Lecture 9

Imperfect Markets

Today we consider wage determination that depends on more than just expected productivity. We develop a model that shows how individuals with the same productivity and the same tastes can end up earning different wages at different firms in a monopsony setting.

Basic Monopsony Model

In a perfectly competitive labour market, the labour supply curve for an individual firm is perfectly inelastic, meaning the labour supply curve is horizontal. Firms must take wages as given. There is no need to raise the wage, because they can hire any number of workers at the going wage, and if they lower the wage (by even a penny) all workers will want to leave. In equilibrium, firms the going wage is driven up to the marginal product of the worker. Thus, in equilibrium, workers in a perfectly competitive market are paid their marginal product: what they contribute to the profit of a firm.

Profit of a firm is: $\pi = Y(L) - wL$, where Y(L) is the production function (with labour as the only input), and w is the wage rate. The first order condition is: $Y'(L^*) = w$. With many firms, the wage is driven down such that $w = \frac{Y(L^*)}{L^*}$, so profits are zero.

A monopsonist firm does not take the wage the wage as given. A monopsonist not only chooses how much labour to hire, but what wage to pay. However, the monopsonist is constrained by the labour supply to the firm. It's just that this labour supply is no longer perfectly inelastic. A firm's labour supply is: $L^s = L^s(w)$, a function of the wage. This is not the same as the individual's overall labour supply. If a firm raises it's wage, relative to the rest of the market, more workers will want to work there: $L^{s}{}'(w) > 0$ (labour supply is no longer perfectly elastic).

Since the firm's labour demand must equal labour supply, $L^s = L^D = L$. Another way of taking into account the firm's labour supply function is that, if a firm wants to hire more labour than it currently has supplied to it, the firm must raise the wage offer: $w = w(L^S) = w(L)$, where $w'(L) > 0$.

Maximizing profit, $\pi = Y(L) - w(L)L$, first order condition is: $Y'(L^*) = w + w'(L^*)L^*$. In deciding how much labour to hire, the monopsonist takes into account, that when he increases one unit of labour hired, not only do costs rise by w, but the wage must be increased for everyone. The most important insight from the monopsony model for our purposes is that profit will no longer be zero. The marginal product from a worker will be greater than the wage paid: $Y'(L^*) = w + w'(L^*)L^* > w$.

The classic example of Monopsony is where there exists a large firm that dominates hiring in a region (such as a mining company in a small town), or an organization that regulates all hiring (such as the nurses association, or the NCAA). But a more general interpretation of monopsony occurs for ANY firm that faces an upward sloping labour supply curve. We will see below a model where there are many firms, but all face an upward sloping supply curve. This leads to a labour market that works strikingly different from the perfectly competitive situation.

Manning's Model of Monopsony

Set-up:

The basic set-up follows Manning's Chapter 2 model of monopsony, which follows Burdett and Mortensen, "Wage differentials, Employer Size, and Unemployment," International Economic Review, 1998.

Workers

There are M_W identical workers, with the same productivity and preferences for leisure. Denote the monetary equivalent of the value attached to leisure by b.

Employers.

There are M_F firms. All firms exhibit constant returns to scale, and the productivity of each worker is p.

Wage Determination

Employers set wages once to maximize steady-state profits. All workers in the firm are paid the same wage. The cumulative distribution function of wages across employers is $F(w)$. This is the probability that a job picked at random pays less than w.

Matching Technology

Each non-employed obtain job offers at a rate λ_u. While working, job offers arrive to an worker at a rate λ_e. I like this feature, because some other models require search to only take place while unemployed. Clearly that is less realistic than this case here. It simplifies the model a lot if we assume both employed and non-employed obtain job offers at the same rate, λ. Burdett and Mortensen consider the case when these rates are different.[1] Job offers are drawn at random from the set of firms (from the distribution $F(w)$). Workers leave for reasons unrelated to getting a higher paying job at rate δ_r, and are replaced by an equal number who initially enter the labor market as unemployed. There is also an exogenous job destruction rate of δ_u. The total rate of exogenous job loss is $\delta = \delta_u + \delta_r$.

Trivial, but important assumption needed so firms can pay workers less than productivity

---

[1] I have notes for Manning's model when λ_e and λ_u are different, if you are interested, or see his Chapter 9.

At least some workers are willing to work for less than their marginal product if the only alternative is unemployment. This occurs when there exists some wage where w<p (the wage paid is less than productivity) and w>b (there are some people in this world where for a given wage lower than productivity, they value leisure less than that wage).

Note, the assumption's we've made so far may seem restrictive, but are relaxed in later chapters of Manning's book: workers have identical productivity, no mobility costs, workers care only about the wage, firms are identical, constant returns, wages set once-and-for-all, job offers drawn at random.

The main difference between this model and one of perfect competition is that a worker does not receive job offers by every firm at once. In this special case when $\lambda \rightarrow \infty$, this model collapses to the perfect competition model, as we will see.

Behavior or workers:

A worker accepts another job if the wage offer is above her current wage. An unemployed worker accepts a job if the wage offer is above her reservation wage, r. Since job offers arrive at the same rate whether employed or non-employed, the decision to accept a current offer has no consequences for future job opportunities. The reservation wage thus does not depend on $F(w)$, but is only equal to b, the value of leisure. As long as the wage offer is higher than this amount, the worker accepts the offer.[2]

[2] If $\lambda_u < \lambda_e$, which is found in most estimates, then the reservation wage will be below b as taking any job improves future employment prospects which means that a worker may want a job even if it makes them worse-off in the short-term.

The employers decision

The employer's decision in this model is to choose a wage to maximize profits: $\pi = (p - w)N(w, F)$, where $N(w, F)$ is the steady-state level of employment in a firm that pays a wage w when the firm distribution of wages as a whole is F. An employer who pays a wage w will recruit workers from the nonemployed (as long as w greater than b), and from workers in other firms which pay less than w. The employer will lose workers who exit to non-employment or leave the labour force or who quit to other firms which pay higher wages. In

With profit maximization, in equilibrium, the distribution of wages does not have any spikes (mass points), if the job offer to loss ratio, λ / δ, is between ∞ and 0.

Proof by contradiction: Suppose a mass of firms try to offer the same wage, w. Consider the perfectly competitive situation where w = p, then all these firms must be making zero profits. If there is some 'stickiness' in the speed at which workers can find jobs paying higher wages (separation rate is finite and recruitment rate is positive), a firm that lowers its wage can make profit. Its quit rate will now be higher and recruitment rate lower, but the recruitment rate is not zero and its quit rate not infinite so it will still have some workers in steady-state. For example, an unemployed worker gets a wage offer from this firm that is above their reservation wage, and decides to accept the offer, since she can always look for a better job while she is working.

Now suppose a mass of firms try to offer the same wage, w, where w < p. If a firm deviates by paying an infinitesimally higher wage (w+e), profit per worker is only inifitesimally reduced but the all workers from the other firms want to work for that firm, (as long as $\lambda > 0$) as recruits now come from workers in all the firms who continue to pay w. Hence, profits must rise and the initial situation could not have been an equilibrium.

This means firms will offer different wages over a distribution between b and p to identical workers. This is a striking feature of the Burdett and Mortensen model. We are left still to solve for the actual distribution of wages that different firms offer, to find the equation for the labour supply function of these workers, and to show that an equilibrium exists.

Approach to solving

We need to solve for the equilibrium distributions of $N(w,F)$ and $F(w)$. In equilibrium,

1) Steady-state unemployment

Define u as the unemployment rate.

$$\dot{u} = \lambda u M_w - \delta(1-u)M_w$$

When $\dot{u} = 0$, inflows of non-employed = outflows of non-employed.

This occurs when $u = \delta/(\delta+\lambda)$.

A note on the interpretation of the rates

$\lambda\Delta$ is the expected number of job offers for a single individual over period Δ. A single offer is distributed randomly across all individuals, so the probability of receiving that offer is $1/M_w$. Let ΔQ be the total number of offers per period. $\Delta\lambda = \Delta Q / M_w$. If we normalize the total number of offers per period to 1, then we can interpret λ as the probability an individual receives an offer. Burdett and Moretensen do this, but it may not be intuitive to think about the model this way with only 'one' firm offer per period.

The analysis is the same regardless of the value for λ if we consider continuous time, as $\Delta \to \infty$. λ and δ may take on any finite value (includes those greater than 1). The values of these parameters determine the highest steady state wage. These values are offer rates and exogenous exit rates. The interpretation is analogous to interest rates in continuous time.

2) Steady state employment for firms paying less than w

Define $G(w,F)$ as the cdf wage distribution of workers (the fraction of workers paid less than w).

The total fraction of persons getting paid w or less is: $G(w,F)M_w(1-u)$.

The job outflow rate for workers paid w or less is: $(\delta+\lambda[1-F(w)])G(w,F)M_w(1-u)$

The number of jobs at firms paying less than w can only increase from hiring unemployed. The inflow of non-employed accepting jobs that pay w or less is $M_w u\lambda F(w)$

When $\dot{G}=0$, inflows to firms paying w or less = outflows of firms paying w or less.

Equating inflows and outflows, we get:

$$G(w,F)=\frac{\delta F(w)}{\delta+\lambda[1-F(w)]}$$

$G(w,F)<F(w)$ for $0<F<1$: workers are concentrated in the better paying jobs, implying that such firms have a higher level of employment. The higher the offer rate, the more concentrated workers are at the high paying firms.

3) Steady state employment for each firm

$N(w,F)$ is the expected number of workers at a firm that pays w.

The separation rate for a firm that pays w is $s(w,F)=\delta+\lambda[1-F(w)]$

So the total outflow of a firm is $s(w,F)N(w,F)$, where if N(w) is the level of employment.

Job offers are spread equally over all firms.

The inflow of new recruits to a firm that pays w from the non-employed is $\frac{M_w}{M_F}u\lambda = \frac{u\lambda}{M}$

λ / M_F is the offer rate from one firm.

The inflow of new recruits to a firm that pays w from the employed is: $\frac{M_w G(w,F)(1-u)\lambda}{M_F} = \frac{G(w,F)(1-u)\lambda}{M}$

So total flow of new recruits to a firm that pays w is:

$$
\begin{aligned}
R(w,F) &= \frac{\lambda}{M}[u+(1-u)G(w,F)] \\
&= \frac{\lambda}{M}[u+(1-u)\frac{\delta F(w)}{\delta+\lambda[1-F(w)]}] \\
&= \frac{\lambda}{M}[\frac{\delta}{\delta+\lambda}+\frac{\lambda}{\delta+\lambda}\frac{\delta F(w)}{\delta+\lambda[1-F(w)]}] \\
&= \frac{\delta\lambda}{M}[\frac{1}{\delta+\lambda[1-F(w)]}]
\end{aligned}
$$

When $\dot{N}=0$, worker inflows of a firm = worker outflows of a firm.

$$
\begin{aligned}
s(w,F)N(w,F) &= R(w,F) \\
(\delta+\lambda[1-F(w)])N(w,F) &= \frac{\delta\lambda}{M}[\frac{1}{\delta+\lambda[1-F(w)]}]
\end{aligned}
$$

$$
N(w,F) = \frac{\delta\lambda}{M(\delta+\lambda[1-F(w)])^2}
$$

This is equilibrium labor supply to a firm paying w. Note that it is the monopsony case where firm labor supply is not infinitely elastic. A firm that pays the lowest wage (the reservation wage) in the distribution has labor supply

$$N(w^L, F) = \frac{1}{M}\frac{\delta\lambda}{(\delta+\lambda)^2}$$

A firm that pays the highest wage has labor supply

$$N(w^H, F) = \frac{1}{M}\frac{\delta\lambda}{\delta^2}$$

4) Steady state profits: All wages offered generate the same expected profits

$$\pi(w, F) = (p-w)N(w, F) = \frac{\delta\lambda(p-w)}{M\left(\delta+\lambda[1-F(w)]\right)^2}$$

The lowest wage offered will be the reservation wage. Suppose not. The lowest paying firm only obtains workers from the non-employed, at rate $\frac{\lambda u}{M}$. Employment is independent of the wage offered for a firm that offers the lowest wage in the distribution. Such a firm can lower the wage (and increase profit per worker) until w=b.

Knowing what the profit is for the lowest paid firm lets us determine the equilibrium profit for all firms:

$$\pi^*(w, F) = \frac{\delta\lambda(p-b)}{M\left(\delta+\lambda[1-F(b)]\right)^2} = \frac{\delta\lambda(p-b)}{M\left(\delta+\lambda\right)^2}$$

5) Solve for equilibrium wage distributions

Wage cdf of firms:

$$\frac{\delta\lambda(p-w)}{M\left(\delta+\lambda[1-F(w)]\right)^2}=\frac{\delta\lambda(p-b)}{M\left(\delta+\lambda\right)^2}$$

$$\frac{\sqrt{p-w}}{\delta+\lambda[1-F(w)]}=\frac{\sqrt{p-b}}{\delta+\lambda}$$

$$\frac{(\delta+\lambda)\sqrt{p-w}}{\lambda\sqrt{p-b}}-\frac{\delta}{\lambda}-1=-F(w)$$

$$\frac{(\delta+\lambda)\sqrt{p-w}}{\lambda\sqrt{p-b}}-\frac{\delta}{\lambda}-1=-F(w)$$

$$F(w)=\frac{(\delta+\lambda)}{\lambda}\left[1-\sqrt{\frac{p-w}{p-b}}\right]$$

From the same equality, the highest wage offer in equilibrium, when F=1, is:

$$w=p-\left(\frac{\delta}{\delta+\lambda}\right)^2(p-b)$$

Wage cdf of workers:

$$G(w,F)=\frac{\delta F(w)}{\delta+\lambda[1-F(w)]}$$

$$G(w)=\frac{\delta}{\lambda}\left[\sqrt{\frac{p-b}{p-w}}-1\right]$$

Wage pdf of workers:

$$g(w)=\frac{\delta}{2\lambda}\left[\frac{\left(p-b\right)^{1/2}}{\left(p-w\right)^{3/2}}\right]$$

B = 0, p = 1, $\lambda = .5$, $\delta = .05$, highest wage: 0.991736

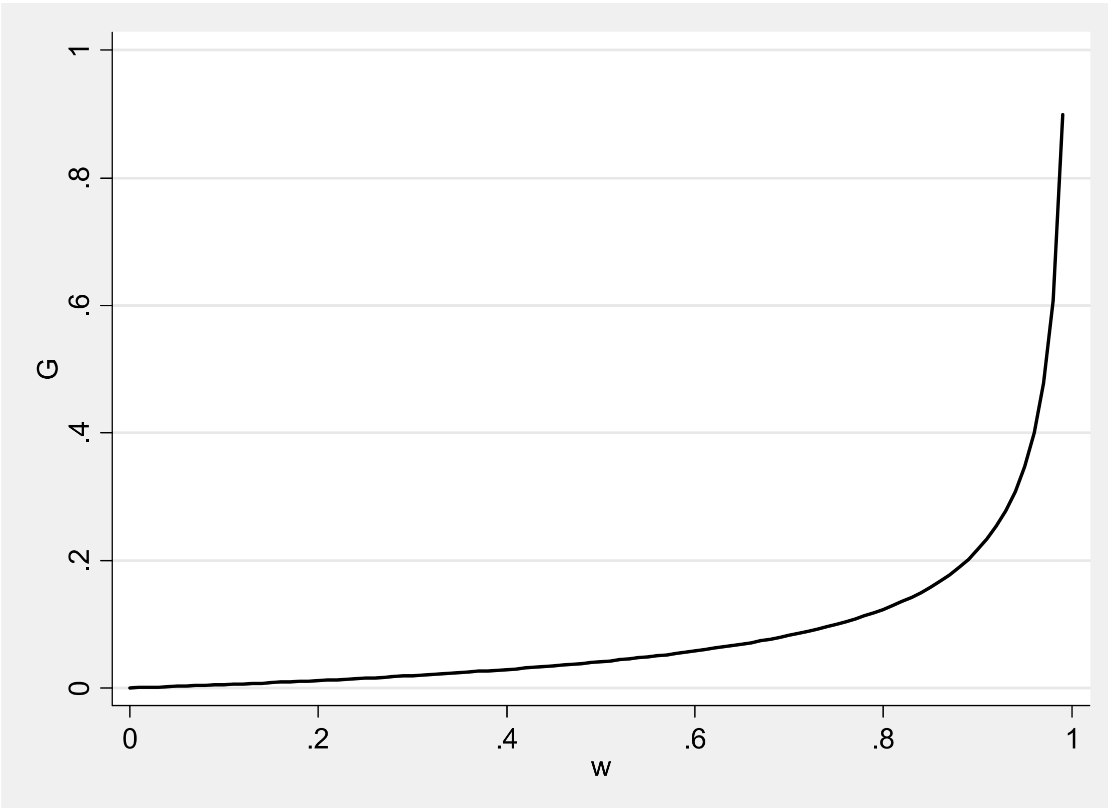

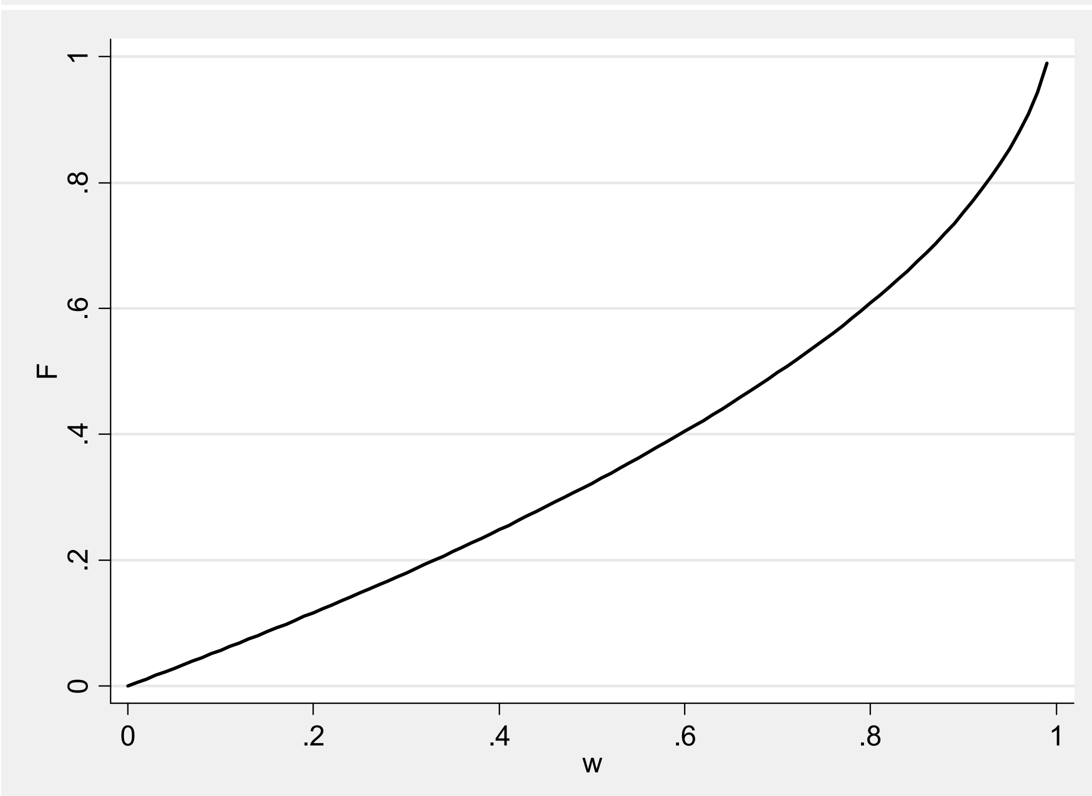

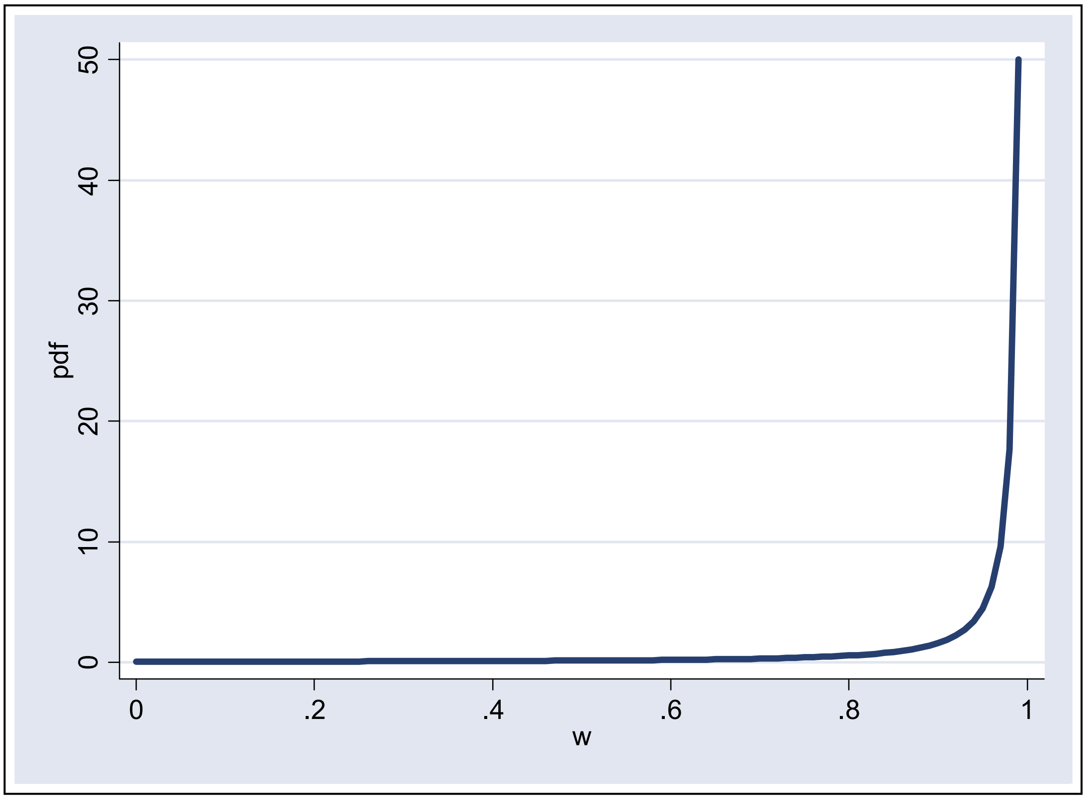

pdf
0
10
20
30
40
50
0
.2
.4
.6
.8
1
w

B = 0, p = 1, $\lambda = .02$, $\delta = .05$, highest wage: 0.489796

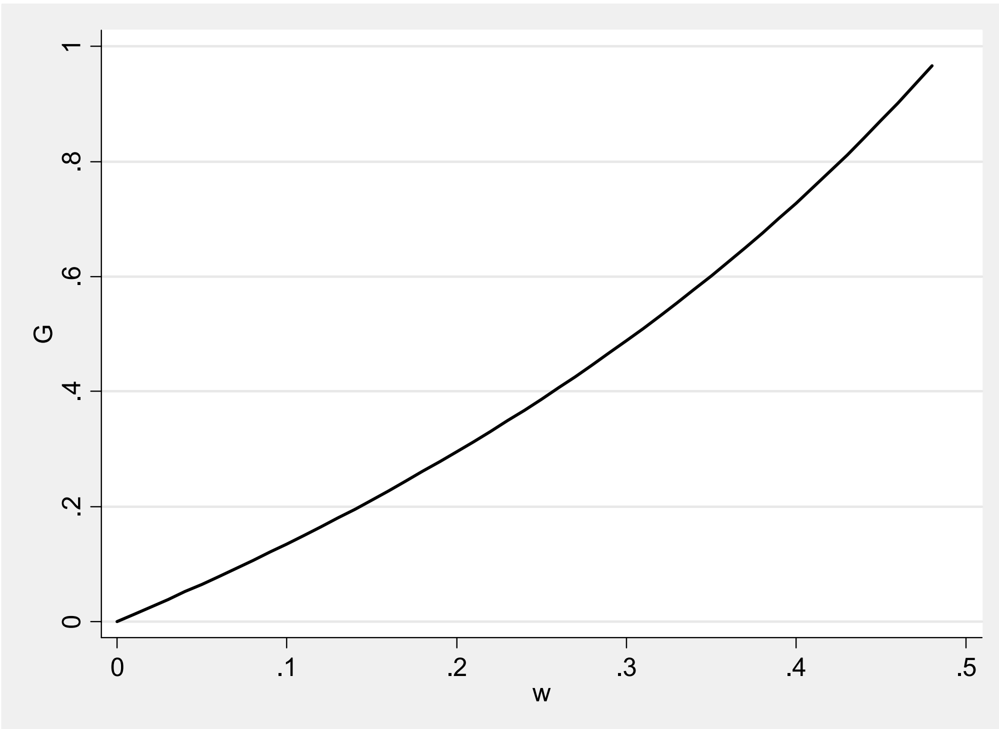

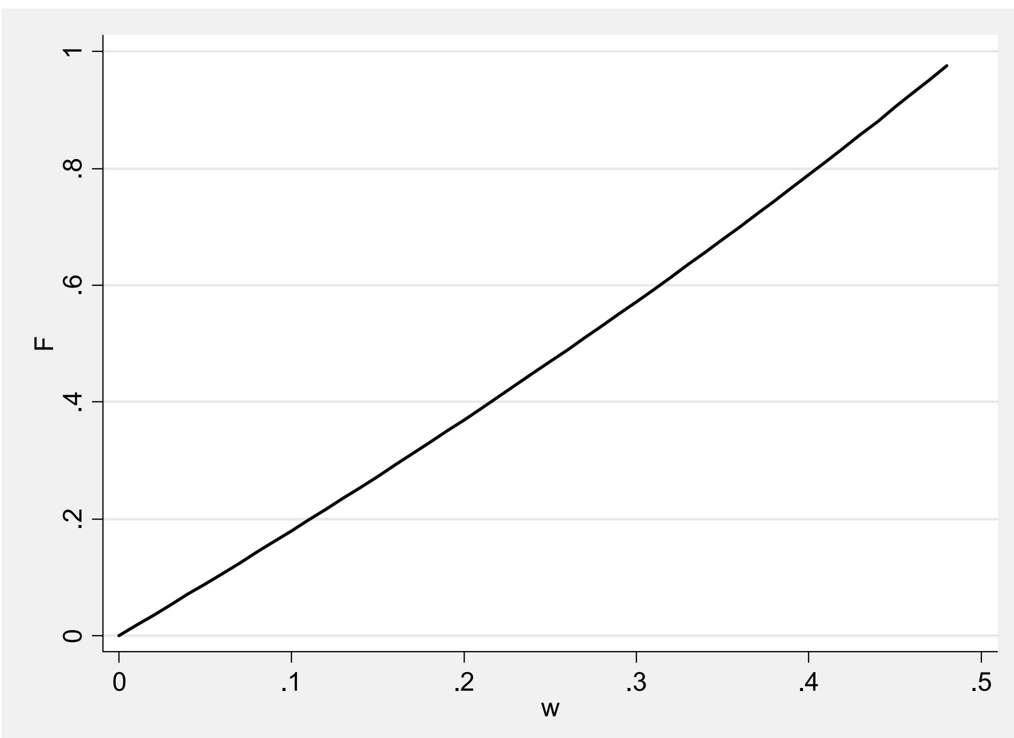

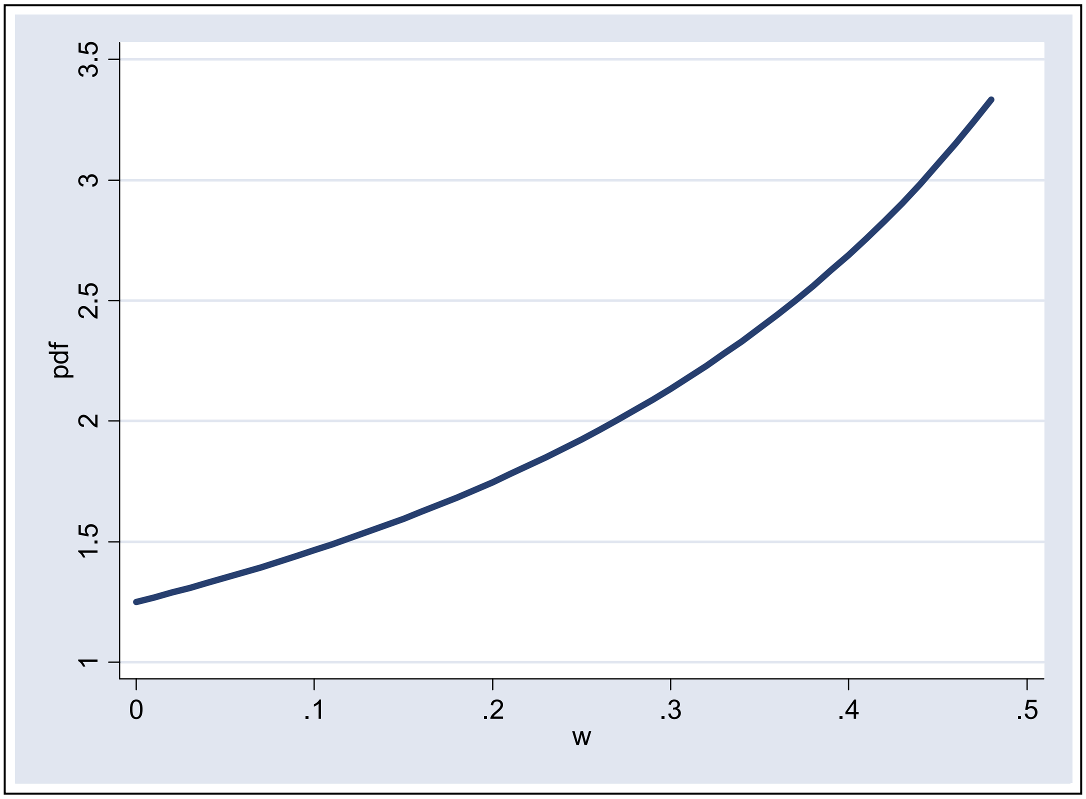
3.5
3
2.5
2
1.5
1
pdf
0
.1
.2
.3
.4
.5
w

B = 0, p = 1, $\lambda = .01$, $\delta = .99$, highest wage: .9999

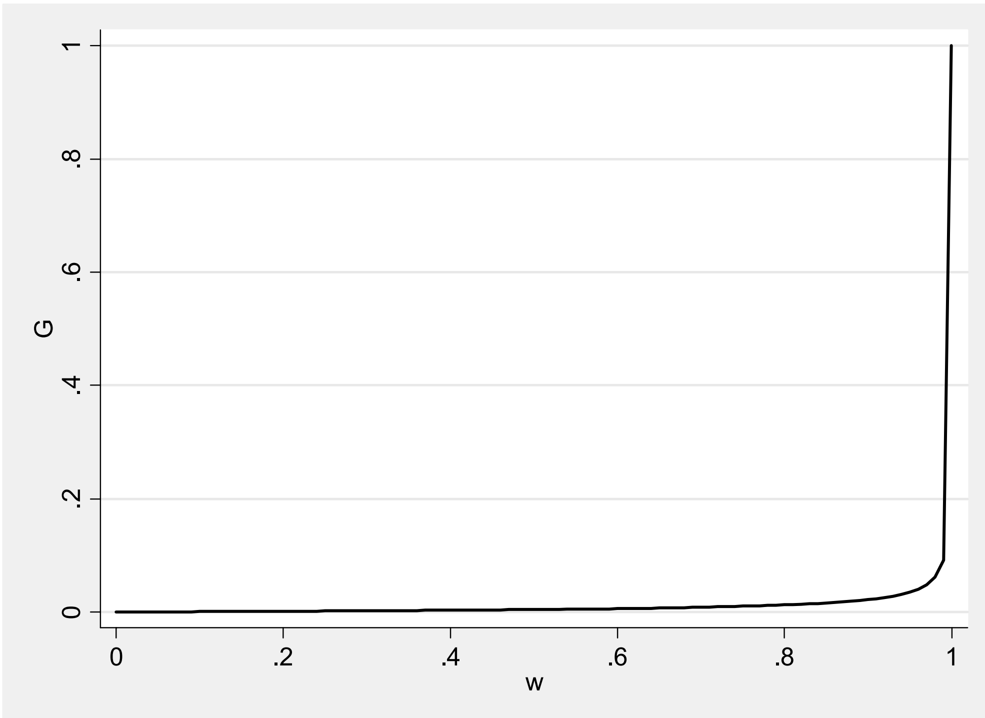

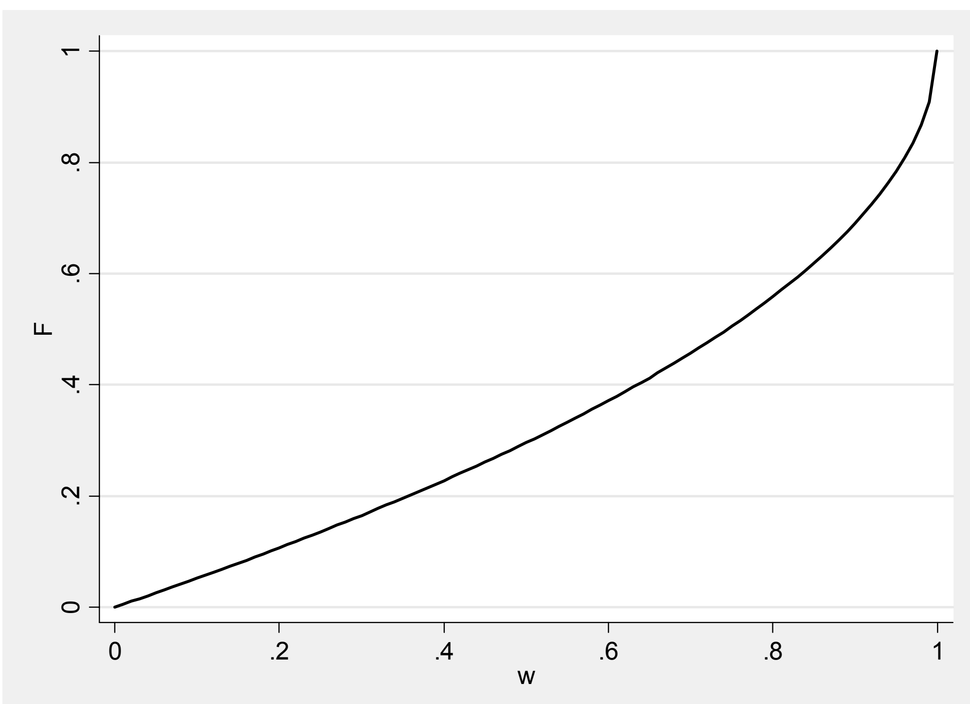

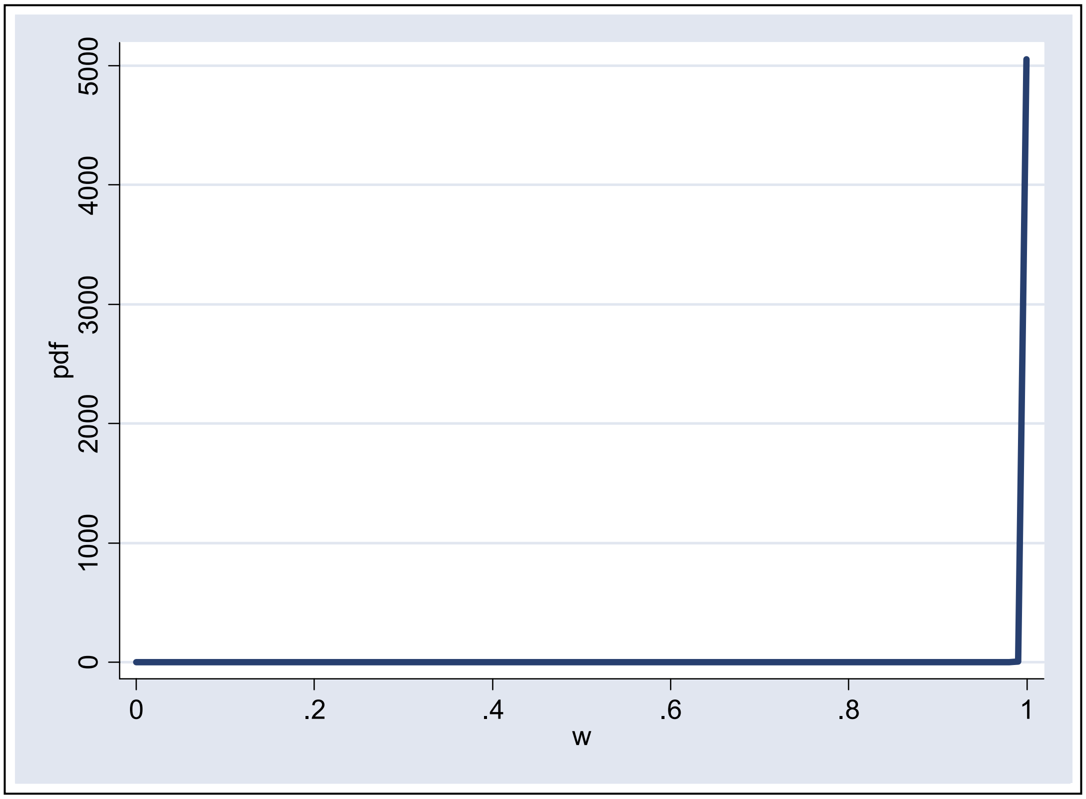
5000
4000
3000
2000
1000
0
pdf
0
.2
.4
.6
.8
1
w

The wage distribution does not look like anything we see in the data. Evidence of the wage distribution for the average worker (across firms) seems follow more of a log-normal distribution. The model presented here, however, is highly stylized model. Allowing for heterogeneous productivity for firms help generate wage distributions that follow more of a hump shape (see, for an introduction to this, Mortensen, cited below).

Now, take a step back to realize what we've done. We've constructed a model where workers with the same productivity and tastes are paid different wages. Why? We've assumed market friction: Workers must search employers for a job, and don't get offers from everyone at once. This causes workers to stay at a firm while searching for a job that pays a higher wage. Each employer chooses a wage given this search behaviour of workers. The labour force available to an employer is constantly changing. There is a nice balancing act, because there is a tradeoff between offering a higher wage and having a more stable labour force (fewer quits) and lower profit.

Note, as the relative rate of new offers, $\frac{\lambda}{\delta}$, becomes infinite, we arrive back at the perfectly competitive situation. Thus, the assumption that not all firms offer jobs to a single worker at once leads to monopsony power, and the general result that wages differ even though everyone is equally productive.

$\frac{\lambda}{\delta}$ is sometimes referred to as the market-friction parameter. It determines the steady state unemployment rate, and helps determine both the location and spread of the wage distributions.

"It is important to correct the impression that those who believe that employers have some market power over workers are extremists – the reality is that those who believe in perfect competition are the fanatics as perfect competition is one point at the edge of the parameter space and every other point in the parameter space gives employers some monopsony power." (from Manning).

Note, another implication, because workers only switch employers in response to a higher wage offer, workers that have been in the labour market longer are more likely to be earning a higher wage (and be at the firm longer).

For more introduction on this elegant model and interesting extensions, see Manning's book and also

Dale T. Moretnsen, 'Wage Dispersion: Why are similar workers paid differently?", MIT press.

Main interesting implications of model:

1) wages rise with experience
2) exogenous (involuntary?) unemployment shocks assumed and necessary for solving model
3) Identical workers at firms with same technology paid differently
4) Firms have monopsony power and make profit

Empirical Evidence?

An implication of the Burdett-Mortensen model is that luck may play a role in labor market success. Identical workers randomly assigned an initial job potentially end up with very different lifetime earnings. It's an empirical question whether this, in fact, occurs, and sceptics don't find the current evidence convincing. The existence of firm fixed effects, controlling for individual fixed effects may reflect an interaction between high productivity workers and high productivity firms, or unobserved changes in individual productivity that are not controlled for by the average individual fixed effects. Francis Kramarz and John Abowd are doing a lot of interesting work in this area to try and uncover the interactions of the firm using French and U.S. panel employer/employee data. These are powerful datasets that open up a lot of possibilities for future research.

Jacobson Lalonde, and Sullivan (AER 93) exploit employer-employee data to examine the wage dynamics from unexpected job loss. The Burdett-Moretensen model predicts job loss lower wages for those that lose their jobs, as they end up trying to climb the wage ladder again from the bottom. In a perfectly competitive setting, workers should be able to obtain similar paying jobs at other firms, and we shouldn't observe a long run impact. There are other reasons to expect a prolonged impact from job loss, especially in the presence of specific firm capital. But the analysis is informative in that mild and temporary impact from job loss shocks would be evidence against the Burdett-Moretensen model.

JLS have quarterly data over 15 years from government administrative records in Pennsylvania.

Sample is: workers at same firm for 6 years between 1974-80.

'Treated' Group: workers that left firm in quarter where at least 30% of firm's labor force left (e.g. mass layoff or firm closure). This is an attempt to identify exogenous job loss.

'Control' Group: The rest of the sample

Note, data limitation is that can't identify workers with zero earnings as those unemployed or those that left the state. JSL drop 25% of the sample of separators that no longer show up in the sample after displacement. They also restrict sample so that every worker in it has at least some wage or salary earnings in one quarter per year. They do this to focus on wage loses do to lower paying jobs rather than to unemployment. They argue this sample selection likely biases the effects of unexpected job loss downwards.

Figure 1

An omitted variables bias concern is that displaced workers are not the same in productivity as retained workers. JSL estimate the following model:

$$w_{it} = \alpha_i + \gamma_t + \beta age_{it} + \sum_{k \geq -m} D_{it}^k \delta_k + e_{it} ,$$

where D_{it}^k is a indicator variable for those in the 'displaced' sample, and δ_k are the coefficients indicating the average effect of displacement k quarters since displacement. With the panel data, we can estimate the effect even prior to displacement actually occurring. The fixed effect captures average earnings for each individual prior to displacement. So to identify this, we need to observe earnings of at least some displaced workers more than m quarters before displacement. JSL has at least 6 years of predisplacement data. In effect, the individual fixed effects are average (age-adjusted) earnings over roughly these six years.

Finally, to account for the possibility that individual productivity is changing over time, and that change is correlated with displacement, JSL add individual linear time trends:

$$w_{it} = \alpha_i + +\omega_i t + \gamma_t + \beta age_{it} + \sum_{k \geq -m} D_{it}^k \delta_k + e_{it}$$

Figure 2 plots the δ_k coefficients.

Even 6 years after displacement, quarterly wages are $1,600 lower than the control group's.

Relative wage differences begin even before displacement, which is a bit disturbing. JSL suggest (exogenous) events leading up to the displacement also lead to lower wages before displacement. They suggest the rate of temporary layoffs may increase prior to mass-layoffs or firm closures. Slower wage growth in these firms may also have this effect. The finding that the wage diversion prior to displacement does not continue past 4 years prior is somewhat assuring.

Figure 3 shows the same, but for workers that leave, but not from a mass-layoff firm. These workers are more likely to be leaving voluntarily, and catch up occurs.

Table
Annual Earnings and Unemployment Insurance Participation of Fathers, 1978 – 1988
Restricted Sample (initial 78-81 firm sample size 5 – 500)

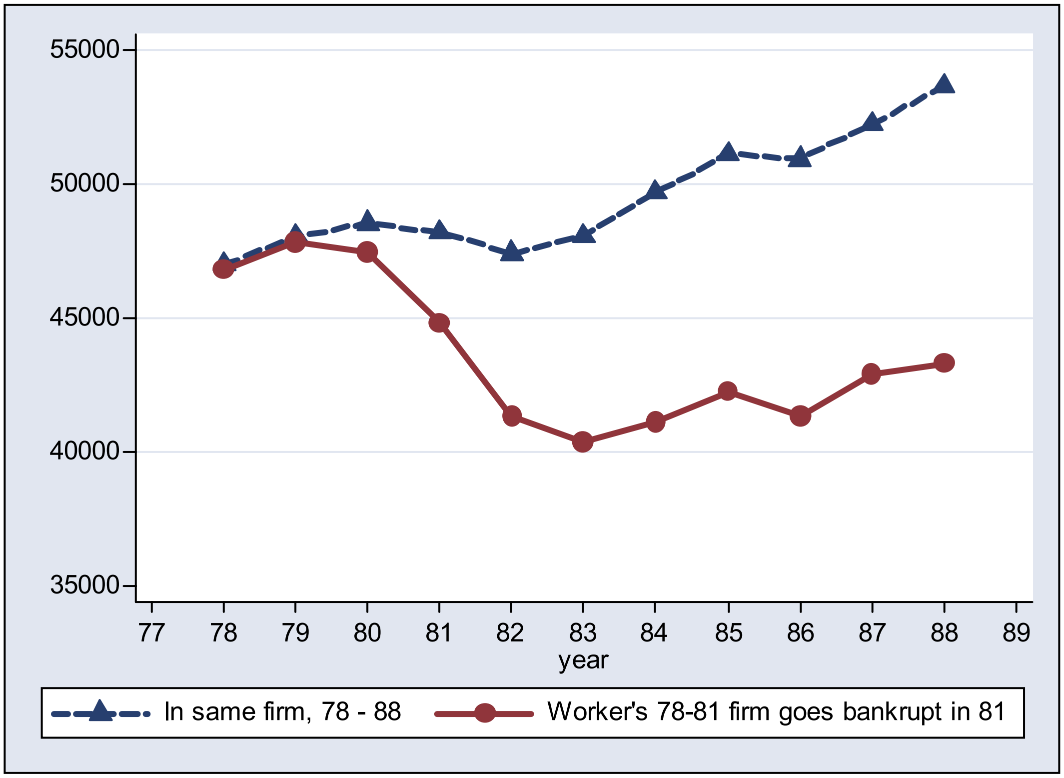

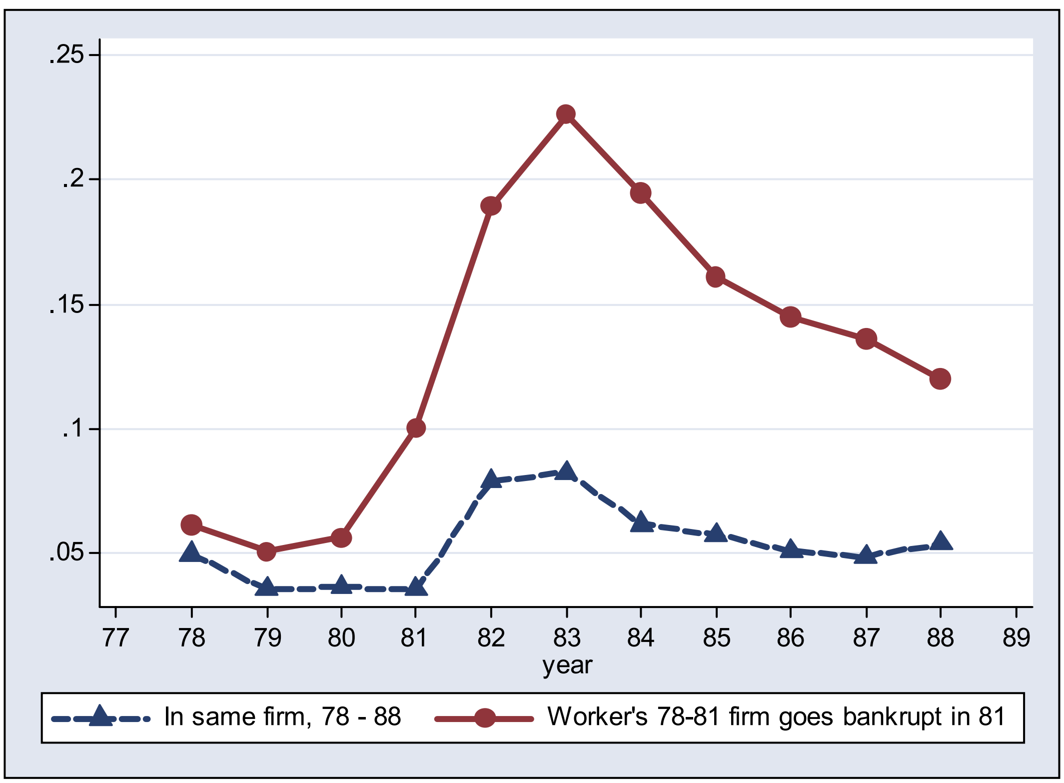

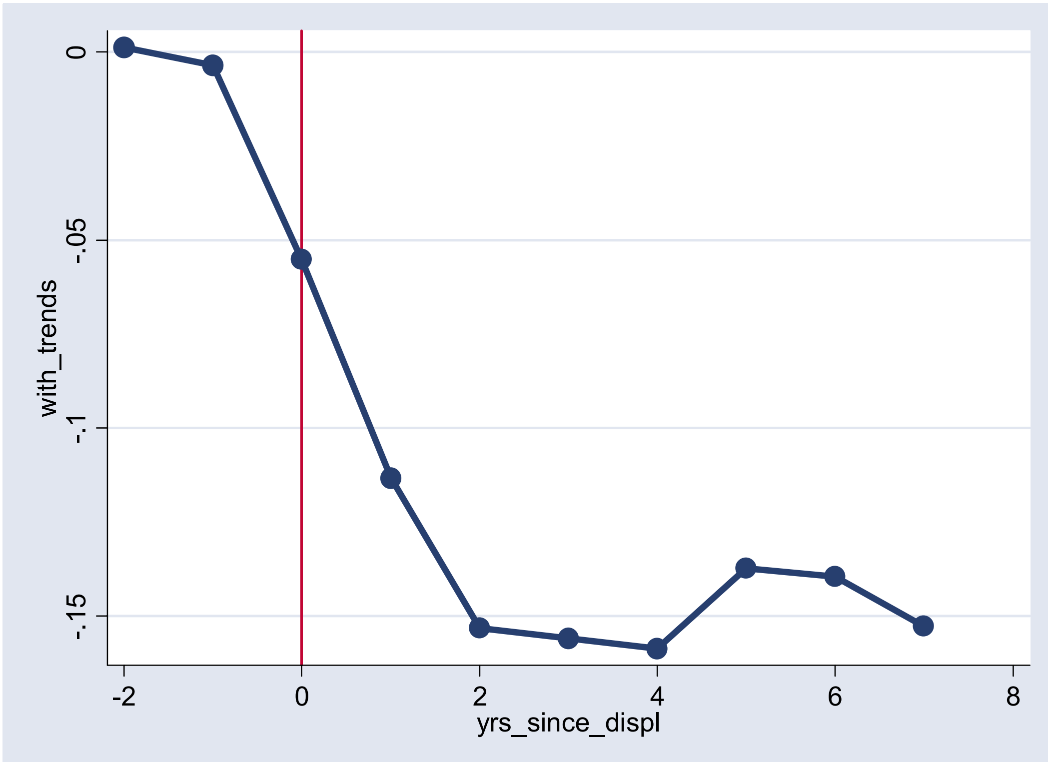

These tables are similar to the JSL analysis, but with a sample of Canadian workers.

The sample is all older workers (around 40) in same firm between 1978-80. The 'treated group' was at a firm that closed in 1980. The control group was never at a firm that experienced a closure the year they were working there, between 78-88.

The data also tracks individuals even if they are not working, but filing taxes (unlike JSL).

Canada experienced a particularly bad recession in 82-83.

The last table shows the coefficients on log earnings marked by the year since firm closure (including individual fixed effects and individual time trends). There is some slippage in the timing because a closure is defined as the firm does not file employment reports after year 0, so the closure could have taken place over a 2 year window.

That wages never catch up is surprising (disturbing?).

This is a sample of fathers. The dataset contains matched children who were 11-14 when father's lost their job. I have earnings data of the children when they are 25-30 and am examining whether there are differences between earnings fathers from the bankruptcy sample to fathers from the 'control'.