wage is significantly positively related to the number of applicants consistent with the monopsony view of the labour market. The second through third rows then investigate the robustness of this inclusion – the second row includes control on the job notably the occupation and whether the job is temporary. The estimated sensitivity of recruits to the wage is higher in this specification. The third and fourth rows then introduce firm controls – this substantially reduces the sample size[5] so the third column reports the estimate of the specification of the second column but on the restricted sample for which firm-level information is available. The coefficient on wages in the fourth row falls but is still very significant and quite large. For this specification we also report the coefficient on the log of establishment size as an important question is whether large firms have an advantage in recruitment (see Manning, 2003; ch 10, for a detailed discussion of this). One might expect that large firms do have an advantage because many workers find out about jobs through personal contacts and large firms are likely to be able to reach out to a larger number of people through the persona contacts of their workers. In the extreme case, known as balanced matching (see Burdett and Vishwanath, 1988) we would find that the coefficient on log employer size is one. As can be seen the estimated coefficient is a long way from this and, although positive, is not statistically significantly different from zero. Large firms do not seem to have a particular advantage in recruiting workers

---

[4] The data set does provide some weights but the way in which these are derived is not entirely clear. The reported results are unweighted but the weighted results are not very different.

[5] Quite why this is so is a mystery on which the documentation of the survey sheds little light. The firm level information is collected in the questionnaire prior to the vacancy information so it seems odd that there should be so much more information on the vacancy events.

One problem with using data on the number of applicants to investigate the extent of choice that employers have is that many of these applicants may be totally unsuitable so should not really be counted in the applicant pool. Of course there is not a precise dividing line between the suitable and the unsuitable but one way of considering this issue is to consider the number of applicants short-listed as, presumably, all of these are deemed potentially suitable. It is likely that this is an under-estimate of the size of the possible applicant pool as the short-list is limited in size but how this varies with the wage is of some interest. The second half of Table 4 changes the dependent variable to the number of short-listed applicants. Again, this is significantly positively related to the wage though the elasticity is smaller than for the number of applicants.

All the discussion so far has assumed that the impact of offering a higher wage is to increase the quantity of applicants to the firm. But it is also important to realise that there is also an impact on the quality of the recruit if there is a larger pool of applicants from which to choose. Evidence from the study of low-wage employers reported in Manning (2000) suggests that very few vacancies go unfilled so that the main impact of a larger applicant pool is unlikely to be on the probability of filling the vacancy but on the quality of the recruit. Once one recognises the importance of labour quality, the distinction between competitive and monopsony models of the labour market becomes a little blurred. For example, the earlier discussion stated that the employer in a perfectly competitive market has no discretion over the wage that it pays: while this is true conditional on the decision to employ a certain quality of worker, the decision on the wage to pay is really a decision about the quality of worker to employ. So, a competitive view of the labour market would expect the quality of recruits to be positively related to the wage just as a monopsony model does.

Table 5 investigates the sensitivity of the quality of recruits to the offered wage. We use three measures of worker quality – whether the recruit has previous relevant experience, whether the recruit was previously employed and the level of reported satisfaction with the recruit. All of these measures are significantly related to the wage whether one includes job or firm controls. Indeed for the last two measures of worker quality we have the information for all recruitment events so can include firm fixed effects. We still find a large significant effect of wages on worker quality. The evidence in Table 5 suggests that the offered wage does affect the quality of recruits.

As the discussion earlier in this section has made clear this evidence is not particularly useful in distinguishing between competitive, monopsony and 'efficiency' wage models of the labour market. But the competitive model suggests a large number of applicants all of the same quality, something that does not seem in line with the perceptions of personnel managers who do think of there being significant variation in labour quality among applicant pools (see Manning, 1994; for more discussion of this).

The analysis of the UK Survey of Employer Recruitment Practices has suggested a rather larger number of applicants for jobs than other studies. But, perhaps this was because it was conducted in a recession. As one final piece of evidence we look at the perspective of job search from the worker and not the employer point of view. If we have information on the number of jobs applied for and the number offered then this gives us (assuming all vacancies are filled all offers accepted) an estimate of the average number of applicants as the number of applications divided by the number of offers. To investigate this we use data from the Survey of Unemployment and Benefits, a survey of those claiming unemployment-related benefits which was done in two waves, one in 1995 and one in 1997. This was a period in which the labour market was doing relatively well and unemployment was falling to levels not seen for a generation. Among many other questions, the respondents were asked how many job applications they had made in the previous 4 weeks, how many interviews they had had and how many offers they had had. Table 6 presents information on the mean and median of the number of applications made, interviews received and offers made for the two waves of the sample. In 1995 the unemployed made an average of 6.5 job applications in the 4-week period receiving 0.22 offers. The ratio is 29.6 suggesting approximately 30 applicants per vacancy. For 1997, when the labour market was stronger both the number of applications and offers are higher but the ratio of the two is very similar.

A problem with using this data for the purpose of inferring the average number of job applicants from this data is that the data refer only to the experience of the unemployed and we know that approximately half of all jobs are filled by someone who already has a job. To see the bias induced by this denote the number of job applications made by the employed as A^e and by the unemployed by A^u. Similarly denote the number of job offers received by the employed by O^e and by the unemployed by O^u. Our estimate of the average number of applicants is then given by:

$$\mu = \frac{A^e + A^u}{O^e + O^u} \tag{16}$$

but all we observe is:

$$\mu^u = \frac{A^u}{O^u} \tag{17}$$

One can write (16) as:

$$\mu = \frac{A^e + A^u}{O^e + O^u} = \alpha\mu^e + (1-\alpha)\mu^u \tag{18}$$

where α is the share of job offers going to the employed. (18) shows that only if the employed and unemployed have similar chances of an application being successful will the application-offer ratio for the unemployed be a good measure of the overall number of applicants per vacancy. However, the data in Manning (2000) suggested that the unemployed do not have a lower 'hit rate' than the employed in converting applications into offers.

The pattern of the data described here is mixed. Previous studies have suggested rather small numbers of applicants for vacancies – less than 10- but both the data sets considered here suggest larger numbers – in the region 15-25. We need more information to be able to resolve this question.

In this section we have argued that information on the number of applicants for vacancies and how responsive is this number to the wage paid is of interest in deciding on the appropriate paradigm for the way the labour market works.

4. An Agenda for Future Research

This paper has identified a number of areas where we do need better information. One is the nature of adjustment costs, in particular whether the per-worker costs are increasing, constant or convex in the scale of recruitment. Although there is an enormous amount of research on dynamic models of labour demand there is remarkably vague knowledge about this question. We also need more information on the process by which employers fill vacancies. Are the wages offered in vacancies so high that employers can easily fill all of their jobs? Or is the wage such that the supply of applicants acts as a constraint on the level of employment of firms?

Conclusion

This paper has argued that, despite initial appearances to the contrary, there is no inevitable contradiction between the view that employers have some monopsony power over their workers and the enormous literature on the economics of labour demand. Indeed, popular models of 'labour demand' with convex costs of adjustment are really better thought of as models of monopsonistic firms. But the 'labour demand' literature could do with some re-focusing as the study of the wage and employment decisions of employers. Among the issues identified by this paper as needing further research are:

- estimates of the labour supply curve to individual employers – labour economics' 'missing literature'

- estimates of the costs of hiring and the shape of these costs

- estimates of the shape of employment adjustment costs

- studies of the process of recruitment

All of these require data on individual employers that has not been readily available in the past but, hopefully, will be increasingly available in the future.

TABLE 1. DIFFERENT PARADIGMS OF THE LABOUR MARKET

	Perfect Competition	Monopsony	Efficiency Wage
Worker Perspective	Jobs always available	Jobs hard to find – 'search model'	Jobs hard to find – 'search model'
Employer Discretion over wages	No	Yes	Yes
Employment always on...	Labour Demand Curve	Labour Supply Curve	Labour Demand Curve

TABLE 2. THE DURATION OF VACANCIES

Occupation	Mean (weeks)	Median (weeks)
Managers/Administrators	4.65	2.91
Professional Occupations	5.15	2.87
Associate Professional/Technical	3.32	0.49
Clerical/Secretarial Occupations	3.01	1.34
Craft/Related Occupations	3.02	1.12
Personal/Protective Service Occupations	3.83	1.63
Sales Occupations	3.32	1.42
Plant/Machine Operatives	2.22	0.55
Other Occupations	2.49	0.81

Notes:

1. Source: NOMIS www.nomisweb.co.uk

2. Data relate to duration of filled vacancies placed with Job Centres.

3. Data are simple averages for period July 1992-July 2000.

TABLE 3. THE NUMBER OF APPLICATIONS

Occupation	Median	Mean	Number of Observations
Routine unskilled	10	26	273
Operative/Assembly	15.5	52	304
Sales	12	31	303
Protective/Personal services	10	35	173
Craft/Skilled	15	50	172
Clerical/Secretarial	20	45	476
Associate professional/technical	18	38	189
Professional	15	36	300
Managerial	24	47	121
Total	15	40	2311

Notes:

1. Data are from the 1992 Employment Recruitment Practices Survey.

TABLE 4. THE INFLUENCE OF WAGES ON THE QUANTITY OF APPLICANTS

Dependent Variable	Log (hourly wage)		Log (total employment)		Job Controls	Firm Controls	Observations
Number of Applicants	0.372	[0.062]			No	No	2067
Number of Applicants	0.467	[0.078]			Yes	No	2063
Number of Applicants	0.496	[0.143]			Yes	No	597
Number of Applicants	0.416	[0.156]	0.051	[0.039]	Yes	Yes	597
Number Short-Listed	0.147	[0.054]			No	No	2265
Number Short-Listed	0.304	[0.062]			Yes	No	2261
Number Short-Listed	0.591	[0.119]			Yes	No	658
Number Short-Listed	0.535	[0.115]	0.009	[0.029]	Yes	Yes	658

Notes:

1. *data are from the 1992 Employment Recruitment Practices Survey.*
2. *Method of estimation is negative binomial model. Standard Errors in parentheses.*
3. *Job controls are 9 occupation dummies and a dummy for whether the job is temporary*
4. *Firm controls are 10 industry dummies, 11 regional dummies and a dummy for public sector.*

TABLE 5. THE INFLUENCE OF WAGES ON THE QUALITY OF RECRUITS

Dependent variable	Log (hourly wage)		Job Controls	Firm Controls	Observations
Previous Experience	-0.786	[0.037]	No	No	4800
Previous Experience	-0.944	[0.053]	Yes	No	4791
Previous Experience	-0.989	[0.093]	Yes	Yes	1307
Previously Employed	0.346	[0.007]	No	No	18653
Previously Employed	0.293	[0.010]	Yes	No	18617
Previously Employed	0.289	[0.014]	Yes	Yes	18617
Satisfaction with Recruit	-0.183	[0.011]	No	No	17434
Satisfaction with Recruit	-0.137	[0.015]	Yes	No	17400
Satisfaction with Recruit	-0.101	[0.019]	Yes	Yes	17400

Notes:

1. data are from the 1992 Employment Recruitment Practices Survey.

2. The variable 'previous experience' is on a five-point scale where 1 means the recruit has experience on similar tasks and 5 means the recruit has no experience. The variable 'previously employed' is a binary variable taking the value one if the recruit was previously employed and zero otherwise. The variable 'satisfaction with recruit' is on a four-point scale where 1 is very satisfactory and 4 is not at all satisfactory.

3. Linear regression models estimated. Standard Errors in parentheses which are clustered on the firm where firm controls are not included.

4. Job controls are 9 occupation dummies, a dummy for whether the job is temporary, and dummy variables for the 'order' of the vacancy.

5. Firm controls are 10 industry dummies, 11 regional dummies and a dummy for public sector for the 'previous experience' variables and firm fixed effects for the other two variables.

TABLE 6. THE HIT RATE FOR JOB APPLICATIONS BY THE UNEMPLOYED

	1995	1997
Average Number of Applications	6.5	8.2
Average Number of Interviews	0.72	0.9
Average Number of Offers	0.22	0.27
Implied Number of applications per vacancy	29.6	30.4
Number of observations	4350	4290

Notes:

1. *Data come from two cohorts of Survey of Unemployment and Benefits.*
2. *Data refer to previous 4-week period*
3. *Data are weighted to be representative of stock of unemployed.*
4. *Implied number of applicants per vacancy is obtained by dividing the average number of applications by the average number of offers.*

FIGURE 1. THE TEXTBOOK MODEL OF MONOPSONY

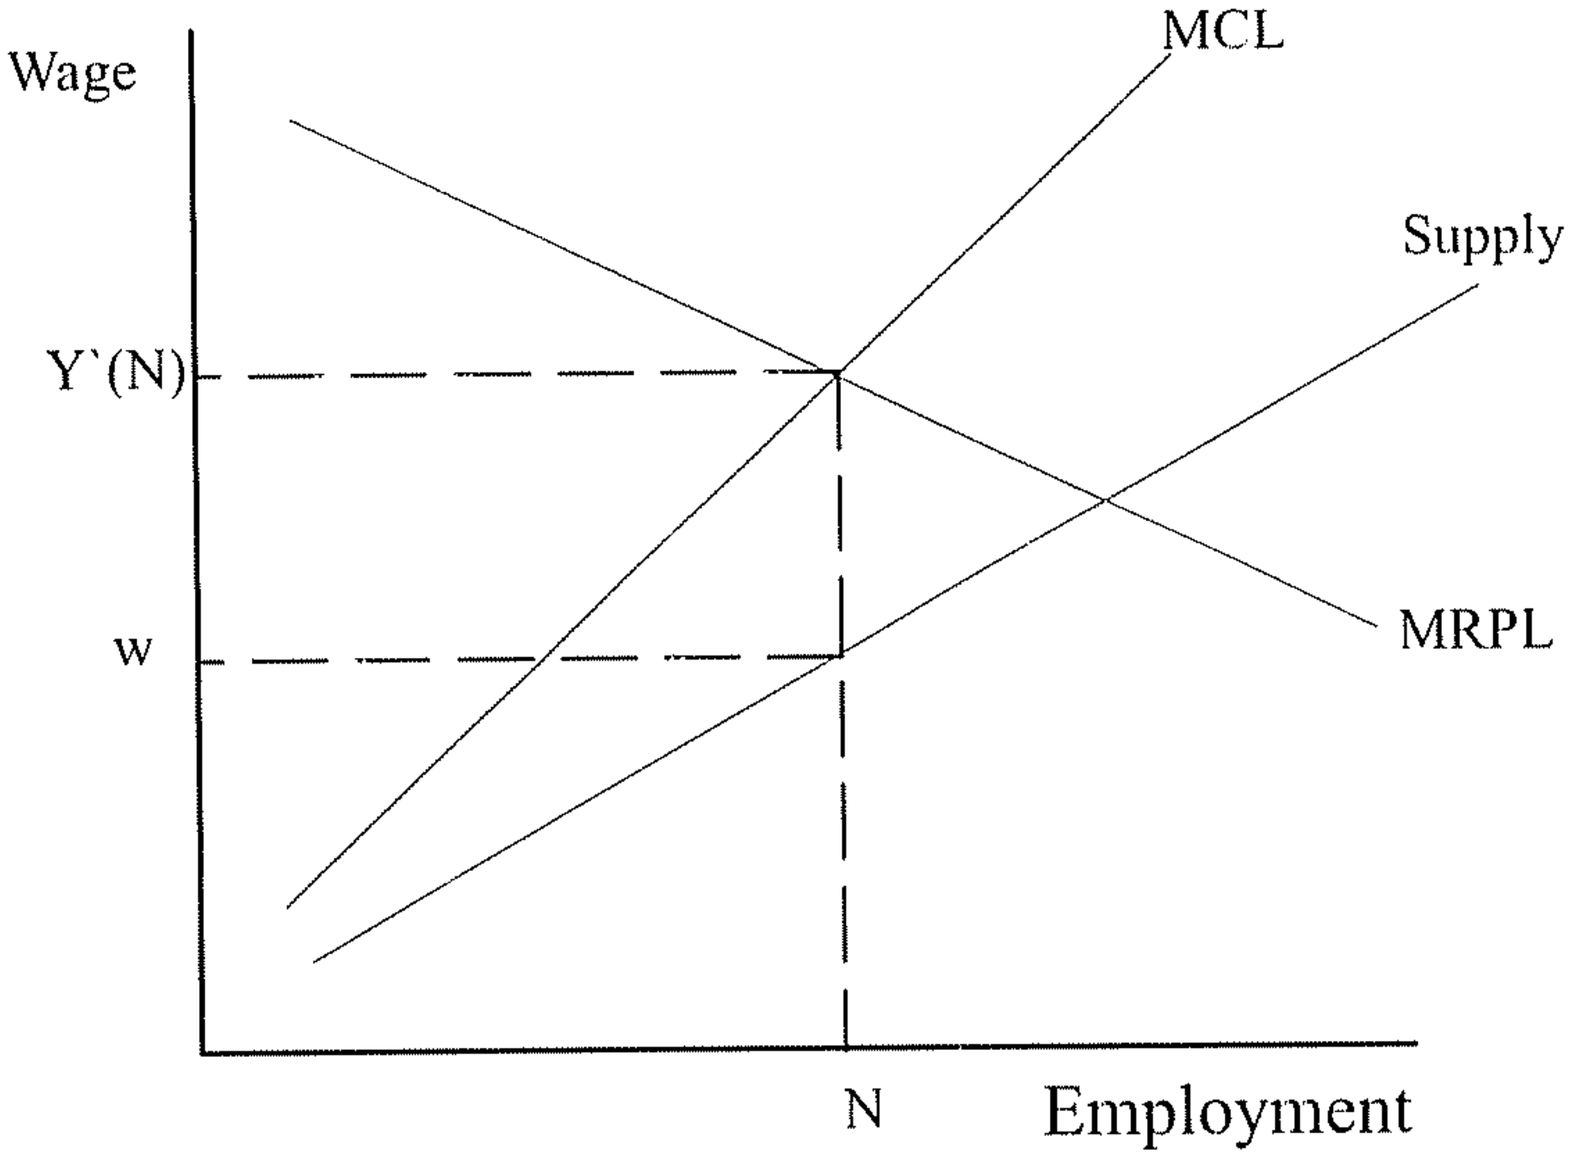

REFERENCES

Akerlof, G. A. and J. L. Yellen, 1986. "Efficiency Wage Models of the Labor Market" Akerlof, George, A., and Janet L. Yellen (eds) *Efficiency wage models of the labor market*. Cambridge: Cambridge University Press, pages 1-21.

Brown, D., R. Dickens, P. Gregg, S. Machin and A. Manning, 2001. *Everything under a Fiver: Recruitment and Retention in lower paying labor markets*, London, Joseph Rowntree Foundation.

Burdett, K., and T. Vishwanath, 1988. "Balanced Matching and Labor Market Equilibrium", *Journal of Political Economy*; 96(5), October, pages 1048-65.

Coles, M. G., and B. Petrongolo, 2002. "A Test Between Unemployment Theories Using Matching Data", CEPR Discussion Paper, No. 3241, March 2002.

Coles, M. G. and E. Smith, 1998. "Marketplaces and Matching" *International Economic Review*; 39(1), February, pages 239-54.

Diamond, P. A., 1981. "Mobility Costs, Frictional Unemployment, and Efficiency", *Journal of Political Economy*; 89(4), August, pages 798-812.

Gregg, P. and B. Petrongolo, 1997. "Random or Non-Random matching? Implications for the use of the UV curve as a measure of matching effectiveness", LSE CEP Discussion Paper, No. 348, June 1997.

Hamermesh, D. S., 1993. *Labor demand*, Princeton: Princeton University Press.

Hamermesh, D. S. and G. A. Pfann, 1996. "Adjustment Costs in Factor Demand", *Journal of Economic Literature*, 34, 1264-1292.

Holzer, H. J., 1994. "Job Vacancy Rates in the Firm: An Empirical Analysis" *Economica*; 61(241), February, pages 17-36.

Holzer, H. J., L. F. Katz and A. B. Krueger, 1991. "Job Queues and Wages", *Quarterly Journal of Economics*; 106(3), August, pages 739-68.

Manning, A., 1994. "Labour Markets with Company Wage Policies", London School of Economics Centre for Economic Performance Discussion Paper No. 214, December.

Manning, A., 1995. "How Do We Know That Real Wages Are Too High?", *Quarterly Journal of Economics*, Vol. 110, pp1111-1125.

Manning, A., 2000. "Pretty Vacant: Recruitment in Low-Wage Labour Markets" *Oxford Bulletin of Economics and Statistics*; vol 62, Special Issue Dec., pages 747-70.

Manning, A., 2005, "A Generalised Model of Monopsony", *Economic Journal*, 116(1), 84-100.

Mortensen, D., 1970. "A Theory of Wage and Employment Dynamics" in E.S. Phelps et al (eds) *Microeconomic Foundations of Employment and Inflation Theory*, New York: Norton.

Pissarides, C. A., 1985. "Short-run Equilibrium Dynamics of Unemployment Vacancies, and Real Wages", *American Economic Review*; 75(4), September, pages 676-90.

Van Ours, J. C. and G. Ridder, 1992. "Vacancies and the Recruitment of New Employees", *Journal of Labor Economics*; 10(2), April, pages 138-55.

Weiss, A. W., 1980. "Job Queues and Layoffs in Labor Markets with Flexible Wages", *Journal of Political Economy*; 88(3), June, pages 526-38.

Weiss, A, W., 1990. *Efficiency wages: Models of unemployment, layoffs, and wage dispersion*, Princeton: Princeton University Press.