14.662 — Spring 2018

Lecture Note: The Canonical Model of Skill Differentials

David Autor, MIT and NBER

February 10, 2018

1 Theory of skill premia

The simplest framework for interpreting skill premia (e.g. returns to schooling or other skills) starts with a competitive supply-demand framework in a simple closed economy setting, where factors are paid their marginal products and the economy operates on its supply and demand curves. The model is a 'workhorse' in several respects: it's as common as livestock; it can carry a big load; and it's often a bit tired and overburdened (much like graduate students). We will enrich this workhorse model in subsequent lectures.

1.1 The constant elasticity of substitution framework

The aggregate production function and the elasticity of substitution

We begin with two types of workers, skilled and unskilled (or high and low education, college and non-college, etc.), who are imperfect substitutes. Imperfect substitutability is crucial for understanding how relative prices affect skills. If workers were perfect substitutes, their wages would always move together (up to a multiplicative constant that reflects relative efficiency units). *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, and this would not change any conclusions). The production function for the aggregate economy takes the constant elasticity of substitution (CES) form:

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

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

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

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

(To derive this, solve for $w_h \equiv \frac{\partial Y}{\partial H}$ and $w_l = \frac{\partial Y}{\partial L}$, divide, reorganize, take logs, and take a derivative to find $\frac{dln(H/L)}{dln(w_l/w_h)}$.) 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—in multiple senses of the word gross.]

Three special cases arise from this model:

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

2. $\sigma \to \infty$ (or $\rho \to 1$). 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. $\sigma \to 1$ (or $\rho \to 0$). 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 (1), 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.

$$Y(t) = \left[(1-b(t))\left[A_l(t)L(t)+B_l(t)\right]^{\rho} + b(t)\left[A_h(t)H(t)+B_h(t)\right]^{\rho}\right]^{1/\rho}. \tag{2}$$

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. (Note that if $\sigma \to 1$, the $b(t)$ terms limit to the exponents in the Cobb-Douglas production function.)

For now, we focus on the production function in (1). This production function 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 third possibility is certainly the most realistic (or least *unrealistic*) but the first is easiest to discuss and we'll use it for convenience.

Wage setting

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

$$w_L = \frac{\partial Y}{\partial L} = A_l^\rho [A_l^\rho + A_h^\rho (H/L)^\rho]^{(1-\rho)/\rho} \tag{3}$$

and similarly

$$w_H = \frac{\partial Y}{\partial H} = A_h^\rho [A_h^\rho + A_l^\rho (H/L)^{-\rho}]^{(1-\rho)/\rho}. \tag{4}$$

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. In other words, 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 π is

$$\pi = \frac{w_H}{w_L} = \left(\frac{A_h}{A_l}\right)^\rho \left(\frac{H}{L}\right)^{-(1-\rho)} = \left(\frac{A_h}{A_l}\right)^{(\sigma-1)/\sigma} \left(\frac{H}{L}\right)^{-1/\sigma}, \tag{5}$$

which can be written more conveniently in logarithmic form:

$$\ln \pi = \left(\frac{\sigma - 1}{\sigma}\right) \ln\left(\frac{A_h}{A_l}\right) - \frac{1}{\sigma} \ln\left(\frac{H}{L}\right). \tag{6}$$

Notice that

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

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 workers, thereby lowering the marginal productivity of high skilled workers and hence their relative wage. 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).

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.[1] The 2005 *QJE* paper by Chad Jones ("the shape of production functions and the direction of technical change") presents a statistical framework for thinking about the properties of aggregate production functions.

Given this uncertainty about σ, 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 $\sigma = 2$.

[1] Concretely, imagine that an auto-maker has two auto plants producing identical cars, one that uses a relatively 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.

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 (6):

$$\frac{\partial \ln \omega}{\partial \ln(A_h/A_l)} = \frac{\sigma - 1}{\sigma}, \tag{7}$$

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 Leontief 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 excess 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 generally thought to be unlikely.

Average wages in the economy will be given by:

$$\overline{w} = \frac{LW_L + HW_H}{L + H} = \frac{[(A_l)^\rho + (A_h H/L)^\rho]^{1/\rho}}{1 + H/L}, \tag{8}$$

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).

Before you get too excited about the beauty of the factor-augmenting CES, we should highlight a property of the model conflicts with an important empirical trend. Note 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. An increase in the productivity of the one is equivalent to an increase in the effective supply (or 'intensity') of that factor, which then boosts the marginal productivity of the other 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 wages of non-college men fell substantially in real terms in the U.S. during the 1980s (and in some other industrialized economies in the 1990s and 2000s). It is difficult for a model of purely factor-augmenting technical change to explain falling wages for a skill group that has declining relative supply. Specifically, we know that non-college workers have become relatively scarce over recent decades. While relative demand for high-skill workers has likely risen (that is A_h/A_l has increased), this should also raise the productivity of L workers through q-complementarity—unless A_l falls even as A_h rises, which is generally a difficult idea to rationalize (i.e., why should there be technological regress?). Thus, in a purely factor augmenting world (and assuming that wages are set competitively), it will be hard to explain the falling wage levels of non-college workers in the U.S.

Other forms of technical change *can* directly lower absolute wages, however. For example, imagine that a machine 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 a case of a skill replacing technical change (see equation (2)). Autor, Levy, and Murnane (2003) present a simple model of skill-replacing technical change that we will discuss. The Acemoglu-Autor (2011) model provides a more flexible and elaborate development of this set of ideas. So while CES is an elegant starting point, it will require some modifications to accurately capture some important trends. We will cover these after we have mastered the mechanics of CES.

Summary of key relationships

In response to an increase in H/L

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 provided the skill premium is positive ($\omega > 1$ or $A_h^\rho (H/L)^\rho - A_l^\rho > 0$).

In response to an increase in A_h, holding A_l and L/H constant

1. $\omega = W_H/W_L$ rises if $\sigma > 1$, falls if $\sigma < 1$, and is unchanged if $\sigma = 1$.
2. Wages of low skill workers rise if $\sigma < \infty$. This follows from q-complementarity, which does not apply if H and L are perfect substitutes.
3. Average wages rise if $\sigma > 0$. (The only case in which a rise in A_l wouldn't increase average wages is if skill groups are perfect complements, $\sigma = 0$, and these groups are already available in the correct Leontief proportions. In that case, a rise in A_h simply creates a surplus of H.)

4. Both W_H and W_L rise if $\sigma \geq 1$.

Note that these results can readily be generalized to a case with capital, i.e., $F(A_l L, A_h H, K)$. But depending on the assumptions on capital supply and capital-skill complementarity, the specific predictions may differ. In a world where capital is equally complementary to both (or all) skill groups, capital supply has no bearing on the *relative* wages of skill groups.

The long term skill bias of technical change

The key result from the above is that as H/L increases, the skill premium (π) falls. But as we know, in every advanced country the supply of educated workers has risen dramatically in the past seven decades, yet relative wages of better educated workers have remained consistently above those of less educated (the degree to which they have done so has varied by decade). So, in the U.S., for example, 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 1915; see the 2007 paper by Katz and Goldin on the *Race Between Education and Technology*).

These observations suggest that 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).

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. Jan Tinbergen (the second Nobel laureate in economics) 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."

Tinbergen's hypothesis provides a useful framework for thinking about the evolution of inequality—or at least the return to skill—is:

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

...and...

- 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 vice versa when educational production outpaces technological advances.

An important caveat to the Tinbergen view is that it implicitly postulates 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 low-skilled labor-biased. These technologies replaced the work of skilled artisans (metal-smiths, carpenters, weavers, etc.) with capital and low-skilled labor (often children working as indentured servants). 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 almost surely 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, the paper lacks direct evidence on unskill-biasedness (it offers evidence only on skill-complementarity).

1.2 Bringing the CES model to the data

From (5), 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 (conditional on taking the rest of the model seriously). 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 pioneered by Katz and Murphy (1992) in their seminal paper.

The Katz-Murphy model (AKA, "The Canonical Model")

Recall from (6) that

$$\ln \omega = \frac{\sigma - 1}{\sigma} \ln \left(\frac{A_h}{A_l} \right) - \frac{1}{\sigma} \ln \left(\frac{H}{L} \right). \tag{9}$$

Let's say we wanted to estimate this model using time series data. To make this operational, we need to add time subscripts to everything in the 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 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:

$$\ln \omega_t = \gamma_0 + \gamma_1 t + \gamma_2 \ln(H/L) + e_t, \tag{10}$$

where $\partial \ln(A_h/A_L)/\partial t$ is linear in time. In estimating this equation, γ_o 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 Current Population Survey data for the years 1963 - 1987, Katz and Murphy's fit this model using a simple OLS regression. They obtain:

$$\ln \omega = \underset{(0.01)}{0.033 \cdot t} \quad \underset{(0.15)}{-0.71 \cdot \ln\left(\tfrac{H}{L}\right)} \quad + \text{constant} \tag{11}$$

Taken literally, 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 $\widehat{\sigma} = -1/0.709 = 1.41$. 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 Katz and Murphy's 1992 paper. Katz-Murphy rewrite the expression for wages above as a relative demand index as follows

$$\ln \frac{W_{Ht}}{W_{Lt}} = \frac{1}{\sigma}\left[D_t - \ln \frac{H_t}{L_t}\right],$$
$$D_t = \ln\left(\frac{W_{Ht}H_t}{W_{Lt}L_t}\right) + (\sigma - 1)\ln\left(\frac{W_{Ht}}{W_{Lt}}\right),$$

where $D_t \equiv (\sigma - 1)\ln\left(\frac{A_{Ht}}{A_{Lt}}\right)$. 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. The only exception is the period from 1975 to 1981, when their model suggests that inequality should have begun to rise earlier than it did. 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 reach 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 *rationalize* 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 implies a dramatic demand shift must have occurred.

A key conclusion of Katz-Murphy '92 is that fluctuations in relative supply overlaid on smoothly rising demand go a long ways towards explaining 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? Definitely not. But it's a good start!

See also the updated estimation of the Katz-Murphy model from the Acemoglu-Autor 2011 Handbook chapter. Projecting the Katz-Murphy estimates forward to 2008 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 the present 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 soon.

1.3 Longer-term evidence

Katz and Murphy showed that one could explain many of the patterns in the data with a very simple, steady increase in demand for educated workers. Is there anything more than fluctuations in supply driving patterns of wage inequality, or have these demand trends shifted over the longer-run? This is what Autor, Katz, Krueger (1998) use a longer time series to address this question. Define the demand index as:

$$D_t = (\sigma - 1)\ln(A_H/A_L) = \ln(w_H H/w_L L) + (\sigma - 1)\ln(w_H/w_L). \tag{12}$$

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 wage-bill shares $(w_H H, w_L L)$

3. Estimates of w_H/w_L

One concern in implementing this framework is that the wage-bill 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

$$\Delta \ln(\text{relative wagebill}) = \Delta \ln(w_H H/w_L L) = \Delta \ln(H/L) + \Delta \ln(w_H/w_L) \qquad (13)$$

Hence, there is both a supply and a price component to the change in the wage-bill. 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:

$$\begin{aligned} \Delta \ln(\text{relative supply}) &= \Delta \ln(\text{relative wagebill}) - \Delta \ln(\text{relative wage}) && (14) \\ &= \Delta \ln(w_H H/w_L L) - \Delta \ln(w_H/w_L). && (15) \end{aligned}$$

This procedure effectively 'subtracts off' the component of wage-bill 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 (updated further in Acemoglu and Autor, 2011), there appears to be some demand *deceleration* after the early 1990s and continuing through the first decade of the 21st century. This is noteworthy and probably unexpected for many versions of SBTC.

So, the key fact that we have so far deduced is the necessity of demand shifts to explain patterns over the longer-run, 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 surprise 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.

But their analysis also points out two periods when the model decidedly does not fit the data. One is the decade of the 1940s, when skill differentials fall much more than the canonical model can rationalize. The second is in the 1970s, when the skill premium declined by more than the model would predict—and then sharply increased after the deep early-1980s recession. Goldin and Katz interpret these deviations as suggesting that institutional forces have a fairly transitory effect on the wage structure; in the long run, supply and demand prevail. Other interpretations are possible.

2 International evidence on wage inequality and the supply and demand of skills

The U.S. has notably higher wage inequality and returns to skills (which are 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 attractive hypothesis, though it's not obvious how to test it. However, a handful of papers have suggested that the supply-demand framework has limited power to explain cross-country differences in skill premia. In this section, we explore whether the supply-demand framework has any explanatory power for understanding cross-country patterns of skill premia and overall inequality.

The first notable entry in this cross-national 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.

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. LOvO use these data to ask how well cross-country differences in supply and demand can explain cross-country differences in skill differentials. (There are a number of other papers that use these data for related exercises, including Blau and Kahn, in *ReStat* and the recent 2015 working paper by Broecke, Quintini, and Vandeweyer, which essentially replicates LOvO using more recent data).

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. Supply is imperfectly measured. Demand is never observable. 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 $k =$\{low, medium, high\} using as cut-points 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:

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

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 *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 (relative to the base country) 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:

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

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). 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. 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 conclusions. It's possible that this produces misleadingly narrow confidence intervals. These 15 separate analyses are not in any sense statistically independent.)

The 2015 working paper by Broecke, Quintini, and Vandeweyer "Wage inequality and cognitive skills: Re-opening the debate," updates the LOvO paper using newer data and a somewhat different set of countries. Their paper adds some nuance to the earlier conclusions. Using a similar framework to LOvO, they conclude that supply/demand can explain some share of the greater upper-tail (90/50) wage inequality in the U.S. than in comparison countries. By contrast, it explains none of the greater lower-tail (50/10) wage inequality in the U.S.. Institutional factors (minimum wages, union coverage) are strong predictors of the cross-country variation in lower-tail inequality, however.

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.

3 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 (or parametric) modeling to the estimation of skill premia.

(b) **Accelerating/decelerating 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. However, history has not been kind to this accelerationist view. Katz-Murphy style evidence suggests a slowdown in the U.S. demand shifts favoring college educated labor starting in the early 1990s. See especially the 2016 *JOLE* paper by Beaudry, Green and Sand ("The Great Reversal in the Demand for Skill and Cognitive Tasks").

2. **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. Whether or not that interpretation is correct, 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. Papers by Markö Tervio and Amanda Pallais explore these ideas theoretically and empirically. Tervio's "mediocrity" paper offers an ingenious alternative (non-benign) interpretation of the 'superstars' idea.

3. **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 Becker and Murphy (1992), Dessain and Santos (2008).

4. **International trade.** 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). For decades, the literature on trade and labor markets had found little compelling evidence that international trade was a major factor in the evolution of wages, skill differentials, or employment rates. But recent evidence has substantially changed the consensus. I will not cover the trade theory this term (another casualty of time constraints), but I will discuss some of the evidence. Interested parties should consider taking 14.581, which provides a deep dive into trade theory.

5. **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 also been quite heated. We will discuss several important institutional explanations. These are not just our grandfather's institutions either (e.g., unions and minimum wages); some are your *great* grandfather's institutions (e.g., the master and servant doctrine).