Abstract

The 'fractal' nature of the rise in earnings dispersion is one of its key features and remains a puzzle. In this paper, we offer a new perspective on the causes of changes in earnings dispersion, focusing on the role of labour reallocation. Once we drop the assumption that all firms pay a given worker the same, the allocation of workers to firms matters for the dispersion of earnings. This perspective highlights two new factors that can affect the dispersion of earnings: rates of job and worker reallocation, and the nature of the process allocating workers to jobs. We set out a framework capturing this idea and quantify the impact of reallocation on earnings dispersion, using a dataset that comprises almost the universe of workers and the universe of employers in Maryland. We show that these factors have potentially large effects in general on earnings dispersion. In the case of Maryland over the period 1985-1994, the changing allocation of workers to jobs played a significant role in explaining movements in the dispersion of earnings.

JEL Codes: J31, D31

Keywords: Earnings inequality, labour reallocation, matched worker and firm panels

Jobs, Workers and Changes in Earnings Dispersion

Simon Burgess, Julia Lane and David Stevens

April 2001

Published by
Centre for Economic Performance
London School of Economics and Political Science
Houghton Street
London WC2A 2AE

ISBN 0 7530 1465 3

Individual copy price: £5

Jobs, Workers and Changes in Earnings Dispersion

Simon Burgess, Julia Lane and David Stevens

The Centre for Economic Performance is financed by the Economic and Social Research Council.

Acknowledgements

This research is partially supported by a grant from the Sloan Foundation through the Urban Institute. The first version of this was presented at the Society of Labor Economists meetings in May 1997. We are grateful to Francis Kramarz for first suggesting the estimation technique to us and to John Abowd for his suggestions. John Haltiwanger also gave us useful comments.

Simon Burgess is a member of the Centre for Economic Performance, London School of Economics and of CEPR. He is a Professor of Economics at the University of Bristol. Contact: simon.burgess@bristol.ac.uk. Julia Lane is at the American University and Census Bureau. Contact: jilane@american.edu. David Stevens is a Professor at the University of Baltimore, Maryland.

1. Introduction

Changes in earnings inequality are a well documented feature of the 1980s in the U.S. and a number of other countries (see OECD, 1996). One of the main features that has been cited as lying behind this is a rise in the return to skill[1]. Beyond this, however, the literature has established that a large component of the level *and growth* in dispersion is within-group (Levy and Murnane, 1992; Moffitt 1990, Burtless 1990). The 'fractal' nature of the change is one of its key features and remains a puzzle. In this paper, we offer a different perspective on the causes of changes in earnings dispersion, focusing on the role of labour reallocation.

We argue that the scope for labour reallocation to influence the earnings distribution is large. Once we drop the assumption that all firms pay a given worker the same, the allocation of workers to firms matters for the dispersion of earnings. The continual re-sorting of workers across different firms paying different wage *premia* means that the economic process matching workers to firms is an important potential determinant of the nature of the earnings distribution. This perspective highlights two new factors that can affect the dispersion of earnings: rates of job and worker reallocation, and the nature of the process allocating workers to jobs. We set out a framework capturing this idea and quantify the impact of reallocation on earnings dispersion, using a dataset which comprises almost the universe of workers and the universe of employers in Maryland.

The rest of the paper is organised as follows: in the next section we briefly and selectively summarise the literature on earnings inequality. Next, in Section 3 we set out the relationship between labour reallocation and earnings dispersion. In Section 4 we describe the data, the selection decisions we made and estimation procedures. Section 5 presents our results and Section 6 concludes and discusses the implications of this approach for understanding the widening inequality of the 1980s.

2. Background

Of the vast literature on the determination of earnings, the overwhelming majority is concerned with modelling the mean. However, interest in the dispersion of earnings has risen alongside the well-established increase in earnings inequality. One substantial factor in this increase is the rise in within-group inequality, groups typically being defined by gender, age and education level. Levy and Murnane (1992) discuss a number of ideas that have been advanced to explain this, but much of the rise remains unexplained.

One prominent strand in this work argues that the return to unobserved skill has risen. One approach is to propose a statistical structure for the underlying unobservable (see for example Juhn, Murphy and Pierce, 1993, and Card and Lemieux, 1996); another is to use a measure of ability such as test scores (see for example Blackburn and Neumark, 1993, and Murnane, Willet and Levy, 1995). Both of these approaches are useful but both face difficulties: the former in the identifying assumptions required, the latter in the sometimes problematic nature of the variables used. In our approach, we take an observable variable, the identity of the worker's employer, and examine the contribution to changes in earnings dispersion by changes in the allocation of workers to firms.

A number of authors have explored the role of establishment effects on wages. Groshen (1991) finds that plant-specific wage differentials explain 27% of residual wage variation. Mitchell (1991) argues that idiosyncratic pay setting practices have become much more important, as the role of unionisation and wage 'contours' in the economy diminishes. Davis and Haltiwanger (1991) use a panel of manufacturing plants to investigate the contribution of plant-specific wage effects to the growth in wage inequality. They find both between- and within-plant dispersion to have changed. Establishment size is a particularly important correlate of the change in wage inequality. Davis and Haltiwanger (1995) have revisited the issue of establishment size and wage inequality, and again show size to be a key factor. The argument of these papers is that both the distribution of size and the impact of size on earnings have changed over time, and that these facts are part of the explanation of rising wage inequality.

In this paper we explore the contribution to earnings dispersion of workers moving between plants that pay differing wage *premia*. The idea that the assignment of workers to jobs matters for the earnings distribution is not new. Roy (1951) proposed a model in which workers had differing talents in different occupations or sectors. Their ability in one sector could be correlated or not with that in other sectors, and they choose which sector to work in simply to maximise income. The overall earnings distribution will depend on the allocation of workers to sectors, which in turn depends through the self-selection mechanism on the means, variances and covariance of the ability distributions. While Roy's presentation is not

[1] See among many others: Bound and Johnson, 1992, Katz and Murphy 1992, Murphy and Welch 1992, and Gottschalk and Moffitt, 1994.

technical, he clearly had a formal model set out; formal treatments are available in Maddala (1983), Heckman and Honore (1990) and Sattinger (1993). This model and related work are referred to as assignment models. The majority of work in this field is theoretical but two empirical implementations of the Roy model are provided by Heckman and Sedlacek (1985, 1990); see also Heckman and Honore (1990) on identification in this model.

Whilst obviously related, our approach has a different econometric set-up. In the Roy model, as in ours, workers will be paid differing amounts as they move between different employers. But in the Roy model, these differing amounts derive from individual-specific ability distributions across jobs. This implies that two workers both re-assigned between the same two employers will in general see their wages change in completely unrelated ways. In our model set out below, wages are decomposed into two components: the standard human capital worker-specific component, common across employers, and an employer-specific component, common across workers[2]. Hence two workers re-assigned between the same two employers will see their wages change by the same (log) amount.

There is another interesting link between our work and models that descend from Roy's original idea. Willis (1986) also focuses on the impact on the earnings distribution of different worker allocation to different occupations where workers possess different individual abilities. He examines the impact of differential sorting of types of individual abilities – positive and negative hierarchical sorting, as well as non-hierarchical sorting. In a related spirit, we examine the impact on the earnings distribution of different types of sorting of workers across firms.

A few recent models have addressed how the changing structure of organisations and the technology of production might affect the matching of worker types and firm types These include Kremer's (1993) O-ring model of production, predicting that with a particular type of production function high-skill workers will be matched together. More recently, Kremer and Maskin (1996) have explored the links between the nature of the production technology and the degree of segregation by skill between firms. The key factors are imperfect substitutability of skills, the existence of different but complementary tasks, and that these tasks are differentially sensitive to skill. Given this, the distribution of skill types available affects the incentives for skill segregation, and hence the effect on earnings inequality. Lindbeck and Snower (1996) also discuss the transformation of the production process and the implications for inequality.

---

[2] Note that this ignores the match-specific component we discuss in the next section; this component is ignored in the empirical work as discussed below.

These changes clearly have implications for earnings inequality. If some firms, as the literature suggests, pay a high premium and others pay a low premium, then the distribution of earnings will depend on how workers are distributed across those firms. Changes in earnings inequality will occur if job reallocation and changes in organisational structure result in changes in the differential attachment of high skill workers to high premium firms and low wage workers to low wage firms. The literature has emphasised that changes in the twin fundamentals of 'technology and trade' can change the skill premium. We show that they can work through other channels: changes in the allocation of workers to firms, and changes in the size and number of high wage firms.

3. The Reallocation of Labour and Earnings Dispersion

In this section we first briefly review models of earnings determination, and secondly show how the process of labour reallocation affects the earnings distribution.

3.1 Determination of earnings

Explanations of change in the dispersion of earnings must start from a model of earnings. The standard approach is a human capital model:

$$\log w_i = b_0 + b_1 S_i + b_2\, sex + b_3\, race + f(t) + b_4\, unemp + e_i \qquad (1)$$

where w_i denotes earnings, S_i refers to years of schooling; $f(t)$ is a function of age, and *unemp* is a measure of aggregate labour market activity. In this framework, a change in the dispersion of individual earnings must stem from changes in one of these components, namely changes in the weights given to individual characteristics (β's) and changes in f(.). This is the approach that has been taken in explaining the rise in earnings dispersion in the U.S. and other countries, with the emphasis being on a rise in the return to skill (b_1 rising over time). As we noted above, however, a large component of the widening of the earnings distribution is accounted for by an increase in within-group dispersion. In terms of (1), this arises from changes in the distribution of e, and remains a puzzle.

The implicit assumption underlying this approach is that all firms[3] pay the same wage to a given worker. However there is a large literature (for example, the efficiency wage literature) which relaxes this by introducing incentive or information problems. We therefore reformulate model (1) to allow for different firm-specific wage *premia*. We think of each individual i as possessing a bundle of skills, $b_t X_{it}$, where the vector X is given by the explanatory variables in (1). The earnings of individual i then depend on her human capital bundle, the earnings mark-up of the firm j she is currently working for, and an error term:

$$w_{ijt} = p_{ijt}.(b_t X_{it}).exp(e_{ij(i)t}) \qquad (2)$$

where p_{ijt} is firm j's valuation of i's skills bundle at t[4]. This may have two components: a firm specific effect denoted a_{jt}, and match specific effect denoted g_{ij}. In principle we can identify a time-varying firm-specific effect, but not a time-varying match specific effect[5]. Taking logs,

$$\log w_{ijt} = a_{jt} + g_{ij} + \log y_{ij(I)} + e_{ij(i)t} \qquad (3)$$

where $a_{jt} = \log a_{jt}$, $g_{ij} = \log g_{ij}$and y_{it} is the index of the human capital bundle, $\psi_{it} = \beta_t X_{it}$. We plan to investigate the role of g_{ij} in future work, but for now we ignore it.

3.2 Labour reallocation and the distribution of earnings

Clearly, reallocation of workers across job slots from one year to the next has no effect on the earnings distribution if all firms pay the same to any given worker. Under this assumption, the earnings distribution is entirely driven by the distribution of human capital variables among workers, and the prices given to those characteristics. The role played by the demand side, by firms, technology, trade and so forth, is purely to determine these prices. This route has been followed by many papers, investigating the impact on earnings of technology adoption and trade penetration among other things.

However, if firms do choose different wage mark-ups, the allocation of workers with

---

[3] In this theoretical discussion, we shall refer to employers as firms. In the empirical work below, note that most employers are single establishment entities; see the Data appendix for details.

[4] A more general model would allow the firm-specific premium to differ for each component of the vector X. Our dataset does not permit us to investigate that.

different human capital bundles to firms paying different *premia* does matter. We index workers by their human capital bundle, y_{it}, and firms by their mark-up, a_{jt}. There is a static element to this and a dynamic element. The former is the nature of this allocation or assignment. Denote the function describing this allocation or 'pairing'[6] as $f_t(a_{jt}|y_{it})$: the probability of finding a worker with index y_{it} in a firm with mark-up a_{jt}. For any given set of worker skill indices and firm mark-ups, this allocation will influence the earnings distribution.

This makes it clear that reallocation potentially affects the residual variation, that is 'within-group' earnings dispersion. For a set of workers defined by a particular value of y, the dispersion of their earnings depends on the distribution of firm effects they are attached to. If the pairing function $f_t(a_{jt}|y_{it})$ changes, then so will this distribution and hence 'within-group' earnings dispersion will change.

Strong positive sorting will clearly lead to a more dispersed distribution than random sorting. The economic behaviour underlying f_t is the decision by firms on the appropriate wage premium and hiring and firing policy (i.e. *who* to hire or fire), and the job acceptance decisions of workers. If all job slots remained owned by the same firms, and if all workers stayed at the same firm, then the earnings distribution would be defined by the distribution of y_{it}, the distribution of a_{jt} and $f_t(a_{jt}|y_{it})$.

But recent research has shown that labour markets are extremely dynamic: there are very high rates of job and worker reallocation (see Davis and Haltiwanger, 1992, Anderson and Meyer, 1994, and Burgess, Lane and Stevens, 2000). This continual resorting of workers to job slots also influences the steady state earnings distribution, as new pairs of worker and firm indices are always being created and old ones destroyed. There are three distinct components to this.

First, we can think of pure job reallocation with no net employment growth (as defined by Davis and Haltiwanger, 1992) as some proportion of a fixed total of job slots changing ownership between firms each period. This changes the distribution of wage premia over the population of job slots (equivalently, the population of workers). A set of workers (defined by their index, y_{it}) are displaced from the declining firms, and a set of worker are hired into the growing firms[7]. Which workers are selected by which firms

---

[5] In practice, data issues force us to assume time-invariant firm-specific effects – see below.

[6] Matching is another appropriate description, but this word is already used in at least two other senses.

[7] In a simple model this would be the same set, but a general model would allow for movement into and out of employment.

(defined by their wage premium, a_{jt}) and which firms hire which workers will both reflect the firms' desired workforce: that is, the allocation function described above.

The second element is turnover of the workforce on top of that arising from job reallocation. We have shown elsewhere (Burgess, Lane and Stevens, 2000) that this worker churning is substantial, and indeed dominates job reallocation as a source of worker turnover. In this context, churning involves a set of workers being reallocated across a fixed set of job slots. High rates of churning mean that a great deal of this reallocation happens each period.

The third element is the *nature* of the reallocation. This element, which is the correlation between the worker type and the firm type, will also influence the distribution of earnings, depending on how workers and firms are paired. We discuss this further below.

We can illustrate these ideas with a brief example, before setting out a more general framework. Assume that there are a fixed number of workers, indexed 1 through 5, and a fixed number of job slots, indexed A through E. Each worker and each job slot has a fixed effect (in this simple example we ignore age, time and labour market conditions). The worker fixed effect is equal to the worker number, and the job fixed effect is indicated in parentheses. Earnings are simply the sum of the worker and job slot fixed effects, so for example, worker 1 in job slot A(10) earns 11 units in t and worker 4 in D(1) earns 5. Workers are allocated to job slots as follows, and then reallocated in t+1:

	Date:	t	t+1
		Workers	
Job Slots	A (10)	1	1
	B (1)	2	2
	C (1)	3	4
	D (1)	4	5
	E (10)	5	3

Note that both job slot and worker fixed effects are constant over time. Nevertheless, he variance of earnings will change over time as a result of worker reshuffling. The effect of such worker reallocation in the example, in t+1 workers 3, 4 and 5 trade places, so worker 5 now earns 6 units and worker 3 earns 13. The variance of earnings changes from 21.4 to 14.2 as a result of the reallocation.

Job reallocation with no net employment growth, or balanced job creation and destruction, can be thought of as job slots swapping between the ownership of different firms. So for example, if in t+2 the firm owning job slot A grew by one, and the firm owning job slot B died, we can think of the fixed effect of that job slot changing from 1 to 10. Suppose also that the firm owning D grew by one, and the firm owning E died. Workers #2 and #3 are displaced and subsequently re-hired into the two new jobs, changing the earnings distribution.

However, this does not isolate the contribution of job creation and destruction on earnings dispersion. The third component in our framework emphasises the importance of pairing: the change in the dispersion depends on how the workers released from the jobs destroyed are paired with the new jobs created. If worker #2 is hired at B and #3 at E, the variance is 13.3. But if they are paired in the opposite manner, the variance is 17.8. Thus the contribution of job reallocation to changes in the earnings distribution is influenced by ϕ_t.

The *scope* for reallocation to affect the variance of earnings depends on the dispersion of worker and firm indices. If either is close to degenerate, reallocation can only have a minor influence. The *actual impact* of reallocation depends on two further things: the amount of labour resorting and the nature of the process matching workers and firms. We know the former is huge; we know very little about the latter. Note that even given constant worker and firm indices *and* constant rates of labour turnover, reallocation can still produce changing earnings dispersion if $f_t(a_{jt}|y_{it})$ changes over time.

To reiterate: the relaxation of the assumption that all firms pay the same introduces a different set of variables that can influence the dispersion of earnings. These include in particular, the degree of job and worker reallocation and the nature of the allocation of workers to firms.

We can use our modified human capital equation, (3) ignoring the g_{ij} term, to investigate changes in the cross sectional variation in log w_{it} over time:

$$\mathrm{var}_t(\log w_{ij}) = \underset{1}{\mathrm{var}_t(y_i)} + \underset{2}{\mathrm{var}_t(a_{jt})} + \underset{3}{2\,\mathrm{cov}(\psi_i, a_{jt})} + \underset{4}{\mathrm{var}_t(e_i)} \tag{4}$$

This decomposition demonstrates that the earnings distribution can change because the distribution of worker indices changes (1), the distribution of firm mark-ups changes (2), the allocation of workers across jobs changes, $f_t(a_{jt}|y_{it})$ (3), or because of further unmodelled changes (4). We have suppressed other terms in this expression: the cross-sectional variance in f(t) and in local labour market conditions are zero in our analysis below, but clearly could be incorporated; covariances with the error term are assumed zero in this theoretical exposition.

We discuss each of these in detail:

(1) The variance of worker indices can change for two reasons. The first of these is changes in the composition of workers, either on a secular or business cycle basis; the second is changes in the indices of a given group of workers - for example, changes in the price of skill.

(2) The variance of firm wage mark-ups can also change for the same two reasons: a change in the composition of firms, or a change in the value of the mark-up among a given set of firms. The former, which essentially refers to changes in the population of job slots, can be characterised as the contribution of job creation and destruction, "industrial restructuring", and the entry and exit of firms. The latter can occur if firms adopt new technologies and personnel practices. It may also include short-term adjustments to the premium to reflect labour market conditions.

(3) The covariance term captures the changing allocation of workers to jobs. This covariance is a measure of the outcome of the underlying 'pairing' function $f_t(a_j|y_i)$, and changes in it therefore reflect changes in the allocation of workers to jobs. The matching of worker human capital indices and firm wage *premia* is unlikely to be random. It arises from the personnel policies of firms and the job acceptance and quitting decisions of workers. The personnel policies involve the joint determination of wage setting, and hiring and firing criteria. The precise nature of these will depend on the specific reason for the failure of the purely competitive model giving rise to firm-specific wage mark-ups. But there is no reason in general to expect that high wage workers (high y) will necessarily be associated with high mark-up firms. High mark-ups may arise from asymmetric information and monitoring problems, and it may be that these are predominant in low skill jobs.

(4) This component captures the unexplained residual due to pure time variation. This may include a systematic component reflecting wage changes over the business cycle, as well as noise and further changes in the within-group residual.

The nature of $f_t(a_{jt}|y_{it})$ is determined by firms' optimal personnel policies and the reallocation process in the labour market. We briefly set out a model to highlight the main points.

A firm faces an environment characterised by a particular set of possible technologies, the nature of its product demand (volatility for example), and the state of the other markets it operates in; denote this set of factors by $\mathbf{M}_j$. This includes worker churning[8]: the firm is aware of the relation between churning and its wage premium. Workers differ in observable skill levels. Part of the optimisation problem the firm faces is to decide its personnel policy. This comprises the joint determination of the optimal wage premium to pay (α^*), and the optimal mix of worker types to employ (π^*, being the proportion of high skill workers employed). Specific versions of this problem have been widely studied in many contexts (see for example, the survey in Parsons, 1986, and more recently and more generally Lazear, 1996); for our purposes, we simply note the general form of the solution:

$$\alpha_j^* = \alpha_j^*(\mathbf{M}_j) \qquad (5)$$
$$\pi_j^* = \pi_j^*(\mathbf{M}_j)$$

This will therefore include the influence of churning.

The environmental factors ($\mathbf{M}_j$) also stochastically influence the firm's growth rate, though the work of Davis and Haltiwanger (1992) also shows the importance of idiosyncratic factors. We suppose that there is a job creation and destruction process that yields a distribution of employment levels at a date, conditional on its size at the previous date, its environment ($\mathbf{M}_j$), and incorporating a stochastic element. This is written as n_{jt} (n_{jt-1}, $\mathbf{M}_j$), interpreted to include births and deaths. For a fixed distribution of $\mathbf{M}_j$, continuous application of this transition process will yield a steady state distribution of firm size, albeit with any individual firm obeying the employment transition process noted above. This steady state distribution and the transition process also define the steady state distribution of $\mathbf{M}$ over jobs, say $\eta(\mathbf{M})$. Finally, given (5) we can also define a steady-state distribution of wage premia and worker types:

$$\alpha^* \sim \omega(\mathbf{M}), \pi^* \sim \varphi(\mathbf{M})$$

Our emphasis in this paper is on the component of earnings dispersion arising from the correlation between worker and firm effects. This is captured by the relationship between α^* (the firm wage effect) and π^* (the worker composition of the firm). This derives from the

[8] Churning is defined as amount of worker turnover net of the amount necessary for job reallocation.

steady state distributions:

$$\alpha^* = \chi(\pi^*) = \omega(\varphi^{-1}(\pi^*))$$

The final stage is to note that this function $\chi(.)$ determines the covariance term (3) above. Thus the functions $\omega(.)$ and $\varphi(.)$ are crucial to the determination of the variance of earnings. Changes in the set of technology options for example, will change these functions and hence the association between α^* and π^*. The personnel policies of firms, and the reallocation of labour between firms, are the main features of our approach to understanding earnings dispersion. We now turn to an empirical implementation of this approach.

4. Data

4.1 Data description

4.1.1 Source

Maryland, like every other state in the U.S. collects quarterly employment and earnings information through its State Employment Security Agency to manage its unemployment compensation program. Each quarter more than 100,000 employers report earnings and employment for over two million employees. Each wage record includes both employee and employer identifiers, enabling us to construct a quarterly longitudinal dataset on employers. The employer's four digit Standard Industrial Classification is then added from another administrative file. Virtually all business employment in Maryland is covered. We use data from 1985:3 to 1994:3. While the data have the advantage of being universal and longitudinal on both employers and workers, there are drawbacks in that we have no data on hours worked, and we have no information on the socioeconomic characteristics of the workers. The data are discussed in more detail in the data appendix.

4.1.2 Construction

We make a series of standard decisions in working with the data. We follow the approach taken by Jacobson, LaLonde and Sullivan (1993), in defining earnings to be the maximum earnings by the individual in a quarter. This picks out a single employer for each individual in any given period and hence ensures that there is a one to one relationship between workers and employers in each quarter. We also follow our previous work and work by Topel and Ward (1992), in that we define employment to be full quarter employment and take only workers whose earnings exceed 70% of the minimum wage during the quarter. Although we do not observe hours or weeks worked in the dataset, making it possible that earnings reported by the employer only reflect partial quarter earnings, our use of full quarter employment avoids this problem. Finally, we could not include all 40 quarters of data, because of the memory constraints mentioned above. Therefore we subset the data one more time to only consist of the third quarter of each year.

However, the type of analysis that is necessary to estimate over 2 million worker and 200,000 employer fixed effects poses a separate challenge. Since worker and employer fixed effects can be correlated in different ways, they cannot be estimated sequentially. Similarly, the usual approach when faced with the estimation of large numbers of dummy variables, deviating from the mean, is not appropriate since there are worker means, employer means, and joint means. Thus the biggest constraint faced in constructing a dataset was the maximum number of variables that could be held in computer memory. This problem is well known, and discussed in some detail by Abowd, Kramarz and Margolis (1999). We approach the estimation problem in two ways: one by directly estimating fixed effects; the second using one of the techniques described by Abowd, Kramarz and Margolis (AKM). We compare the outcomes below.

a) Basic approach

In order to make the analysis manageable, we took a random subset of 4,000 workers who were employed in each of the 10 quarters. The length of the panel is then sufficient to estimate parameters with some degree of accuracy. To identify worker effects we need inter-firm movement: of the 4,000 individuals who are the focus of the analysis, 29% are movers in the sense that their firm identifier changes at least once: 19.5% have 1, 6.5% have 2 and 3% have 3 or more changes.

To identify firm effects, we need intra-firm variation. To accomplish this, we then identified all the 2,426 employers who employed the 4,000 workers during the decade, and identified all other workers who worked for those employers in the same quarter as our 4,000 workers. Those 4,724,499 workers were then included in the estimation sample as long as they continued working at that firm. Thus, while we have 10 quarters of data on each of the 4,000 workers, we do not necessarily have 10 quarters of data on the others. Table 1(a) describes the distribution of quarters worked for the other employees in the dataset. This approach effectively denotes the employers who employed the 4,000 workers during this period as the universe: we include everyone who ever worked for those employers, and no-one who didn't. This is, on a smaller scale, similar to defining geographical boundaries on datasets.

Finally, the distribution of employers is described in Table 1(b). The drawback to this approach is that, while it is possible to accurately estimate the fixed effects of both individuals and employers, the choice of the sample is not completely representative of the typical *employer*. Our employers are, in general, larger and more heavily concentrated in manufacturing than employers in general. However, the choice of sample is quite close to the distribution experienced by the typical *worker* as is evident when the industrial distribution is weighted by employment. The restriction to continuously employed workers also means that this sample is not representative. Both these factors are necessary for data reasons, and not uncommon in this literature. We provide some checks on the likely effects of this approach by comparing results from a random sample of workers using the AKM approach.

Our approach is to specify the individual worker wage index and the firm wage premium as dummy variables (one for each worker and one for each firm) and to estimate these fixed effects from the data. Given the size of the dataset, this is not straightforward. An alternative, which we also estimate, has been suggested by AKM.

b) AKM approach

AKM suggest that one way of addressing the size constraint is to choose a matrix, Z, so that the individual effects, β are orthogonal to the projection of the firm effects, α, onto Z. Once Z is chosen, they offer two approaches: an order independent and an order dependent method of recovering the fixed effects.

The clear advantage to either method is that the choice of the sample is less constrained: one can choose a subset of firms and then from that a subset of workers in those firms. AKM use readily available proxies for Z: firm size, its square and 10 industry dummies, as do we. However their method also relies on some within-firm variation: they also interact individual characteristics (age at the end of schooling and labor force experience) with firm characteristics. We are unable to do this, since we have no information on the former and labor force experience in our dataset is severely left truncated.

We do follow the spirit of their approach, however, with two distinct purposes in mind. First, to compare our approach to theirs, we re-estimate the results for a smaller dataset (based on 2,000 workers and 2,729097 observations) using their framework. Second, to explore the sensitivity of the results to the sample selections we are forced to make (by our method but not theirs), we then chose a new random sample of 100 employers, and re-estimate the model using their framework for all workers ever employed by those employers. This addresses the problem that by sampling *workers* first and then their employers in our basic approach, we end up over-sampling large employers.

4.1.3 Estimation

The unit of observation is the earnings of an individual in a quarter. The equation we estimate is:

$$\log w_{ijt} = a_j + b_i + d_1.t + d_2.g + e_{it} \tag{6}$$

Comparing this with (3) derived above, a number of simplifications forced by the data can be noted. First, the firm mark-up is assumed to be time-invariant. Second, individual characteristics that in general do not vary with time (education level, ability, gender and race for example) are collapsed into an individual fixed effect, b_i. This is simply because we do not have data on these features. Note that this assumes that the value of these characteristics does not change over time[9], mirroring the time invariance assumption for firms' fixed effects. Third, the polynomial in age has to be assumed linear, and is split into the age at which individuals are first seen in this window (which goes into the fixed effect) and then a common linear component thereafter. Fourthly, the local labour market conditions are measured by including the Maryland aggregate employment growth rate, g.

[9] This is the converse of the usual practice: data constraints mean that we cannot look at temporal variations in the return to skill but can look at reallocation, whereas the previous literature had the data to look at the former but not the latter.

a) Basic approach

The focus of attention is estimating fixed effects for the 4,000 individuals and the 2,426 employers who employed them, using the data on the 4.7 million observations. However, computational limitations meant that coefficients on only five variables at a time could be estimated. We thus estimated the model in two steps. We first created dummy variables for five individuals at a time, and created a separate dummy for the other 3,995 workers. We chose those workers who were outside the sample of 4,000 to be the reference group and used the absorb command in the SAS GLM procedure to control for firm fixed effects. We repeated this until each of the 4,000 coefficients were estimated. We then created dummy variables for five employers at a time, created a separate dummy for the balance of the employers, chose one employer to be the reference group, and used the absorb command in the SAS GLM procedure to control for individual fixed effects. The interpretation of the coefficients, then, should be as relative, rather than absolute, fixed effects. The fixed effects on individuals are earnings of these 4,000 over and above the earnings of the group in the omitted category.

There are two technical issues with this approach. The first is whether it is possible to identify the fixed effects. Table 1(a) demonstrates that there is enough movement of both groups of workers (sample and non-sample) across the "universe" of 2426 employers to permit the estimation of separate employer fixed effects. Indeed, over 90% of the employer fixed effects were significantly different from zero, suggesting that the data were sufficient to separate these effects out. The second issue is the possibility of correlation between the employer premium and the error term. Clearly if the reallocation of workers across employers is random, then this is not an issue. Although moves are unlikely to be random, we are unable to find satisfactory instruments, and simply note the problem and proceed with the analysis.

b) AKM approach

We follow the AKM approach as follows. For our 2,000 person dataset we estimate the person-effects while controlling for the firm proxies for Z (in the manner described above), and then estimate the firm effects by directly substituting in the directly estimated individual effects into the basic regression (4). We then compare the results from our approach and the AKM approach and find the correlation between them. We also estimate firm and individual fixed effects for our sample of 100 firms and the 198,667 workers who ever worked for those firms, and report the results together with those for our main dataset.

4.2 Trends in earnings inequality in this dataset

In our (complete) base dataset, there is also only a slight change in overall earnings inequality for all full quarter workers in Maryland over this period: see Table 2. The 4% overall increase in the 90/10 ratio hides a great deal of variation by industry sector however: it is evident from the table that earnings inequality actually fell in five sectors of the economy (agriculture, manufacturing, transportation communication and utilities, finance insurance and real estate and government) while rising in others (wholesale trade, retail trade, professional services and other services).

In our restricted sample of 4000 workers, earnings inequality is (not surprisingly) lower than for the sample at large. The 90/10 ratio shows dispersion falling from 4.17 in 1985 to 3.80 by 1994. The disparity between this and the slight rise in the overall figure is likely to arise from the fact that this sample necessarily has high labour force attachment, and is likely to have above average job tenure.

5. Results

The purpose of this empirical analysis is to quantify the impact of labour reallocation on earnings dispersion using the results from the estimation just described. This will answer the question as to whether the new set of factors this perspective introduces as *potential* influences on dispersion are worth further investigation. We argue that they are. We also look at the *actual* contribution of labour reallocation to changes in the earnings dispersion in Maryland in our decade of data. We cannot specifically investigate the capability of our model to explain events in the U.S. as a whole in the 1980s and 1990s.

We first present the main results quantifying the potential role of labour reallocation on the earnings distribution. We then turn to examine the actual contribution in Maryland over the period 1985 to 1994.

5.1 Potential impact of labour reallocation on earnings dispersion.

We noted in section 3 that the scope for reallocation to influence the distribution of earnings depends first on the distribution of worker and firm fixed effects. The estimation yielded distributions of these that are not degenerate: see Figure 2. This shows that the range of values is far from negligible, compared to mean log earnings in the dataset of around 5.

As a first step, we look at the impact on the variance of earnings if we resorted *all* these individuals across the job slots in a perfectly systematic manner. The results of this are in Table 3 and reveal a very strong effect on dispersion. Column 1 shows a random assignment of workers to firms, and can be compared to a perfect positive sort in column 2[10] and a perfect negative sort in column 3. The numbers show that such complete resorting can double or eliminate the variance derived from a random matching.

However, this is not a realistic exercise, and it is more useful to look at the impact using data on labour turnover that approximates U.S. labour markets. To do this we set up a simple simulation of labour reallocation. There is a fixed population of workers and an equal number of job slots; these belong to a smaller number of firms. Each worker is paired with a job slot and the wage is simply the sum of a time-invariant worker component and a time-invariant firm effect. We use our estimated distributions for these fixed effects[11]. Initially, workers are randomly assigned to job slots. Each year, a given fraction of randomly chosen job matches split up[12]. These newly unattached workers and jobs re-match in a systematic way. We investigate the impact of the nature and extent of this re-matching on the cross-sectional variance of wages as this mechanism operates repeatedly.

The matching process contains the economics of the problem (discussed above), but for the purposes of this section we simply treat it as a black box and parameterise it as follows. We rank the unattached workers in terms of their fixed effect; similarly for the unattached job slots. We then perturb the worker ranking by adding $k*\varepsilon$, where ε is a standard normal random variable and k is the parameter we vary[13]. Workers and firms are then re-matched on the basis of this new ranking. At $k = 0$, the two sides of the market are perfectly matched, injecting a high degree of correlation[14] into the market as a whole. As k is

---

[10] That is, the highest ranked worker is allocated to the highest ranked firm, and so on.
[11] Replicated tenfold to give a bigger sample for the randomisation discussed below to work on.
[12] Clearly, having some systematic component in separations that was consistent with the systematic component in hiring would simply make our point more strongly.
[13] Obviously there are many ways of adding noise to the ranking.
[14] It is the *values* of the fixed effects (not the rank) that matter for the variance of wages; these will be highly but not perfectly correlated.

increased, more noise is added to the ranking and consequently the correlation between the fixed effects of the new hires and their firms is lower. Clearly, as this separating and resorting process is repeated, the correlation of the fixed effects across the market as a whole converges on that of the matching process.

The base case for these simulations is a steady state of random allocation and reallocation. This produces a variance of wages of 0.43 (the sum of the variance of individual fixed effects of 0.175 and the variance of firm fixed effects of 0.260). The main experiment is to show the impact of different reallocation functions, see Panel A of Figure 1. This is based on taking a reallocation rate of 40% per year, and changing k to achieve a steady state covariance between the fixed effects of 0.02, 0.04 and 0.08. The effect on the variance of wages is easily deduced from equation (1): the change in the variance will be equal to twice the change in this covariance. So relative to the initial value, small changes in the covariance of the fixed effects produce relatively large changes in the variance. The new steady state is achieved after about 10 years, by when the market correlation fully reflects the reallocation function.

The second experiment is to see how the rate of reallocation, coupled with a particular reallocation function, affects the fixed effect covariance and hence the variance of wages. We chose to use the rates of 20%, 40% and 60%. The rate of reallocation will affect the speed with which the steady state distribution is approached, but also may affect the steady state covariance itself. Examples are shown in panels B and C of Figure 1. Panel B is drawn with $k = 0$, i.e. perfect ranking, and shows that the paths for different reallocation rates converge at different speeds on the same steady state. Given the time span (the reallocation rates approximate those of U.S. labour markets in an annual time frame), this suggests that differences in labour reallocation rate will affect the cross-sectional variance of wages in the transition from one steady state to another. The result also suggests that differences between industries in labour reallocation rates will be reflected in differences in variance of wages during transitions between steady states. Panel C shows a different case: at a covariance of 0.04, the reallocation rate influences the steady state variance. This is perhaps a less expected result, but can be understood as follows. A higher reallocation rate yields a bigger pool of fixed effects that can be re-matched, hence a higher chance of picking the (absolutely) larger values. Consequently, combining these even with a low correlation will imply a higher variance. Conversely, a low reallocation rate will have only a low probability of picking some non-central values and hence will produce a lower variance. When the correlation is very high (as in Panel B), this effect is blunted by the fact that the whole population rapidly becomes highly sorted.

To summarise: these simple simulations, using estimated labour reallocation rates and fixed effect distributions, show that changes in the process of labour reallocation can have substantial effects on the dispersion of earnings.

5.2 Actual contribution of changing allocation to changes in earnings dispersion

We now turn to consider the actual contribution of changes in the allocation process to changes in the dispersion of earnings in our dataset. Over the period as a whole the variance of worker fixed effects accounted for some 55% of the variance of log wages, and firm effects around 35%. The remainder is mostly accounted for by the error, with all covariances being generally small.

It is the size of the correlation between firm and worker effects that we are most concerned with here. This is small and negative in this dataset. Figure 2 cross-plots the individual and worker fixed effects at one particular date (1994), and shows clearly that there is no strong pattern. Figure 3 compares changes in the workers' employers in terms of their fixed effects over the decade. Most workers stayed with the same employer, many had moved (not necessarily directly, and not necessarily voluntarily) to higher fixed effect firms and some had moved to lower fixed effect firms.

We are interested in changes in the dispersion of earnings over time and the contribution to that of changes in the allocation of workers to firms. We showed above that relatively small changes in this allocation could have substantial impacts on earnings dispersion. Figure 4 plots the variance of log wages and the correlation between worker and firm fixed effects over the ten years of our data. We see that the pattern in its change over time is very closely mirrored by the change in the variance[15]. This provides some evidence that changes in the allocation of workers to job slots are associated with changes in the dispersion of earnings.

However, there are other components of the variance of log earnings and we must consider those too. Since we have imposed in our estimation the assumption that the worker fixed effects are constant over time and taken a fixed set of continuously employed workers, clearly the cross-sectional variation in these fixed effects is the same in each year. However, while the fixed effect of any particular firm is fixed, since workers move among the

[15] This is not an artefact of using the variance: Appendix figure 1 shows that the 90/10 ratio of earnings follows the same pattern too.

population of firms, the cross-sectional variance of wage *premia* of firms currently occupied by our sample of workers is not fixed. However, in practice, variation in this is negligible.

The simulations suggest that the allocation function, the job reallocation rate and the churning flow rate may have an impact on earnings dispersion. We attempt to separate out the contribution of these factors in a very simple way by regressing var_t (log w_i) (the cross-sectional variance of log earnings at t) on cov_t (y_i, a_j) (the covariance of firm and worker effects at t), and JRR (job reallocation rate) and CFR (churning flow rate)[16]. The results of this regression on the 10 data points for 1985 - 1994 are in panel A of Table 4. They support the idea that there is a strong association between the changing pattern of worker/firm allocation and changing earnings dispersion. This exercise is simply designed to summarise the relationship between the time series change in the cross-sectional variation in earnings and the time series change in the covariance of worker and firm fixed effects. To re-emphasise, we are saying that changing allocation appears to be playing an important role in affecting the earnings distribution; we are unable to investigate (because of data constraints) whether this contribution remains important once we allow for changing worker fixed effects.

If $\phi_t(a_{jt}|\ y_{it})$ depends on production and organisational technology, it seems likely that this varies between industry. Our final piece of analysis exploits this idea and examines earnings dispersion and the covariance of worker and firm effects by major industry group. Figure 5 plots out the estimated correlation over time for the 9 major groups. We see that in the Retail Trade and Other Services industries (groups 5 and 7), the correlation is significantly negative and falls quite sharply, while in FIRE (group 6) it rises over the decade. This suggests quite plausibly that the nature of the optimal personnel decisions may differ between industries and may change in different ways in response to a changing environment. In Figure 6, we match this up with the movement in var_t(log w_i) within each industry; again, there is some evidence that the two move together. We test this using regression analysis, in Table 4, panel B. These results use the 81 observations available from 9 years times 9 industries. The results confirm that the covariance is a significant factor in explaining earnings dispersion. They also show that the industry-specific job reallocation and churning rates do not appear to matter.

[16] Note that this is not a variance-decomposition exercise on equation (4), nor is (4) an identity: in the data, all the other covariances may be non-zero, and the error variance may not be constant.

5.3 Comparisons using the AKM approach and random sample.

We performed two major sensitivity checks on our results. First, we took a dataset based around 2,000 workers chosen according to the sample selection criteria noted above; this produced a total of 2,729,097 observations on 468,549 individuals. We estimated worker and firm fixed effects on this dataset using both our approach and the AKM approach, and compared the outcomes. The correlation between the worker fixed effects is very high, 0.82. The correlation between the firm fixed effects using the two approaches is substantially lower, 0.45. We ascribe most of this difference to the problem with finding satisfactory instruments for Z. Although our approach directly estimates the employer and worker fixed effects, it is clear that our sample selection criteria, while delivering a sharp investigation of our approach, do not pretend to deliver a purely random sample of workers. But we want to consider whether it seriously biases the results. So we also re-estimated using a sample of 100 randomly chosen employers and their employees using the AKM approach. Our previous analysis suggests, however, that the gain in representativeness comes at the cost of noisier employer fixed effect estimates. Taking the resulting estimates of worker and firm fixed effects, we re-ran the regressions reported in Table 4. The key covariance term remained significant, indicating that our main results are not an artefact of our sample selection criteria.

6. Conclusions

This paper has provided a different perspective on changes in the dispersion of earnings, particularly changes in within-group inequality, focusing on the influence of the reallocation of labour. We argue that this introduces a new set of factors that determine the steady state distribution of earnings. These are principally the rates of job reallocation and worker reallocation, and the nature of the economic behaviour allocating particular workers to particular jobs. We show that these can have a quantitatively large impact on earnings dispersion. We also show that the changes we estimate in the outcome of this allocation mechanism play a significant role in explaining movements in the variance of log earnings in our dataset.

We can finally speculate about the implications of this for the changes in earnings inequality experienced in the U.S. during the 1980s. It is possible, for example, that the change in inequality during the 1980s was the result of a change from one steady state to another. This may have manifested itself in different ways. The literature has exhaustively investigated potential changes in the skill premium itself. We explore different possibilities. For example, changes in technology may have resulted in a change in the pairing between workers and firms - in our notation, a different $f(.)$ - as has been suggested by Kremer and Maskin (1996). Or alternatively, changes in trade may have resulted changes in the sizes of firms with a given wage premium. Finally, there may have been a change in the amount of labour reallocation.

The contribution of this paper has been to raise each of these possibilities, and to establish that, at least in Maryland, labour reallocation and rematching played a significant role in explaining changes in earnings dispersion. This result is robust to different sample choices and different estimation techniques.

Table 1(a): Characteristics of Workers

	Sample Dataset	4,000 individuals
Median Earnings	4812	7180
Number of employers:		
1	78.5	70.5
2	16.4	19.6
3+	5.1	7.9

Source: authors' data

Table 1(b): Characteristics of Employers

	The 2426 Employers		All Employers	
Industry distribution (%):	Unweighted	Weighted by Employment	Unweighted	Weighted by Employment
Ag., Min	10.3	9.1	15.1	12.8
Manufacturing	11.1	9.9	3.3	9.7
TCU	4.2	6.4	3.7	5.8
Wholesale Trade	11.0	3.4	9.9	6.5
Retail Trade	14.2	9.3	18.8	15.5
FIRE	9.4	6.3	7.3	6.2
Other Services	17.6	34.8	18.8	23.8
Professional Services	13.0	6.0	21.3	12.6
Government	9.2	14.7	1.9	7.0
Average Employment	262.15		18.69	

Source: authors' data

Table 2: Earnings Inequality in Maryland

90/10 ratio	1985	1994	Change
Overall	9.33	9.70	4.06%
Ag,Min	7.00	6.49	-7.33%
Manufacturing	4.53	4.43	-2.18%
TCU	8.99	8.12	-9.69%
Wholesale Trade	6.12	6.56	7.27%
Retail Trade	10.68	11.17	4.68%
FIRE	5.49	5.24	-4.47%
Prof. Services	9.85	11.03	12.00%
Other Services	12.63	13.10	3.70%
Government	4.27	3.67	-14.17%

Source: authors' data.

Table 3: Potential Impact of Reallocation on Inequality

	Simulated Matching		
	Random	Perfect positive Sorting	Perfect negative Sorting
var ($b_i + a_j$)	0.257	0.501	0.009
cov (b_i , a_j)	0	0.122	-0.124
Mean ($b_i + a_j$)	0.713	0.713	0.713

Table 4: Regressions for Cross-Sectional Variance of Log Earnings

Sample: 1985 - 1994
Dependent variable: cross-sectional variance of log earnings

AGGREGATE
Unit = year
Obs = 10

	(1)	(2)	(3)	(4)
Covariance	8.965 (7.3)	13.481 (6.7)	13.227 (4.5)	8.239 (5.5)
JRR	-	0.467 (0.5)	0.520 (0.5)	-
CFR	-	-0.281 (0.9)	-0.300 (0.8)	-
g	-	-	0.105 (0.1)	0.556 (0.9)
Adj. R^2	0.853	0.859	0.824	0.848

DISAGGREGATE
Unit = industry-year
Obs = 81
Regression includes industry dummies

	(1)	(2)	(3)	(4)
Covariance	1.594 (3.8)	1.327 (2.8)	1.193 (2.5)	1.406 (3.2)
JRR	-	-0.031 (0.1)	-0.062 (0.2)	-
CFR	-	-0.022 (0.2)	-0.074 (0.7)	-
g	-	-	0.857 (1.4)	0.830 (1.5)
Adj. R^2	0.920	0.920	0.921	0.921

Notes: Tables give coefficient and t-statistic
Excluding industry group 0, because Figure 6 shows it to be an outlier; including it simply strengthens our results.
Cov = Covariance of worker and firm fixed effect
JRR = Job Reallocation Rate
CFR = Churning Flow Rate
g = Aggregate (Maryland) employment growth rate

Figure 1: Simulated Changes in Earnings Inequality

Panel A: Differences in the Matching of Workers and Firms

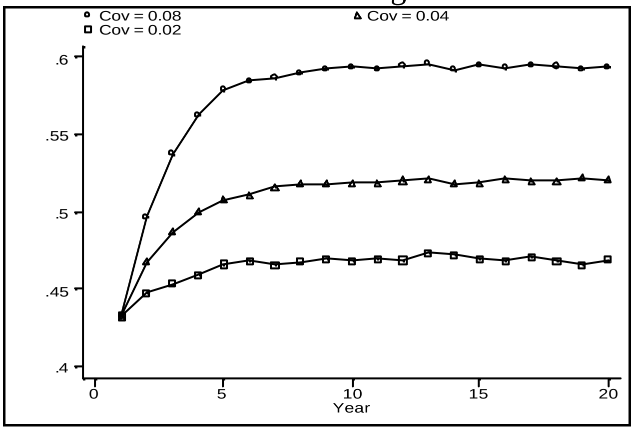

Panel B: Differences in the Reallocation Rate with perfect matching of Workers and Firms

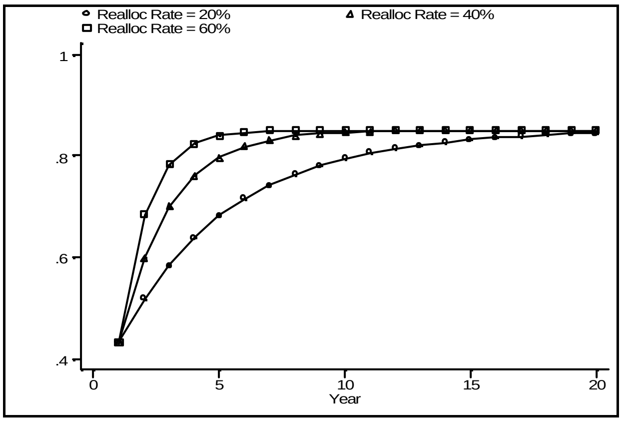

Panel C: Differences in the Reallocation Rate with matching of Workers and Firms at Cov = 0.04

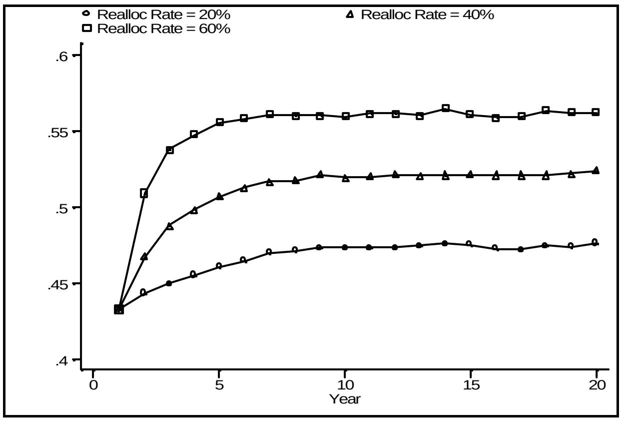

Figure 2: Worker and Firm Effects in 1994

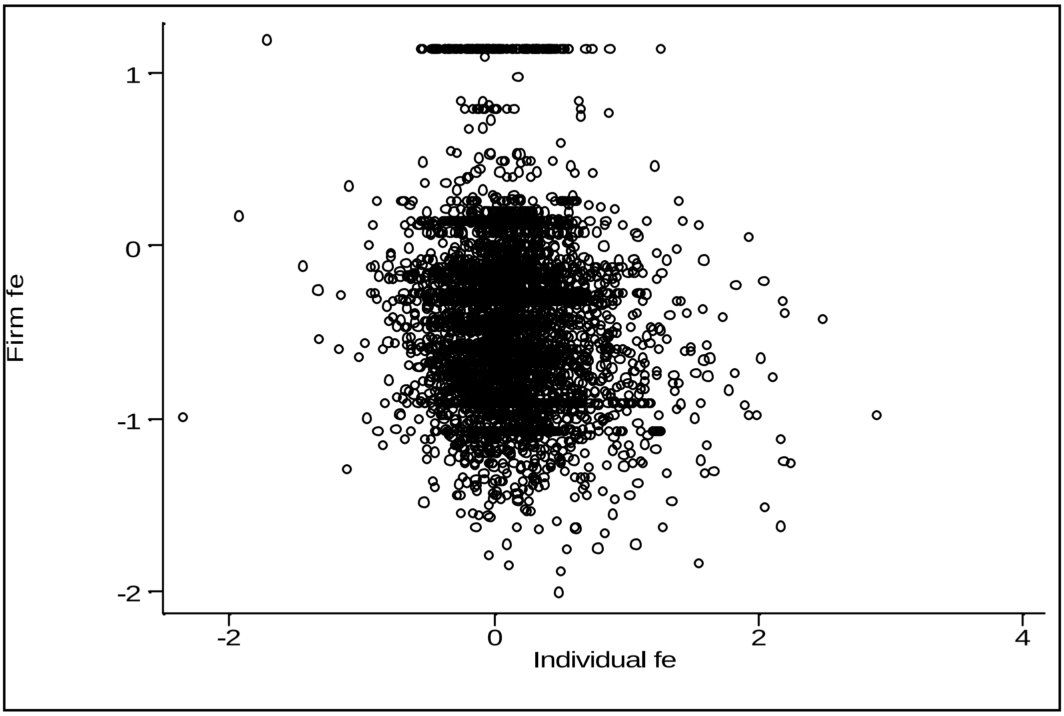

Note: each observation is a worker and the coordinates are given by that worker's individual fixed effect and the fixed effect of that worker's employer in 1994.

Figure 3: Change in Firm Fixed Effects between 1985 and 1994

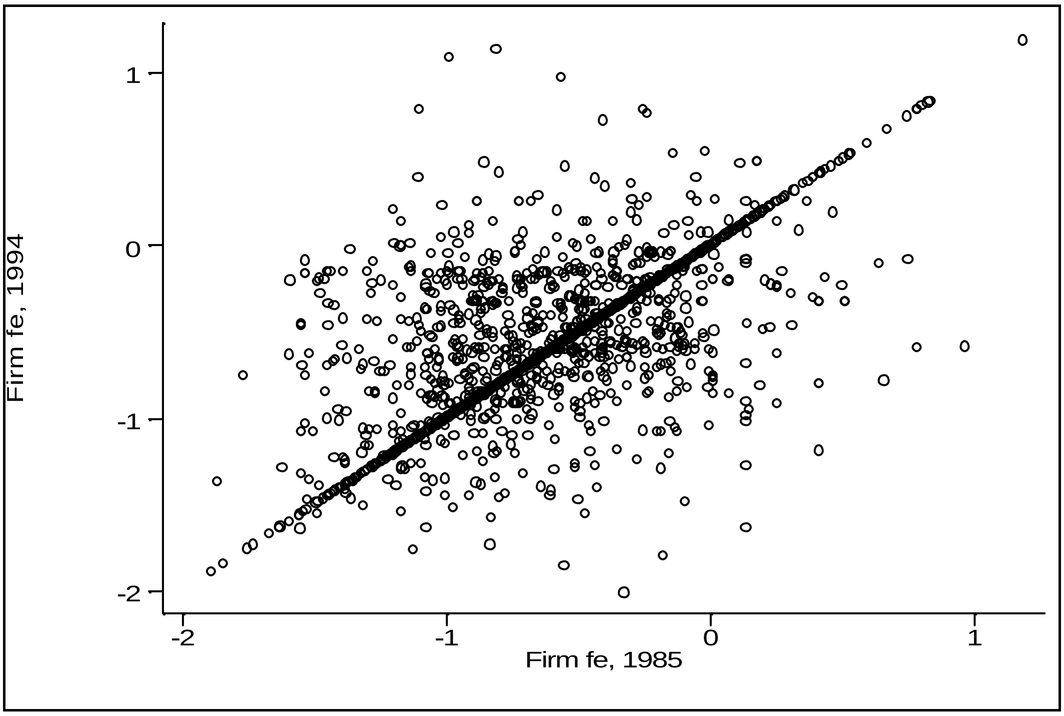

Note: each observation is a worker, and the coordinates are given by the fixed effect of that worker's employer in 1985 and the fixed effect of that worker's employer in 1994. Points above a 45^o line therefore represent workers moving to a higher fixed effect firm.

Figure 4: Variance of log earnings and Correlation of firm and worker fixed effects

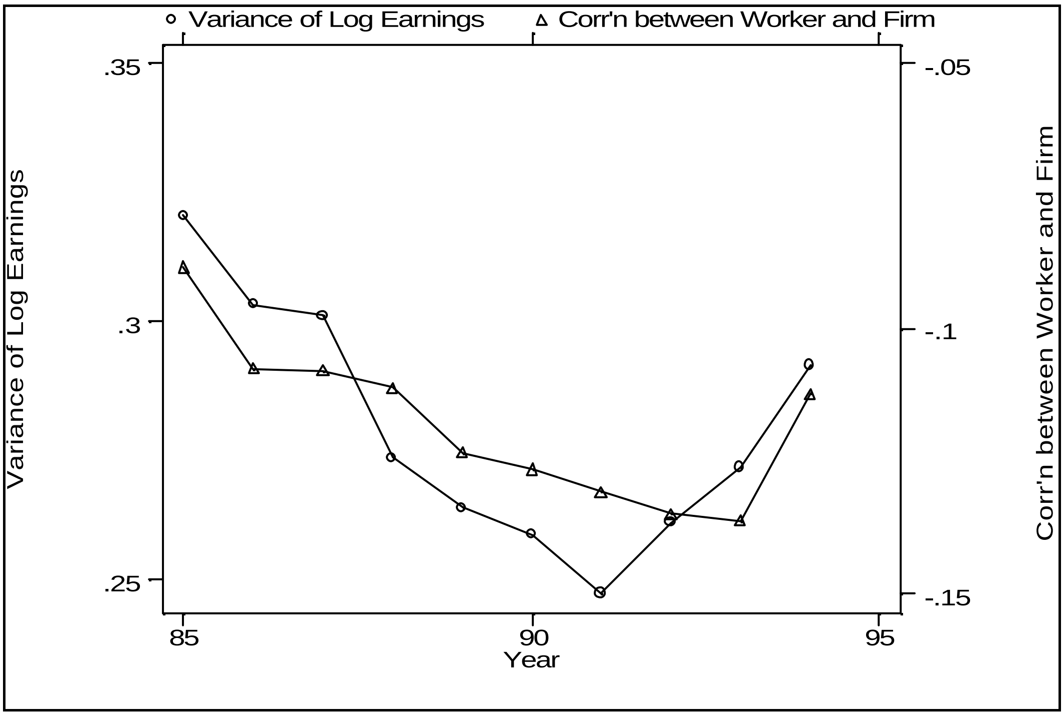

Note the split scale: variance of log wages is the left scale, covariance of firm and worker fixed effects is the right scale.

Figure 5: Correlation of Worker and Firm Effects by Industry Group

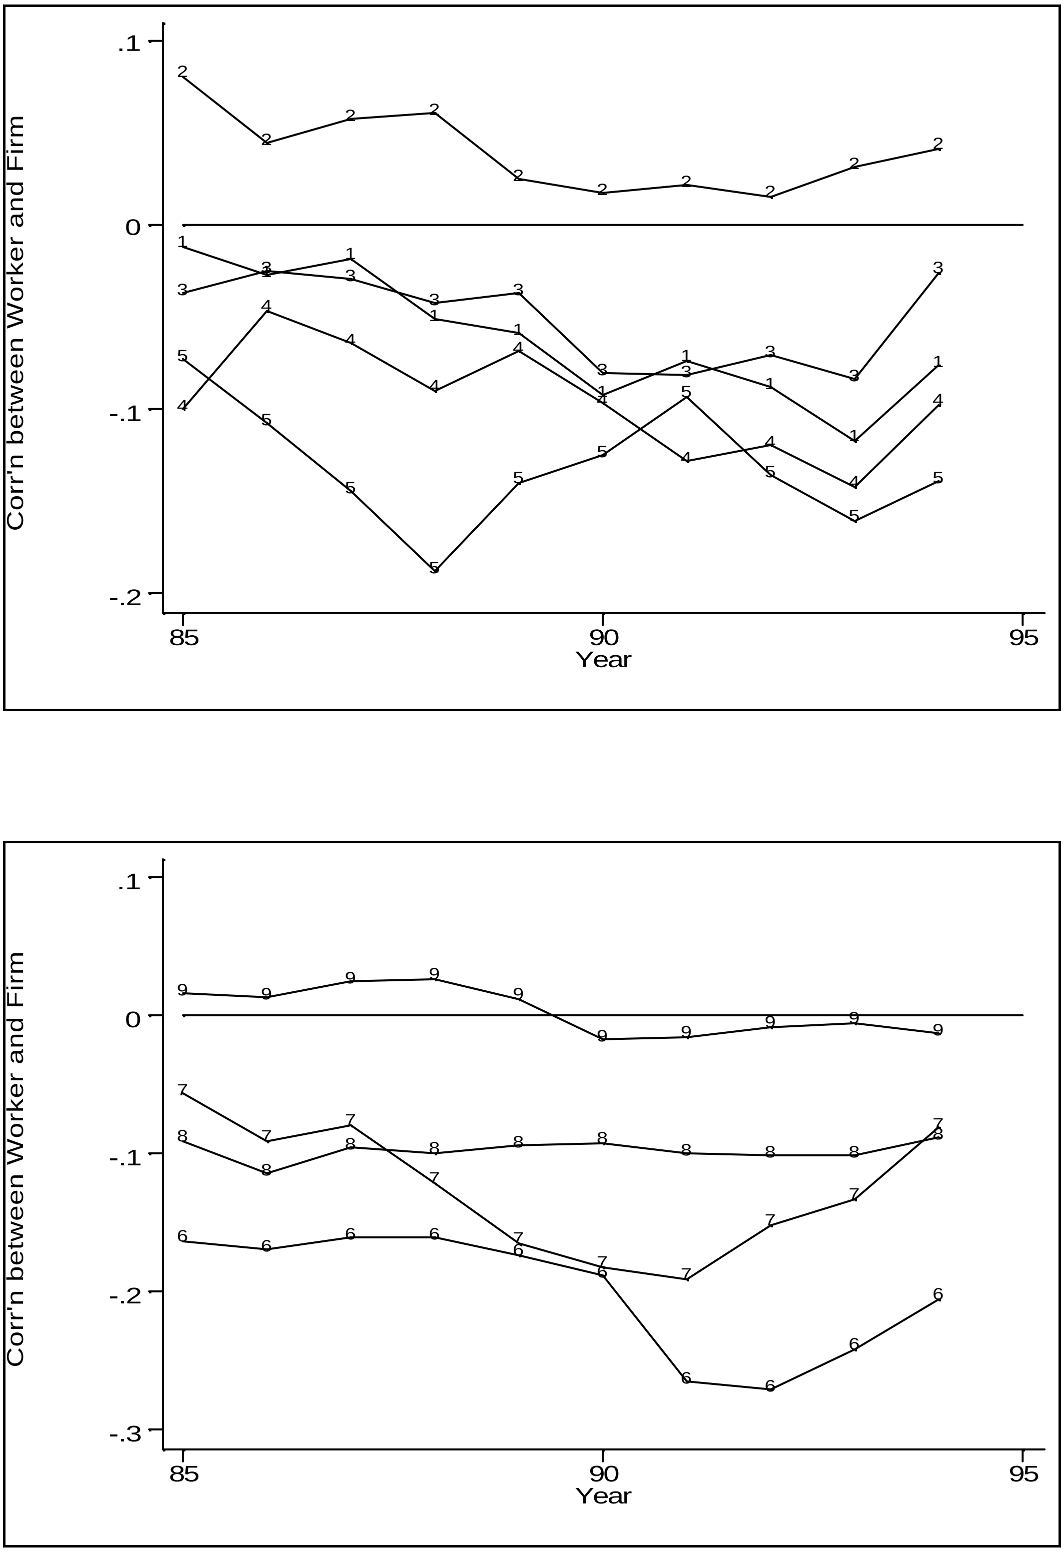

Figure 6: Earnings Variance and Correlation of Worker and Firm Effects by Industry

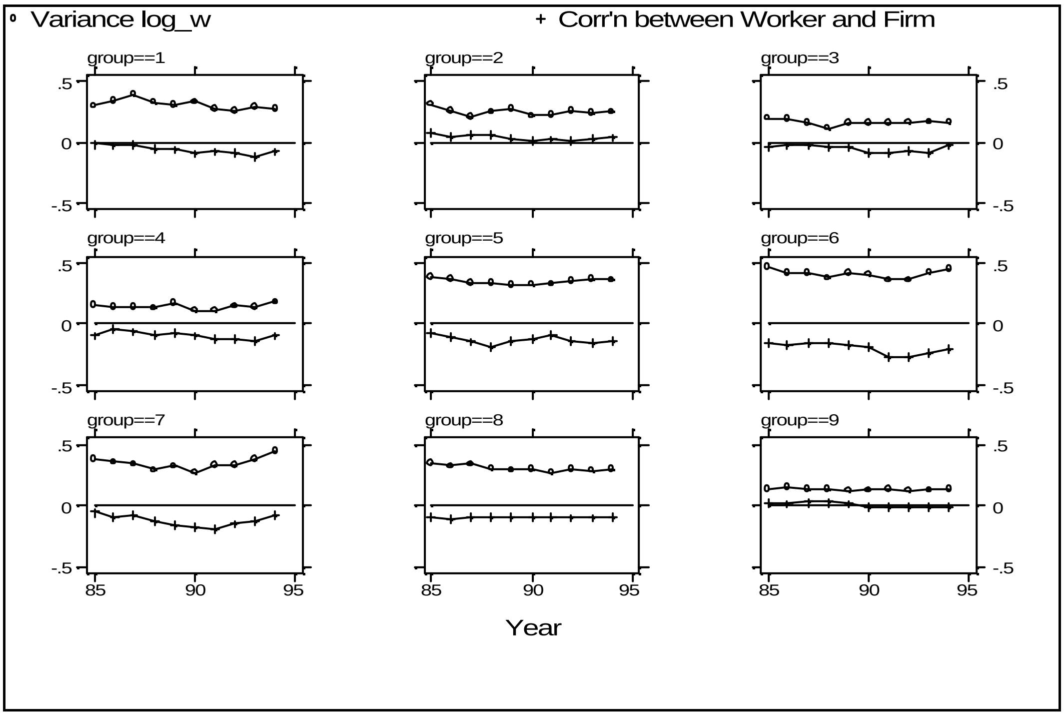

On a different scale:

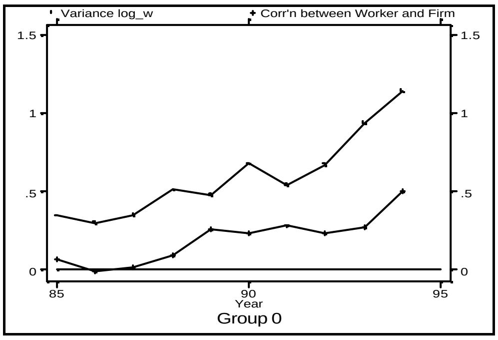

Data Appendix

Maryland, like every other state except New York, collects quarterly employment and earnings information in order for the State Employment Security Agency to manage the state unemployment compensation program. These UI wage records cover virtually all (over 90%) employment in the labor market.[17] Each wage-record includes an employee's social security number, a unique employer identification number and the employee's total earnings during the reference quarter. The employer's total employment level and four-digit Standard Industrial Classification code are acquired from administrative units within the state and merged with the wage data elements. All private sector employers who employ one or more paid employees are required to file quarterly reports, as are state and local government agencies. Only federal civilian and military personnel, employees of the U.S. Postal Service, railroad employees, employees of religious and philanthropic organizations, self-employed individuals, those who receive only commissions and some agricultural employees are not covered.

Total wages, including tips, commissions and bonuses are covered, up to a ceiling of $100,000 per quarter. In other words, virtually all business employment is covered and virtually all earnings are subject to required reporting. These administrative records are confidential. Public release of the identity of any individual or reporting business is strictly prohibited. The employer reporting unit is usually the firm, but also government and nonprofit organizations. Over 90% of organizations are single establishment employers. Of the remaining organizations, the reporting unit sometimes encompasses all establishments and sometimes only some establishments.

The data in this paper are subject to several limitations. They are not directly comparable with earnings inequality measures derived from national household surveys. Since the data in this paper deal with employers only in Maryland, the results may not necessarily extrapolate to the country as a whole. In the case of the largest organizations, reporting units are not firms. Another limitation is the absence of data on hours worked: the reports provide information only on quarterly earnings. As a result, some workers may have low earnings not because of their low hourly earnings but because they leave the state or the labor force or they spend time frictionally unemployed moving from one job to another during a quarter. In order to reduce the possibility that we are capturing sequential rather

[17] The wage records are archived in a secure environment for research purposes through an agreement between

than concurrent multiple job holding, we follow Topel and Ward in only looking at full quarter job spells[18].

The employer data do not include information on worker characteristics. The data do, however, have a number of advantages over household survey data. Recall error is less likely when earnings and industry affiliation are reported by employers, rather than by households. Second, the focus on the earnings in jobs rather than of individuals who possibly fill multiple jobs, documents another dimension of earnings inequality.

Maryland's Department of Economic and Employment and University of Baltimore.

[18] In other words, for a job spell to be included in this analysis, a worker must have had positive earnings with the employer in both the quarter before and the quarter after the quarter in question. This means that we minimize the possibility of having partial quarter earnings and employment spells.

References

Abowd, J., Kramarz, F. and Margolis, S. (1999), 'High Wage Workers and High Wage Firms', Econometrica, 67(2), pp. 251-333.

Anderson, P. and Meyer, B. (1994), 'The Extent and Consequences of Job Turnover', Brookings Papers on Economic Activity.

Blackburn, M. and Neumark, D. (1993), 'Omitted-Ability Bias and the Increase in the Return to Schooling', Journal of Labor Economics, vol. 11, pp. 521-544.

Bound, J. and Johnson, J. (1992), 'Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations', American Economic Review, vol. 82, no. 3, pp. 371-392.

Burgess, S. M., Lane, J. I. and Stevens, D. (2000), 'Job Flows, Worker Flows and Churning', forthcoming, Journal of Labor Economics.

Burtless, G. (1990), *A Future of Lousy jobs?: The Changing Structure of U.S. Wages*, Brookings Institution, Washington, D.C.

Card, D. and Lemieux, T. (1996), 'Wage Dispersion, Returns to Skill and Black-White Wage Differentials', Journal of Econometrics, vol. 74, pp. 319-361.

Davis, S. and Haltiwanger, J. (1991), 'Wage Dispersion Between and Within U.S. Manufacturing Plants, 1963-1986', Brooking Papers on Economic Activity, pp. 115-180

Davis, S. and Haltiwanger, J. (1992), 'Gross Job Creation, Gross Job Destruction and Employment Reallocation', Quarterly Journal of Economics, vol. 107, pp. 819-863.

Davis, S. and Haltiwanger, J. (1995), 'Employer Size and the Wage Structure in U.S. Manufacturing', NBER Working Paper No. 5393.

Gottschalk, P. and Moffitt, R. (1994), 'The Growth in Earnings Instability in the U.S. Labor Market', Brooking Papers on Economic Activity, pp. 217-254.

Groshen, E. (1991), 'Five Reasons Why Wages Vary Among Employers', Industrial Relations, pp. 350-381.

Heckman, J. J. and Sedlacek, G. L. (1985), 'Heterogeneity, Aggregation and Market Wage Functions: An Empirical Model of Self-Selection in the Labor Market', Journal of Political Economy, vol. 93, no. 6, pp. 1077-1125.

Heckman, J. J. and Sedlacek, G. L. (1990), 'Self-Selection and the Distribution of Hourly Wages', Journal of Labor Economics, vol. 8, no. 1, Part 2, pp. S329-S363.

Heckman, J. J. and Honore, B. E. (1990), 'The Empirical Content of the Roy Model', Econometrica, vol. 58, no. 5, pp. 1121-1149.

Jacobson, L., LaLonde, R., and Sullivan, D. (1993), 'Earnings Losses of Displaced Workers', American Economic Review, vol. 89, pp. 685-709.

Juhn, C., Murphy, K. and Pierce, B. (1993), 'Wage Inequality and the Rise in the Returns to Skill', Journal of Political Economy, vol. 101, pp. 410-442.

Katz, L., and K. Murphy (1992), 'Changes in Relative Wages, 1963-1987: Supply and Demand Factors', Quarterly Journal of Economics, vol. 107, pp. 35-78.

Kremer, M. (1993), 'The O-Ring Theory of Economic Development', Quarterly Journal of Economics, vol. 108, pp. 551-575.

Kremer, M. and Maskin, E. (1996), 'Wage Inequality and Segregation by Skill', NBER Working Paper no. 5718.

Lane, J., Stevens, D. and Burgess, S. M. (1996), 'Worker and Job Flows', Economics Letters, vol. 51, pp. 109-113.

Lazear, E. (1996), *Personnel Economics*, MIT Press: Cambridge, Mass. and London, Eng.

Lindbeck, A and Snower, D. (1996), 'Reorganization of Firms and Labor-Market Inequality', American Economic Review, May, vol. 86, pp. 315-321.

Levy, F. and Murnane, R. (1992), 'U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations', Journal of Economic Literature, pp. 1333-1381.

Maddala, G. S. (1983), *Limited-Dependent and Qualitative Variables in Econometrics*, CUP: Cambridge.

Mitchell, D. (1991), 'Wage Pressures and Labor Shortages: the 1960s and 1980s', Brooking Papers on Economic Activity, pp. 191-231.

Moffitt, R. (1990), 'The Distribution of Earnings and the Welfare State', in Burtless, G. (ed), *A Future of Lousy Jobs?: the Changing Structure of U.S. Wages*.

Murnane, R., Willet, J. and Levy, F. (1995), 'The Growing Importance of Cognitive Skills in Wage Determination', Review of Economics and Statistics, vol. 77, pp. 251-266.

Murphy, K. and Welch, F. (1992), 'The Structure of Wages', Quarterly Journal of Economics, vol. 107, pp. 215-226.

OECD (1996), 'Earnings Inequality, Low-Paid Employment and Earnings Mobility', Employment Outlook, pp. 59-107.

Parsons, D. (1986), 'The Employment Relationship: Job Attachment, Work Effort and the Nature of Contracts', in O. Ashenfelter and R. Layard (eds), *Handbook of Labour Economics* V*ol. 2,* North Holland: Amsterdam.

Roy, A. D. (1951), 'Some Thoughts on the Distribution of Earnings', Oxford Economic Papers, pp. 135 – 146.

Sattinger, M. (1993), 'Assignment Models of the Distribution of Earnings', Journal of Economic Literature, vol. 31, no. 2, pp. 831-880.

Topel, R. H. and Ward, M. P. (1992), 'Job Mobility and the Careers of Young Men', Quarterly Journal of Economics, pp. 439-478.

Willis, R. (1986), 'Wage Determinants: A survey and Reinterpretation of Human Capital Earnings Functions', in O. Ashenfelter and R. Layard (eds), *Handbook of Labour Economics*, North Holland: Amsterdam.

CENTRE FOR ECONOMIC PERFORMANCE
Recent Discussion Papers

490	S. Burgess S. Profit	Externalities in the Matching of Workers and Firms in Britain
489	S. Nickell G. Quintini	Nominal Wage Rigidity and the Rate of Inflation
488	S. Nickell J. Van Reenen	Technological Innovation and Performance in the United Kingdom
487	M. M. Tudela	Explaining Currency Crises: A Duration Model Approach
486	D. Sturm	Product Standards, Trade Disputes and Protectionism
485	G. Duranton V. Monastiriotis	Mind the Gaps: The Evolution of Regional Inequalities in the UK 1982-1997
484	H. G. Overman Y. Ioannides	Zipfs Law for Cities: An Empirical Examination
483	H. G. Overman Y. Ioannides	Cross Sectional Evolution of the US City Size Distribution
482	Y. Ioannides H. G. Overman	Spatial Evolution of the US Urban System
481	H. G. Overman	Neighbourhood Effects in Small Neighbourhoods
480	S. Gomulka	Pension Problems and Reforms in the Czech Republic, Hungary, Poland and Romania
479	S. Nickell T. Jones G. Quintini	A Picture of the Job Insecurity Facing British Men
478	C. Dougherty	Numeracy, Literacy and Earnings: Evidence from the National Longitudinal Survey of Youth
477	P. Willman	The Viability of Trade Union Organisation: A Bargaining Unit Analysis
476	D. Marsden S. French K. Kubo	Why Does Performance Pay De-Motivate? Financial Incentives versus Performance Appraisal

475	S. Gomulka	Macroeconomic Policies and Achievements in Transition Economies, 1989-1999
474	S. Burgess H. Turon	Unemployment Dynamics, Duration and Equilibrium: Evidence from Britain
473	D. Robertson J. Symons	Factor Residuals in SUR Regressions: Estimating Panels Allowing for Cross Sectional Correlation
472	B. Bell S. Nickell G. Quintini	Wage Equations, Wage Curves and All That
471	M. Dabrowski S. Gomulka J. Rostowski	Whence Reform? A Critique of the Stiglitz Perspective
470	B. Petrongolo C. A. Pissarides	Looking Into the Black Box: A Survey of the Matching Function
469	W. H. Buiter	Monetary Misconceptions
468	A. S. Litwin	Trade Unions and Industrial Injury in Great Britain
467	P. B. Kenen	Currency Areas, Policy Domains and the Institutionalization of Fixed Exchange Rates
466	S. Gomulka J. Lane	A Simple Model of the Transformational Recession Under a Limited Mobility Constraint
465	F. Green S. McIntosh	Working on the Chain Gang? An Examination of Rising Effort Levels in Europe in the 1990s
464	J. P. Neary	R&D in Developing Countries: What Should Governments Do?
463	M. Güell	Employment Protection and Unemployment in an Efficiency Wage Model
462	W. H. Buiter	Optimal Currency Areas: Why Does the Exchange Rate Regime Matter?
461	M. Güell	Fixed-Term Contracts and Unemployment: An Efficiency Wage Analysis

To order a discussion paper, please contact the Publications Unit
Tel 020 7955 7673 Fax 020 7955 7595 Email info@cep.lse.ac.uk
Web site http://cep.lse.ac.uk