



Munich Personal RePEc Archive

Directed technical change: A macro perspective on life cycle earnings profiles

Cragun, Randy and Tamura, Robert and Jerzmanowski, Michal

Niagara University, Clemson University

20 November 2017

Online at <https://mpra.ub.uni-muenchen.de/81830/>
MPRA Paper No. 81830, posted 21 Nov 2017 16:43 UTC

Directed Technical Change: a Macro Perspective on Life Cycle Earnings Profiles

Randy Cragun^{*}, Robert Tamura^{**}, and Michal Jerzmanowski^{**}

^{*}*Department of Economics and Finance, Niagara University. Contact: rragun@niagara.edu*

^{**}*Department of Economics, Clemson University*

November 20, 2017

First version: 2014

Abstract

We propose a new macroeconomic mechanism for generating patterns in age-earnings profiles based on directed technical change. The mechanism does not depend on changes in the human capital of the individual; rather differences in the human capital shares of age groups affect the profitability of developing age-specific technologies, biasing innovation toward improving the productivity of workers whose cohorts provide large shares of the labor force's human capital. The theory should be taken as supplemental to (rather than replacing) human-capital-based theories of age-earnings profiles. Using recently developed data and human capital estimates, we simulate reductions in wages due to age-biased directed technical change over workers' lives for most of the world's nations over two centuries. Our simulations indicate that age-specific directed technical change contributes to wage concavity for the average country \times cohort group. Because younger workers make up a larger share of human capital in most years and countries, most cohorts in most countries experienced wage gains early in life and losses later in life from age-specific directed technology, making life-cycle earnings profiles flatter. The late-life losses are larger than the early-life gains, increasing concavity in life-cycle earnings profiles. We also calculate the present value of lifetime wage gains from age-specific endogenous technology, and find particularly large and economically-significant gains for baby boomers and losses for cohorts born during low population growth periods.

1 Introduction

This paper builds on the endogenous directed technology theory developed by Acemoglu (1998, 2002). In those models, workers with different skill sets are imperfect substitutes and technologies are skill-specific; that is they raise the productivity only of the workers for whom they have been designed. Which technologies get developed in equilibrium—and as a result which groups of workers become more productive and more highly compensated—depends crucially on the relative supply of workers in the different groups. Those groups that are large ensure a greater market for innovators, allowing them to more easily recoup the fixed cost of research and thus, *ceteris paribus*, provide a stronger incentive to direct innovation at them.

Our model extends this logic to different age groups of the labor force; technologies are age-specific in the sense that they raise the productivity of workers of only the specific age cohort for which they have been designed. We think this is a natural extension because older and younger worker are indeed different and many technological improvements do seem inherently tied to specific ages in the labor force. For instance, a technology like eye glasses or

hearing aids, while obviously used by all age groups, improves the productivity of older workers in a disproportionate manner. This extends beyond technologies that overcome ageing of human bodies. For example, a remote office technology enhances the productivity of parents (especially women) with young kids and has a relatively smaller impact on middle-age and older workers.

We write down an age-specific directed technology model and calibrate its parameters in order to back out the age-specific productivity levels for a large sample of countries during the period 1820–2010. One of the crucial parameters of this exercise is the elasticity of substitution between workers of different ages. This is a parameter without a consensus value, and our model sheds some new light on how we should interpret its existing estimates. We also produce our own estimates of this elasticity, which is an additional contribution of the paper. With parameters calibrated, we back out the age-specific productivity levels and the wages they imply. We compare the simulated wages to wages from a similar model in which workers of different ages are imperfect substitutes but where technologies are *not* age-specific and find that endogenously directed age-specific technology induces economically-significant concavity in age-earnings profiles for the average birth cohort (and also for most cohorts in most countries over the 20th century) and increases the present value of lifetime earnings for baby boom cohorts while decreasing the present value of lifetime earnings for cohorts in East Asia (other than in Japan) over the last half century).

2 A new mechanism driving age-earnings profiles

The general shape of the life-cycle earnings profile is well-known: wages grow quickly early in the career and then slowly late in life (Becker, 1964, 1967; Ben-Porath, 1967; Casanova, 2013; Lagakos et al., 2016; Mincer, 1958, 1974, 1997; Murphy and Welch, 1990, 1992; Psacharopoulos, 1994; Rosen, 1976). Typically economists explain this with individual human-capital accumulation (Becker, 1964, 1967; Ben-Porath, 1967; Mincer, 1974, 1997; Rosen, 1976).

Our contribution is to show that it is not necessary for the individual's human capital to decline in order for her wage to grow slowly relative to others' wages. It is enough that the human capital possessed by her age group represent a smaller share of the total human capital in the market because this shrinks the relative profits for innovators producing technology for her age group. Changes in human capital shares occur over the life cycle for various reasons:

- The size of a cohort usually declines over time as members die, retire, or reduce their hours of labor supplied (immigration could reverse this, but immigrants are usually young [CITE]).
- In many places new generations are larger than the previous generation because of high fertility (it takes low fertility by historical standards for this to not be true [CITE]).
- Younger generations usually have more average human capital than older generations both (for instance because high human capital reduces the costs of human capital accumulation and because expectations of better techno-

logies and lower mortality rates in the future raise the expected return to investing in your children's human capital).

- Workers gain skills throughout their careers through on-the-job training and learning by doing (e.g. Mincer, 1974 and Ben-Porath, 1967, 1970). Although this contributes to an individual's wage through that worker's individual human capital, it also contributes to increasing the wages of other workers around the same age by increasing the share of human capital supplied by that age group.

If the first three conditions dominate (as we find), then the market size for technologies directed at younger workers will be greater than the market size for technologies directed at older workers, and workers will see their wages stagnate as a result as they age. On the other hand, we could imagine a scenario in which the first three effects were small, and then workers' experience-derived human capital would determine most of the changes in the share of human capital supplied by a cohort throughout its life, and the market size for technologies directed at older workers would be greater than the market size for technologies directed at younger workers, causing workers to see their wages grow more as they age. These comparative statics must be understood within the context of general technology growth driving up wages over time (so that wage losses from being from a low-human-capital-share age group are deviations from an upward trend).

3 Model of age-specific directed technical change

Here we develop a model of directed technical change to study differences in wages by age group. We use ten-year age bins (15–24, 25–34, 35–44, 45–54, and 65–54), and study how the age-group-specific technologies and wages depend on the age-group-specific labor supplies¹ and human capital.²

The model follows closely with Acemoglu (2002) with two major differences. We use the total human capital provided to the market by a given age group as a productive input rather than using the quantity of labor supplied by people of different levels of human capital. This causes little difference in solving the model. However, because we care about the experiences of workers rather than the abstract idea of the wage of a unit of human capital (which would have an equivalent function in the model to wages per worker in Acemoglu, 2002), we solve the model for the wage per unit of labor. Because we include multiple countries (and because country populations grew substantially over the last two centuries), we modify the innovation possibilities frontier to limit scale effects and to allow countries to have different costs of innovation (as in Jerzmanowski and Tamura, 2017). Although economies are assumed to evolve in continuous time, we leave off time subscripts where this does not cause confusion.

¹We actually use the age structure of the population as the assumed age structure of the work force. This may be a better assumption earlier in our sample because in a world where education ends by age 10 (or even as old as 14), the age structure of the 15–64 population is likely similar to the age structure of the work force.

²Human capital is constructed following Tamura et al., 2016 and includes schooling, experience, intergenerational transfers, and international spillovers. See Section 6.3.2 for details.

We solve two versions of the model: one with age-specific technologies and one where all innovation complements every worker age group. In both cases worker age groups are imperfect substitutes.

3.1 Consumption Demand

For simplicity, assume a continuum of measure³ 1 of infinitely-lived representative consumers with CRRA preferences over a single consumption good over time. Consumers can borrow and lend at the interest rate r and have an internal rate of discount ρ . The representative consumers own a balanced portfolio of physical capital and technology patent rights and from those collect income equal to the aggregate output of final goods net of taxes and spend it on consumption C , investment in capital I , and research and development efforts $R\&D$. With preferences ordered by

$$U = \int_0^{\infty} e^{-\rho t} \frac{C_t^{1-\theta} - 1}{1-\theta} dt \quad (1)$$

consumption growth will be

$$\frac{\dot{C}}{C} = \frac{1}{\theta} (r - \rho) \quad (2)$$

3.2 Final Goods

Aggregate output Y is produced competitively with intermediate goods Y_a , where a indexes age groups, according to⁴

$$Y = \left[\sum Y_a^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \quad (3)$$

The representative final good producer purchases intermediate goods at prices P_a and faces the problem

$$\max_{Y_a} \{ Y - \sum P_a Y_a \} \quad (4)$$

which implies that

$$\frac{P_a}{P_b} = \left(\frac{Y_a}{Y_b} \right)^{-\frac{1}{\varepsilon}} \quad (5)$$

where b indicates a second age group. Given relative prices, a zero profit condition, and data on GDP, we can solve

³The choice of size of the consumer base does not impact the results. This normalization is simply convenient for talking about a representative consumer.

⁴A more general version of this production function would seem to be one that allows for different weights on the intermediate goods:

$$Y = \left[\sum \gamma_a Y_a^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}$$

However, these weights are not separately identified from the levels of production of the intermediate goods (which are not observed) and the technology levels that lead to that production. Thus we normalize $\gamma_a = 1$ and allow the production of intermediate goods to capture all aspects of the value of constituent markets in final good production.

for the intermediate goods prices:

$$P_a = Y Y_a^{-\frac{1}{\varepsilon}} \left(\sum_b Y_b^{\frac{\varepsilon-1}{\varepsilon}} \right)^{-1} \quad (6)$$

Final good output is the numeraire, so

$$\sum P_a^{1-\varepsilon} = 1 \quad (7)$$

3.3 Intermediate Goods

Intermediate goods are produced competitively using machines κ_a , labor L_a , and human capital h_a (which we construct for each cohort and country from average schooling levels and using an intergenerational transfer model from TDDB and Tamura et al., 2012, as described in Section 6.3.2) that are specific to the intermediate good:

$$Y_a = \frac{1}{1-\beta} (h_a L_a)^\beta \int_0^{X_a} \kappa_a(x)^{1-\beta} dx \quad (8)$$

In this particular application, the labor types are age groups, so readers can think of these intermediate goods industries indexed by a as composite sets of services offered to the market by workers of different ages. There is a continuum of length X_a of machine types, representing the level of technology in age group (sector) a ; Y_a sells at price P_a ; machines of type x can be rented at the price $p_a(x)$; and workers of age group a get paid wage w_a . Thus the problem for the intermediate good producer is

$$\max_{\kappa, L} \left[\frac{P_a}{1-\beta} (h_a L_a)^\beta \int_0^{X_a} \kappa_a(x)^{1-\beta} dx - \int_0^{X_a} p_a(x) \kappa_a(x) dx - w_a L_a \right]$$

3.4 Demand for machines

Assuming that the first order conditions hold with equality, the profit maximization condition with respect to machines for age group a is that

$$P_a (h_a L_a)^\beta \kappa_a(x)^{-\beta} = p_a(x) \quad (9)$$

The blueprint for each line of machines is controlled by a monopolist who produces and maintains each unit of the machine with a unit of capital ⁶ rented at price R . Thus the stock of capital is always equal to the stock of machines. Each monopolist has a competitor who can produce the same machine at cost νR , so the monopolist sets the limit

⁵In the case without age-specific innovation, the equation is the same with the condition that $X = X_a \forall a$ while $\kappa_a(x)$ is still drawn from a separate pool of machines for each age group.

⁶This is like assuming that one unit of capital can produce one unit of machine per unit of time and then machines depreciate fully but capital does not. The reader can think of machines as capital that has been temporarily repurposed for a specific method of production. As Acemoglu (2002) shows, relaxing the assumption of a 100% depreciation rate on machines does not change the balanced growth path.

price $p_a(x) = vR$. The demand for a machine line for age group a for a given capital rental rate is then

$$\kappa_a(x) = \left(\frac{P_a}{vR} \right)^{\frac{1}{\beta}} h_a L_a \quad (10)$$

which is independent of the type x of machine. Production of intermediate goods then takes the form

$$Y_a = \frac{1}{1-\beta} \left(\frac{P_a}{vR} \right)^{\frac{1-\beta}{\beta}} X_a h_a L_a \quad (11)$$

Combined with equation 5, this means that

$$\frac{P_a}{P_b} = \left(\frac{X_a h_a L_a}{X_b h_b L_b} \right)^{-\frac{\beta}{\sigma}} \quad (12)$$

where $\sigma = 1 + (\varepsilon - 1)\beta$ is the elasticity of substitution between worker age groups conditional on fixed technology.

3.5 Monopoly profits for machine blueprint owners

The stream of profits per unit of κ_a is $(v-1)R$, so the stream of profits per machine blueprint is

$$\pi_a = (v-1)R\kappa_a = \frac{v-1}{v} P_a^{\frac{1}{\beta}} h_a L_a (vR)^{\frac{\beta-1}{\beta}} \quad (13)$$

3.6 Demand for labor

The profit maximization condition with respect to labor is that

$$w_a = \frac{\beta}{1-\beta} P_a (h_a L_a)^{\beta} L_a^{-1} \int_0^{X_a} \kappa_a(x)^{1-\beta} dx$$

Combined with the condition for machines, this gives us that

$$w_a = \frac{\beta}{1-\beta} P_a^{\frac{1}{\beta}} (vR)^{-\frac{1-\beta}{\beta}} X_a h_a \quad (14)$$

which means that the relative wages are

$$\frac{w_a}{w_b} = \frac{X_a}{X_b} \left(\frac{P_a}{P_b} \right)^{\frac{1}{\beta}} \frac{h_a}{h_b} \quad (15)$$

Combining 15 with 12 gives

$$\frac{w_a}{w_b} = \left(\frac{X_a h_a}{X_b h_b} \right)^{\frac{\sigma-1}{\sigma}} \left(\frac{L_a}{L_b} \right)^{-\frac{1}{\sigma}} \quad (16)$$

3.7 Allocation of capital

Because one unit of capital produces one machine and machine varieties are symmetric within age groups, $K_a(x) = \kappa_a(x) = \kappa_a(s)$ for any machine types x and s , and the total capital used in an intermediate good market is

$$K_a = \int_0^{X_a} \kappa_a(x) dx = X_a \kappa_a(x) = X_a \left(\frac{P_a}{vR} \right)^{\frac{1}{\beta}} h_a L_a \quad (17)$$

Foregone consumption produces capital K , which is allocated between the intermediate goods markets to produce machines. The relative capital allocations (using the relative prices already derived) are

$$\frac{K_a}{K_b} = \left(\frac{X_a h_a L_a}{X_b h_b L_b} \right)^{\frac{\sigma-1}{\sigma}} \quad (18)$$

Given data on the aggregate capital stock, we can find the values of capital allocated to each age group by considering that $K = \sum_a K_a$, where K is the size of the aggregate capital stock.

3.8 Interest rates

In order to get the rental rate, write the intermediate good output as

$$Y_a = \frac{1}{1-\beta} K_a^{1-\beta} (X_a h_a L_a)^\beta \quad (19)$$

and note that the value marginal product of capital in this market is

$$P_a MPK_a = (1-\beta) P_a \frac{Y_a}{K_a} \quad (20)$$

Capital market equilibrium requires that $vR = P_a MPK_a$, so

$$P_a Y_a = R K_a \frac{v}{1-\beta} \quad (21)$$

The intermediate goods markets adjust so that there are zero profits in the final good market, so the sum of the age-group-specific versions of Equation 21 must equal Y , which requires that

$$R = \frac{1-\beta}{v} \frac{Y}{K} \quad (22)$$

The interest rate in this economy must be the rate that could be earned by producing capital:

$$r = (1 - \tau)R - \delta = (1 - \tau) \frac{1 - \beta}{v} \frac{Y}{K} - \delta \quad (23)$$

where τ is the tax rate on capital income and δ is the rate of capital depreciation. This condition is not strictly in terms of observables for the purposes of this paper because we do not observe τ (although we assume that firms and consumers take it as exogenous).

3.9 Technological innovation with age-specific technologies

There are multiple countries that build their own stock of technology with $R\&D$ expenditures (foregone consumption) with the following innovation possibilities frontier (for the version without age-specific technologies, see Section 3.11):

$$\dot{X}_a = \frac{1}{\zeta} \frac{R\&D_a}{L} \quad (24)$$

where $R\&D_a$ is age- a -specific spending on innovation-producing activities, $\sum_a R\&D_a = R\&D$, L is the size of the labor force, and ζ is a country-specific innovation entry barrier. $\frac{R\&D_a}{\zeta}$ should be thought of as the effective expenditure on $R\&D$, with differences between countries presumably due to less pro-growth institutions or environments. This form is different from that used by Acemoglu because $R\&D$ expenditures are scaled by the size of the labor force to reduce scale effects. We justify this by recognizing that people compete to produce new technologies by partially duplicating efforts (as, for example, in Klenow and Rodríguez-Clare, 1997a).

Although the owners of machine blueprints produce their machines with some market power once they have the blueprint (e.g. because of intellectual property protections), there is free entry into innovation. One unit of $R\&D$ expenditure produces a stream of $\frac{1}{\zeta L}$ new blueprints, each with a present value V_a , so as long as there is any innovation in sector a ,⁷ free entry requires

$$V_a = \zeta L = V_b \quad (26)$$

and thus the value of a line of machines specific to any age group must be equal to the value for other age groups:

3.10 Characterizing the balanced growth path

No arbitrage requires that

$$rV_a = \pi_a + \dot{V}_a \quad (27)$$

⁷More generally, the conditions for equilibrium that allow for corner solutions are

$$\left(\frac{1}{\zeta L} V_a - 1 \right) R\&D_a = 0, \frac{1}{\zeta L} V_a \leq 1, \text{ and } R\&D_a \geq 0 \quad (25)$$

This application deals only with the balanced growth path⁸. Differentiating Equation 26 with respect to time (taking labor supplies as constant) gives, $\dot{V}_a = 0$, so $rV_a = \pi_a$ for any equilibrium in which innovation occurs in all sectors. Using the value for instantaneous profits from Equation 13, the BGP values of machine blueprints must satisfy

$$V_a = \frac{\pi_a}{r} = \frac{\nu - 1}{\nu r} P_a^{\frac{1}{\beta}} h_a L_a (\nu R)^{\frac{\beta-1}{\beta}} \quad (28)$$

Equation 28 implies

$$\frac{V_a}{V_b} = \left(\frac{P_a}{P_b} \right)^{\frac{1}{\beta}} \frac{h_a L_a}{h_b L_b} = \left(\frac{X_a}{X_b} \right)^{-\frac{1}{\sigma}} \left(\frac{h_a L_a}{h_b L_b} \right)^{\frac{\sigma-1}{\sigma}}$$

using Equation 12 to give us the relative prices. Combining this with the equilibrium condition in 26 shows how relative technology between age groups varies with relative human capital:

$$\frac{X_a}{X_b} = \left(\frac{h_a L_a}{h_b L_b} \right)^{\sigma-1} \quad (29)$$

Given these relative technologies, we can get relative wages with Equation 16:

$$\frac{w_a}{w_b} = \left(\frac{h_a}{h_b} \right)^{\sigma-1} \left(\frac{L_a}{L_b} \right)^{\sigma-2} \quad (30)$$

Knowing the relative technologies (Equation 29) and using Equation 19 and the final good production function, we can determine the levels of technology:

$$X_b = \left[\frac{Y(1-\beta)}{\left(\sum_a \left[\left(h_a L_a \frac{X_a}{X_b} \right)^\beta K_a^{1-\beta} \right]^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}} \right]^{\frac{1}{\beta}} \quad (31)$$

where the values of K_a above are those determined by Equation 18 and data on aggregate capital.

Combining the condition on relative technologies in Equation 29 with what we know about relative intermediate goods prices from Equation 12, we find that

$$\frac{P_a}{P_b} = \left(\frac{h_a L_a}{h_b L_b} \right)^{-\beta} \quad (32)$$

which depends only on observables. Using the price normalization from Equation 7 we find a form for the levels of

⁸Although Acemoglu (2007) showed that off-BGP trends may be needed to match historical changes in skill premia because the price effect dominates the short run and the market size effect often dominates in the long run, we do not expect this to be a substantial weakness of our work. The reason is that when workers age, they are simply moving to a different but established role in the economy rather than changing the expected distribution of workers across ages. If there were instead a sudden, unexpected age demographic shock, then this could cause problems for this analysis. The baby boom is one potential shock, but innovators had over an entire decade to prepare for this shock (average fertility rates are not private information). Other potential shocks are wars, which may differentially kill young, male workers) and diseases, which often differentially kill older workers).

intermediate goods prices that also depends only on observed values. They must satisfy

$$P_a = \left[\sum_b \left(\frac{P_b}{P_a} \right)^{1-\varepsilon} \right]^{\frac{-1}{1-\varepsilon}} \quad (33)$$

3.11 Technological innovation without age-specific technologies

We also calculate balanced growth paths without age-specific technologies for comparative statics. The utility, final goods production function, and intermediate goods production functions are the same as before, so workers of different ages are still imperfect substitutes. The innovation possibilities frontier in this case was chosen to be as comparable as possible to the case with age-specific technology:

$$\dot{X} = \frac{\eta}{\zeta} \frac{R\&D}{L} \quad (34)$$

where $X = X_a = X_b \ \forall (a, b)$, $R\&D$ gives total spending on innovation-producing activities, and η is chosen to allow ζ to have the same value as in the case with age-specific endogenous technologies.

The model solutions turn out to be the same as before but with the relative technologies fixed at 1. In particular Equations 16, 18, 12, 7, 31, and 14 continue to hold (but with relative technologies set to 1).

4 Why might age-specific directed technical change drive down wage growth rates late in life?

“Weak bias” is defined to mean that when the relative supply of a factor increases its relative technology also increases. This is equivalent to saying that demand for the factor increases when its supply increases. “Strong bias” occurs when an increase in the relative supply of a factor leads to an increase in its relative wage.⁹

Strong bias requires greater substitutability of workers than weak bias does. Using Equation 30 we can see that an increase in the relative average human capital of a particular age of workers will always increase their relative wage when $\sigma > 1$, but an increase in the labor share supplied by a particular age group will lead to an increase in its relative

⁹While the wage used in Equation 30 is per worker, firms care about the total human capital provided to them and hence about the wage per unit of human capital. To look at the typical story of technical bias (e.g. Acemoglu, 2002) using this model, we would ask how changes in the total human capital provided to a sector impact the wage earned (or technology used) by a unit of human capital in that sector:

$$\frac{w_a/h_a}{w_b/h_b} = \left(\frac{h_a L_a}{h_b L_b} \right)^{\sigma-2} \quad (35)$$

Because we are taking average human capital as exogenous, and producers care only about the total human capital they hire, the relative wage per unit of human capital depends only on the relative total human capital between age groups. We have strong technical bias as long as $\sigma > 2$. It is not surprising that this condition is the same as the one for strong bias using just the relative supplies of labor. The reason is that the total human capital can be increased either by increasing the quantity of labor supplied or by increasing the average human capital. Thus the condition for strong bias from increasing total human capital will be the stronger of the conditions for getting strong bias from either of these. This result for bias in the relative wages paid to human capital is the same as for those paid to labor in Acemoglu (2002).

wage only when $\sigma > 2$.

This report, however, is interested primarily in weak bias, and it is clear from Equation 29 that we have weak bias as long as $\sigma > 1$. That is: any increase in the relative total human capital of a particular age group induces a relative increase in the technology used by that age group as long as workers of different ages are moderately good substitutes for each other regardless of whether the increase in relative human capital occurs because the quantity of labor supplied increases or because the average human capital of workers increases. The existence of weak bias matters for this report because we are concerned with how wages differ because of one particular mechanism; we are interested in how wages would be different without age-specific endogenous technology. When an age group supplies a large share of the human capital in a country, technology will heavily complement that age group, and this effect will increase their wages *relative to what their wages would have been without DTC*. Their wages could fall as a result of their higher labor share (when there is not strong bias), but the wages fall less than they would without the DTC mechanism.

A little algebra gives the technology levels in terms of human capital shares:

$$X_a = \frac{[(1-\beta)Y]^{\frac{1}{\beta}}}{HK^{\frac{1-\beta}{\beta}} (\sum_b \lambda_b^{\sigma-1})^{\frac{\sigma}{\sigma-1}}} \lambda_a^{\sigma-1} \quad (36)$$

where λ_a is share of total human capital possessed by age group a . An increase in the human capital share of an age group usually increases the technology *level* available to it for reasonable values of σ . For instance, when $\sigma = 2$ (similar to our estimates in Table 1), the index of human capital shares in the denominator is 1, and the technology is proportional to the cohort's human capital share (with the same proportion for every cohort). With $\sigma = 2.2$ (a high value given our estimates in Table 1), if we simulate 100,000 years of random human capital shares (from a uniform density) and then reconstruct technologies from Equation 36 (assuming that everything other than human capital shares is constant), the correlation between an age group's human capital share and its technology is .997. Thus DTC will contribute to wage losses later in life if human capital shares tend to fall later in life.

Cohort human capital shares fall as people age because their cohort usually shrinks while young cohorts get larger and larger, because younger cohorts tend to be better educated, and because the benefits of human capital accumulation fall and the costs rise for most individuals in the cohort as the cohort ages (e.g. (Becker, 1964, 1967; Ben-Porath, 1967; Mincer, 1974, 1997; Rosen, 1976)). Figure 1 shows that human capital shares (which we construct with average lifetime schooling attainment, synthetic work experience, and intergenerational transfers using the methodology and schooling data from TDDDB) for younger age groups are indeed higher than for older age groups in various countries. The same graph for other countries typically produces the same result. We also see this pattern even though many countries are aging, as seen in Figure 2 for the US. This occurs because of increases in education over time and intergenerational transfers of skills.



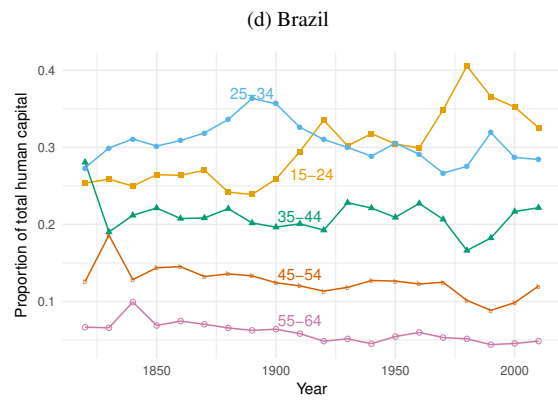


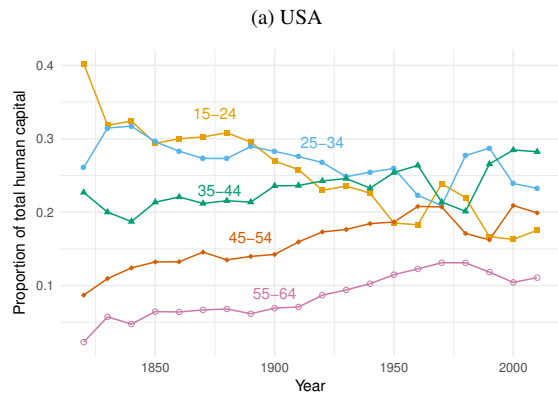


Figure 1: Age distribution of human capital in various countries




Figure 2: Age distribution of labor in the US

If we split these data up into time series for each cohort, the pattern becomes even more clear. Figure 3 does this with human capital shares. Each line is one cohort and starts when the cohort is approximately 20 years old and ends when the cohort is approximately 60 years old. Because of increased schooling over time and inter-generational transfers, cohorts almost always see their human capital shrink later in their careers and thus will see the technology devoted to them shrink relative to what it would have been without age-specific endogenous technology.

It is important to remember that our choice of σ as the border between strong and weak bias (see Section 5) does not mean that larger cohorts do not benefit from DTC for fixed levels of human capital per worker. Rather larger cohorts receive more *R&D* even if they do not have more human capital per worker, and the wage benefit from this DTC (which we will plot in a later section—see Figure 13) is not sufficient to offset losses from competing with a larger cohort. Thus when we say that wage curvature is induced by DTC, this is not the same as saying that wage curvature is induced by having a smaller cohort when older. We are not arguing that a higher human capital share necessarily induced higher average wages; we are arguing instead that a higher human capital share induces captures more *R&D* and that that is one factor contributing to wage profiles. We are interpreting effects of human capital on technology as the mechanism for DTC effecting wages.

5 Estimating the elasticity of substitution, σ

Because the elasticity of substitution between workers by age group is so important for the strength of technical bias, this section presents estimates for that elasticity. We consider the role of directed technology in generating shocks




Figure 3: Cohort human capital shares in the US. Each line is a separate 10-year birth cohort with human capital estimates starting at age 15–24 and going to age 55–64.

to the relative demand for workers, incorporate an instrument for labor supply by age group, and reconcile multiple estimates.

Card and Lemieux (2001) (C&L) estimate the elasticity of substitution between age groups of workers with the same education level in the US, UK, and Canada, with typical estimates ranging from 4 to 6. Our σ is calculated for human capital based on the average schooling attainment in a cohort, whereas C&L calculate the elasticity for specific schooling levels, but the values should be made more comparable by the fact that the elasticity in C&L is assumed to be the same within different education groups. The values also differ in that ours is a measure of the substitutability of effective units of labor, whereas C&L are looking at the substitutability of workers, but the elasticity of substitution between labor provided by different age groups is the same in our model as the elasticity of substitution between human capital provided by different age groups (provided that average human capital is not permitted to respond to changes in labor supply).

5.1 The estimation problem and the trouble with ignoring age-specific technical change

Estimates of the elasticity of substitution between worker groups are concerned with measuring the shape of the relative demand curve for a worker group by employing (hopefully) exogenous shifts in relative supply. Unfortunately for

these estimates, shifts in relative factor demand from increased relative technology will bias estimates of the elasticity upward so long as wages do not increase from the new technology and could create a positive relationship between prices and quantities if the increase in technology is large, falsely indicating a negative elasticity of substitution.

Equations 30 and 16 highlight the difficulty of estimating the relevant demand elasticities. We have recreated them below to help illustrate the problem.

$$\frac{w_a}{w_b} = \left(\frac{X_a}{X_b}\right)^{\frac{\sigma-1}{\sigma}} \left(\frac{h_a}{h_b}\right)^{\frac{\sigma-1}{\sigma}} \left(\frac{L_a}{L_b}\right)^{-\frac{1}{\sigma}} \quad (16 \text{ revisited})$$

$$\frac{w_a}{w_b} = \left(\frac{h_a}{h_b}\right)^{\sigma-1} \left(\frac{L_a}{L_b}\right)^{\sigma-2} \quad (30 \text{ revisited})$$

Equation 16 is an analog to the C&L regressions, and Equation 30 is the source of our regression model. Linear regressions of the logs of relative wages on the logs of relative labor supply, as in C&L, produce coefficients whose meaning is ambiguous. Suppose the regression coefficient is γ . If technology does not adjust to a supply shock, then transforming γ into σ as in Equation 16 ($\sigma = -\frac{1}{\gamma}$) would be justified. However, if technology has been able to adjust close to the BGP, we would want to transform γ into σ as implied by Equation 30 ($\sigma = \gamma + 2$). This is not a minor difference.

Consider the results of using the wrong transformation. Figure 4 plots point estimates for σ that we would get from various plausible regression coefficients using either the methodology from C&L or incorporating endogenous directed technology. The C&L estimates are unstable near the threshold for strong technical bias (where the coefficient is zero and our methodology gives $\sigma = 2$) even if the coefficient is estimated precisely. Estimates by C&L using wage and labor supply data by five-year age group (and split by college education versus high school education) from the 1960 US Census and the 1970-1997 CPS fall within the highlighted box, corresponding to an actual elasticity of substitution slightly below 2 if they did not properly control for technology changes and if our methodology is correct. We will shortly produce similar estimates and demonstrate the volatility in the C&L methodology by producing large, negative values with the C&L methodology using only a slightly different sample.

C&L would possibly argue that their estimation method might capture changing technologies. They estimate models with time fixed effects to capture changes in relative productivities of worker groups (here called “technologies”). However, as they point out, “the relative supplies of different age groups are not all trending at the same rate”; this will necessarily induce technology changes that are not captured by the time trends that they use if the directed technology theory is applicable.

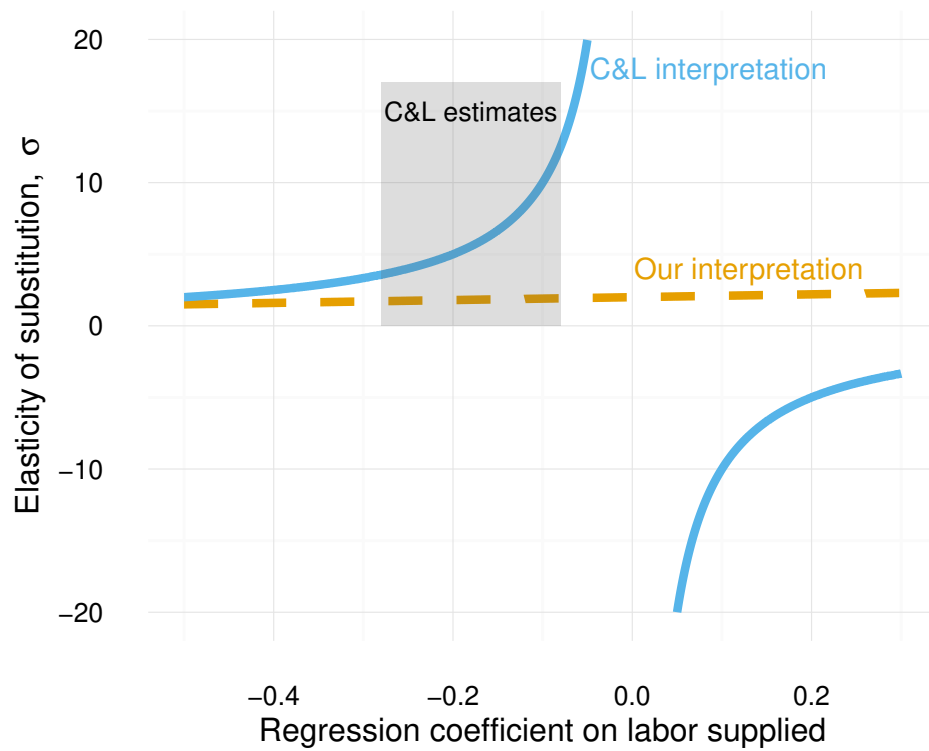


Figure 4: Estimates of σ derived by C&L's methodology and our methodology

Note: Estimates by Card and Lemieux (2001) lie within the highlighted box

5.2 Regression model and assumptions

We assume that measures of the marginal products of labor and of wages are noisy¹⁰ and assume that the noise can be captured by a full set of age group indicators and a full set of year indicators¹¹. Then taking the log of Equation 30, adding time subscripts, and allowing for some measurement error in wages, human capital, and labor supplied, we write our regression model as

$$\ln\left(\frac{w_{at}}{w_{20t}}\right) = (\sigma - 1)\ln\left(\frac{h_{at}}{h_{20t}}\right) + (\sigma - 2)\ln\left(\frac{L_{at}}{L_{20t}}\right) + \alpha_a + \delta_t + \xi_{a,t} \quad (37)$$

where ξ has conditional mean zero and we have indicated that the base category is age 15–24 with subscript “20”. The reader should note the implied restriction on the coefficients on labor supplied and human capital.

5.3 Instrument for supply shifts

Estimating these demand models is problematic because it requires assuming that fluctuations in relative labor supplied are exogenous. Contrarily, they are likely sometimes responses to the relative wages. Because higher relative wages should drive up relative labor supply, a simple fit of these demand functions to the data will overestimate the coefficient on relative labor supplied on average. In order to limit this endogeneity problem, we propose using relative cohort sizes to instrument for supply shifts. It seems unlikely that past population measures would respond much to wages.

A potential weakness of this instrument is that immigration could respond to wage differentials between age groups, so we use relative population values lagged by eleven¹² years from the sample month, whereas wages and labor supplied are calculated for work done in the year leading up to the sample month. Even when using lagged population values, earning potential and expected wages could drive both populations and wages wages independently of labor supplies because because of death (poorer workers may be more likely to die), measurement error (work choices could influence the likelihood that a worker would be counted by the Census), or because innovators might see the population size—rather than labor supplied—as the potential market for their innovations. However, we use the lagged population of the cohort rather than the lagged population of the same age group as our instrument. There probably are cohort-specific technologies, but we focus on age-group-specific technologies in this paper, and there is no overlap in the members of the cohort that is currently 25–34 and the people who were 25–34 eleven years ago (except to the extent that people misreport their ages).

¹⁰The data we use to estimate σ are estimates of wages, labor supplied, and human capital from the Current Population Survey of the United States, so even if Equation 30 holds exactly, there will still be a stochastic component in the regression using the available data.

¹¹This is in line with Murphy and Welch (1992), Card and Lemieux (2001), and Jerzmanowski and Tamura (2017). Alternative specifications with linear and quadratic time trends produce almost identical results.

¹²We use eleven rather than ten because the wages and labor supplied are from the year leading up to the sample.

5.4 Data for estimating σ

5.4.1 US wages and labor supplies

We construct wages and labor supplied for individuals and age groups (15–24, 25–34, 35–44, 45–54, and 55–64) in the United States using wage micro data from the Annual Social and Economic Supplement (ASEC) of the Current Population Survey (CPS) from the US Census from 1968 to 2017 (Ruggles et al., 2010).¹³ We use only workers who worked at least 40 weeks in the previous year.

We attempt to add up all hours of civilian labor supplied by people aged 15 to 64. To get the hours of labor supplied by each individual for samples from 1976 on, we multiply the number of weeks worked in the last year by the usual hours worked per week in the last year. Usual hours worked per week are not available before the 1976 ASEC, so for earlier years, we multiply the number of weeks worked in the last year by the number of hours worked in the previous week. We report results both with and without these earlier, noisy years. We then calculate totals by 10-year age group and year by multiplying the individual hours supplied by weights from the CPS that indicate how many people at the national level the sampled individual represents. Because the ASEC is early in the calendar year (March), we assign labor hours to the age group the person would have been in one year earlier. The smallest year by age group cell contains over 5,500 observed non-institutionalized workers, and the median cell size is 14,667.

For wages, we use only people with reported positive wage and salary income, giving a minimum cell size of 4,435 and a median cell size of 9,981. We calculate a person’s hourly wage by dividing their wage and salary income over the previous year by the number of hours they worked in the previous year. We then calculate weighted means of these hourly wages within age groups and years, where the weights are the same as those used to add up labor hours. As with the labor hours, we assign wages to the age group the person would have been in one year earlier. Wages are deflated by the GDP deflator.¹⁴ The wage and salary data are top-coded, and the strategy for dealing with top-codes changed significantly over the sample period. We use a set of corrected wage and salary data distributed by the Census Bureau (acquired through IPUMS) that apply the 2012 rank proximity swap rule to top-coded incomes in all years from 1975 on.

Figure 5 shows the total hours of labor supplied by age group and Figure 6 shows the average hourly wage by age group. The baby boom is clearly visible in the labor supplied graphs, and as others have pointed out (e.g. Murphy and Welch, 1992, and Card and Lemieux, 2001), this set of demographic changes had a big impact on “the structure of wages” (Murphy and Welch, 1992). The instrument for labor supply this report employs (lagged population) has the benefit of capturing changes due to these demographic shifts.

¹³We use only the 5/8 sample for 2014

¹⁴Others (e.g. Murphy and Welch, 1990; Murphy and Welch, 1992; or Card and Lemieux, 2001) use the CPI, but since this report is interested primarily in the shape of the demand curve and the strength of incentives for firms to innovate rather than in explaining workers’ purchasing power, the GDP deflator seems more appropriate.




Figure 5: Total hours of labor supplied by each age group




Figure 6: Average hourly wage of each age group

5.4.2 Population instrument

We construct the measures of population by age group from the CPS (ASEC) by adding up the weights for all non-institutionalized persons. The population values are lagged by eleven years. For the age group 35–44 in 1992, for example, we find the population aged 25–34 in the 1982 ASEC by adding up all the weights for people with those ages and divide that by the population aged 15–24 in that sample. This gets combined with wages and labor supplied by people aged 36–45 in the ASEC sample from the subsequent year, 1993 (again, because labor hours and wages were worked and earned over the year leading up to the sample). The earliest ASEC sample is from 1962, so the earliest wage and labor supply data we can use with this instrument are from the 1973 sample.

5.4.3 Human capital

We construct data on human capital following the methodology from TDDDB, who borrow their measure of the returns to experience from macro estimates from Bils and Klenow (1996, 2000), Klenow and Rodríguez-Clare, 1997b, and Hall and Jones (1999) and their measure of returns to schooling from Banerjee and Duflo (1995). The human capital of person i in age group a in year t is

$$h_{ait} = \exp\left(\frac{\text{SCHOOL}_{ait}}{10} + .0495\text{XPR}_{ait} - \frac{.0495}{67}\text{XPR}_{ait}^2\right) \quad (38)$$

where SCHOOL is the number of years of school the person completed and XPR is a measure of synthetic years of work experience for that individual, equal to $\max\{0, \text{AGE}_{ait} - \text{SCHOOL}_{ait} - 5.5\}$, where AGE is again the age the person would have been one year before the survey. We use 5.5 rather than the standard 6 to get the experience at the midpoint of the previous year rather than the experience at the beginning of the previous year. Thus someone reported as aged 30 in the survey with 11 years of schooling would be counted as having $29 - 11 - 5.5 = 12.5$ years of experience during the year of work leading up to the survey.

We average human capital within age groups using the person weights from the CPS and using the same sample of workers that were used for the average wage calculations. Figure 7 shows the average human capital of each age group over time. The differences between age groups (driven mainly by higher age producing higher implied experience) dominate the differences within age groups over time for younger workers. Older workers are further in their careers, so the negative quadratic term in experience has a greater influence on them and makes the relative human capital levels between age groups of older workers closer to 1.




Figure 7: Average human capital of each age group

5.5 OLS estimates of σ

There are five age groups for labor, so there will be four separate age-specific intercepts and time trends. We take 15–24 as the base category in all regressions and indicate this group with the subscript “20”.

Table 1 presents estimates of the model in Equation 37. At the bottom of the table, we present the estimates for σ from transforming the coefficients. Columns 1 and 2 present OLS estimates for 1967–2016 and for 1975–2016 (respectively) and suggest an elasticity of substitution a little above 2. In these data, unlike in Card and Lemieux (2001), increases in the labor supply are associated with rising relative wages. Although the coefficient estimate we get is close to theirs, the point estimate we calculate for σ with their methodology (-31.02) is far from their estimates (typically less than 10 and greater than 4) and not even the expected sign. We remind the reader that this instability occurs because they transform the coefficient on labor supplied into the elasticity with the relationship $COEFFICIENT = -\frac{1}{\sigma}$.

5.6 IV estimates of σ

We attempt to correct for potential endogeneity in labor supplied by using lagged relative population levels by birth cohort as an instrument for labor supply. The IV estimates in Columns 3–6 of Table 1 are comparable to the OLS estimates.

Table 1: Least squares estimates of σ (data from CPS ASEC from 1968 on)

Dependent Variable:	OLS		IV			
	1967–2016	1975–2016	1972–2016		1975–2016	
	$\ln\left(\frac{w_{at}}{w_{20t}}\right)$	$\ln\left(\frac{w_{at}}{w_{20t}}\right)$	1 st Stage $\ln\left(\frac{L_{at}}{L_{20t}}\right)$	2 nd Stage $\ln\left(\frac{w_{at}}{w_{20t}}\right)$	1 st Stage $\ln\left(\frac{L_{at}}{L_{20t}}\right)$	2 nd Stage $\ln\left(\frac{w_{at}}{w_{20t}}\right)$
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln\left(\frac{pop_{at}}{pop_{20t}}\right)$			0.590 (0.0664)		0.592 (0.0654)	
$\ln\left(\frac{h_{at}}{h_{20t}}\right)$	1.031 (0.0112)	1.024 (0.0113)	2.877 (0.3071)	1.084 (0.0284)	3.209 (0.2990)	1.093 (0.0287)
$\ln\left(\frac{L_{at}}{L_{20t}}\right)$	0.031 (0.0112)	0.024 (0.0113)		0.084 (0.0284)		0.093 (0.0287)
Age × Year Groups	206	206	206	206	206	206
Age group FE	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES
R^2	0.921	0.899	0.760		0.769	
σ	2.031	2.024		2.084		2.093
C&L σ	-32.20	-42.00		-11.92		-10.72

– Standard errors are in parentheses and are bootstrap estimates from sampling data de-meaned by age group
– R^2 and standard errors calculated using data de-meaned by age group
– Confidence intervals are bootstrap estimates

The main take-away from these estimates should be that the elasticity is close to and slightly above the boundary between strong and weak bias (2 in our model). Please see Appendix A for further discussion of these estimates and additional estimates including possible non-linearities and structural breaks.

6 Simulating model solutions for technology and wages

We assume that the data are observations from balanced growth paths (BGP). For each decade, the relative technologies and wages and the technology and wage levels from a new BGP for each country can be calculated by supplying values of the parameters σ , β , and ν and supplying data on GDP per worker, total capital stock, and labor supplied and average human capital of each age group.

6.1 Parameters: β and ν

We choose $\beta = 2/3$ in line with a long tradition of estimates on the share of income going to capital (Gollin, 2002, Hall and Jones, 1996, 1999; Klenow and Rodríguez-Clare, 1997b; Mankiw, Romer, and Weil, 1992; and Caselli, 2005, among many others). We use a value of 1.4 for the markup ν in line with the work of Ramey and Nekarda (2013), Basu (1996), Norbin (1993), Comín (2004), and Jones and Williams (2000) and consistent with similar work by Jerzmanowski and Tamura (2017).

6.2 Parameters: elasticity of substitution, σ

Card and Lemieux (2001) estimate the elasticity of substitution between age groups of workers with the same education level in the US, UK, and Canada, with typical estimates ranging from 4 to 6. Section 5 shows that the Card and Lemieux estimates may be inappropriate in the presence of directed technology and estimated σ with the model setup used here with estimates at approximately 2 (and a lower bound around 1.9 from Appendix A). $\sigma = 2$ is the boundary between strong and weak bias and is consistent with estimates from section 5, so we report estimates with this value.

6.3 Data for constructing experience-earnings profiles

6.3.1 Income, labor force, and physical capital

The data on income, physical capital, populations by age group, labor force size, and human capital are from Tamura et al. (2016) (TDDB), who get their income data mainly from The Maddison-Project (2013). The data include observations on 168 countries over the years 1790 to 2010. Early in the sample, the number of countries represented is small relative to late in the sample (one country in 1790, 58 in 1820, and 168 in 1990). We use only the data from 1820 on because there are so few countries present before 1820. The data from TDDB are available at variable frequencies. We construct decadal data by interpolating¹⁵ the values for countries whose data are for years other than multiples of 10. All values are in year 2000 PPP dollars.

Because these data do not include explicit information on the amount of labor supplied from particular age groups, we assume that persons from every (working age) age group are equally likely to supply labor. Thus if L is the size of the labor force and s_{1524} is the proportion of the working aged population¹⁶ that is aged 15–24, then the size of the labor force for ages 15–24 is taken to be $s_{1524}L$. We also assume that members of the labor force have the same average human capital as the population in their age group.

6.3.2 Human capital

The algorithm for calculating the human capital data is taken from TDDB. It includes four major components:

- Higher average schooling raises the cohort's human capital.
- Higher work experience raises the cohort's human capital (where the size of the effect of experience is based on estimates from Klenow and Rodríguez-Clare, 1997b).
- Intergenerational transfers raise the human capital of a cohort if the previous cohort in that nation had high human capital.

¹⁵Most data years are multiples of 10, so this only matters in a few cases.

¹⁶Working age is 15–64 in this case. Some members of the labor force are outside of this age range, and the validity of this approach depends on the number of such workers being small.

- There are positive externalities on human capital accumulation from the world supply of human capital.

This form reflects the fact that parents with more human capital can invest more in their children’s human capital through efforts outside of the labor market and schooling. The international spillovers reflect the possibility that more investment in human capital in one place may improve the available technology for human capital production.

6.3.3 Work experience

We construct each age group’s synthetic work experience as the midpoint of that age group’s current age bin minus their average schooling minus 6. This produces some negative levels of experience. In order to deal with this when squaring experience, we use instead $[\text{Experience} - \min(\text{Experience})]^2$, where the minimum is taken over all countries and years.

7 Fit of the model

In this section, we provide evidence on the quality of the fit of the model with actual data. For comparison, we look at macro estimates of experience-earnings profiles from Bils and Klenow, 2000 and calculate wage profiles from the Annual Social and Economic Supplement of the Current Population Survey of the United States.

Figure 8 shows our simulated life-cycle wage trends for the United States, France, China, Brazil, Russia, and Jordan (the same patterns for all countries can be obtained from the authors or online¹⁷), which are representative of the countries observed. Each line is the life path of one birth cohort. For instance, the highest curve in the graph for France tells us how the wages constructed for French people born between 1946 and 1955 (marked as “1950”) grow throughout their lives. 1950 is the last cohort that has completed its normal working age. Later cohorts are generally higher in the graph because average national incomes around the world grew over our sample period. Maximum experience falls for later cohorts because they spend more time in school on average. For most cohorts, the model’s predicted earnings rise when the cohort is young and then plateau (and sometimes fall) as the cohort ages, consistent with past estimates of experience-earnings profiles.

It is interesting to note that the analysis does not predict concave experience-earnings profiles at all times in all countries. For instance, the 1950-born and 1930-born cohorts in China see rapid growth of predicted wages throughout their lives as they worked through a period when total factor productivity exploded. This can occur if GDP growth is very high during the cohort’s lifetime. There is nothing inherent in the theory that requires quadratic experience-earnings profiles; they are instead a product of the labor demographics, barriers to innovation, and taxes on capital.

¹⁷<http://randycragun.com/Research/DirectedTech/wexp.html>

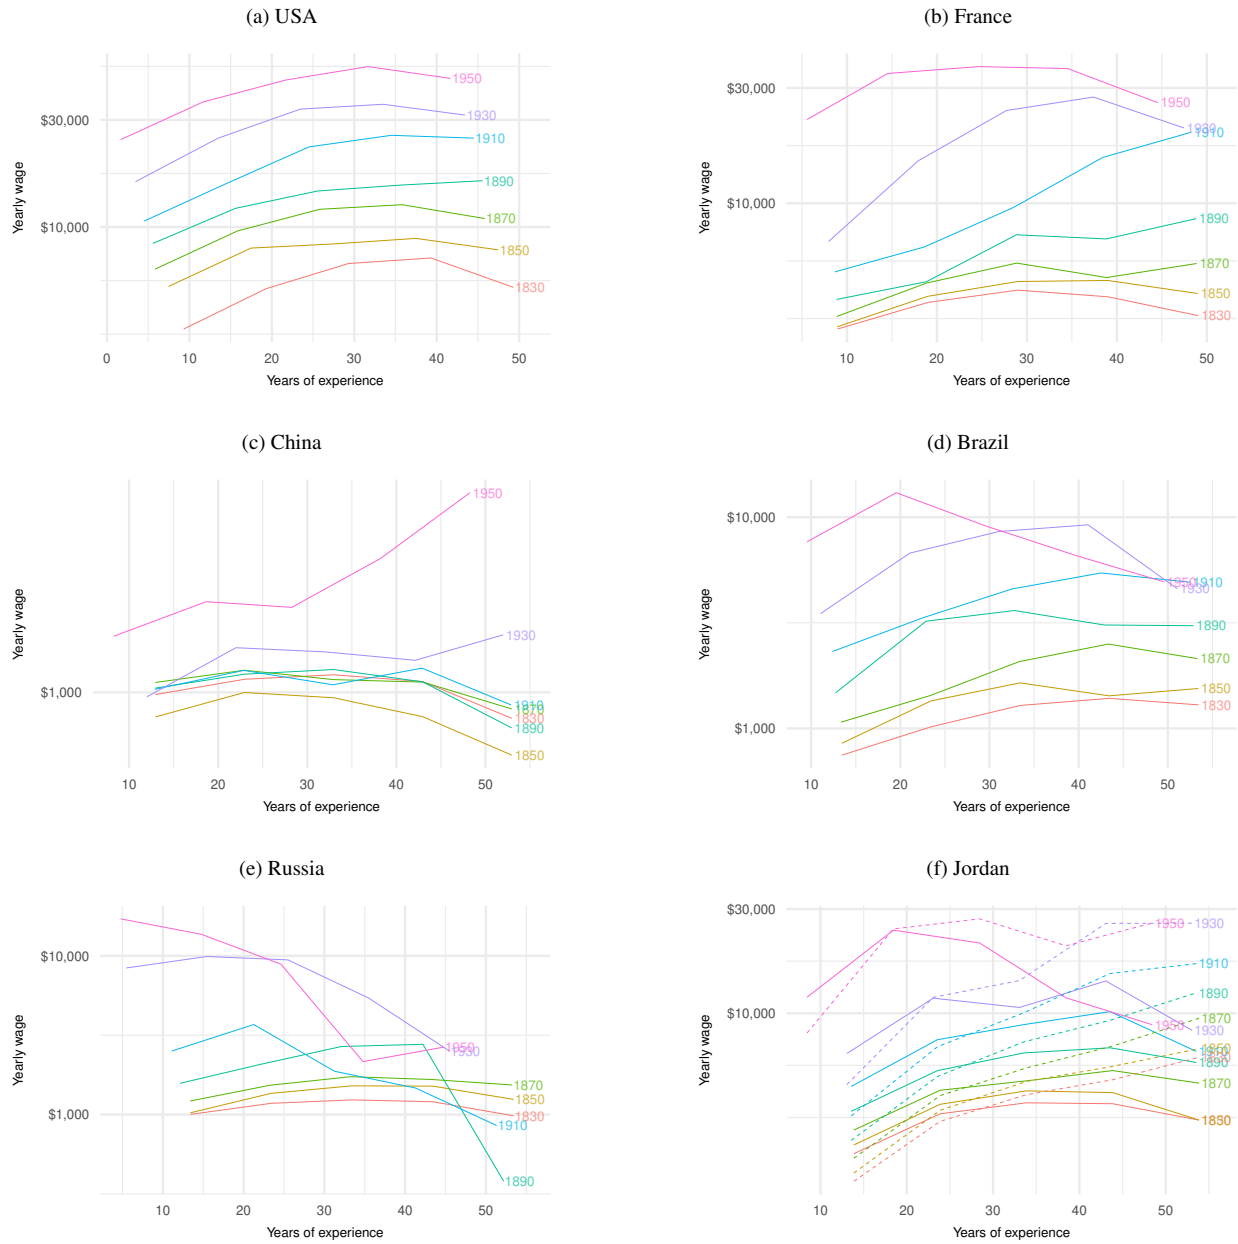


Figure 8: Predicted wages for each cohort throughout their lifetimes as they gain experience. Dashed lines are simulated counter-factual wages without age-specific directed technology.

- Notes: More recent cohorts are typically higher.
- The vertical axis has a log scale.
- The axes do not have the same scale in each subfigure.

We estimate the average experience-earnings profile with a fixed effects estimator of

$$\ln w_{c,t-a,t} = \alpha_{c,t-a} + \gamma_1 \text{XPR}_{c,t-a,t} + \gamma_2 \text{XPR}_{c,t-a,t}^2 + \mu_{c,t-a,t} \quad (39)$$

where c indexes countries and $t - a$ indexes a birth cohort (indicated by its approximate birth year, year t minus age a). Thus α picks up fixed effects for each birth cohort by country group. Table 2 contains the estimates.

Regression (1) of Table 2 contains estimates using the full sample of countries and years from our simulated wages. Our estimates are similar to the estimates of .512 and -.071 for the linear and quadratic coefficients (respectively) using international macro data from Bils and Klenow (2000), which are shown in column (2) and preliminary estimates from an early (1996) copy of that same paper of .0495 and -.0007, which were used in Klenow and Rodríguez-Clare (1997b), Hall and Jones (1999), and TDDDB. Thus our model does a decent job of reproducing empirical experience-earnings profiles.

Table 2: Simulated experience-earnings and empirical experience-wage profiles

Dependent Variable:	Log of wage (hourly or yearly)				
	Data source:	Simulation (full sample)	Bils & Klenow, 2000	Simulation (US only)	ASEC
		(1)	(2)	(3)	(4)
Decades of Experience		0.413 (0.011)	0.512	0.515 (0.014)	0.601 (0.052)
Decades of Experience ²		-0.060 (0.002)	-0.071	-0.078 (0.003)	-0.081 (0.011)
Cohort × Country × Year Groups		10,075		25	25
Cohort × Country FE		YES		YES	YES
R ²		0.172		0.993	0.948

– Standard errors are in parentheses

Figure 9 shows estimates from Regression (3) beside Bils and Klenow’s (2000) values. Note that the heights of the curves are not meaningful here because the heights depend on the values of the covariates observed in the data, which we have simply assumed to be values that place the curves close together in the 10–30 years of experience range in order to draw the figure. Our coefficient on the linear term is slightly smaller than theirs and our quadratic term is slightly closer to zero, making the curve from the estimates in this paper slightly flatter with a later peak.¹⁸

In addition to simulating wages using the TDDDB data, we construct wages and labor supplied for individuals and age groups (15–24, 25–34, 35–44, 45–54, and 55–64) in the United States using wage and salary micro data from the

¹⁸Unlike findings by Murphy and Welch (1990), Murphy and Welch (1992), and Casanova (2013) that older workers do not experience wage decreases as long as they remain in full-time jobs, our estimates from regression 1 place the peak before the end of the usual working life. This might be a weakness of our fit if our goal were to estimate the effect of gaining experience for an individual of average skill who plans to keep working full time until retirement, but we are looking at the average wage by age group rather than the wage that a typical worker from a given cohort would experience. These differ because older workers tend to work part time at disproportionately high rates, and those who choose to work part time or retire early are likely systematically different from those who do not (and working part time reduces wages).




Figure 9: Predicted log wage from Regression (3) in Table 2 and from Bils and Klenow (2000). The heights of the curves are based on arbitrary choices of fixed effects.

Annual Social and Economic Supplement (ASEC) of the Current Population Survey (CPS) from the US Census from 1968 to 2017 (Ruggles et al., 2010).¹⁹ To calculate the wages, we use the same methods as those used to get wages for estimating sigma except that we pool sample years within two years of a decade (in other words, wages from 1988, 1989, 1990, 1991, and 1992 all get counted as wages in 1990) and only use these pools centered at the decades. This allows us to compare the ASEC wages to the wages we simulate at 10-year-intervals.

Figure 10 presents our simulated experience-earnings profiles for a few birth cohorts in the US alongside the empirical experience-wage profiles for the same cohorts estimated from the ASEC. These cohorts are the ones that are observed in the ASEC for at least three decades during working age. The simulated wages match fairly well with the shapes of the estimated ASEC wage profiles but seem to overestimate the increases in levels of wages over time. In this paper, we are interested in how DTC affects the shape of the wage profile within cohorts' lives, so we are not very concerned with the poor fit to wage growth between cohorts.

We estimate the average profiles for these same cohorts with a fixed effects estimator of Equation 39 limited only to the United States so that α picks up fixed effects for each birth cohort. Regression (3) of Table 2 is estimated with the ASEC data, and Regression (4) is estimated with the simulated wages for the US. These coefficients tell us nothing about the levels of wages but do show that the shapes of the estimated and simulated wage profiles are similar.

¹⁹We use only the 5/8 sample for 2014.




Figure 10: Wages for each cohort throughout their lifetimes as they gain experience. Solid lines are our simulated wages, and dashed lines are estimated wages from the CPS ASEC.

Notes: More recent cohorts are typically higher.
The vertical axis has a log scale.

8 Life-cycle earnings with and without DTC

We use the model to produce wages for each age group in each country in each decade both with and without age-group-specific technology. This section presents comparisons of the earnings profiles with age-specific technology to the simulated counterfactual earnings without age-specific technology. We start by showing how the profiles change for select cohorts and countries. Then we estimate the change in the average profile. Finally, we construct the present value of the wage gains from DTC for each cohort and discuss which cohorts win and lose.

8.1 Life-cycle earnings profiles by country and cohort with and without DTC

Figure 11 takes the simulated wages in Figure 8 and adds the simulated counterfactual wages without age-specific technology (the same patterns for all countries can be obtained from the authors or online²⁰). The solid lines are wages simulated with age-specific directed technical change, and the dashed lines eliminate age-specific technology. For most cohorts, age-specific directed technology increases concavity in the experience-earnings profile and particularly hurts the elderly while usually helping the young.

8.2 Average life-cycle earnings profiles with and without DTC

To estimate the impact of DTC on the average experience-earnings profile, consider regressions of the form

$$\ln w_{c,t-a,t} = T_t + \alpha_{c,t-a,\text{NoDTC}} + \sum_{i=1}^4 (\gamma_i \text{XPR}_{c,t-a,t}^i) + \text{NoDTC} \sum_{i=1}^4 (\xi_i \text{XPR}_{c,t-a,t}^i) + \mu_{c,t-a,t} \quad (40)$$

where $t - a$ indexes a birth cohort (indicated by its approximate birth year) in country c in year t and NoDTC is 1 if the wage was simulated with age-specific technology turned off and 0 otherwise. Thus α picks up fixed effects for each birth cohort in each country in each version of the simulation. We estimate both quadratic and quartic versions of this equation, following Murphy and Welch (1990), who show that the quadratic version does a poor job of fitting observed experience-earnings profiles.

We consider specifications both with and without year fixed effects (imposing $T_t = 0 \forall t$). The case with the year fixed effects likely gets closer to a causal effect of experience on wages (rather than letting experience capture generic income growth over time, for instance²¹), but the purpose of this exercise is not to estimate a causal effect of experience on wages. Rather the goal in estimating this model is to determine how age-specific directed technology affects the path of wages over the lifetime. Thus the main interest is in the case without year fixed effects (e.g. $T_t = 0 \forall t$), which helps us to produce life cycle earnings profiles both with and without age-specific directed technology.

²⁰<http://randycragun.com/Research/DirectedTech/wexp.html>

²¹We might also want to condition the regression on age-specific capital per worker if we were trying to show how gaining experience would impact an individual's wage on average. Schooling attainment is generally picked up by the cohort fixed effects.

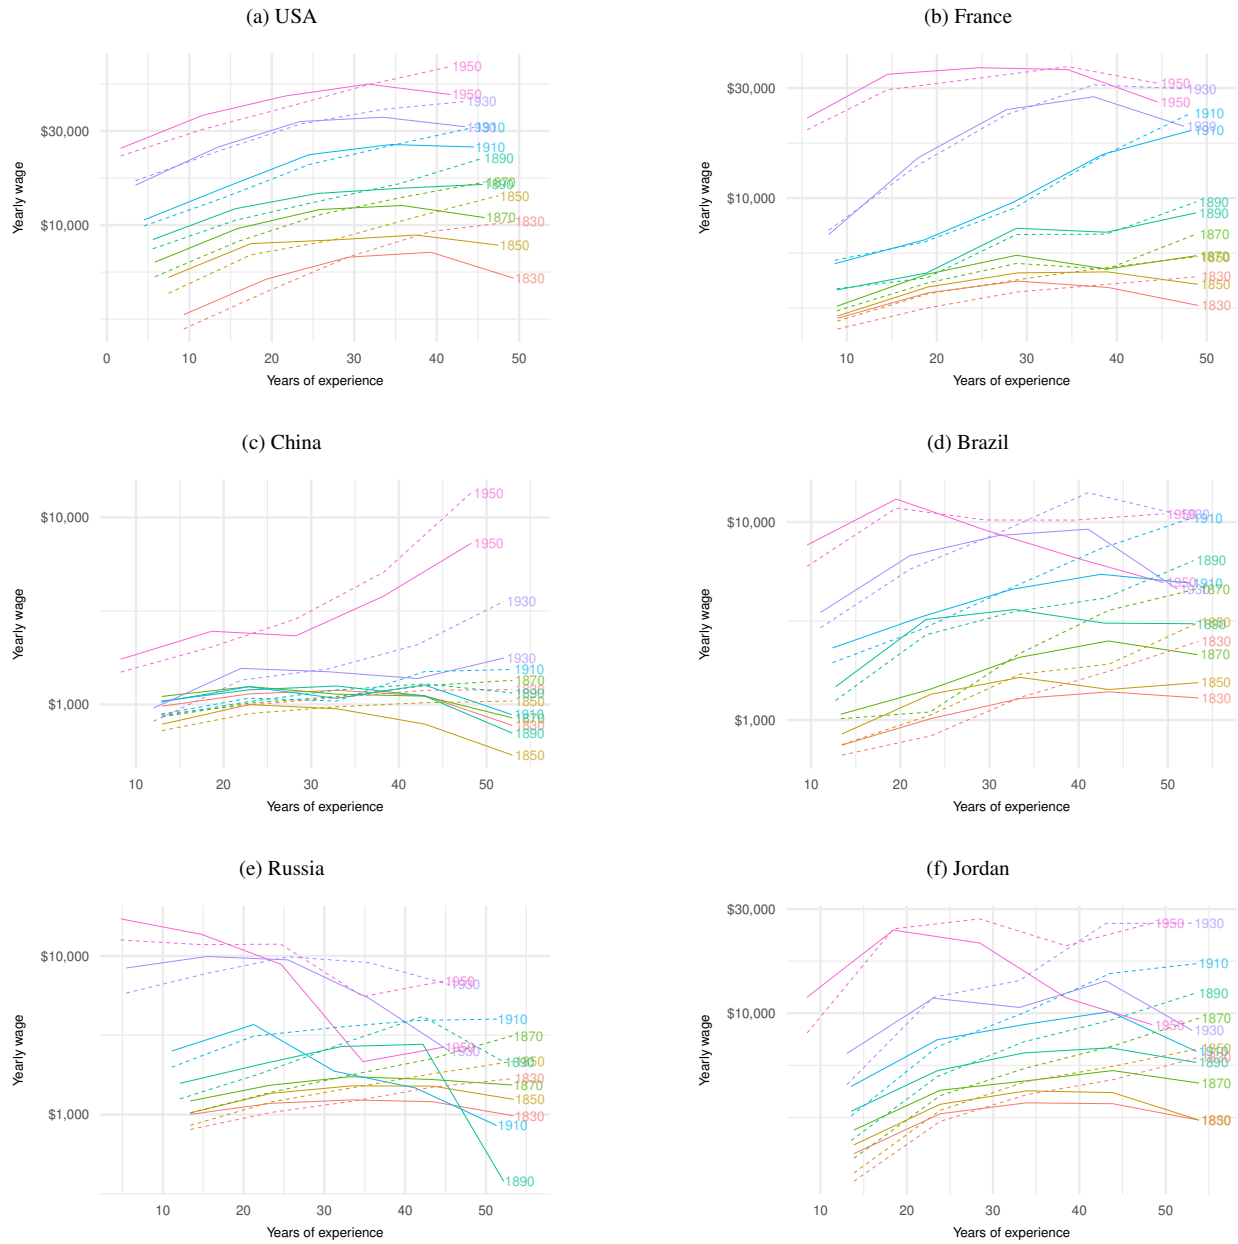


Figure 11: Predicted wages for each cohort throughout their lifetimes as they gain experience. Dashed lines are simulated counter-factual wages without age-specific directed technology.

Notes: More recent cohorts are typically higher.
 The vertical axis has a log scale.
 The axes do not have the same scale in each subfigure.

Estimates of Equation 40 are in Table 3. Regression (1) gives the estimates of the quadratic version without year fixed effects and shows that on average a decade increase in work experience is associated with a proportional increase in real wage of approximately $.405 - .112(\text{Decades of Experience})$ or that a year increase in work experience is associated with a proportional increase in real wage of approximately $.0405 - .00112(\text{Years of Experience})$. Thus workers early in their career can expect to earn approximately 4.05% more after a year, whereas after 40 years in the labor force they could expect a wage loss of .43% raise over the next year. Regression (1) also shows the wage profile using the wages simulated without age-specific technical development. The average profile is less steep and has less curvature, with a slope of $.0305 - .0002(\text{Years of Experience})$. This gets at our primary results: age-specific directed technical change on average shifts earnings earlier in life and helps explain concavity of age-earnings profiles. Figure 12 panel (a) plots the estimates from Regression (1), showing how turning off the DTC mechanism eliminates the wage profile curvature.

The lower coefficients on Experience in regression (2) in the presence of year fixed effects show that generic growth of worker productivity over time accounts for the positive slopes of the curves in Figure 11. However, concavity of the experience-earnings profile remains despite controlling for year effects, and the major result that age-specific directed technical change increases concavity of the profile remains.




Figure 12: Predicted log wage from Regressions (1) and (3) in Table 3. The heights of the curves are based on arbitrary choices of fixed effects.

Estimates of the quartic version of Equation 40 are in Regressions (3) and (4) of Table 3, and we plot the predicted values from Regression (3) in Table 3 against synthetic work experience in Figure 12 panel (b). We again see wage losses later in life and gains earlier in life and an increase in concavity of the profile later in life.

Table 3: Difference between wages with and without DTC as a function of synthetic work experience

	Log of Hourly Wage			
	(1)	(2)	(3)	(4)
Decades of Experience	0.405*** (0.010)	-1.304*** (0.222)	-0.504*** (0.067)	-1.909*** (0.229)
Decades of Experience ²	-0.056*** (0.002)	-0.066*** (0.001)	0.479*** (0.041)	0.292*** (0.038)
Decades of Experience ³			-0.122*** (0.010)	-0.082*** (0.009)
Decades of Experience ⁴			0.009*** (0.001)	0.006*** (0.001)
Decades of Experience (no DTC)	-0.100*** (0.014)	-0.100*** (0.012)	0.183* (0.095)	0.183** (0.085)
Decades of Experience ² (no DTC)	0.046*** (0.002)	0.046*** (0.002)	-0.100* (0.058)	-0.100* (0.052)
Decades of Experience ³ (no DTC)			0.029** (0.014)	0.029** (0.013)
Decades of Experience ⁴ (no DTC)			-0.002* (0.001)	-0.002* (0.001)
Year FE	<i>No</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
Cohort × Country × DTC condition FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Cohort × Country Observations	10,075	10,075	10,075	10,075
R ²	0.469	0.317	0.479	0.324
Adjusted R ²	0.276	0.068	0.290	0.076

Notes:

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

8.3 Which cohorts win and lose from DTC?

The results from the last two sections show how life-cycle wage profiles change due to DTC, but they do not address the value of those changes to the cohorts that actually earn those wages. Workers may be willing to take the wage increases when young at the expense of wage losses when older.

To address this issue, we calculate the present value of lifetime wage gains due to age-specific endogenous technology for each birth cohort for each country by region of the world²² and plot the results in Figure 13. We see the big impact of the baby boom on wages in Western Europe and its offshoots as well as in Japan and some of the rest of Europe, with large baby boom cohorts getting a great deal of technology directed at them in their youth, while we estimate that inter-war cohorts suffered economically-significant losses. Recent cohorts in the East Asian miracles other than Japan did not benefit from age-specific DTC because low population growth produced only small wage gains from DTC in their youth while standard human capital accumulation incentives still mean that cohort human capital accumulation slowed down later in their lives.

9 Discussion and conclusion

This work has attempted to show that directed technical change shapes life cycle earnings profiles in salient ways and can account for some of the concavity in life cycle earnings. This effect occurs because of technology changes due to the relative human capital of cohorts, supplementing other researchers' theories regarding the effects of individual human capital choices on life cycle earnings profiles. We should reiterate here that this report has not shown that directed technology replaces the other theories; what we have shown is that the framework can help explain the data. Some or all of the proposed theories could contribute to observed profiles.

This work has three major implications for estimations of human capital and returns to schooling:

- Because changes in experience are typically negatively correlated with changes in cohort human capital relative to the workforce in late life (and, therefore, with changes in age-specific productive technology), causal estimates of the effect of late-life increases in experience on wages will be downward biased.
- Because individual investments in human capital are correlated with cohort investments in human capital (and, therefore, with changes in productive technology), we may observe an illusory correlation of individual human capital investments with individual wages.
- Individual investments in human capital have external effects on other workers' productivity through the directed technology mechanism, and estimates of returns to education or other human capital investments should include this effect.

²²We again assume an intertemporal discount rate of 4% (see Section 6.1), but increasing the value to 20% does not alter the major patterns.

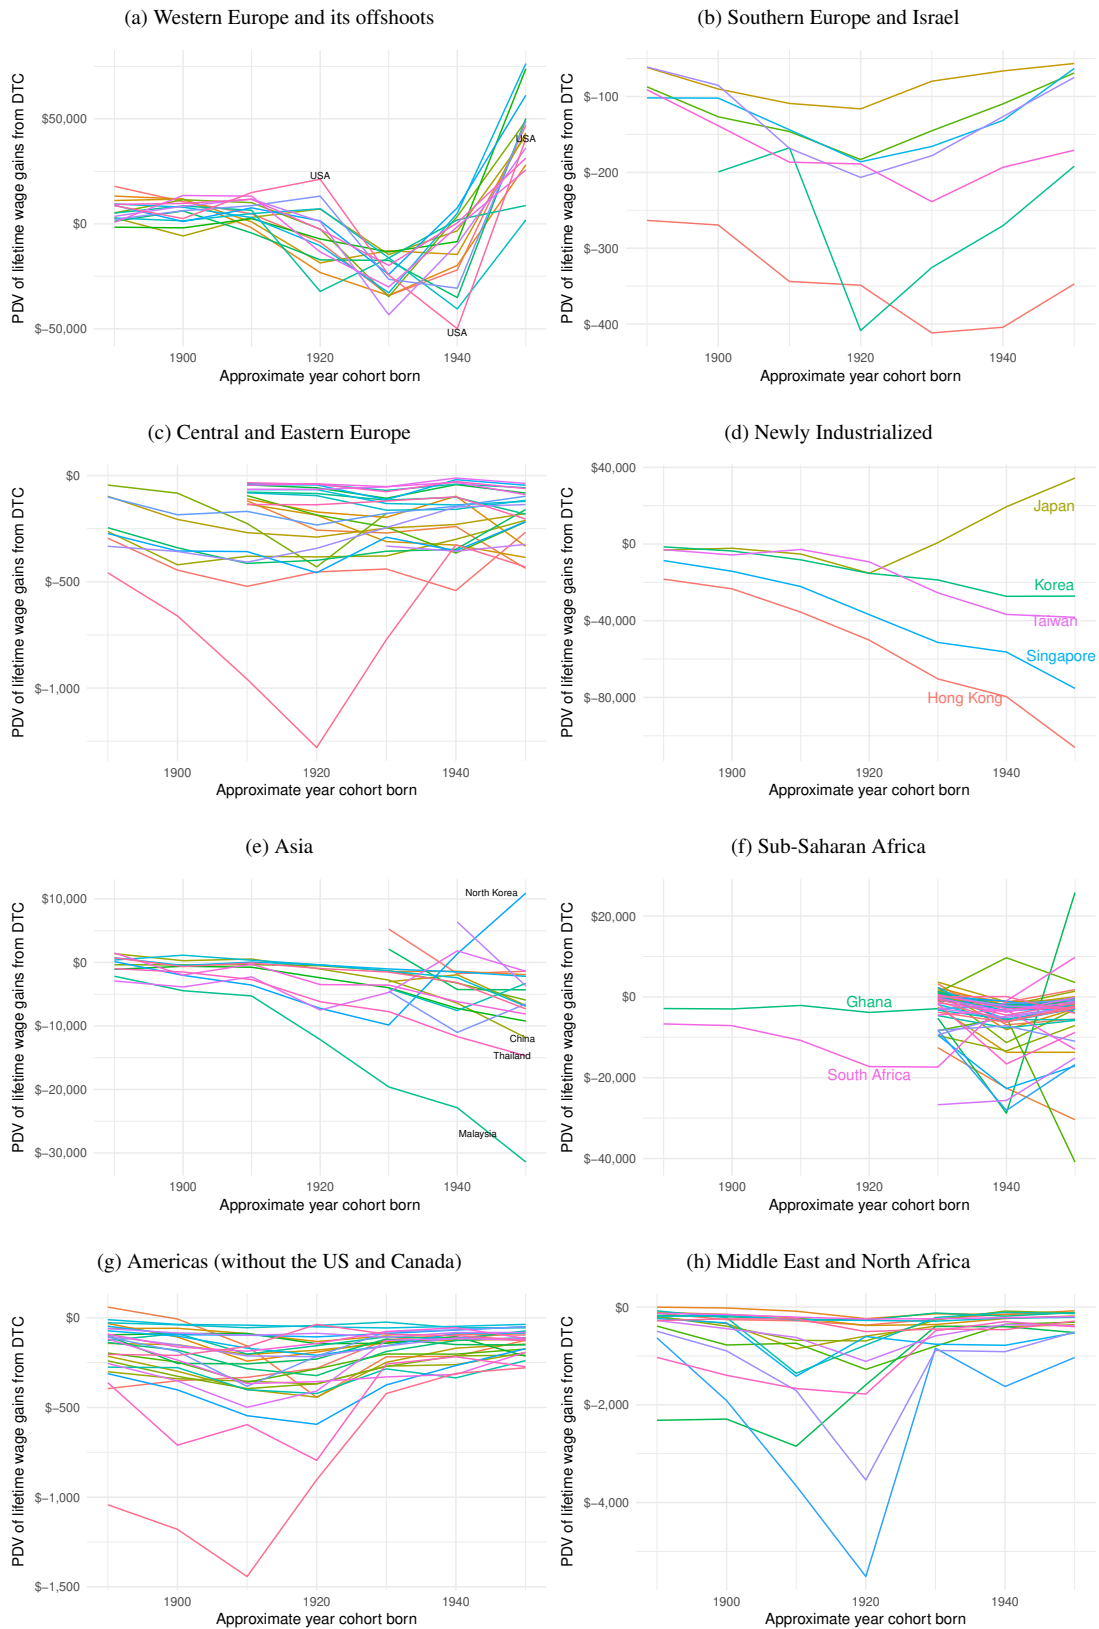


Figure 13: PDV of lifetime wage gains from DTC by birth cohort. Each line is one country.

Consider estimates of human capital derived from fitting life cycle earnings trends. Usually human capital estimates are required to account for almost all of the variation in the mean of wages, but allowing directed technical development to explain some of the variance in wages will produce human capital estimates that change less over the lifetime. This is a problem for the methodology as in TDDB and Klenow and Rodríguez-Clare (1997b), but it is not a problem where entire age-earnings profiles of persons with the *same* age are compared (see Psacharopoulos 1994, for examples).

Another important aspect of this result is in what it says about estimates of the returns to investment in human capital. Past research (following specifications by Mincer, 1974) has attempted to measure the cost of investments at different periods of a worker's career and estimated the returns based on their increased earnings. The actual private returns to investment are the wage gains that would occur with the effects of directed technology (and general improvements in productivity) removed. Because the direction of technical bias is likely positively correlated with rates of human capital accumulation, the results in this report suggest that this past research underestimates the returns to investments in skill acquisition late in life (and overestimates the value early in life). Additionally, the size of this bias will not be constant over cohorts but will depend on heterogeneous demographic differences. Research comparing the lifetime earnings paths of similar workers that do and do not invest does not face this problem (again see Psacharopoulos 1994, for examples) and could help to measure the size of the curvature due to technology effects. Getting these estimates right could be quite beneficial to older workers considering learning new skills.

This research suggests an additional measurement error in past estimates of returns to human capital investment. While researchers have rightly emphasized the importance of the social costs and benefits of education and even of externalities of education on one's cohort (Psacharopoulos, 1994), those estimates often will not sufficiently account for the external effect a worker has specifically on the technology available to her cohort when she invests in human capital.

This work also provides a guide for other researchers wishing to treat the theory of directed technical change as a viable representation of wages and factor productivity for broad classes of productive inputs, and we hope that we have inspired others to demand generality from this theory and others.

References

- Acemoglu, Daron (1998). “Why do new technologies complement skills? Directed Technical change and wage inequality”. In: *Quarterly Journal of Economics* 113, pp. 1055–1090.
- (2002). “Directed technical change”. In: *Review of Economic Studies* 69, pp. 781–809.
- (2007). “Equilibrium bias of technology”. In: *Econometrica* 75, pp. 1371–1410.
- Banerjee, Abhijit and Esther Duflo (1995). “Growth Theory through the Lens of Development Economics”. In: *Handbook of Economic Growth*. Ed. by Philippe Aghion and Steven Durlauf. 1st ed. Vol. 1. 1. Elsevier.
- Barro, Robert and Xavier Sala-i-Martin (1991). “Convergence across States and Regions”. In: *Brookings Papers on Economic Activity* 1, pp. 107–182.
- (1992). “Convergence”. In: *Journal of Political Economy* 100, pp. 223–251.
- (2004). *Economic Growth*. Cambridge, MA: MIT Press, pp. 107–182.
- Basu, Susanto (1996). “Procyclical productivity: Increasing returns or cyclical utilization?” In: *Quarterly Journal of Economics* 111 (3), pp. 719–751.
- Becker, Gary (1964). *Human Capital*. New York: NBER.
- (1967). “Human Capital and the Personal Distribution of Income”. In: *Woytinsky Lecture No. 1*.
- Ben-Porath, Yoram (1967). “The Production of Human Capital and the Life-Cycle of Earnings”. In: *Journal of Political Economy* 75 (4), pp. 352–365.
- (1970). “The Production of Human Capital Over Time”. In: *Education, Income, and Human Capital*. Ed. by W. L. Hansen. Columbia University Press for NBER.
- Bils, Mark and Peter J. Klenow (1996). “Does Schooling Cause Growth or the Other Way Around?” University of Chicago. Mimeo.
- (2000). “Does Schooling Cause Growth?” In: *American Economic Review* 90 (5), pp. 1160–1183.
- Card, David and Thomas Lemieux (2001). “Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis”. In: *The Quarterly Journal of Economics* 116 (2), pp. 705–746.
- Casanova, María (2013). “Revisiting the Hump-Shaped Wage Profile”.
- Caselli, Francesco (2005). “Accounting for Cross-Country Income Differences”. In: *Handbook of Economic Growth*. Ed. by Philippe Aghion and Steven N. Durlauf. Amsterdam: North-Holland.
- Caselli, Francesco and Wilbur John Coleman (2006). “The World Technology Frontier”. In: *The American Economic Review* 96 (3), pp. 499–522.
- Chui, Michael, Paul Levine, and Joseph Pearlman (2001). “Winners and Losers in a North-South Model of Growth, Innovation and Product Cycles”. In: *Journal of Development Economics* 65, pp. 333–365.
- Comín, Diego (2004). “R&D: A small contribution to productivity growth”. In: *Journal of Economic Growth* 9, pp. 391–421.
- Gollin, Douglas (2002). “Getting Income Shares Right”. In: *Journal of Political Economy* 110 (2), pp. 458–474.
- Hall, Robert E. and Charles I. Jones (1996). “The Productivity of Nations”. NBER Working Paper No. w5812.
- (1999). “Why Do Some Countries Produce So Much More Output per Worker than Others?” In: *Quarterly Journal of Economics* 114, pp. 83–116.
- Jerzmanowski, Michal and Robert Tamura (2017). “Directed Technological Change & Cross-Country Income Differences: A Quantitative Analysis”. Working paper.
- Jones, Charles I. and John C. Williams (2000). “Too much of a good thing? The economics of investment in R&D”. In: *Journal of Economic Growth* 5, pp. 65–85.
- Klenow, Peter J. and Andrés Rodríguez-Clare (1997a). “Externalities and Growth”. In: *NBER Macroeconomics Annual 1997* 12, pp. 73–103.
- (1997b). “The Neoclassical Revival in Growth Economics: Has It Gone Too Far?” In: *NBER Macroeconomics Annual 1997* 12, pp. 73–103.
- Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman (2016). “Life-Cycle Wage Growth Across Countries”. In: *Journal of Political Economy*. Forthcoming.
- The Maddison-Project (2013). URL: <http://www.ggd.net/maddison/maddison-project/home.htm>.
- Mankiw, N. Gregory, David Romer, and David N. Weil (1992). “A Contribution to the Empirics of Economic Growth”. In: *The Quarterly Journal of Economics* 107 (2), pp. 407–437.
- Mincer, Jacob (1958). “Investment in Human Capital and Personal Income Distribution”. In: *Journal of Political Economy* 66 (4), pp. 281–302.
- (1974). *Schooling, Experience, and Earnings*. New York: Columbia University Press.

- Mincer, Jacob (1997). “The Production of Human Capital and the Life Cycle of Earnings: Variations on a Theme”. In: *Journal of Labor Economics* 15 (1), S26–S47.
- Murphy, Kevin and Finis Welch (1990). “Empirical Age-Earnings Profiles”. In: *Journal of Labor Economics* 8 (2), pp. 202–229.
- (1992). “The Structure of Wages”. In: *The Quarterly Journal of Economics* 107 (1), pp. 285–326.
- Norbin, Stefan C. (1993). “The relation between price and marginal cost in U.S. industry: A contradiction”. In: *Journal of Political Economy* 101 (6), pp. 1149–1164.
- Psacharopoulos, George (1994). “Returns to Investment in Education: A Global Update”. In: *World Development* 22 (9), pp. 1325–1343.
- Ramey, Valerie A. and Christopher J. Nekarda (2013). “The Cyclical Behavior of the Price- Cost Markup”. In: University of California at San Diego working paper.
- Rosen, Sherwin (1976). “A Theory of Life Earnings”. In: *Journal of Political Economy* 84 (4), S45–S67.
- Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek (2010). *Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]*. Minneapolis: University of Minnesota.
- Tamura, Robert, Gerald Dwyer, John Devereux, and Scott Baier (2012). *Data appendix for economic growth in the long run*. MPRA Paper. University Library of Munich, Germany. URL: <http://EconPapers.repec.org/RePEc:pra:mprapa:41325>.
- (2016). “Economic growth in the long run”. Clemson University working paper.

A The wage-to-labor-supply relationship in the ASEC data

A.1 Non-linearities

It is valuable to graph the data to see what we might expect from the relationship between labor supplied and wages. Figure 14 splits up the data by age group and shows average wages (relative to 15–24-year-olds) against total hours of labor supplied (again relative to 15–24-year-olds) for each sample year. The sample year is represented with color intensity so that later years are lighter colors. If there is an overall relationship between the two variables, it is that there is a slight positive correlation, but we might wonder if the left and right sides of the older workers' data come from different data generating processes.

(Acemoglu, 2007) discusses the dynamics of directed technology and shows that there was an early period of falling skill premiums as the relative supply of skilled workers rose that was followed by a period of positive correlation between the relative supply and relative wage. We may be seeing a similar pattern here. Early samples for some age groups seem to exhibit a negative relationship between relative quantities supplied and relative wages (consistent with movement along a labor demand curve), while later samples are inconsistent with movements along a demand curve. As pointed out by (Acemoglu, 2007), it could take some time before technology changes respond to supply shocks enough to shift out demand, creating an essentially null relationship in later years. These results may weaken the argument of this paper, which relies on the assumption that technology adjusts to its BGP every year. Alternatively, there could be something correlated with labor supplies by age group in later years (e.g. labor supplies by schooling group) that is driving the technical changes. This is a concern that warrants further investigation.

It is also important to note that because Card and Lemieux, 2001 (and Murphy and Welch, 1990, 1992) do not use the later data that we use, that they will see a stronger negative relationships between wages and labor supplies. Hence C&L estimate a positive elasticity of substitution, whereas applying their methodology to the data used here sometimes produces negative estimates for the elasticity of substitution. We leave out the pre-1976 data because using them requires making very strong assumptions about hours of work over the previous year, and attempts to model labor supplied based on the information observed in the pre-1976 samples (similar to in Murphy and Welch, 1990, 1992) produces low-quality estimates. Cursory analysis suggests that including pre-1976 data would lower the estimated elasticities if the pre-1976 data followed a similar trend to the early data presented here.

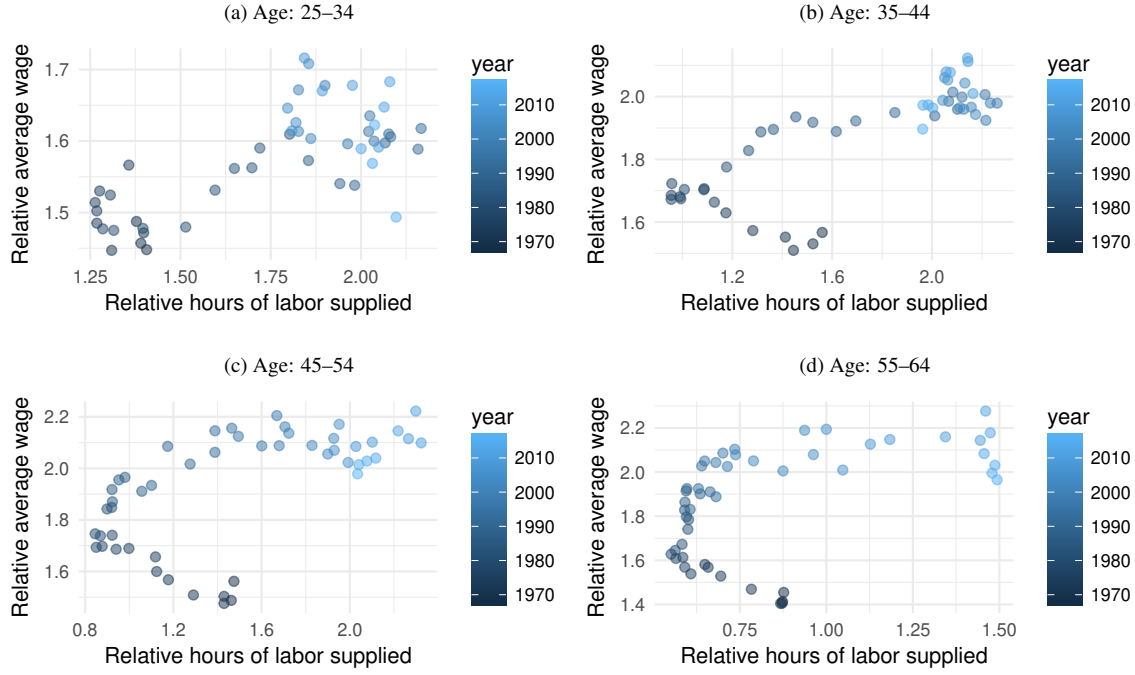


Figure 14: Labor supply and wage data relative to 15–24-year-olds

A.2 Estimates with a structural break

One potential interpretation of the evidence in Figure 14 is that there was a structural break in the data generating process in the late 1980s or early 1990s. This could be due to demographic shifts resulting from the baby boom or increased educational attainment for post-baby-boom cohorts. If we know the timing of the break, we can estimate (where I_{early} is an indicator for being before the break and I_{late} is an indicator for being after the break):²³

$$\begin{aligned}
 \ln\left(\frac{w_{at}}{w_{20r}}\right) &= I_{early} \left[(\sigma_{early} - 1) \ln\left(\frac{h_{at}}{h_{20r}}\right) + (\sigma_{early} - 2) \ln\left(\frac{L_{at}}{L_{20r}}\right) \right] + \\
 & I_{late} \left[(\sigma_{late} - 1) \ln\left(\frac{h_{at}}{h_{20r}}\right) + (\sigma_{late} - 2) \ln\left(\frac{L_{at}}{L_{20r}}\right) \right] + \\
 & I_{early} \alpha_a + I_{late} \alpha_a + \delta_t + \xi_{a,t}
 \end{aligned} \tag{41}$$

Estimating this model by searching over years for the break indicates that if there is a break, it is around 1992, the year that maximizes the R^2 of the model estimates.

Estimates for the model with a break in 1992 are in Table 4. The restricted estimates in Regressions 1 and 4 suggest a negative relationship (“weak bias”) between relative hours worked and relative wages conditional on supply shocks and demand shifts due to directed technology. The unrestricted estimates in Regressions 2 and 5 might suggest strong

²³The age group fixed effects differ in the early and late periods because these are a function of both σ and the productivity of innovation for each age group (see Equation eq:etas).

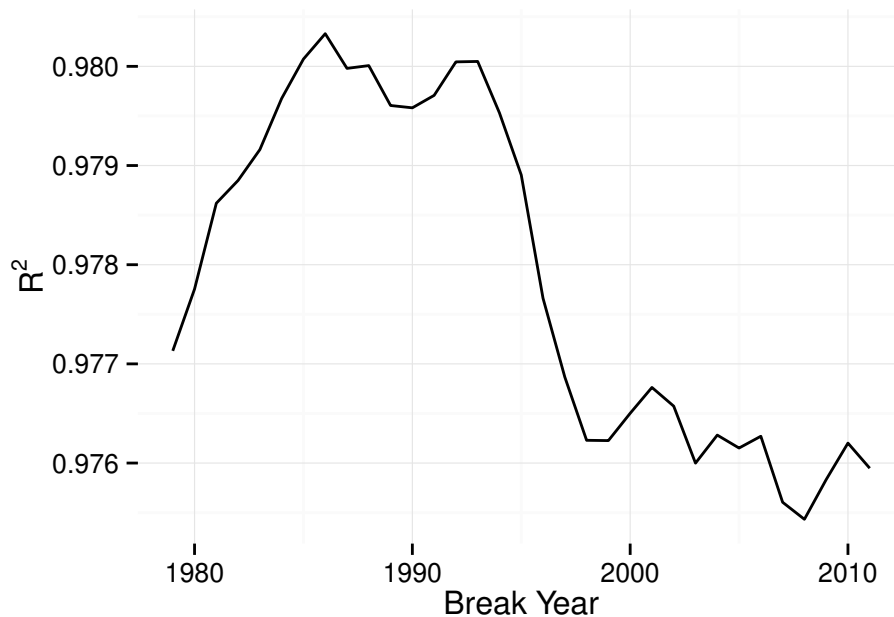


Figure 15: R^2 from estimating the model allowing σ to differ before and after the indicated year

Note: R^2 measures include variation within age groups because changes in the fit of the model within age groups are an important component of the effect of changing the location of the break

bias as in the pooled regressions but only for the coefficients on human capital and only for the early period. However, the main purpose of the unrestricted regressions is to test the restriction on the coefficients on human capital and labor. There seems to be evidence against the restriction (an assumed difference of 1) in both the early and later periods, but the error is in opposite directions between these periods.

A.3 Be cautious in applying point estimates

The estimates of the elasticity of substitution between workers by age group in the report are consistently near 2. The higher values in the regressions with earlier and later data pooled imply “strong bias” under the directed technology framework. This means that as relative labor supply for a group increased, the increase in technology specific to that group was so great that relative demand shifted out enough to raise their wages (again in relative terms).

Our results suggest caution in applying the estimates by Card and Lemieux (2001), since their methodology led to unstable estimates for the elasticity even when the variance of the regression coefficients was small. An alternative interpretation, of course, is that their methodology is sound but that the elasticity of substitution varies greatly (and is sometimes negative) over time and with slight changes in specifications of worker groups.

Readers should also be cautious about applying the estimates here in the wrong context. The regression coefficients should be generally applicable, but how we should interpret those coefficients depends on what we know about the

Table 4: Least squares estimates of σ with a structural break

Dependent Variable:	OLS		1 st Stage	IV 2 nd Stage	
	$\ln\left(\frac{w_{at}}{w_{20t}}\right)$		$\ln\left(\frac{L_{at}}{L_{20t}}\right)$	$\ln\left(\frac{w_{at}}{w_{20t}}\right)$	
	Restricted (1)	(2)	(3)	Restricted (4)	(5)
$\ln\left(\frac{pop_{at}}{pop_{20t}}\right)$ (Early)			1.131 (0.0444)		
$\ln\left(\frac{pop_{at}}{pop_{20t}}\right)$ (Late)			1.180 (0.0222)		
$\ln\left(\frac{h_{at}}{h_{20t}}\right)$ (Early)	0.916 (0.0165)	1.409 (0.0948)	-0.593 (0.2299)	0.932 (0.0148)	1.430 (0.0978)
$\ln\left(\frac{h_{at}}{h_{20t}}\right)$ (Late)	0.934 (0.0207)	0.487 (0.1999)	1.064 (0.1565)	0.934 (0.0206)	0.506 (0.2055)
$\ln\left(\frac{L_{at}}{L_{20t}}\right)$ (Early)	-0.084 (0.0165)	-0.052 (0.0139)		-0.068 (0.0148)	-0.045 (0.0140)
$\ln\left(\frac{L_{at}}{L_{20t}}\right)$ (Late)	-0.066 (0.0207)	-0.039 (0.0213)		-0.066 (0.0206)	-0.044 (0.0220)
Observations	200	200	200	200	200
Age group \times early/late FE	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES
R^2	0.911	0.921	0.335		
σ_h (Early)	1.916	2.409		1.932	2.430
σ_L (Early)	1.916	1.948		1.932	1.955
C&L σ_h (Early)	11.95	-2.44		14.80	-2.32
C&L σ_L (Early)	11.95	19.23		14.80	22.40
σ_h (Late)	1.934	1.487		1.934	1.506
σ_L (Late)	1.934	1.961		1.934	1.956
C&L σ_h (Late)	15.12	1.95		15.07	2.02
C&L σ_L (Late)	15.12	25.70		15.07	22.96
95% CI for diff btwn h and L coefs (Early)		(1.291, 1.643)			(1.301, 1.663)
99% CI for diff btwn h and L coefs (Early)		(1.238, 1.717)			(1.251, 1.742)
99% CI for diff btwn h and L coefs (Late)		(-0.009, 1.067)			(-0.002, 1.105)

– Standard errors are in parentheses and are bootstrap estimates from sampling data de-means by age group
– R^2 and standard errors calculated using data de-means by age group
– Confidence intervals are bootstrap estimates

scenario at hand. Consider, for instance, that the innovation possibilities frontier we proposed may be too simplistic. No nation is totally insulated from the supply of technologies throughout the rest of the world. Suppose that we modified Equation 24 to account for international spillovers of technology:

$$\dot{X}_a = \frac{1}{\zeta} \left(\frac{X_a^W}{X_a} \right)^\phi \frac{R\&D_a}{L} \quad (42)$$

where X_a^W is some function of world technology (possibly the maximum or some average of the highest values for nearby countries). The ratio $\left(\frac{X_a^W}{X_a} \right)$ is included to account for the fact that it is easier for lower-income countries to catch up to the frontier (e.g. Barro and Sala-i-Martin, 1991, 1992, 2004; Caselli, 2005; Caselli and Coleman, 2006; Chui, Levine, and Pearlman, 2001). In this case relative endogenous technologies will take a slightly different form:

$$\frac{X_a}{X_b} = \left(\frac{h_a L_a}{h_b L_b} \right)^{\frac{\sigma-1}{1+\sigma\phi}} \left(\frac{X_a^W}{X_b^W} \right)^{\frac{\sigma\phi}{1+\sigma\phi}} \quad (43)$$

Relative wages will similarly depend on the world technologies and the rate of international spillovers (ϕ):

$$\frac{w_a}{w_b} = \left(\frac{h_a}{h_b} \right)^{\frac{\sigma-\phi-2}{1+\sigma\phi}+1} \left(\frac{L_a}{L_b} \right)^{\frac{\sigma-\phi-2}{1+\sigma\phi}} \left(\frac{X_a^W}{X_b^W} \right)^{\frac{\phi(\sigma-1)}{1+\sigma\phi}} \quad (44)$$

This new form very clearly changes the way we would interpret regression coefficients, so that letting γ_L be the (log-log) regression coefficient on labor supply, the elasticity of substitution, σ , can be recovered as $\sigma = \frac{\gamma_L + \phi + 2}{1 - \gamma_L \phi}$, which is the same as the interpretation in the main body of this report only when $\phi = 0$.

This discussion might seem to imply that we cannot be confident about the elasticity of substitution, and to some degree that interpretation is correct. However, we can be confident in the conclusions we make based on the regression coefficient itself. For instance, the condition on σ that generates strong technical bias when there are no international technical spillovers is $\sigma > 2$. In the presence of international technical spillovers, this condition changes into $\sigma > 2 + \phi$. These inequalities correspond to exactly the same set of regression coefficients. Thus we can say (based on the regressions presented above) that there is evidence that within the US over the last four decades there has been technical bias between age groups that is near the border between strong and weak bias.