

MIT 14.662 Graduate Labor Economics II Spring 2009
Lecture Note 1: Earnings Inequality: Facts, Models and
Hypotheses

David Autor
MIT and NBER

April 2, 2009

1 PART I: LAYING OUT FACTS AND CLEARING EMPIRICAL UNDERBRUSH

It is useful to start with some key facts on changes in the wage structure in the U.S. and internationally that we will use to frame the discussion. Before discussing these facts, you may ask: why bother studying wage inequality at all? Substantively, this topic—the rapid, unanticipated growth of earnings inequality in many developed economies during the 1980s and 1990s, most notably in the U.S. and U.K. and West Germany (though the latter was only recently discovered)—has animated an outsized share of research in labor, development and international economics for a decade. This research has helped to set the agenda in macroeconomics, trade, theory and development, with many useful feedbacks. Because of the breadth of the topic and the sustained interest of the profession, researchers have brought a very large number of tools to bear on this general equilibrium phenomenon. Hence, the set of hypotheses and techniques is extremely rich. Using inequality as an organizing principal, one can tackle a lot of theory, econometrics and substantive knowledge that will continue to be useful to you even when the professional interest in inequality has died down.

1.1 KEY FACTS

1. Returns to education in the U.S. fell during the 1970s when there was a very sharp increase in the supply of educated workers. Returns to education then began a sharp rise in the 1980s. This rise slowed in the 1990s, though it did not reverse direction.
 - (a) This conclusion is robust to many sensible ways of measuring education returns.
 - (b) In a standard Mincer wage equation, return to year of education rose from about 7.5 percent in 1980 to 10.0 percent in 1990. But largest increase is between college and HS grads.
 - (c) Also an increase in the estimated return to experience for younger cohorts. Note: Since ‘returns to experience’ are typically estimated cross-sectionally, this pattern probably has no meaningful predictive power for lifecycle earnings.
2. Many other economies experienced a rise in earnings differentials in the 1980s, but only in the UK was the rise as pronounced as in the U.S. The entire distribution of outcomes was shifted far to the right, but experiences varied from large falls to very large growth:

- (a) Very large increases: US and UK
 - (b) Modest increases: Australia, Canada, Japan, Spain, and Sweden
 - (c) No noticeable changes: France and Italy
 - (d) Modest falls: Netherlands
 - (e) Large falls: South Korea
 - (f) Due to recent work by Dustmann, Ludsteck and Schonberg using Germany social security records, we now know that German earnings inequality increased rapidly during the 1980s and 1990s. Previous evidence from the GSEOP (a standard German data set) had indicated quiescent earnings inequality during this period.
3. Note that all of the countries above saw large falls in differentials in the 1970s (except perhaps Korea, little data.)
4. Overall earnings inequality as measured by the 90-10 rose sharply in the U.S. starting in the early 1980s. In many other countries, this began later. This is visible in almost any measure of inequality
- (a) Hourly and weekly earnings
 - (b) Reinforced by including non-wage compensation
 - (c) Looking at the very top - CEO's, top 1, 5, 10 percent of earners reinforces this picture (though .
 - (d) Also true in markets for workplace disamenities: non-standard work hours, safety in manufacturing.
5. After the 1990s, this monotone rise in overall inequality in the U.S. gave way to a 'polarization'—with wage inequality rising in the upper-half of the wage distribution but flattening out or declining in the lower half.
6. Some share of the rise in overall inequality is due to a rise in 'residual' inequality—that is, the inequality remaining after parceling out the estimated contribution of observables (such as education and potential experience). The exact decomposition of overall inequality into 'between-group' and 'residual' components is inherently arbitrary—it depends on

what X 's are included in the regression model used to explain the between-group components. However, it is still meaningful to talk about a rise in residual inequality if we hold the conditioning set of X 's fixed across time.

7. There is some disagreement about the timing of the rise in overall versus residual inequality. DiNardo, Fortin and Lemieux and Card/DiNardo find that inequality does not begin to grow until 1979. Other analyses suggest the rise starting earlier, perhaps in 1973, though all agree that it was much more rapid in the early 1980s. The answer is sensitive to the choice of data set (May/MORG CPS, versus March CPS, versus Census of populations).
8. A 2006 *AER* paper by Lemieux (2006) also argues that a large part of the rise in residual inequality is explained by mechanical effects of labor market composition. Because the U.S. workforce became simultaneously older and more educated during the 1980s and 1990s, and because the earnings of older, educated workers are typically more dispersed (as predicted by the lifecycle human capital model), this phenomenon gives rise to an increase in earnings dispersion without any change in the underlying factors affecting wage distribution. We'll discuss this hypothesis in more detail in a later lecture.
9. In the U.S., average and median wages stagnated after 1973 and fell considerably in absolute terms for low wage workers. This trend was reversed only after 1995. Other OECD countries did not experience this pattern of declining absolute wages, though did experience slowing growth.
10. The stagnation in wages corresponds to a post-1973 decline in the growth of Total Factor Productivity that had risen rapidly during the post-War 'golden age.' This experience of sharply slowing TFP growth was shared in all developed economies. TFP began to rebound in the mid to late 1990s in the U.S. and has continued to grow rapidly during the first half of the decade of the 2000's despite a recession and slow employment growth.
11. The wage gap between males and females closed considerably. This was particularly noteworthy given rapidly rising relative female labor supply since 1970. Most other advanced economies also saw a declining gender gap in the 1980s, though the U.S. again stands apart in that the trend change was quite sudden after 1979.

12. Notably, the gender gap *stopped* closing in the early 1990s, at a time when female labor supply also stopped rising. This appears a striking coincidence and we will discuss several papers examining the link between the gender gap, female labor supply and overall inequality.
13. The black-white wage gap, which closed rapidly in the 1970s, stagnated and/or expanded in the 1980s. The declining labor force participation of black males (combined with a severe rise in the rate of incarceration) probably masks an even larger decline in earnings capacity, as emphasized in a 2003 NBER paper by Amitabh Chandra. See also the startling work by Bruce Western and Becky Petit on the incarceration rates of black males.
14. Changes in the supply of skills. Changes in the growth rate of supply of skills will be an important explanatory variable for understanding wage structure changes.
 - (a) There has been remarkable growth in the supply of skills among all advanced economies.
 - (b) But there are also very large cross-country differences in the timing of acceleration and slowdown in production of skills
 - (c) U.S. had particularly rapid rise in 70s (partly due to the Vietnam war), slowdown in 80s. UK experienced the slowdown later. Netherlands and North Korea had extremely rapid supply growth during the 1980s, producing a rapid fall in earnings differentials.
 - (d) Slowdown in UK and Canada came later and was not as severe in the 1980s
 - (e) In the Netherlands and Korea, supply actually grew faster in the 1980s than the 1970s—and skill premia declined.

1.2 CLEARING SOME EMPIRICAL UNDERBRUSH [FOR SELF-STUDY: NOT COVERED IN CLASS]

Before formally developing the theory of skill premia, it is useful to develop a few tools that allow us to better describe the facts and dispense with some hypotheses. These tools are:

- The relationship between observed and unobserved skills and ‘between group’ and ‘within group’ (residual) inequality.

- Distinguishing age, time and cohort effects in skill quantities and prices.
- The relationship between permanent and transitory components of earnings inequality.

1.2.1 INEQUALITY: OBSERVED AND UNOBSERVED SKILLS.

It's helpful to start by asking what the relationship between overall and residual inequality 'should' be. The simplest model of residual inequality is a single index model, in which there is only one type of skill, and this skill is imperfectly approximated by education (or experience). This idea can also be stated as 'observed and unobserved skills are perfect substitutes.'

Say there are two skill levels, H and L . A college graduate has probability ϕ_c of being high skilled. A non-college worker has probability ϕ_n of being high skilled, with $\phi_c > \phi_n$. Suppose the skill premium is $\omega = w_H/w_L$. The observed college premium will be

$$\omega^c = \frac{w_C}{w_N} = \frac{\phi_c w_H + (1 - \phi_c)w_L}{\phi_n w_H + (1 - \phi_n)w_L} = \frac{\phi_c \omega + (1 - \phi_c)}{\phi_n \omega + (1 - \phi_n)}. \quad (1)$$

In this setup, there is both between and within group inequality. Between group inequality is the component explained by observables (education). Residual inequality is inequality among those with equivalent education. Total inequality is the sum of these two components. In this model, 'within group' inequality among those with equivalent education, equal to the ratio of high to low earnings within an education group, will equal ω . Between group inequality will equal ω^c .

A key observation from this model is that a change in the skill premium (reflected in a change in ω) will always raise or lower between, within, and total inequality together (provided that ϕ_H, ϕ_L are fixed; you can work out the derivatives).

This single-index model immediately makes a strong prediction: the timing of residual inequality growth should match the timing of overall and between group inequality growth. Between group inequality:

$$\frac{\partial \omega^c}{\partial \omega} = \frac{\phi_c - \phi_n}{(\phi_n (\omega - 1) + 1)^2} > 0. \quad (2)$$

Within group inequality: the relative wage of high skill to low skill college graduates and high to low skill non-college graduates will rise:

$$\frac{\partial [(w_C|H) / (w_C|L)]}{\partial \omega} = \frac{\partial [(w_N|H) / (w_N|L)]}{\partial \omega} = 1. \quad (3)$$

If this does not occur—that is, the time trends in between and within group inequality do not coincide—there are (at least) two potential explanations. One, we need more ‘skill indexes.’ Two, there have been changes in ϕ_H, ϕ_L such that they vary by cohort. If for example, in some cohorts the match between education and skill were near-perfect ($\phi_H \approx 1, \phi_L \approx 0$), these cohorts would have much less residual inequality. Many people posit in the U.S. that the quality of education has declined for recent cohorts. In this case, the correspondence between skill and education might have weakened, potentially changing the observed education premium ω^c and giving rise to more residual inequality.

These observations motivate two questions: 1) Are the growth of between and within group inequality contemporaneous, as the single index model predicts?; and 2) Are there ‘cohort effects’ in inequality? We’ll discuss both in turn.

1.2.2 COMPARING THE TIMING OF BETWEEN VERSUS WITHIN GROUP INEQUALITY: A DATA CONUNDRUM

Start with the simple regression model

$$\ln w_{it} = \alpha + \gamma_t S_i^* + \varepsilon_{it},$$

where S_i^* is the human capital of i . In this model, the variance of wages is composed of two terms

$$V(\ln w_t) = V(\gamma_t S^*) + V(\varepsilon_t),$$

where the first term is ‘between group’ inequality and the latter term is ‘within group’ or residual inequality.

In this model, a rise in the return to ability (that is a rise in γ_t), will raise between-group inequality but will not *not* raise residual inequality.

Now, assume that S^* is not observed. Instead we proxy for S_i^* using a measure of human capital such as schooling, S . Assume that $S_i - S_i^* = \nu_i$, where ν_i is an iid error term. If we estimate the model

$$\ln w_{it} = \alpha + \gamma_t S_i + \varepsilon_{it},$$

the variance of this expression will be

$$V(\ln w_t) = V(\gamma_t S) + V(\gamma_t \nu) + V(\varepsilon_t),$$

where the residual is now $V(\gamma_t\nu) + V(\varepsilon_t)$. Thus, if ability is imperfectly measured, a rise in the returns to ability will cause both between and within-group inequality to rise.

A rise in residual inequality corresponds to a rise in $V(e_t)$.

Under the single index assumption (as above), these two terms ought to move together. Did this occur? Unfortunately, there is not a universally accepted answer to this question, and the reason is that the key data sources do not agree. Figures 2a, 5a, and 6a and Tables 2 and 3 from Lemieux (2003) reveal the crux of the problem:

- There are two annual earnings series that are the source for much of what we know about the U.S. wage structure: the March Current Population Survey Annual Demographic File ('the March') and the May and (later) Monthly Outgoing Rotation Groups of the CPS. Although both are components of the Current Population Survey, they measure different earnings constructs. The March survey collects data on annual income whereas the May/MORG collects data on weekly or hourly income.
- During the 1970s, 1980s and 1990s, both surveys agree that between group inequality first fell slightly in the 1970s, rose rapidly in the 1980s, and rose much more slowly in the 1990s. See Figure 5 of Lemieux.
- However, the two surveys do not agree on trends in *overall* inequality in the 1970s. As Figure 2B of Lemieux shows, in the March CPS overall inequality either remained flat (women) or rose (men) in the 1970s. But in the May/MORG, overall inequality *fell* for both genders in the same decade.
- Given that both surveys agree on the trend in between group inequality in the 1970s, the implication is that they must disagree on residual (within-group) inequality. As Figure 6a of Lemieux shows, the March survey shows a rise in residual inequality during the 1970s while the May/MORG does not.
- The regressions in Table 3 test the 'single index' hypothesis for the consistency of within and between-group trends using both May/MORG and March files. As implied by the discussion above, the May/MORG accepts this hypothesis and the March data reject it.
- This leads to an unfortunate state of affairs: we cannot interpret the facts if we don't know what they are.

Lemieux’s paper makes the argument that one should prefer the May/MORG estimates to the March estimates since the March estimates are noisier. His paper does offer a convincing case that self-reported hourly or weekly wages from the May/MORG have higher precision than hourly wages calculated from the March files using annual earnings divided by annual hours. However, this observation it is not a particularly compelling argument for putting greater weight on trends in the May/MORG relative to March data. Why?

1. While May/MORG data appears less noisy than the March data in levels, this observation has no implications for the relative accuracy of trends in the two data sources—which is what the debate is about. There is no evidence suggesting that the March data suddenly (in the 1970s) got noisier or the May/MORG suddenly became less noisy. Hence, there is no reason to discount trends in one or the other.
2. Substantively, it’s possible that trends in both data reflect real phenomena (since the two surveys measure different earnings construct). It may be that inequality in annual incomes within-group earnings rose during the 1970s while inequality in hourly incomes did not. (Note that as far as we know, this is not simply due to greater variance of hours.) This is slightly ad hoc, but not more so than discarding the findings from the March data.
3. The Census files from 1970, 1980, and 1990 also appear to confirm the trends in the March data. Notably, the Census data also record annual earnings. This lends weight to the idea that the difference between the May/MORG and March series is substantive rather than mechanical.

Whatever conclusions one draws from this debate, it is indisputable that the oft-discussed trends in residual inequality are less robust across data sources than trends in between-group inequality. Hence, hypotheses for the growth of inequality that hinge critically on the timing of residual versus between-group inequality are also somewhat fragile.

1.2.3 COMPOSITION/COHORT EFFECTS IN INEQUALITY

We noted above that between or within group inequality could rise if there are changes in the skill level of college relative to high school graduates. This would be a *composition* effect as opposed to a *price* effect. In the extreme case, an increase in the ‘returns to education’ could

simply be a result of more able people going to college. To get a true measure of the return to education, we need to control for possible composition bias.

Consider a model with two education levels, high $h = 1$ and low $h = 0$. Suppose wages are given by

$$\ln w_{cit} = A_c^h + \gamma_t h_{ic} + e_{cit}, \quad (4)$$

where i indexes individuals, t indexes the year of the observation, and c denotes a cohort, i.e., a group of individuals who are born in the same year (or alternatively have entered the market in the same year), and $A_c^h \equiv E(a_{ic}|h_i = h)$, which is the average ability of members of education group h in cohort c . Here, the true return to education is γ_t , but unless we can control directly for A_c^h , we cannot measure this return unless $A_c^1 = A_c^0$, which seems unlikely. Note that this model is restrictive in that we assume that returns to skill are the same for all cohorts and ages; they only vary by time, γ_t . We'll relax this in a moment.

At the moment, the cohort specific education premium is:

$$\ln w_{ct} \equiv E(\ln w_{cit}|h_i = 1) - E(\ln w_{cit}|h_i = 0) = \gamma_t + A_c^1 - A_c^0, \quad (5)$$

under the additional assumption that ability (or education) is fixed for a cohort. In this simple model, the true 'skill price' γ_t cannot be identified unless we can directly observe A_c^1, A_c^0 . But, the change in the skill price can easily be measured:

$$\Delta \ln \omega_{c,t_1-t_0} \equiv \ln \omega_{ct_1} - \ln \omega_{ct_0} = \gamma_{t_1} - \gamma_{t_0}, \quad (6)$$

Hence, we estimate the true change in the return to education ($\Delta\gamma = \gamma_{t_1} - \gamma_{t_0}$) using within cohort, over time variation to difference out the cohort effect.

The assumption that returns to skills A_c^1, A_c^0 are constant over the lifetime of an individual may be too restrictive, however. There are quite different age-earnings profiles by education. Effectively, the model above allows for cohort effects (A_{1c}, A_{0c}) and time effects (γ_t) but no age effects. So, let's enrich the model slightly to allow for age effects:

$$\ln w_{cit} = A_{ic} + \delta_a + \gamma_t h_{ic} + e_{cit} \quad (7)$$

where the δ_a 's represent additive age effects. This model is now relatively general: earnings depend upon both cohort effects, age effects, and a time-varying return to education.

Unfortunately, we cannot estimate this model: cohorts are a linear combination of age and time; if you know someone’s age, and you know the current year t , you know their cohort. This means it’s not possible to separately identify age, cohort, and time effects.

But there is a workaround in some cases. Rewriting (6), we have:

$$\Delta \ln \omega_{c,t_1-t_0} \equiv \ln \omega_{ct_1} - \ln \omega_{ct_0} = \gamma_{t_1} - \gamma_{t_0} + \delta_{a_1=c+t_1} - \delta_{a_0=c+t_0}. \quad (8)$$

So now, the first-difference gives a combination of time effects in the return to education ($\gamma_1 - \gamma_0$), which is what we want to isolate, and age effects ($\delta_{a_1=c+t_1} - \delta_{a_0=c+t_0}$), which are a confounding source of variation.

Now consider a 2nd cohort c_1 that is age a_1 in year t_0 and age a in year t_{-1} . For this cohort, we estimate:

$$\Delta \ln \omega_{c_1,t_0-t_{-1}} \equiv \ln \omega_{c_1 t_0} - \ln \omega_{c_1 t_{-1}} = \gamma_{t_0} - \gamma_{t_{-1}} + \delta_{a_1} - \delta_{a_0} \quad (9)$$

Taking the double-difference, we obtain:

$$\begin{aligned} \Delta^2 \ln \omega &\equiv \Delta \ln \omega_{c,t_1-t_0} - \Delta \ln \omega_{c_1,t_0-t_{-1}} \\ &= (\gamma_{t_1} - \gamma_{t_0} + \delta_{a_1} - \delta_{a_0}) - (\gamma_{t_0} - \gamma_{t_{-1}} + \delta_{a_1} - \delta_{a_0}) = (\gamma_{t_1} - \gamma_{t_0}) - (\gamma_{t_0} - \gamma_{t_{-1}}) \end{aligned} \quad (10)$$

So, this expression lets us isolate the double-difference in the time effect between two intervals, that is, the *change* in inequality growth between two adjacent time periods. This is the approach taken by Juhn, Murphy and Pierce (1993) in Table 3 (a famous table). The clear finding from their analysis is that the growth in inequality does not appear to be a cohort-specific phenomenon. Instead, growth in overall and residual inequality appears to be ‘explained’ by a time effect, with the average change growing after 1970. Hence, their results seem to rule out an entire class of explanations for rising inequality based on changes in sorting/composition/cohort quality. (However, these conclusions here have not fully withstood the test of time. Later research has found that the return to education rose more for younger than older workers, suggesting a type of ‘cohort’ effect, which is studied by Card and Lemieux, 2001).

1.2.4 PERMANENT VERSUS TRANSITORY COMPONENTS OF INEQUALITY

One question that these cross-sectional analyses cannot answer is whether the observed growth of inequality is due to increased ‘churning’ in the earnings distribution—that is, the same people

have more variability in wages over time—or increased ‘spread’ among people residing at stable points in the distribution—that is, the gaps among neighbors just widen.

These two hypothesis have distinct welfare consequences. In theory, one could purchase insurance against transitory earnings shifts. And, on the positive side, a rise in transitory inequality implies increased earnings mobility: the chances that an individual will rise (or fall) in the earnings distribution improves when earnings fluctuate. On the other hand, increases in the spread of earnings without any increase in ‘churning’ implies greater permanent differences in earnings among individuals. Permanent wealth changes are typically not insurable and hence the welfare consequences of the latter scenario are generally worse.

To see this more formally, consider the earnings model

$$\ln w_{it} = \rho_t \alpha_i + \gamma_t e_{it}, \tag{11}$$

where α_i is an individual’s ‘permanent’ skill endowment, e_{it} is an iid, mean zero ‘skill’ or ‘luck’ shock that varies in each period, and ρ_t, γ_t are the time varying ‘prices’ associated with each of these components. Since $V(\ln w_{it}) = \rho_t^2 V(\alpha) + \gamma_t^2 V(e)$, it is clear that an increase in $V(\ln w_{it})$ can be caused by an increase in either price component ρ_t, γ_t . To tell these hypotheses apart, we need *time series data* on individuals’ earnings.

Define earnings mobility as

$$M_t = \frac{\gamma_t^2 V(e)}{\rho_t^2 V(\alpha) + \gamma_t^2 V(e)}, \tag{12}$$

which is simply the share of transitory variance in total variance. Notice that if the numerator and denominator rise proportionately—that is, the share of transitory variance in overall variance remain the same—then earnings inequality rises but earnings mobility is unchanged.

The question of permanent versus transitory inequality is studied by Gottschalk and Moffit (1994) who estimate a very simple model of earnings for 1970 to 1988: first, regressing out the age-experience profile, next doing a simple variance decomposition of the residuals into permanent and transitory components. The interesting result of their analysis is that permanent and transitory components have grown roughly proportionally—in other words, earnings mobility has been essentially unaffected by the growth of overall inequality.

This is neither especially good nor especially bad news. But it again rules out a class of explanations having to do with increased ‘chaos’ in the market of various forms. In fact,

there is a variety of evidence that would have led one to suspect that ‘churning’ was not the predominant explanation for rising inequality. Many studies that have looked for a decline in job stability in the U.S. labor market have found limited evidence of this (except for white collar managers). Nor is there much evidence of an increase in job reallocation.

1.2.5 TENTATIVE CONCLUSIONS

The evidence from study of observed/unobserved skills, cohort effects, and permanent/transitory earnings motivates the following tentative conclusions about the rise of earnings inequality in the United States:

1. The timing of the growth of residual and between group inequality may be distinct, at least in the U.S. This suggests that we should not think of ‘the rise in inequality’ as necessarily being a single phenomenon. However, this claim is challenged by the Lemieux (2003) paper and hence we should maintain disbelief on this point.
2. The rise in income inequality was not primarily a cohort specific phenomenon, meaning that it is unlikely to be due to differences in sorting or skill composition by different cohorts of labor market entrants. However, new evidence on this point is provided by Card and Lemieux (2001) in a paper that we will discuss in detail. (More precisely, C-L argue that there *are* cohort effects in the returns to education, and that these can be usefully interpreted in a supply-demand framework.)
3. The rise in inequality is not primarily driven by increased churning in the labor market. This last observation focuses attention on explanations affecting the prices (or returns to) skills rather than ‘instability’ per se. Unlike the first two tentative conclusions, this one is not currently in dispute.

To interpret these facts, we need a model of the determinants of skill premiums. The next section develops this model.

2 THEORY OF SKILL PREMIA

The simplest framework for interpreting skill premia (that is returns to schooling and other skills) starts with a competitive supply-demand framework. We begin with a simple closed

economy, aggregate production framework where factors are paid their marginal products and the economy operates on its supply and demand curves.

2.1 THE CONSTANT ELASTICITY OF SUBSTITUTION FRAMEWORK

2.1.1 THE AGGREGATE PRODUCTION FUNCTION AND THE ELASTICITY OF SUBSTITUTION

Begin with two types of workers, skilled and unskilled (or high and low education, or college and non-college, etc.), who are imperfect substitutes. Imperfect substitutability is crucial for understanding how relative prices affect skills. If workers were instead perfect substitutes, their wages would always move together up to a multiplicative constant (reflecting relative efficiency units)—so relative wages would not depend upon relative supplies, only on the return to the single factor of skill.

Suppose that there are $L(t)$ unskilled workers and $H(t)$ skilled workers supplying labor inelastically at time t . (It's a small matter to add elastic labor supply, but this would not change any conclusions and so we won't bother for now). The production function for the aggregate economy takes the constant elasticity of substitution (CES) form:

$$Y(t) = [(A_l(t)L(t))^\rho + (A_h(t)H(t))^\rho]^{1/\rho} \quad (13)$$

where $\rho \leq 1$. For now, we ignore capital and drop time subscripts where possible.

In this model, the elasticity of substitution between skilled and unskilled workers (that is, the percentage change in relative demand for low (high) skill workers for a percentage change in the relative price of high (low) skill workers) is given by

$$\sigma \equiv 1/(1 - \rho), \quad \rho \in (-\infty, 1)$$

Skilled and unskilled workers are 'gross substitutes' when the elasticity of substitution $\sigma > 1$ (or $\rho > 0$) and 'gross complements' when $\sigma < 1$ (or $\rho < 0$). When two productive inputs are gross substitutes, a reduction in supply of one creates added demand for the other. When these inputs are gross complements, a reduction in supply of one reduces demand for the other. [Hot dogs and buns are gross complements. Butter and margarine are gross substitutes.] Three special cases arises from this model:

1. $\sigma \rightarrow 0$ (or $\rho \rightarrow -\infty$). In this case, skilled and unskilled workers are Leontif, and output can only be produced using skilled and unskilled workers in fixed proportions. This is a case of 'perfect complements.'

2. $\sigma \rightarrow \infty$ (or $\rho \rightarrow 1$). Skilled and unskilled workers are perfect substitutes. Relative supplies of each do not affect relative wages. Changes in aggregate supplies will affect wages by affecting the price of skill overall. But the relative wage of skilled vs. unskilled (w_H/w_L) will be constant.
3. $\sigma \rightarrow 1$ (or $\rho \rightarrow 0$). The production function is Cobb Douglas, with fixed shares paid to each factor

In the CES framework, the value of σ plays a critical role because it determines how changes in either technology (given by the A 's) or supplies (L 's) affects demand and wages.

Note that in the variant of the CES function written in equation (13), there are no directly skill replacing technologies. Technologies in this equation are factor augmenting in that they augment the productivity of skilled or unskilled workers by raising A_l or A_h .

It is also possible to write a more general production function that has skill replacing technologies.

$$Y(t) = [(1 - b_t) [A_l(t)L(t) + B_l(t)]^\rho + b_t [A_h(t)H(t) + B_h(t)]^\rho]^{1/\rho}. \quad (14)$$

In this specification, B_l and B_h are directly-skill replacing technologies (they are perfect substitutes for the respective skill groups), while b_t corresponds to technology or organization factors that shift the distribution of tasks between skill groups (e.g., a new machine that ‘deskills’ a previously skilled task). The role of b_t in this expression is sometimes called ‘extensive’ technical change—a technology that shifts the allocation of tasks among factors—while the A_l , A_h terms are termed ‘intensive’ technical change, things that alter the productivity of factors.

If we were to drop $B_l(t)$, $B_h(t)$ from this specification, however, the addition of the extensive technical change term, b_t , would be redundant, though perhaps notionally useful. The reason is that one could always re-normalize A_l and A_h to $A'_l(t) = (1 - b_t) A_l(t)$ and $A'_h(t) = b_t A_h(t)$. So, the conceptual distinction between intensive and extensive technical change in the CES model is often not operationally useful.

The production function in (13) admits three interpretations:

1. There is only one good and skilled and unskilled workers are imperfect substitutes in its production.

2. The production function is equivalent to an economy where consumers have utility function $[Y_l^\rho + Y_h^\rho]^{1/\rho}$ defined over two goods. Good Y_h is produced with $Y_h = A_h H$ and good Y_l is produced with $Y_l = A_l L$ (hence, they have linear, single-factor technology). The parameter σ measures the elasticity of substitution between these goods in consumption.
3. A mixture of the two whereby two different sectors produce goods that are imperfect substitutes, and high and low education workers are employed in all sectors.

The 3rd possibility is the most realistic but the 1st is easiest to discuss and we'll use it for convenience.

2.1.2 WAGE SETTING

Given competitive labor markets, wages are set according to marginal products. The unskilled wage is given by

$$w_L = \frac{\partial Y}{\partial L} = A_l^\rho [A_l^\rho + A_h^\rho (H/L)^\rho]^{(1-\rho)/\rho} \quad (15)$$

and similarly

$$w_H = \frac{\partial Y}{\partial H} = A_h^\rho [A_h^\rho + A_l^\rho (H/L)^{-\rho}]^{(1-\rho)/\rho}. \quad (16)$$

Two important results follow from these equations.

1. First $\partial W_H / \partial (H/L) < 0$. The own labor demand curve is downward sloping.
2. Second $\partial W_L / \partial (H/L) > 0$. Everything else equal, as the fraction of skilled workers in the labor force increases, the wages of unskilled workers should increase. Hence, skilled and unskilled workers are 'Q-complements,' a greater quantity of the one increases the marginal product of the other. (This seems more natural if you think of the two inputs as capital and labor; more intensive use of capital raises the marginal productivity of labor and vice versa).

Combining these two equations, the skill premium is

$$\omega = \frac{w_H}{w_L} = \left(\frac{A_h}{A_l} \right)^\rho \left(\frac{H}{L} \right)^{-(1-\rho)} = \left(\frac{A_h}{A_l} \right)^{(\sigma-1)/\sigma} \left(\frac{H}{L} \right)^{-1/\sigma}, \quad (17)$$

which can be written more conveniently in logarithmic form:

$$\ln \omega = \left(\frac{\sigma - 1}{\sigma} \right) \ln \left(\frac{A_h}{A_l} \right) - \frac{1}{\sigma} \ln \left(\frac{H}{L} \right). \quad (18)$$

Notice that

$$\frac{\partial \ln \omega}{\partial \ln(H/L)} = -\frac{1}{\sigma} < 0,$$

the *relative* demand curve for high versus low skilled workers is downward sloping (recall that $\sigma \geq 0$). That is for given ‘skill bias,’ A_h/A_l , an increase in relative supplies H/L lowers relative wages with elasticity σ .

You can think of this substitution as occurring through two channels. If all workers are producing the same good, then an increase in the relative supply of high skilled worker will cause firms to reassign some ‘tasks’ performed by low skilled workers to high skilled, thereby lowering the marginal productivity and hence the wages of high skilled. If they are instead producing different goods, then output of the high-skilled good will rise, increasing consumption of this good but lowering consumers’ marginal utility of consuming it and hence its price (due to diminishing marginal rate of substitution).

2.1.3 RELATIVE SUPPLY OF SKILLS AND THE ELASTICITY OF SUBSTITUTION

Although σ is the crucial parameter of this model, it is difficult to know what its value is in reality since it combines substitution in production and consumption across consumers, across industries etc. It’s important to stress that this aggregate production function is an abstraction and is *not* intended to correspond to the production function of any given firm. We would generally expect factors to be less substitutable at the firm level than at the aggregate level. [For example, you would expect the ability of a single firm to substitute among skill groups to be lower (less elastic) than for a group of firms or for an industry. Concretely, imagine that an auto-maker has two auto plants producing identical cars, one that uses a relative skill intensive technology and the other using an unskilled-intensive technology. The technology is fixed in the short run so that neither plant can adjust its factor input ratios. In response to a decline in the price of skill, the manufacturer can shift production towards the skill-intensive plant. In the long run, it can close the unskilled-intensive plant altogether.] The 2005 *QJE* paper by Chad Jones presents an intriguing framework for thinking about the properties of aggregate production functions.

Given this uncertainty, there is a surprising consensus across estimates for the U.S. that $\sigma \approx 2$, with the most commonly used estimate of $\sigma = 1.4$. Also, Angrist’s (1995) *AER* paper on the demand for Palestinian labor, which uses a nice natural experimental design, finds an

implied elasticity of substitution between Palestinians of 16 years of schooling and those with less than 12 of approximately $\sigma = 2$.

2.1.4 TECHNICAL CHANGE AND THE SKILL PREMIUM

In this model, the A 's are so called efficiency parameters, and a rise in either is referred to as factor augmenting technical change. How does the skill premium respond to a shift in A_h/A_l ? The result depends upon the elasticity of substitution. From equation (18):

$$\frac{\partial \ln \omega}{\partial \ln(A_h/A_l)} = \frac{\sigma - 1}{\sigma}, \quad (19)$$

the sign of which depends upon $\sigma \gtrless 1$. Many people find this result counter-intuitive. Concretely, why would an increase in the productivity of more skilled workers, that is a rise in A_h/A_l , cause their wages to fall (when $\sigma < 1$)? An intuitive way to see this is to consider a Leontif production function where high and low skilled workers are used in constant proportions in a competitive market. An increase in the supply of high skilled workers in this setting effectively creates “excess supply” for a given number of unskilled workers. The extra skilled workers will either bid down wages of other high skilled workers or will become unemployed (lowering average wages for skilled workers if zeros are counted). Since the broad consensus is that $\sigma > 1$, this case is not considered likely.

Average wages in the economy will be given by:

$$\bar{w} = \frac{LW_L + HW_H}{L + H} = \frac{[(A_l L)^\rho + (A_h H)^\rho]^{1/\rho}}{1 + H/L}, \quad (20)$$

which is also increasing in H/L provided the skill premium is positive ($\omega > 1$ or $A_h^\rho(H/L)^\rho - A_l^\rho > 0$). Hence, when the skill composition of the labor force rises, wages increase. (When the wages of skilled are below unskilled, this effectively implies that unskilled are scarcer than skilled; additional skilled workers effectively lower the skill composition of the labor force).

Another important observation is that if A_h or A_l rises with $\sigma > 1$, wages should rise for all workers, both skilled and unskilled (though inequality may increase). This comes from the fact that skilled and unskilled labor are q-complements in the production function (think capital and labor). An increase in the productivity of the one is like an increase in the effective supply (or ‘intensity’) of that factor, which then boosts the marginal productivity of the complementary factor.

Factor augmenting technical change always raises societal wealth since we can get more output for a given set of inputs. This observation is important to bear in mind since male wages below the median fell substantially in real terms in the U.S. during the 1980s (though not in other countries). This suggests that this model will be unable to explain falling real wages.

Note that there are other forms of technical change that can directly lower absolute wages. For example, imagine that there is a machine that is invented that is perfectly substitutable for high skill workers. In this case, the wages paid to high skilled workers cannot exceed the rental price (per efficiency unit) of the machine, and declines in the price of the machine (or increases in its efficiency) will lower the price of skilled workers. This is a case of a skill replacing technical change (see equation (14)). Autor, Levy, and Murnane (2003) present a simple model of skill-replacing technical change that we will discuss.

2.1.5 SUMMARY

In response to an increase in H/L (and assuming the skill premium is positive):

1. The skill premium $\omega = W_H/W_L$ falls.
2. Wages of unskilled workers rise.
3. Wages of skilled workers decrease.
4. Average wages rise.

These results can readily be generalized to a case with capital, i.e., $F(A_l L, A_h H, K)$, and they will generally go through.

2.1.6 THE LONG TERM SKILL BIAS OF TECHNICAL CHANGE

The key result from the above is that as H/L increases, ω falls. But as we know, in every advanced country the supply of educated workers has risen dramatically in the past 6 decades but relative wages of better educated workers have remained consistently above those of less educated, though the degree to which they have done so has varied by decade. So, in the U.S., the college educated share rose from 6.4 to 29.7 percent of the workforce from 1940 to 2000, whereas the those with less than 12th grade declined from 68 to 9 percent of the workforce.

Yet, the skill premium in 2000 (measured in a variety of ways) was at or above that of in 1940 (though not above that in 1910; see the 1999 paper by Katz and Goldin on Returns to Skill in the United States across the Twentieth Century). Hence, the relative demand for skilled workers *must* have risen practically everywhere. This is not a surprising conclusion: it is hard to think of a modern economy that isn't utterly reliant on the literacy and numeracy of its workforce. The contemporary structure of production could clearly not have come into being and could not be sustained without a massive accumulation of human capital (i.e., an educated workforce). But this observation—that skill demands have been rising secularly for decades and hence skill-biased technical change is not exclusively a recent phenomenon—is often overlooked in otherwise intelligent discussions of SBTC.

The pattern of generally rising returns to education across the developed world does not imply that either (1) technical change has always outstripped the growth in supply of educated workers or; (2) that the rate of technical change is constant. Tinbergen advanced a hypothesis in 1975 that is useful for thinking about the demand for skills throughout this century:

“The two preponderant forces at work are *technological development*, which made for a relative increase in demand and hence in the income ratio... and *increased access to schooling*, which made for a relative decrease.”

Hence, a useful framework for thinking about the evolution of inequality—or at least the return to skill—is:

1. Long term trend increases towards greater relative demand *and* greater supply of skilled workers

...and...

1. Bursts of supply and/or technologically-induced demand accelerations/decelerations that cause demand to temporarily move out more rapidly than supply or vice versa in some eras.

Under the ‘education race’ view of Tinbergen, skill returns will rise when the rate of technological development outpaces the production of new human capital (that is, the growth in

education of the workforce) and v.v. when educational production outpaces technological advances.

An important caveat to the Tinbergen view is that it appears to take as given that technical change is always and everywhere skill-biased. Most economic historians would dispute this view. Some of the great technological innovations of the nineteenth century, in particular the “factory system,” which gave rise to mass production and the interchangeable parts revolution—were probably unskill-biased. These technologies replaced the work of skilled artisans (metal-smiths, carpenters, weavers, etc.) with capital and low-skilled labor. The Luddite rebellion of the 19th century is a case in point. The Luddites were skilled weavers who rebelled against capitalists by destroying power-looms because they feared that this machinery would devalue their skill—allowing unskilled workers to accomplish the artisanal tasks that their livelihoods depended upon. Their fears were justified; automation substantially devalued their skills. The 1998 paper by Katz and Goldin on “The Origins of Capital Skill Complementarity,” discusses when and why mass-production technology, which was initially unskill-biased, became skill-biased. Unfortunately, their paper doesn’t have any direct evidence on unskill-biasedness (only skill-biasedness).

2.2 BRINGING THE CES MODEL TO THE DATA

From (17), the relative productivity of skilled workers is given by $(A_h/A_l)^{(\sigma-1)/\sigma}$, and all long term evidence implies it must have increased considerably since the first time period in which we have consistent measures, which is 1939 (from the 1940 U.S. Census). How much has A_h/A_l increased? Deducing the answer to this question implicitly depends upon knowing σ . If we ‘know,’ assume, or estimate a value of σ , we can back out the implied change in relative demand for high versus low-skilled workers. This approach was popularized by Katz and Murphy (1992).

2.2.1 THE ‘KATZ-MURPHY’ MODEL

Recall from (18) that

$$\ln \omega = \frac{\sigma - 1}{\sigma} \ln \left(\frac{A_h}{A_l} \right) - \frac{1}{\sigma} \ln \left(\frac{H}{L} \right). \quad (21)$$

Let’s say we wanted to estimate this model using time series data. We need to add time subscripts to everything in this equation (save for σ , which we assume to be fixed), so that the A ’s vary by year, as do the supplies of skilled and unskilled workers and the skilled/unskilled

wage ratio. Of course, we observe supplies and we can estimate the wage premium. Hence, the unknowns are σ and the A_h/A_L . Our hypothesis is that $\partial \ln(A_h/A_L)/\partial t > 0$, the relative productivity of skilled workers is rising over time. So we can estimate this model as:

$$\ln \omega_t = \gamma_0 + \gamma_1 t + \gamma_2 \ln(H/L) + e_t, \quad (22)$$

where t is a linear function of time. In estimating this equation, γ_0 is a constant, γ_1 gives the time trend on $(\frac{\sigma-1}{\sigma}) \ln(A_{ht}/A_{Lt})$, and $\hat{\gamma}_2$ is an estimate of $1/\sigma$. Using CPS data from for 1963 - 1987, Katz and Murphy's estimate this model in a simple OLS regression:

$$\ln \omega = \begin{matrix} 0.033 \cdot t \\ (0.01) \end{matrix} \quad \begin{matrix} -0.71 \cdot \ln\left(\frac{H}{L}\right) \\ (0.15) \end{matrix} + \text{constant} \quad (23)$$

This estimate suggests two things: 1) there has been a trend increase in the relative demand for skilled workers; 2) the elasticity of substitution between them $\hat{\sigma} = -1/0.709 = 1.41$. The K-M regression must be treated with care: there are only 25 data points, and they are highly serially correlated. But this model appears surprisingly informative.

The operation of this simple model can be seen in Figure IV of the Katz-Murphy paper. Several observations:

1. The skilled wage differential increases extremely rapidly after 1979. (Panel A)
2. This jump coincides with a very rapid deceleration in the trend growth rate of college educated workers (Panel B). As many have noted (and Card and Lemieux have written), part of this decline appears due to the end of the Vietnam war.
3. The model with $\sigma = 1.4$ fits the data well—except for the period from 1975 to 1981 when their model suggests that inequality should have begun to rise. This did not occur until the start of the very deep U.S. recession in 1980. (Panel C) (This unexpected drop in inequality—or more accurately, its failure to rise—during 1975 - 1981 is also visible in the UK, as we'll see later.)

Another interesting exercise is seen in Panel D of the figure. If we assume different values of σ , we get different conclusions about the behavior of relative demand which Katz-Murphy index in Panel D as $(\sigma - 1) \ln(A_h/A_l)$. For higher values of σ , there appears to be a rapid acceleration in relative demand in the 1980s. The reason for this inference is that the higher

the elasticity of substitution between H/L , the greater the demand shift required to induce a rise in relative wages of given magnitude. We know relative wages rose considerably in the 1980s. If we believe that factors are highly substitutable, this requires a dramatic demand shift to rationalize.

Hence, a key conclusion of Katz-Murphy is that fluctuations in relative supply overlaid on smoothly rising demand might be sufficient to explain trends in relative wages in the U.S. for 1967 to 1987 (though they are agnostic on this point). Is this the end of the story? Not necessarily.

See also the updated estimation of the Katz-Murphy model from the Autor-Katz-Kearney 2008 ‘Revisionists’ paper. Projecting the Katz-Murphy estimates forward to 2005 and using the observed changes in skill supplies in each year, AKK show that the Katz-Murphy model continues to fit the aggregate data extremely well to 1992, which is five years beyond the data available to K-M at the time of their writing. But the model goes somewhat awry after that. In particular, it predicts a substantially greater increase in the skill premium between 1992 and 2005 than is observed in the data. Assuming σ is constant, the K-M model therefore implies that demand growth *decelerates* (but does not plateau or reverse) after 1992. We’ll have more to say on this topic later in the semester.

2.3 LONGER-TERM EVIDENCE

Katz and Murphy showed that one could explain many of the patterns in the data with a very simple steady demand hypothesis. Is there anything more than fluctuations in supply driving patterns of wage inequality? To get some evidence on this point, one needs a longer time series. This is what Autor, Katz, Krueger (1998) provide. Define the demand index as:

$$D_t = (\sigma - 1) \ln(A_H/A_L) = \ln(w_H H/w_L L) + (\sigma - 1) \ln(w_H/w_L). \quad (24)$$

We’d like to ask whether there has been an acceleration in the rate of change in D_t , in other words is $\Delta D_t^{70-99} > \Delta D_t^{40-69}$. To perform this test, we need:

1. A consistent series for wages and employment.
2. Estimates of wagebill shares ($w_H H, w_L L$)
3. Estimates of w_H/w_L

One concern in implementing this framework is that the wagebill shares may confound changes in prices (w_H/w_L) with changes in the composition labor if the quality of high and low skill workers varies with time. Here's a simple fix. Note that

$$\Delta \ln(\text{relative wagebill}) = \Delta \ln(w_H H/w_L L) = \Delta \ln(H/L) + \Delta \ln(w_H/w_L) \quad (25)$$

Hence, there is both a supply and a price component to the change in the wagebill. We want to isolate the supply component, which we can do by backing it out from the total observed change in the relative wage bill:

$$\Delta \ln(\text{relative supply}) = \Delta \ln(\text{relative wagebill}) - \Delta \ln(\text{relative wage}) \quad (26)$$

$$= \Delta \ln(w_H H/w_L L) - \Delta \ln(w_H/w_L). \quad (27)$$

This procedure effectively ‘subtracts off’ the component of wagebill share change due to pure price changes (estimated from a regression) and hence calculates the effective supply change as a residual.

Table II of AKK (updated to 2000 here) makes this set of calculations:

1. There is evidence of growing relative demand for skilled workers in every decade except the 1940s (when it is believed war-era industrialization dramatically raised relative demand for less-skilled workers).
2. There is clear evidence that net demand changes were larger over 1970 - 2000 (i.e., the most recent three decades relative to the prior three). This is ‘gross’ evidence of acceleration, but note that it does depend heavily on including the 1940s.
3. Whether demand accelerated in the 1970s or 1980s relative to the prior decades depends sensitively on the assumed elasticity. For higher elasticities, demand appears to have accelerated in the 1980s. For lower elasticities, it accelerated in the 1970s. This is itself an interesting finding: a demand acceleration in the 1970s may have been masked by a simultaneous supply acceleration. If this inference is correct, inequality would have grown in the 1970s had supply not suddenly jumped.
4. As per the figure from Autor-Katz-Kearney, there appears to be some demand *deceleration* in the 1990s. This is noteworthy and probably unexpected for many versions of SBTC (though not all).

So, the key fact that we have so far deduced is the necessity of demand shifts, with some evidence of their acceleration in the 1970s or 1980s and some evidence of deceleration in the 1990s.

Katz and Goldin (2007) present a more up to date version of this analysis, using data from 1915 to 2005. A surprising result of their analysis is just how far one can get by fitting this two-factor model to the data for a very extended time period and assuming an acceleration from 1949 forward and a deceleration after 1992.

3 INTERNATIONAL EVIDENCE ON WAGE INEQUALITY AND THE SUPPLY AND DEMAND OF SKILLS

The U.S. has notably higher wage inequality and returns to skills (not necessarily one and the same) than do almost all other advanced countries. See, for example, Figures 2 and 3 of the 2005 *ReStat* article by Blau and Kahn. Two major hypotheses have been advanced to explain these cross-country differences. One is that the effective ratio of relative skill demand to relative skill supply is higher in the U.S. than in most other countries. Stated differently: the U.S. has an abundant supply of low-skilled workers relative to other advanced nations. A second explanation attributes international differences in wage inequality across skill groups to differences in labour market institutions. In this view high minimum wages, employment protection and labor unions are responsible for the relatively higher wages of low skilled workers in continental Europe. (A number of papers by Emmanuel Saez and co-authors attribute cross-country inequality differences to ‘social norms.’ This is an interesting idea but the authors do not explain how this view can be tested.)

A notable entry in this debate is the 1996 paper by Blau and Kahn in the *JPE*. B&K found that labor market institutions (in particular, labor unions and wage centralization) are much better predictors of cross-country wage inequality than are supply and demand indices constructed using education (for supply) and industrial and occupational composition (for demand). In fact, they find *no evidence* that supply and demand can explain any of the cross-country differences in the relative earnings of the less-skilled. B&K’s rejection of the supply and demand hypothesis was sufficiently spectacular that the *JPE* saw fit to publish it, in spite of (perhaps because of) the journal’s “Chicago view” of the world. I will not devote class time to the B&K paper, but it is worth reading.

The 2004 *Economic Journal* paper by Leuven, Oosterbeek and van Ophem (‘LOvO’) rejoins this debate with the aid of much better data and perhaps a better-formulated supply and demand framework. They draw on the International Adult Literacy Survey (IALS), which was designed to provide internationally comparable measures of cognitive skills across 20 advanced countries. (There are a number of other papers that use these data for related exercises, including Blau and Kahn, in *ReStat* and Freeman and Devroye 2001 (NBER WP). The LOvO paper is, in my view, the most persuasive.) LOvO use these data to ask how well cross-country differences in supply and demand can explain cross-country differences in skill differentials.

Figure 1 of LOvO demonstrates the surprisingly weak link between average years of schooling and average IALS score across countries. For example, the U.S. has by far the highest average years of completed schooling of the 15 countries included in their sample yet ranks only 9 of 15 on the average test score. This suggests that using years of schooling to compare cross-country skill levels could generate misleading inferences (if wages primarily reflect cognitive ability rather than years of schooling).

The next step to the analysis applies an internationally comparable demand and supply index to the data to ask if wage differentials are relatively lower where the ratio of supply to demand is relatively greater. There is no perfect way to do this. LOvO use a modified Blau and Kahn (1996) procedure.

They choose a baseline country b , and group workers in *all* countries into three skill groups $k = \{\text{low, medium, high}\}$ using as cutpoints the values in country b that break the skill distribution (proxied by IALS scores) into three even parts. For each country $j \neq b$, LOvO form a relative skill supply index of:

$$s_{kj} = \ln(E_{kj}/E_{kb}),$$

where E_{kj}, E_{kb} are the shares of total labor input supplied by skill group k in countries j and b respectively (the latter being equal to $\frac{1}{3}$ by construction).

To form a relative demand index for each skill group k , they use a rough ‘skill requirements’ index. Skill input in the base country b is measured as the share of each skill group k employed in industry-occupation cells o , c_{ok} . LOvO form a relative demand index for other countries by contrasting the employment shares in industry-occupation cells o in countries j to the skill

input in country b . Specifically, the demand index is:

$$d_{kj} = \ln \left(1 + \sum_0 c_{ok} \frac{\Delta E_{oj}}{E_{kb}} \right),$$

where ΔE_{oj} is the country j minus country b difference in employment shares in industry-occupation o , and E_{kb} is employment share of skill group k in country b (again equal to $\frac{1}{3}$). If country j is relatively concentrated in ind-occs in which country b has relatively high intensity of skill group k (relative to its own endowment of k), country j will be said to have relatively high demand for skill group k . The results of this exercise could depend heavily on the choice of the baseline country b . Hence, LOvO do the analysis 15 times, once using each country as a baseline.

Define w_{kj} as the mean relative wage of skill group k relative to a base skill group in country j . LOvO estimate the following model for relative wages:

$$(w_{kj} - w_{kb}) = \alpha + \beta [(s_{kj} - d_{kj}) - (s_{kb} - d_{kb})] + \varepsilon_j,$$

The key prediction from the supply-demand framework is that $\beta < 0$, relative prices and relative net supplies negatively covary. This is pretty much what LOvO find (Figure 3 and Table 4). Perhaps most importantly, the explanatory power of the model appears best for the wages of low-skilled workers. Interestingly, when LOvO use the same empirical tools as Blau and Kahn 1996—in particular, forming skill groups using education rather than IALS scores—they find very weak (and insignificant) evidence that supply-demand is an important explanation for differences in skill differentials across countries (see Figure 2 and Table 4).

This article makes a valuable contribution to the cross-country inequality debate, which has been dominated by the view that institutional differences are the predominant factor explaining cross-country wage structure differences. It must be stressed, however, that: 1) LOvO do not attempt to assess the role of institutional differences, which might also appear important in their data; 2) the LOvO analysis is only cross-sectional, presumably because there has only been one round of the IALS study. It would be valuable to be able to use equally good data to ask whether international differences in skill demand and supply could also explain the differential changes in wage structure across the U.S., U.K., and Europe over the last three decades. One suspects that a cross-country panel analysis using such data would not provide nearly as clear-cut conclusions. (In addition, it's unclear how the re-normalization of the supply-measures by

skill and employment levels in 15 separate countries affects the veracity of the conclusions. It's possible that this produces misleadingly narrow confidence intervals.)

A related 2003 paper by Acemoglu in the *Economic Journal* poses an interesting but as yet untested hypothesis for *why* inequality rose so much more in the U.S. and U.K. than continental Europe. (Acemoglu takes as a starting point that differences in skill supplies do not entirely explain these differences.) His hypothesis is that institutional factors that compress wages (and in particular, prevent low skill workers from receiving very low wages) spur firms to endogenously adopt technologies that raise the productivity of the low-skilled. The mechanism here is similar in spirit to the Acemoglu and Pischke papers (*QJE 1998* and *JPE 1999*), who consider why wage compression may lead firms to invest in workers' general skills training: by raising worker productivity above an artificially imposed floor, investments in workers' productivity allow firms to capture the difference between workers' value marginal product and their outside wage, which is pinned down by exogenous institutional forces. The Acemoglu hypothesis awaits empirical testing for labor markets.

4 KEY HYPOTHESES FOR RISING INEQUALITY

After this quick overview of simple supply-demand stories, we are now ready to think more rigorously about the major explanations for rising inequality:

1. Supply and demand shifts

- (a) **Steady demand.** One possibility motivated by the Tinbergen framework (and partly affirmed for the 1960s-1980s by Katz-Murphy) is that for unspecified (and perhaps exogenous) reasons, there has been a steady rise in demand for skills throughout the century. Hence, movements in the wage premium reflect changes in the trend growth of supply—when supply lags demand, the premium rises (and vice versa). This is a rather pedestrian explanation but certainly a plausible one, and it may explain much of the data if not all of it. Richer versions of the K-M model (e.g., Card and Lemieux 2001) may do a better job of explaining the data using this hypothesis than does the original K-M paper. We'll spend a lecture on the Card-Lemieux paper because it is an unusually rich, sophisticated and successful application of 'structural modeling' to the estimation of skill premia.

- (b) **Accelerating demand.** This hypothesis posits a discontinuous increase in the trend rate of demand growth, perhaps occurring in the 1970s or 1980s, that, *coupled with the slowdown in supply*, caused inequality to rise. This is also a reasonable hypothesis *a priori*: why should the rate of movement of the relative demand curve be steady across periods? What gives this hypothesis added plausibility to many economists is the coincidence of the ‘computer revolution’ with the rise in inequality in advanced countries. Generally, hypotheses in this vein (accelerating demand) will specify a variety of reasons why demand will have accelerated, and provide evidence for these causes. We’ll look at a number of these. All of the key SBTC hypotheses fall into this category.
2. **Changes in the organization of production.** Technical change is often conceived of as improvements in capital. But changes in work organization (such as the factory system) can potentially effect skill demand even without a corresponding advance in physical capital (though some types of capital and organizational structures may be complementary). A number of papers present theory and some evidence for this type of organizational change story. These include Acemoglu 1999, Beaudry and Green 2003, Bartel, Ichniowski and Shaw 2004, Caroli and Van Reenan 2001, Bresnahan, Brynjolfsson and Hitt 2002, and Autor, Levy and Murnane 2003. We will discuss two of these papers (Autor-Levy-Murnane and Beuadry-Green) in detail.
3. **Market structure and returns to talent.** The 1981 paper by Sherwin Rosen on “The Economics of Superstars” is often cited as prescient harbinger of the rise in returns to skills experienced by many developed economies in the subsequent decades. In fact, the superstars explanation probably has little to do with it. However, the paper offers a fascinating insight that has considerable currency with many economists as an explanation for why wages of CEOs, entertainers and athletes are incomparably higher than for other occupations. The ingenious working paper by MIT graduate (currently at Haas) Markö Tervio offers a far less benign explanation for the same phenomena. Time-permitting, we will briefly discuss both papers.
4. **International trade.** There was substantial growth in world trade flows in the United States, especially in the United States. (However, the most rapid growth is during the

1970s, not the 1980s.) Trade between countries with different factor endowments will change relative prices and will therefore raise or lower inequality among owners of those factors (depending on whether your country has more of the relatively scarce or abundant factor after trade opening). This hypothesis has numerous testable implications that we'll look at. Here, you'll learn just enough trade theory to be dangerous (to yourself only).

5. **International outsourcing.** This is subtly different from trade. Rather than opening factor markets to trade, you simply purchase certain factor-intensive inputs from overseas and turn them into final products in your own country. Observationally, this can look a lot like SBTC. Feenstra and Hanson [2001], on the syllabus, provide theory and evidence on this hypothesis. I won't spend class time on this topic because I'm not convinced that it's genuinely important, yet...
6. **Institutional changes.** Declining union penetration and falling minimum wages are a major feature of the U.S. and U.K. labor markets during the period of inequality growth. In other countries, these institutional changes have been far more moderate. A number of authors have argued that these institutional changes explain the observed changes in wage setting rather than the forces of supply and demand. The debate about the role of the U.S. minimum wage has been the most heated, and we'll spend some time on the leading papers on this topic.
7. **Labor force composition.** Although we often think of prices and quantities only interacting through supply and demand, an interesting debate has opened on the role of prices versus quantities in explaining the growth of residual earnings inequality. The question is whether the dispersion of wages can in part be explained by the dispersion of skill characteristics of the labor force. Juhn, Murphy and Pierce (1993) posed this question and empirically dismissed it—that is, they concluded that dispersion of quantities was not an important factor in explaining the growing dispersion of wages. Lemieux (2006 *AER*) revisits this topic and reaches a dramatically different conclusion. Autor-Katz-Kearney (2005 NBER 11628) conclude something different again. We'll spend some time on this debate because it has fostered the development of some interesting tools for empirical work, particularly the use of kernel reweighting (popularized by DiNardo, Fortin and Lemieux 1996) and quantile regression (used by Autor-Katz-Kearney paper).