Graduate Labor Economics
ECN 250A, Spring 2019

Professor Brendan Price*
University of California, Davis

This is a closed-book exam. Answer all questions briefly but clearly. The exam is graded out of 75 points; each subquestion is worth 5 points.

I Short-answer questions (30 points)

(a) In contrast to the classical intuition that there should be a "law of one wage" for skill, there is substantial empirical evidence that equally skilled workers may be paid different (hourly) wages depending on which firm happens to employ them. Give two possible explanations for why such wage differences may occur.

Among many possible explanations, these wage differences may reflect compensating differences for job amenities (or disamenities); monopsony power, whereby firms offer different wages to different workers with different perceived elasticities of supply; search frictions, which can generate equilibrium dispersion in match quality across employer-employee pairings; wage premiums paid by unionized firms; or efficiency-wage motives that are stronger for some firms than for others (for example, firms that invest heavily in on-the-job training may offer elevated wages to deter costly turnover).

(b) Consider a simple model in which aggregate output depends on a combination of routine and non-routine tasks (as in Autor, Levy, and Murnane 2003). How would you expect a decline in the price of computing power to affect *employment* in routine-intensive occupations? How would you expect it to affect *wages* in routine-intensive occupations, relative to wages in non-routine occupations?

The basic idea of the ALM model (and related models) is that computer capital substitutes for human labor in the execution of routine tasks and complements human labor in the execution of (some) non-routine tasks. In these models, a reduction in the cost of computing will reduce demand for human workers in routine occupations. This will result in lower routine employment (and typically higher non-routine employment), but the effect on observed wages is ambiguous: the wage per efficiency unit in routine tasks will decline, but depending on which workers self-select from routine into non-routine occupations, it's possible that observed hourly wages might rise if the composition effect dominates the task-price effect.

(c) Retail pharmacies like CVS rely on a mixture of human cashiers and automated self-checkout kiosks. Suppose that the price of an automated kiosk falls by 50% in the next 5 years. In a two-factor model, how will this price change affect (i) cashier employment and (ii) the number of kiosks used?

Declines in the price of a kiosk have theoretically ambiguous effects on cashier employment. The substitution effect says that pharmacies will substitute away from human cashiers and towards kiosks,

 Given my subsequent affiliation at the Federal Reserve Board, I note that any views presented here are my own and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff.

but the scale effect says that it will be cheaper to operate a pharmacy establishment, so companies may open up new pharmacy locations. The net effect is ambiguous and it depends on the elasticity of substitution between cashiers and kiosks. The number of kiosks will rise, since both scale and substitution effects favor installing more kiosks.

(d) Whereas many studies based on changes in state-level minimum wages have tended to find at most modest disemployment effects from minimum-wage increases, the Jardim et al. (2018) study of Seattle's $13 minimum wage found substantial negative employment effects. Provide a possible explanation for why this study found larger effects than most prior work.

Here are three possible explanations. First, Jardim et al. studied a larger wage increase than are typically studied in the state-level literature; if there are nonlinear effects, then the marginal disemployment effect of raising the minimum wage may be larger around $13 than at lower levels. Second, employers may find it easier to shift employment to other locations in response to a city-level minimum wage than to state-level minimum wages. Third, Jardim et al. are able to identify low-wage jobs, rather than relying on teenagers or restaurant workers as imperfect proxies for the set of jobs that are likely to be affected by minimum wage laws.

(e) Commuting zones are constructed by merging adjacent counties linked by dense commuting flows. Explain why researchers are increasingly using CZs, rather than counties, to define local labor markets.

In response to a localized shock, jobs may be lost in one county while other jobs are created in a nearby county. If these counties share dense commuting flows, then there may be no net change in workers' access to jobs. Aggregating to the CZ level ensures that our measure of a local labor market encompasses these local spillovers. Another reason to use CZs is that people may live and work in different counties, in which case an adverse shock to labor demand *in* a particular county may be a poor measure of the shock to labor demand for residents *of* that county. Aggregating sidesteps this problem as well.

(f) Describe a method for identifying seasonal workers in high-frequency panel data.

Workers who are employed in the same kind of seasonal job for multiple consecutive years will tend to lose their jobs around the same time each year. Coglianese and Price (2019) find that a disproportionately large share of job losses are spaced exactly 12 months apart, and they show how machine-learning tools can be used to pinpoint workers who (based on their observable characteristics at the time of an initial job loss) are excessively likely to separate again 12 months later. Such workers are likely to be seasonally employed.

II Job-finding (10 points)

If we examine the hazard rate of reemployment among a group of new job-losers, we often see that this hazard rate initially rises with duration since job loss, peaks a few months later, and then steadily declines.

(a) Provide an explanation for why the job-finding hazard tends to rise in the initial months after job loss.

The hazard rate may initially rise because workers are being recalled to previous employers, or because many workers lose their jobs at the onset of seasonal downturns (e.g., construction or retail in winter) and are able to find work again once their industry rebounds from its seasonal decline. Another possibility is that the application and hiring process takes some time, so that even if a worker submits an application the day she is laid off, it may take a couple of months before she is hired and shows up for her first day of work.

(b) Offer two possible explanations for why job-finding hazards tend to decline after the first few months.

Four possible explanations are skill depreciation, employer discrimination against the long-term unemployed, discouragement causing workers to search less, and dynamic selection whereby the workers who are most prone to find jobs find jobs faster and hence enter the risk set for the hazard rate.

(c) Job-finding hazards sometimes exhibit a "spike" at a particular duration after job loss. Why?

Such spikes typically reflect increased job-finding at the moment when workers exhaust unemployment insurance benefits. (Other phenomena such as recalls to past employers or the end of the "off-season" can also cause upticks in the job-finding hazard at certain durations, but those increases are generally not as sharply concentrated as the "spike" caused by UI exhaustion.)

III Displaced workers (15 points)

The figure below, taken from Jacobson, Lalonde, and Sullivan (1993, hereafter "JLS"), plots estimated differences in earnings between Pennsylvania workers displaced in mass layoffs during 1980–1986 and a comparison group of non-separators, partialing out worker and time fixed effects. (Differences between the "without trends" and "with trends" series are not important for answering this question.)

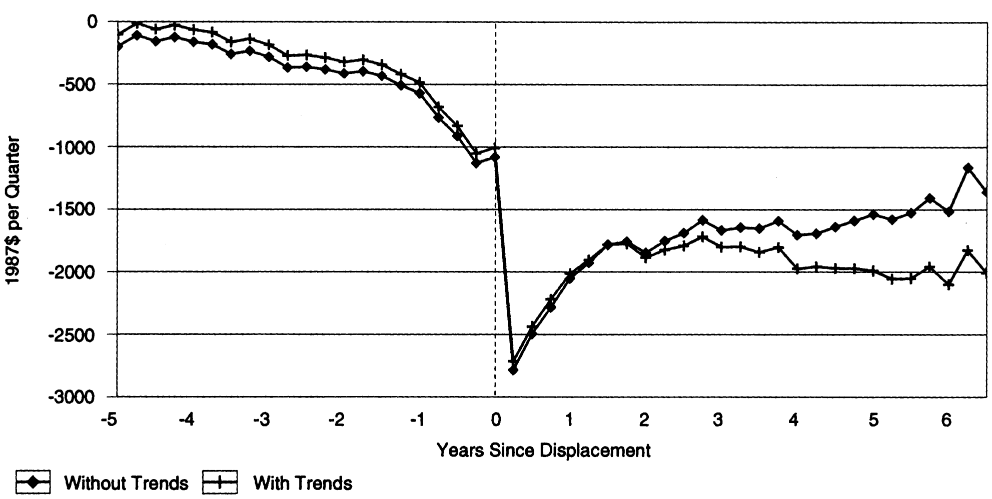

(a) Why do JLS focus primarily on workers who lose their jobs during mass layoffs, rather than looking at all job-separators?

Two reasons. First, it is typically hard to distinguish voluntary quitters from involuntary job-losers in administrative data like those that JLS use. By focusing on workers who separate in mass layoffs, we can be reasonably confident that most of the observed separations reflect involuntary job loss. Second, workers who lose their jobs separately from mass layoffs may be adversely selected on earnings potential (due, e.g., to public learning about adverse job performance, an adverse health shock, or another life event that would result in earnings losses even in the absence of job loss). Focusing on mass layoffs avoids this selection concern.

(b) Provide a possible explanation for why displaced workers' earnings begin to decline several years before they are actually displaced.

Firms may experience prolonged periods of financial distress before eventually conducting mass layoffs. Therefore, in the lead-up to a layoff event, workers at such firms may miss out on wage growth they would normally have experienced, or they may be subject to an increased frequency of temporary layoffs that lead to partial reductions in earnings even before a permanent separation.

(c) Provide a possible explanation for why displaced workers experience long-term earnings losses. Explain why such losses are sometimes interpreted as evidence of imperfect competition in the labor market.

There are many possible explanations. Since mass layoffs typically represent substantial reductions in labor demand within particular industries, displaced workers may be unable to find new jobs in their previous occupation or industry, in which case they will no longer be able to use their previously accumulated occupation or industry-specific human capital. Displaced workers may also lose rents or "match capital" accumulated through job-to-job transitions earlier in their careers. The stress from job loss may contribute to longer-term health problems that reduce productivity. Long-term earnings losses suggest that displaced workers may have been receiving substantial rents in their pre-separation jobs, which in turn is consistent with imperfect competition.

(d) JLS also report results from a specification that includes interactions between time effects and a dummy variable for each 1979 (pre-separation) firm, so that displaced workers' earnings are compared to those of non-separators from the same firm. Explain why non-separators from the same firm might be a *better* control group than non-separators in general. Explain why they may be a *worse* control group.

A displaced worker's former coworkers are likely to be especially similar to him or her in terms of unobserved ability and other attributes, so that coworkers are more likely to be on similar earnings trends in the absence of a layoff event. On the other hand, these coworkers may also experience depressed earnings because of the same financial difficulties that prompted the mass layoff in the first place, in which case using coworkers as a control group is likely to understate the true magnitude of the earnings losses. Relying only on within-firm variation may also amplify any selection concerns relating to which particular workers are laid off.

(e) Because workers who move out of state disappear from the dataset, JLS restrict their sample to workers who are observed with positive earnings at some point following their job loss. Given this sample restriction, are the true mean earnings losses experienced by displaced workers likely to be larger or smaller than the estimates shown above suggest? Justify your answer.

Excluding the "zeroes" presumably excludes some workers who remain in-state but drop out of the labor force. If this is the case, then true earnings losses are likely to be even bigger than the (large) effects reported by JLS. On the other hand, it's possible that the most capable workers move out of state after job displacement and find high-paying jobs in other states, in which case the true earnings losses could be smaller than those reported by JLS.

IV Regional evolutions (10 points)

The figure below (from Amior and Manning, 2018) plots 2010 emp-pop rates and 1980–2010 population growth against 1980 emp-pop rates for the 50 largest US commuting zones (as ranked in 1980).

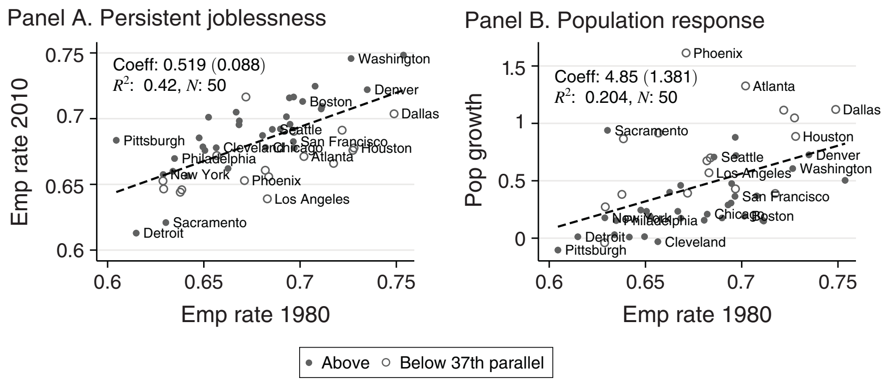

(a) Imagine that the right panel showed a *flat* (horizontal) relationship between baseline employment rates and subsequent population growth. Under this alternative fact pattern, offer two explanations for the persistence of local employment rates between 1980 and 2010.

If the right panel showed a flat relationship, it would suggest that there is no systematic migration from low-employment areas into high-employment areas, so that utility is equalized for the marginal resident of each location. Here are two explanations for equalized utility. First, there may be sorting on the basis of skill or preferences. For example, less-educated individuals might choose to live in rural areas whereas more-educated individuals choose to live in cities, or locations may differ in female labor force participation. Second, there may be differences in local amenities that compensate workers for geographic variation in job opportunities.

(b) Why is it surprising to see upward slopes in both of these relationships? What is a possible explanation for this fact pattern? (Amior and Manning offer one reconciliation, but feel free to propose another.)

There is directed migration from low-employment areas to high-employment areas. If labor demand is downward sloping in each location, we would expect this to put downward pressure on relative wages and employment rates in high-employment areas, so that regions gradually converge. What's surprising is that joblessness persists over such long periods of time. Amior and Manning argue that these facts can be explained by serially correlated industry shocks: places hit by a negative shock today are likely to get hit by another negative shock tomorrow, so that population adjustments fail to keep up with adjustments in labor demand. An alternative possibility is that migration causes general equilibrium effects (e.g., through aggregate demand or through knowledge spillovers) that increase labor demand in destination markets while lowering labor demand in origin markets, which would again cause employment differences to persist even in the face of a migration response.

(c) Amior and Manning instrument for employment growth using a classic Bartik IV. Write the formula for such an instrument (defining your notation), and describe one possible concern about using it.

There are multiple, subtly different ways of writing a shift-share instrument, but a typical representation is along the lines of

$$B_{it} = \sum_j \frac{L_{ij0}}{L_{i0}} \,\mathrm{dlog}\, L_{-i,jt}$$

where i indexes locations, j indexes sectors, t indexes time, L_{ij0} is baseline employment in a given sector in a given location, L_{i0} is total employment in that location, and $\text{dlog}\, L_{i,jt}$ is the log change in employment at the national level (excluding location i itself). One possible concern is that there may be a mechanical bias if we include location i in the calculation of the national shock, but this is addressed by the "leave-one-out" adjustment. A bigger concern is that the variation in the instrument may be largely driven by a few large or heavily shocked industries, in which case spurious unobserved industry-level shocks may lead to inconsistent estimates (e.g., industries' susceptibility to automation may be correlated with changes in Chinese import penetration).

V Reemployment bonuses (10 points)

Consider a continuous-time search model in which a worker enters unemployment at date $t = 0$. She can generate job offers at flow rate $s \geq 0$ by incurring search cost $\psi(s) = \frac{1}{2}cs^2$ (with $c > 0$). While unemployed, she receives flow utility b; once reemployed, she receives a flow wage $w > b$ forever. Her discount rate is δ.

In an effort to encourage job search, the government will pay this worker a one-time lump sum bonus valued at $\overline{L}$ if she is reemployed by some deadline date $T > 0$. The associated Bellman equation is

$$\delta U(t) = \max_{s \geq 0}\ b - \frac{1}{2}cs^2 + s(J + L(t) - U(t)) + \dot{U}(t)$$

where $U(t)$ is the value of unemployment at date t, J is the value of a job, $\dot{U}(t)$ is a time derivative, and

$$L(t) = \begin{cases} \overline{L} & \text{if } t \leq T \\ 0 & \text{if } t > T \end{cases}$$

(a) Suppose we want to study this model in discrete time. Write the associated discrete-time Bellman equation, with discount factor $\beta < 1$. State any assumptions you are making about timing conventions.

Let's assume that a worker who searches for (and finds) a job this period will start that job next period, and that the bonus payment will be paid next period as well. In this case, the Bellman is

$$U(t) = \max_{s}\ b - \frac{1}{2}cs^2 + \beta(s(J + L(t)) + (1 - s)U(t + 1))$$

(b) Returning to continuous time, prove that $s^*(t)$ is a decreasing function of b for all values of t.

The FOC for search effort implies that $s^* = \frac{1}{c}(J + L(t) - U(t))$. Since J and $L(t)$ don't depend on b, it suffices to show that $U(t)$ is increasing in b at all durations. We can prove this using a revealed preference argument: starting from a baseline value of b and the associated optimal search policy $s^*(t)$, an increase in b will increase the present discounted value of future utility flows enjoyed by an unemployed worker even if this worker does not reoptimize in light of the higher benefit level, so she is clearly better off. This implies that $U(t)$ is increasing in b at any duration. (An analytic argument based on the envelope theorem works too.)

(c) The bonus policy induces a discontinuous change in search intensity at date T. Compute the quantity

$$\Delta s \equiv \lim_{t \to T^+} s^*(t) - \lim_{t \to T^-} s^*(t)$$

as a function of exogenous parameters.

Let's compare search intensity at times $T-\varepsilon$ and $T+\varepsilon$ for ε very small. These are

$$s^*(T-\varepsilon) = \frac{1}{c}(J + \overline{L} - U(T-\varepsilon))$$
$$s^*(T+\varepsilon) = \frac{1}{c}(J - U(T+\varepsilon))$$

so that $\Delta s = \frac{1}{c}(\overline{L} - (U(T-\varepsilon) - U(T+\varepsilon))$. The trick is to recognize that $U(t)$ is continuous at date T, since the probability of finding a job (and hence getting the bonus) within a time interval of length 2ε goes to zero as $\varepsilon \to 0$. So this simplifies to just $\Delta s = \frac{\overline{L}}{c}$.