

Experience-biased Technical Change

Francesco Caselli¹

Frist Draft: August 2014; This Draft: July 2015

¹London School of Economics, Centre for Macroeconomics, BREAD, CEP, CEPR, and NBER. Email: f.caselli@lse.ac.uk. The material in this paper was part of my discussion of Guvenen and Kuruscu (2009) at the 2009 NBER Macroeconomics Annual conference, but was not included in the published volume. I thank Georg Graetz, Guy Michaels, and Christian Siegel for helpful discussions, and Vincenzo Scrutinio for excellent research assistance.

Abstract

The baby-boom cycle has caused very large swings in the relative supply of experienced workers (first a large decline, and then a large increase). Yet, the experience premium has failed to decline markedly in the period where the supply of experience has increased. I develop a methodology to estimate the increase in the relative demand for experience that is required to reconcile the behavior of prices and quantities, and show this to have been large - a phenomenon I dub experience-biased technical change. I conjecture that one of the drivers of experience-biased technical change is a decline in the relative demand for physical strength. In support this conjecture, I show that occupations requiring high or moderate physical strength have accounted for a declining share of weeks worked in the economy, with sedentary occupations experiencing a corresponding increase. I also confirm that older workers have a comparative disadvantage in occupations requiring physical strength.




Figure 1: Relative prices and quantities of experience, by educational attainment.

1 Introduction

The baby-boom cycle has caused major swings in the relative supply of experienced workers. In the left panel of Figure 1 I plot a measure of the efficiency units supplied by high-school-educated workers with more than 20 years of experience, relative to high-school-educated workers with less than 20 years of experience (in logs). There is a vast decline in the relative supply of experienced high-school-educated workers between 1960 and the late 1980s (reflecting the entrance of the baby-boom generation in the labor market), and an equivalently vast increase from the late 1980s to 2010 (as the baby boomers age). In the right panel, I plot the relative supply of experience for college graduates. The trends for this group are even more pronounced.

As noted by, among others, Katz and Murphy (1992) and Card and Lemieux (2001), workers with different experience are likely to be imperfect substitutes, leading us to expect these vast swings in the relative supply of experience to be mirrored in changes in the relative wage of experienced workers, or “experience premium.” Indeed, Figure 1 confirms that – within both skill groups – the experience premium has increased during the period of falling relative supply of experience, and fallen as the relative supply of experience has rebounded.

Yet, it is also clear from the figure that relative wages have failed to fully respond

to the rebound in relative supplies in the second half of the sample. For the high-school educated, the rebound in supply is of the same magnitude as the original decline, and yet relative wages only drop half-way towards their original level. For the college graduates, the rebound in supply ends up vastly overshooting the original level, and yet the experience premium by the end of the sample is still far above what it was at the beginning.

Clearly supply changes alone are insufficient to explain the joint dynamics of prices and quantities. I suggest, therefore, that the relative demand for experienced workers must have risen since the 1980s. I attribute this increase to *experience-biased technical change*. Experience-biased technical change is likely to have occurred within industries and occupations, as automation (and other advances) have diminished the requirements for physical strength and stamina, which the youth (inexperienced) possess in greater quantity than older (experienced) workers. At the aggregate level, experience-biased technical change may also be the reduced-form implication of structural changes that have diminished the weight of sectors where workers perform physical tasks, for which younger workers are typically more suitable, such as manufacturing.

I propose a simple analytical framework where experience premia respond both to changes in relative supply and to changes in the experience bias of technology. I use this framework to estimate the elasticity of substitution between experienced and inexperienced workers, and to back out the experience bias. Consistent with the observations above, I find that the experience bias has increased markedly both for high-school and college-educated workers since the 1960s, with a sharp acceleration since the late 1980s.

I complement this analysis with new evidence on the evolution of the physical strength required by occupations in the economy. Adapting a methodology first pioneered by Autor, Levy, and Murnane (2003) for a different set of occupational characteristics, I show that there has been a steady decline in weeks accounted for by occupations requiring heavy and moderate physical strength, and a corresponding rise in weeks accounted for by sedentary occupations. Finally, I show that, as conjectured above, older workers do indeed seem to have a comparative disadvantage in occupations requiring physical strength.

The idea of inferring biases in technology from relative wage changes that are unexplained by relative supply changes has famously been pioneered by Katz and Murphy (1992), who also presented an application to the experience premium for the period 1963-

1987. Like me, they found that shifts in relative supplies alone could not fully account for the rise in the price of experience in the 1980s. However, they attributed this to labor-market institutions that shielded older uneducated workers from relative changes in the demand for skills. In other words, they concluded that changes in the price of experience were adequately explained by a combination of changes in the relative supply of experience and skill-biased technical change - hence, with no need for experience-biased technical change. In contrast, I work with a more general formal framework that separately identifies skill-bias technical change and experience-biased technical change, and still find a role for experience-biased technical change.

As mentioned, another important antecedent is Card and Lemieux (2011), who highlight the changes in the relative supply of experience. My analytical framework is very similar to theirs. However, Card and Lemieux do not allow for changes in the experience bias. Moreover, the focus of Card and Lemieux is on differences in the evolution of the college premium for workers with different experience, while my focus is on the behavior of the experience premium. Guvenen and Kuruscu (2009), and Acemoglu and Autor (2011) also investigate differences in the behavior of college premia for younger and older workers.

More recently, Jeong et al. (2012) conclude that there is no need of demand shifts to explain changes over time in the “price of experience” in the US. However, their conceptual framework is very different from mine and their notion of the price of experience does not match well with the experience premium analyzed here.^{1,2}

The paper also relates to (mostly theoretical) contributions on changes in the relative importance of “brain v. brawn.” For example, Galor and Weil (1996) argue that a decline

¹Jeong et al. postulate an aggregate production function defined on two inputs: a “pure labor” input, and an “experience” input. The price of experience is the relative price between these two. In my framework the production function is defined over four inputs: experienced/inexperienced high school/college graduates. The experience premia are the relative wages of experienced workers. It is perhaps possible to argue that my framework, being defined in terms of bodies, poses fewer measurement challenges than the one founded on the abstract notions of the overall supply of “pure labor” and “experience”. Interpretation is also perhaps a bit more straightforward.

²Boehm and Siegel (2014) combine Jeong et al.’s framework with a panel-IV strategy. Unlike Jeong et al. their preliminary results do show a significant role for demand shifts.

in the relative importance of brawn has contributed to the decline in the gender wage gap. Here I suggest that a decline in the importance of brawn may also contribute to increasing experience premia, as older workers are typically not as strong as younger ones..

2 Analytical Framework

My analytical framework is one in which output is produced using a composite labor input, as well as other inputs that do not need to be specified. I will work with the following representation of the composite labor input (which is essentially Card and Lemieux's)

$$\tilde{L}_t = \left\{ [UI_t^{\eta_U} + A_{UEt}UE_t^{\eta_U}]^{\frac{\rho}{\eta_U}} + A_{St} [SI_t^{\eta_S} + A_{SEt}SE_t^{\eta_S}]^{\frac{\rho}{\eta_S}} \right\}^{1/\rho}, \quad (1)$$

where UI , UE , SI , and SE denote the quantities of high-school educated, inexperienced inputs, high-school educated, experienced inputs, college-educated, inexperienced inputs, and college-educated experienced inputs, respectively. For economy, from now on I will refer to high-school (college) educated workers as unskilled (skilled), hence the U (S) notation.

The time-invariant coefficients η_U and η_S govern the elasticity of substitution between unskilled inexperienced and unskilled experienced workers, and skilled inexperienced and skilled experienced ones, respectively. The parameter ρ governs the elasticity of substitution between unskilled and skilled workers. Finally, the time-varying coefficients A identify non-neutralities in technological change: A_{UEt} , and A_{SEt} capture the “experience bias” within the unskilled and the skilled group, respectively; A_{St} captures the skill bias. The goal is to characterize the time-series behavior of these A s, and particularly of the experience biases, as the skill bias is already the focus of a huge literature.

Assuming perfectly competitive labor markets, we can derive the following formulas for the *experience premia*:

$$\frac{w_{UEt}}{w_{UIt}} = A_{UEt} \left(\frac{UE_t}{UI_t} \right)^{\eta_U - 1}, \quad (2)$$

$$\frac{w_{SEt}}{w_{SI_t}} = A_{SEt} \left(\frac{SE_t}{SI_t} \right)^{\eta_S - 1}. \quad (3)$$

It is clear from these equations that – with data on relative wages and relative quantities (i.e. the ones plotted in Figure 1) – we can back out the time-series path of the experience

biases if we have estimates of the elasticities of substitutions. In particular, we will observe a rising experience bias if the experience premium drops “too slowly” in the face of a rapid rise in the relative supply of experience.

Card and Lemieux (2011) essentially estimate η_U and η_S from (2) and (3) by taking logs and assuming that A_{UEt} and A_{SEt} are constant. Clearly this is not an option if we want to allow for changes over time in the experience bias. I therefore propose the following alternative. Begin by assuming

$$A_{UEt} = \chi A_{SEt} \omega_t,$$

where ω_t is i.i.d. In other words, assume that the experience biases have a common trend for skilled and unskilled workers. With this assumption, we can combine the two expressions for the experience premium into a Difference in Difference specification:

$$\log \frac{w_{SEt}}{w_{SI_t}} - \log \frac{w_{UEt}}{w_{UI_t}} = \alpha + (\eta_S - 1) \log \frac{SE_t}{SI_t} - (\eta_U - 1) \log \frac{UE_t}{UI_t} + \varepsilon_t, \quad (4)$$

which can be estimated by OLS.³ With estimates of the elasticities η_S and η_U at hand, we can return to equations (2) and (3) and solve them for the experience biases A_{UEt} and A_{SEt} .⁴

While A_{UEt} and A_{SEt} are the key focus, for completeness we can also use the framework to generate new estimates of the time series of the skill (or college) premium. For inexperienced and experienced workers the college premium is given by

$$\frac{w_{SI_t}}{w_{UI_t}} = A_{St} \frac{[SI_t^{\eta_S} + A_{SEt} SE_t^{\eta_S}]^{\frac{\rho}{\eta_S} - 1} SI_t^{\eta_S - 1}}{[UI_t^{\eta_U} + A_{UEt} UE_t^{\eta_U}]^{\frac{\rho}{\eta_U} - 1} UI_t^{\eta_U - 1}}, \quad (5)$$

$$\frac{w_{SEt}}{w_{UEt}} = A_{St} \frac{[SI_t^{\eta_S} + A_{SEt} SE_t^{\eta_S}]^{\frac{\rho}{\eta_S} - 1} A_{SEt} SE_t^{\eta_S - 1}}{[UI_t^{\eta_U} + A_{UEt} UE_t^{\eta_U}]^{\frac{\rho}{\eta_U} - 1} A_{UEt} UE_t^{\eta_U - 1}}. \quad (6)$$

Hence, we can back out the skill bias A_S either from data on college premia among experienced workers, coupled with data on relative supplies *augmented* with our estimates of

³The error term in (4) contains temporary deviations of A_{UE} from the common trend with A_{SE} . Given the annual nature of the data, it does not appear plausible that the relative supplies on the right-hand side can respond endogenously to such deviations.

⁴I run equation (4) with and without the inclusion of a time trend, and find that the coefficients are virtually identical. The coefficient on the time trend is close to zero and insignificant.

the (skilled) experience bias, or from data on college premia among inexperienced workers. The only additional input required is an estimate of the elasticity of substitution between skilled and unskilled workers, $1/(1 - \sigma)$. There is now a large body of work focusing on this coefficient. A particularly compelling estimate is due to Ciccone and Peri (2005), who are able to deploy an instrumental-variable approach, and find a value close to 1.4 (in line with the vast majority of estimates).

As noted by Card and Lemieux (2001), the skill premium for a given experience group depends on (i) the overall skill bias in technology, (ii) the overall relative supply of skills, and (iii) the relative supply of skills specific to the given experience group. What I have added here is that technology can have an experience bias, while Card and Lemieux only allow for a skill bias. In this sense, my approach is closer to Katz and Murphy's (1992) strategy to back out the rate of *skill-biased* technical change from the behavior of the *college* premium. However, Katz and Murphy simultaneously estimate the elasticity of substitution between skilled and unskilled workers and the rate of change in the skill bias, by exploiting an assumption that the latter is roughly constant over time. Figure 1 suggests that the rate of change in the experience bias has been time varying. This motivates my two-step procedure where I first estimate elasticities of substitution between experienced and inexperienced workers, and then back out the relative efficiency of experienced workers in each year of the sample. The ability of getting year-by-year estimates of the technology bias makes the paper methodologically closer to Caselli and Coleman (2002) (who focus on the skill bias).⁵

3 Data

I need time series data on the labor supplies UI_t , UE_t , SE_t , SI_t , and the corresponding wages w_{UI_t} , w_{UE_t} , w_{SE_t} , w_{SI_t} . I construct these series from data developed by Acemoglu and Autor (2011), henceforth AA, using the 1963-2008 March CPS samples.

AA make available a variable measuring total annual hours of labor by gender, 5 educa-

⁵In turn, the method in Caselli and Coleman (2002) is a time-series application of the cross-country studies in Caselli and Coleman (2006) and Caselli (2005).

tion categories, and 48 experience categories (i.e. from 0 to 48 years of experience). I define “inexperienced” all workers with 19 or less years of experience, and “experienced” those with 20-to-48 years of experience.⁶ I further define as “unskilled” all high-school dropouts, high-school graduates, and workers with incomplete college (education categories 1-3). The “skilled” are those with a college or a post-graduate degree (education categories 4-5).

AA also compute the average weekly full-time equivalent earnings within each of these gender-education-experience cells. I pick male, high-school graduates with 10 years of experience as the reference group for the unskilled-inexperienced category, male high-school graduates with 30 years of experience as benchmark for the unskilled-experienced group, and male college graduates with 10 and 30 years of experience as baseline for skilled-inexperienced and skilled-experienced. Then, for each gender-education-years of experience cell I construct a fixed weight given by mean earnings in that cell relative to the relevant benchmark mean earnings, averaged over the sample period. The idea of these weights is that they represent an efficiency-unit conversion factor to express hours supplied by a given cell into hours supplied by the reference cell within the education-experience category. Using these weights, I construct UI_t , UE_t , SE_t , SI_t as weighted sums of the hours supplied by each gender-education-years of experience cell within each of the four broad education-experience categories.

For each annual data set, AA also regress individual log weekly wages on 5 dummies corresponding to the five levels of educational attainment, a quartic in experience, a gender dummy, a race dummy, and several interactions of these variables. They then construct predicted real log weekly wage series for white workers by gender, 5 educational attainment categories, and 5 levels of experience, namely 5, 15, 25, 35, and 45 years. I simply take the predicted wage series for high-school graduates (college graduates) as representative of the unskilled-category, and the series for workers with 5 (25) years of experience as representative of the experienced category. This gives me (logs of) w_{UI_t} , w_{UE_t} , w_{SE_t} , w_{SI_t} .

⁶19 years of experience is the (hours of labor supply weighted) average over time in the CPS (the unweighted average is 22). Corresponding medians are 18 and 22.




Figure 2: Experience-biased technical change

4 Results

The OLS coefficients (standard errors) from estimating regression (4) are $(\eta_S - 1) = -.342$ (.046) and $(\eta_U - 1) = -.303$ (.054). These imply that the elasticities of substitutions are

$$\frac{1}{1 - \eta_U} = 3.3 \text{ and } \frac{1}{1 - \eta_S} = 2.9,$$

with standard errors 0.586 and 0.392, respectively.⁷

Given these estimates, and the relative quantities and relative wages displayed in Figure 1, the experience biases A_{UEt} and A_{SEt} are plotted (in logs) in figure 2. Interestingly, the figure reveals a positive experience bias through most of the sample period, though clearly there is an acceleration since 1980.

As we have seen in the Introduction, since 1980 or so, the relative supply of experience has increased very markedly (in both skill groups), and yet experience premia have declined little. Even with our relatively large estimated elasticities of substitution between skilled and unskilled workers, the experience premium relative stability in the face of a large increase in the relative supply of experience must imply that technological change has been experience biased. Before 1980, the relative supply of skills was declining in both

⁷The fit of the regression is reasonable, with an R-squared of 0.57 and a mean square error of 0.05.

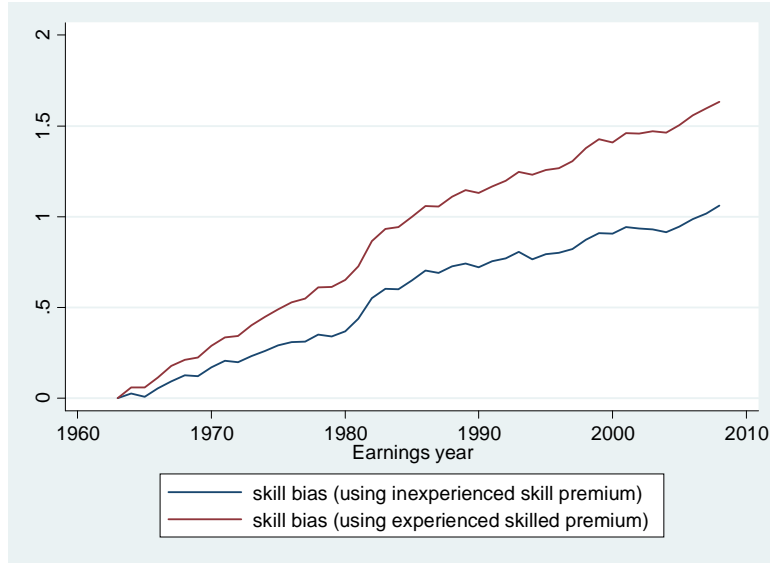


Figure 3: Skill-biased technical change

groups, and both groups duly experienced an increase in the experience premium. For the unskilled, the increase in the premium was roughly what one would expect given the estimate of η_U , so the experience bias is relatively flat. For the skilled, the increase in the premium is actually *greater* than what one would expect given η_S , leading to a positive trend in the skilled experience bias for the early sub-period as well.

Figure 3 plots the two variants of A_s (in logs, after normalizing to one in the first period) backed out from equations (5) and (6), respectively. The skill bias implied by the experienced skill premium shows a larger increase than the skill bias implied by the inexperienced skill premium (which clearly implies that the model does not fit the data quite perfectly), but clearly in both cases there is a pronounced upward trend, confirming the conclusions of the large SBTC literature. Indeed, not surprisingly given the attention that SBTC has received, quantitatively SBTC is one order of magnitude larger than experience-biased technical change.

5 The Demand for Physical Strength

As mentioned in the introduction, a possible mechanism for experience-biased technical change is that the kind of tasks workers are performing are becoming less demanding in terms of physical strength. Because the young are physically stronger and have greater

endurance, this would result in a decline of their marginal productivity relative to older workers. In this section, I document changes in the demand for physical strength in the US economy.

In a pioneering study, Autor, Levy, and Murnane (2003) have documented significant changes in the composition of demand between "routine manual," "non-routine manual," "routine cognitive," "non-routine analytical", and "non-routine interpersonal" tasks. It is difficult to use their results to infer the extent to which the demand for physical strength has declined. The input of analytical and interpersonal tasks has increased, and these likely demand little strength, but so do routine cognitive tasks, whose input has declined [see also Autor and Price (2013)]. And it is hard to tell how the strength content of manual tasks has changed over time. Hence, a separate look at changes in the "strength-content" of the tasks performed by US workers seems warranted.

The Dictionary of Occupational Titles (1977) [DOT] contains a variable that describes, for each occupation, the amount of physical strength required of the worker. This is a categorical variable taking 5 values: sedentary, low, medium, heavy, very heavy. In the appendix I describe how, using a cross-walk from the Current Population Survey, I am able to map the list of occupations in the DOT into the list of occupations in the US Census. Hence, I am able to place each respondent in each census year in one of the five DOT strength-intensity categories. For each of the five categories I can then construct, for each year, the share of full-time equivalent weeks worked by workers in the relevant category in the total number of weeks worked by US workers. The result is displayed in Figure 4.

The most apparent trend is the substantial increase in the fraction of weeks in sedentary occupations, which has increased by about 50%, with a corresponding decline in weeks accounted for by medium and high physical-demand occupations.

It seems plausible that the decline in the demand for strength will have favored older, less physically strong workers. I check on the assumption that older workers have a comparative disadvantage at tasks requiring physical strength in Table 1. The table presents a battery of ordered probit regressions where the dependent variable is the five-valued indicator of physical-strength requirement in the worker's occupation, and explanatory variables include worker's age, gender, and controls for race and education. Both when all census years are pooled together (with and without year effects), and in each individual

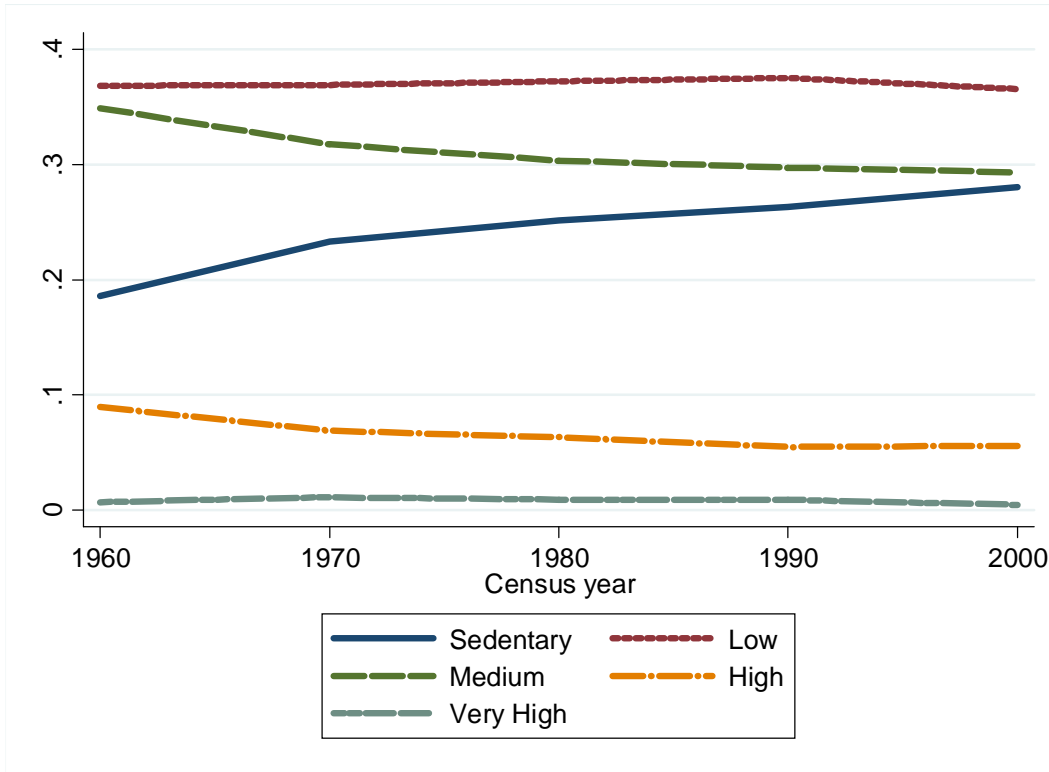


Figure 4: Changes in the composition of tasks by physical demand.

census year, age is a consistently negative predictor of the physical demands of a worker’s occupation. The results also confirm the widely-held presumption that women have a comparative disadvantage at tasks involving heavy physical requirements.

To quantify the effect of age on the physical requirements of tasks performed, Figure 5 plots the probability that a white, male high school graduates (left panel) and a white, male, college graduate has an occupation in each of the five strength-requirement classes, as implied by the coefficients in Column 3 of Table 1 (results for high-school dropouts and post-graduates, as well as women and non-whites, are qualitatively similar, and available upon request). The probabilities are normalized to those applying to 20 years old. Both high-school and college graduates are 40% less likely to be performing a job with “very high” physical demands when age 60 than when age 20. 60 year-old high-school (college) graduates are also 40% (30%) more likely to be in a sedentary occupation.

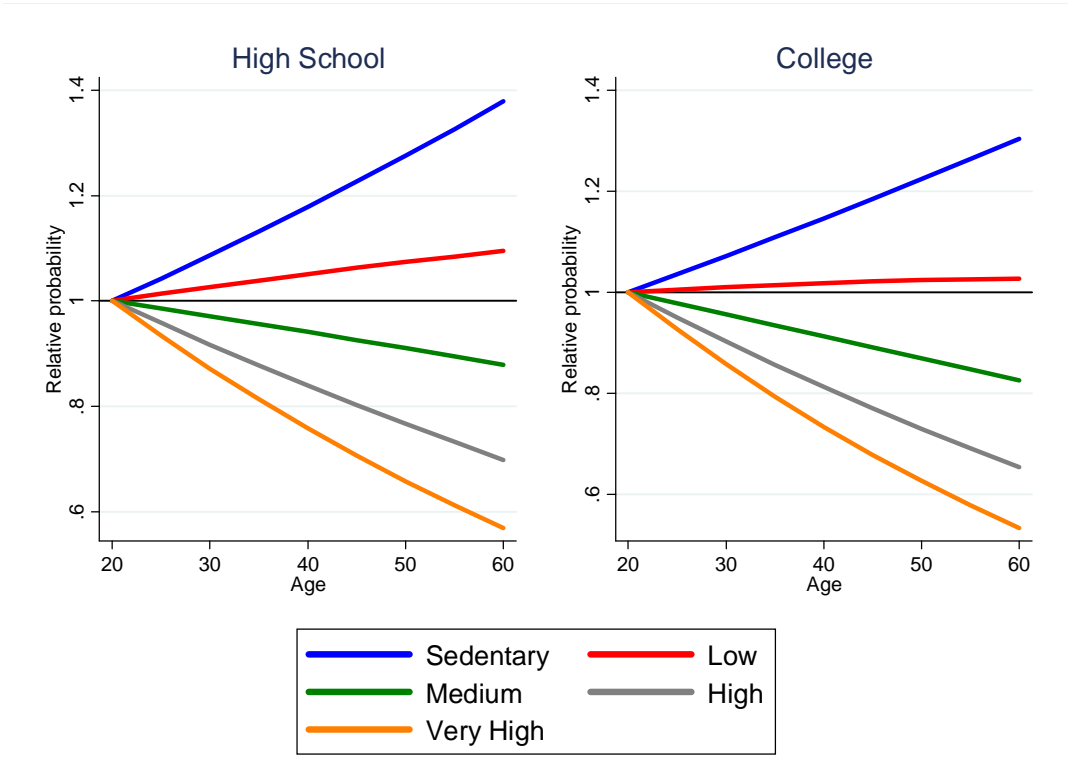


Figure 5: Probability of having occupations with different strength requirements as a function of age

Table 1: Age and Physical Demands of Occupation

	Census Years							
	All	All	All	1960	1970	1980	1990	2000
age	-0.00223*** (4.42e-05)	-0.00485*** (4.54e-05)	-0.00485*** (4.54e-05)	-0.00566*** (0.000126)	-0.00517*** (0.000106)	-0.00770*** (9.90e-05)	-0.00615*** (9.76e-05)	-0.00156*** (9.26e-05)
female	-0.639*** (0.00116)	-0.639*** (0.00116)	-0.768*** (0.00333)	-0.737*** (0.00285)	-0.674*** (0.00258)	-0.643*** (0.00237)	-0.550*** (0.00225)	
Obs.	3,821,722	3,821,722	3,821,722	519,740	661,129	770,882	895,880	974,091
Controls	NO	YES	YES	YES	YES	YES	YES	YES
Year FE	NO	NO	YES	YES	YES	YES	YES	YES

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Note: Individual controls include a dummies for education (high school dropout or less, high school, college, higher than college) and race dummies

6 Conclusions

In sum, I confirm many previous findings of a significant skill bias in technology change over the last 50 years. In addition, I present (novel, I believe) evidence of an experience bias in technological change over roughly the same period, but especially since the 1980s. Further, I present novel evidence of the decline of weeks devoted to tasks requiring physical strength, and a corresponding increase in sedentary work. Finally, I show that older workers have a comparative advantage in occupations requiring little physical strength.

It is of course quite possible that the simultaneous occurrence of experience-biased technical change and increased relative supply of experience since the 1980s is coincidental. However, it is also possible to advance a causal explanation. In models of endogenous technology choice (e.g. the model in Caselli and Coleman, 2006), or directed technical change (e.g. Acemoglu, 1998, 2002), technology choices are biased towards the relatively abundant factor whenever factors of production are good substitutes with each other. The good-substitutes assumption is clearly satisfied here, so one could argue that the experience-biased in technology is an endogenous response to the aging of the baby boom generation.

It is true that the supply of experience declined in the 1960s and 1970s, and we do not

observe a corresponding decline in the relative efficiency of experienced labor over this earlier period. To the contrary, the experience bias in technology appears to have begun before the turnaround in relative supplies, albeit at a slower pace than in the post-1980s period. One possible explanation, though, is that the subsequent, demographic driven, reversal in relative supplies was predictable. Firms that were aware of the coming acceleration in the relative supply of experience would probably not have wanted to temporarily switch to inexperienced-biased technologies. Indeed, they may have wanted to begin readying themselves for the impending aging of the labor force.

7 References

Acemoglu, Daron (1998): “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality.” *Quarterly Journal of Economics*, 1998, 113 (4), pp. 1055–89.

Acemoglu, Daron (2002). “Directed Technical Change.” *Review of Economic Studies* , 2002, 69 (4), pp.781–809.

Acemoglu and Autor (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” *Handbook of Labor Economics*, Volume 4, Orley Ashenfelter and David E. Card (eds.), Amsterdam: Elsevier.

Autor, David H., Levy, Frank, and Richard J. Murnane (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration,” *Quarterly Journal of Economics*, 1279-1333.

Autor, David H., and Brendan Price (2013): “The Changing Task Composition of the US Labor Market: An Update of Autor, Levy, and Murnane (2003), unpublished, MIT.

Boehm, Michael, and Christian Siegel (2014): “The Race between the Demand and Supply of Experience,” PDF slides.

Card, David, and Thomas Lemieux (2011): “Can Falling Supply Explain the Rising Return to College of Younger Men? A Cohort-Based Analysis,” *Quarterly Journal of Economics*, May, 705-46.

Caselli, Francesco (2005): “Accounting for Cross-Country Income Differences,” *Handbook of Economic Growth*, 679-741.

Caselli, Francesco, and John Coleman (2002): “The US Technology Frontier,” *American Economic Review P&P*, May.

Caselli, Francesco, and John Coleman (2006): “The World Technology Frontier,” *American Economic Review*, 93, 3, June, 499-522.

Ciccone, Antonio and Peri, Giovanni: “Long-Run Substitutability between More and Less Educated Workers: Evidence from U.S. States 1950 –1990,” *Review of Economics and Statistics*, 2005, 87 (4), pp. 652– 63

Galor, Oded, and David Weil (1999): “The Gender Gap, Fertility, and Growth,” *American Economic Review*, 86, 3, 374-387.

Guvenen, Fatih and Burnahettin Kuruscu (2009): “A Quantitative Analysis of the Evolution of the U.S. Wage Distribution: 1970-2000,” *NBER Macroeconomics Annual*, 24, 227-276.

Katz, Lawrence F. and Murphy, Kevin M. “Changes in Relative Wages, 1963–1987: Supply and Demand Factors.” *Quarterly Journal of Economics*, 1992, 107(1), pp. 35–78.

APPENDIX

The construction of our strength-content variables combines information from two publicly-available data sets:

1. The *Current Population Survey, April 1971, Augmented With DOT Characteristics* available at <http://www.icpsr.umich.edu/icpsrweb/ICPSR/studies/7845>.
2. The *Census 1pc Sample* for the years 1960-2000, available at <https://usa.ipums.org/usa/index.shtml>.

The first data set contains data from the CPS, together with data from the Dictionary of Occupational Titles. The key variables in this data set are two variables coding occupation, one according to the DOT (henceforth "DOT occupation"), and one according to the list of occupations used in the 1970 census (henceforth "1970 census occupation"); the variable for "Physical Demand", which is a 5-value categorical indicator of strength required in each of the DOT occupations; a measure of hours worked in the previous week; and CPS person weights.

Each census sample contains a measure of weeks worked in the previous year, a variable storing occupation according to a census-year specific list of occupations, a variable storing occupation according to the 1990 census list of occupations (henceforth "1990 occupation"), and census person weights.

We begin by using the 1970 census sample to create a cross-walk from the "1970 census occupation" to the "1990 census occupation". Each 1970 occupation maps into one 1990 occupation so there are no 1970 occupations that we cannot turn into 1990 occupations.

After importing this cross-walk into the DOT-augmented CPS file, we begin the task of assigning each 1990 census occupation to one of the five levels of physical demand. Because the list of "DOT occupations" is much finer than the list of "1990 census occupations", DOT occupations with different strength requirement map into the same 1990 census occupation. To solve this problem, for each level of physical demand I construct the share of (CPS-weighted) hours accounted for by DOT occupations characterized by that level of physical demand, in the total (CPS-weighted) hours of each 1990 census occupation. Then, the physical demand of a 1990 census occupation is the level of physical demand that accounts for the largest fraction of hours in that occupation. Since the distribution of DOT occupations within a 1990 census occupation varies significantly for men and women,

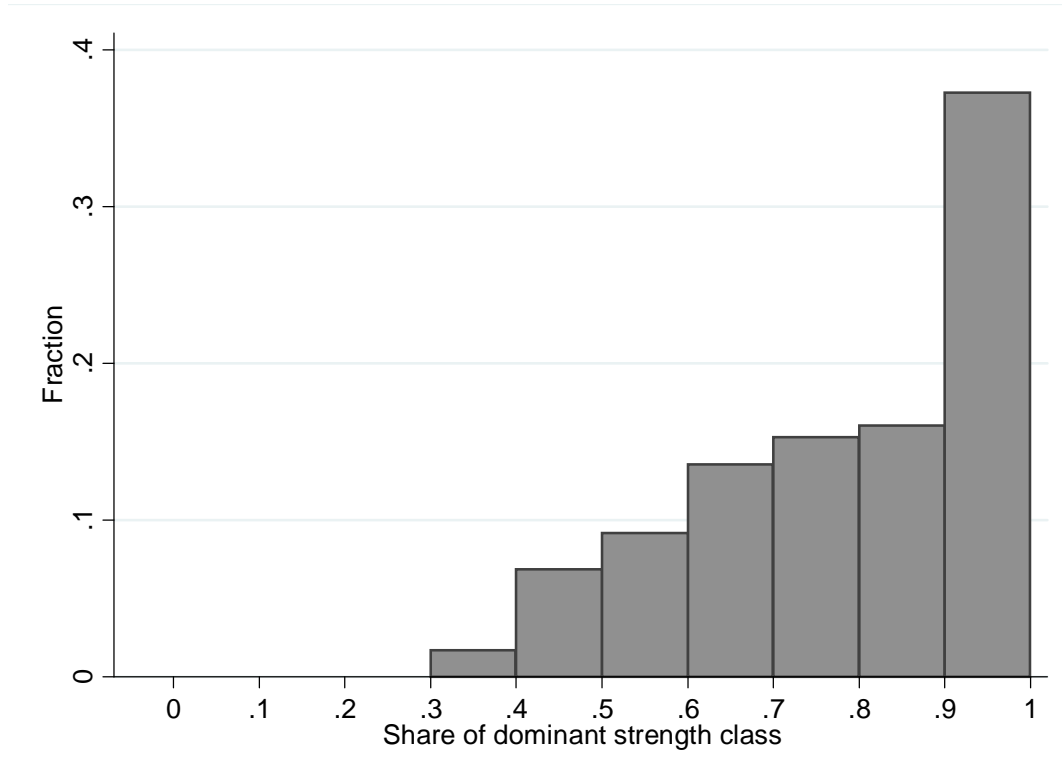


Figure 6: Distribution of the share of hours of the dominant strength-requirement level by occupation

I do this by gender. In other words for each 1990 census occupation there will be one level of physical demand for men and one for women.

Figure 6 shows the distribution of the largest share of physical demand across 1990 census occupation-gender cells. It looks like for the vast majority of cells the dominant level of physical demand accounted for a very large fraction of total hours.

Having assigned a level of physical demand to each 1990 census occupation-gender combination, I can add this information to the Census samples. For each of the five indicators of strength content I then simply compute the total number of weeks worked in the previous year (weighted by census person weights) and divide by total hours worked by all (again weighted by person weights). Since weeks worked in the previous year is a categorical variable, I use the mid-point of each interval.

In performing these calculations I restrict all samples to workers between 18 and 64 years of age, working in the non-institutionalized sector.

One complication induced by the procedure outlined above is that not all 1990 cen-




Figure 7: Fraction of census observations matched to a level of Physical Demand

sus occupations have a corresponding 1970 census occupation (although all 1970 census occupations have a corresponding 1990 census occupation). I am thus unable to assign a level of physical demand to those 1990 occupations which do not have a corresponding 1970 occupation. Figure 7 shows the fraction of observations in each census year for which we are able to assign a level of physical demand. Not surprisingly the fraction of matches decays as we move away from 1970, but remains relatively elevated even in 2000.⁸

⁸The less than 100% match in 1970 is due to the fact that a few (21 out of 440) of the 1970 census occupations in the census files do not have a match in the 1970 census occupations in the CPS files.