NBER WORKING PAPER SERIES

# A QUANTITATIVE ANALYSIS OF THE EVOLUTION OF THE U.S. WAGE DISTRIBUTION: 1970-2000

Fatih Guvenen
Burhanettin Kuruscu

Working Paper 13095
http://www.nber.org/papers/w13095

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2007

For helpful discussions, we thank Orazio Attanasio, Richard Blundell, Satyajit Chatterjee, Dean Corbae, Gueorgui Kambourov, Larry Katz, Per Krusell, Edward Prescott, Tom Sargent, Robert Shimer, Kjetil Storesletten, Aleh Tsyvinski, Amir Yaron, and especially Daron Acemoglu, David Autor and Steve Davis, as well as the seminar participants at Carnegie Mellon University, European University Institute, Georgetown University, Goethe University Frankfurt, HEC Lausanne, Humboldt University, LSE, NYU, Rice University, Stony Brook, University College London, University of Michigan, University of Rochester, University of Toulouse, UT-Austin, University of Zurich, Yale University, York University, the Federal Reserve Banks of Atlanta, Cleveland and Philadelphia, NBER's Economic Fluctuations and Growth Meeting in Cambridge, the Minnesota Macro Workshop, the "Labor Supply and Life-cycle Productivity" Workshop organized by the Bank of Canada, the 20th Canadian Macro Study Group Workshop, and the SED conference. Guvenen acknowledges financial support from the National Science Foundation (SES-0649437). All remaining errors are our own. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

A Quantitative Analysis of the Evolution of the U.S. Wage Distribution: 1970-2000
Fatih Guvenen and Burhanettin Kuruscu
NBER Working Paper No. 13095
May 2007
JEL No. E21,E25,J24,J31

## ABSTRACT

In this paper, we construct a parsimonious overlapping generations model of human capital accumulation, and study its quantitative implications for the evolution of the U.S. wage distribution from 1970 to 2000. One of the key features of the model is that individuals differ in their ability to accumulate human capital, which is the main source of wage inequality in this model. We examine the response of this model to skill-biased technical change (SBTC), which is modeled as an increase in the trend growth rate of the price of human capital starting in early 1970's. Due to the heterogeneity in ability and age, the responses of different individuals to SBTC are systematically different from each other, generating rich behavior in the evolution of relative wages. We consider different scenarios regarding how individuals' expectations evolve during SBTC. Specifically, we study the case where individuals immediately realize the advent of SBTC (perfect foresight); and the case where they initially underestimate the future growth of the price of human capital (pessimistic priors), but learn the truth in a Bayesian fashion over time. Lack of perfect foresight appears to have little effect on the main results of the paper. The model is quantitatively consistent with several trends including the rise in overall wage inequality; the fall and rise in the college premium; the rise in within-group inequality; the stagnation in median wage growth, and the small rise in consumption inequality despite the large rise in wage inequality. Overall, the model shows promise for explaining disparate trends in the evolution of the wage distribution in a unifying human capital framework.

Fatih Guvenen
Department of Economics, BRB 3.118
University of Texas
1 University Station, C3100
Austin, TX 78712
and NBER
guvenen@eco.utexas.edu

Burhanettin Kuruscu
University of Texas at Austin
Deparment of Economics
Office BRB 3.122
Austin, Texas 78712
kuruscu@gmail.com

# 1 Introduction

The U.S. economy, and the labor market in particular, has experienced substantial changes since the early 1970's. Among the most notable of the trends during this period was a rise in overall wage inequality that started slowly in the 70's, but accelerated substantially in the 80's. Perhaps surprisingly, however, wage inequality *between* education groups (i.e., the college premium) followed a different pattern: it fell precipitously during the 70's, but then also rose substantially in most of the subsequent two decades. Despite these big changes in cross-sectional measures of wage dispersion, the rise in consumption inequality (which is a proxy for the inequality in lifetime wage incomes) was very much muted. Finally, in addition to these trends in inequality, macroeconomists have documented a sharp slowdown in aggregate labor productivity growth, and labor economists have found a parallel stagnation in median wage growth, which both started around 1973 and lasted until the mid-90's.[1]

While these trends have typically been documented by economists in different fields, studying different questions, they in fact all point to changes in the moments of the U.S. wage distribution during this time. Motivated by decades of research emphasizing the central role of human capital accumulation for the determination of wages, it seems natural to wonder whether the human capital theory can shed light on these developments. Therefore, the goal of this paper is to investigate how much mileage one can get towards explaining these phenomena using a parsimonious overlapping generations model of human capital accumulation.

The model we construct has the following features. Individuals are born with a fixed endowment of "raw labor" (health, strength, etc.), but are able to accumulate "human capital" (skill, knowledge, etc.) over the life cycle. Raw labor and human capital earn separate wages in the labor market and each individual supplies both of these factors of production at competitively determined wages. In a given period an individual is either employed full-time or is enrolled in school (accumulating human capital full time). However, while employed, an individual can choose to allocate any fraction of his time—subject to an upper bound—to human capital accumulation. We assume that skills are general and the labor market is competitive. As a result, the cost of this on-the-job investment will be completely borne by workers, and firms will adjust the hourly wage rate downward by the fraction of time workers spend learning new skills (Becker (1964)). Thus, the cost of human capital investment is given by these forgone earnings.

The model described so far extends the classic Ben-Porath (1967) framework only by introducing raw labor as a second factor of production. While this is a seemingly simple addition, it serves a key role in the results of this paper. To see this, notice that one element that is missing from the Ben-Porath model is a well-defined notion of "returns-to-skill" (which is essential for studying *skill*-biased technical change as we do below). As we elaborate in Section 2.4, while the Ben-Porath model has a price per unit of human capital, in a stationary

[1] For extensive documentation of these trends, see Bound and Johnson (1992), Katz and Murphy (1992), Juhn, Murphy, and Pierce (1993), Card and Lemieux (2001), Autor, Katz, and Kearney (2005a,b), Krueger and Perri (2006), and the surveys by Katz and Autor (1999), and Acemoglu (2002).

world, a higher price level (or wage) affects the cost and benefit of investment in exactly the same way, leaving the trade-off, and therefore the investment decision, unaffected.[2] Instead, in the present model, the marginal cost of investment (i.e., forgone earnings) is proportional to the prices of both human capital *and* raw labor, whereas the marginal benefit of investment is proportional only to the former. Therefore, a high price of human capital increases the benefit more than the cost, resulting in higher investment.

A second key feature we introduce into this framework is heterogeneity in the ability to accumulate human capital. As a result, individuals differ systematically in the amount of investment they undertake, and consequently, in the growth rate of their wages over the life cycle. Although the idea that individuals may differ in their ability is not new, here we are motivated by recent microeconometric evidence which finds that such heterogeneity is substantial (Baker (1997), Guvenen (2005, 2006), and Huggett, Ventura, and Yaron (2006a,b)). For the parameters chosen, the model generates a large rise in within-cohort wage inequality over the *life cycle*, quantitatively consistent with the values reported in these studies. This is different, for example, than the strategy followed by Heckman, Lochner and Taber (1998) who also allow for ability differences (in the standard Ben-Porath model) but proxy learning ability with the Armed Forces Qualification Test (AFQT) score, which results in very small differences in cross-sectional investment and wage profiles.

The production side of the economy is modeled as an aggregate CES technology that takes raw labor and human capital as inputs. The third key feature in the model, and the driving force behind the non-stationary changes during this period, is skill-biased technical change (SBTC)—modeled here as a rise in the price of human capital relative to raw labor—that occurs starting in the early 1970's.[3] In the baseline model we assume that individuals do not anticipate SBTC before it happens, but have perfect foresight about the future once it starts. (We relax this assumption later). It is important to point out that the specification of the production function here departs from the existing literature, which typically assumes a CES production function that takes the labor supplied by workers with high- and low-education as its two inputs (cf., Katz and Murphy (1992), Juhn, Murphy and Pierce (1993)). Furthermore, because SBTC is typically modeled in these studies as an increase in the relative demand for educated workers, it creates variation in wages between education groups, but does not have a differential impact on individuals within each group. Instead, in the present framework all individuals supply both factors of production, and therefore, simultaneously lose from the fall in the price of raw labor and gain from the increase in the price of human capital. Moreover, since individuals differ in both age and ability level, these gains and losses are distributed differently across the population (including within each education group), which creates rich dynamics in the evolution of the wage distribution. This feature allows us to study both between-group and within-group wage inequality in a single framework.

[2] To increase investment incentives in the Ben-Porath model one would need to assume an acceleration in the growth rate of the price of human capital. But in a model with a single factor of production this would also mean an acceleration in TFP growth rates during the period since 1970's, which is clearly counterfactual.

[3] For empirical evidence on SBTC, see, for example, Berman, Bound, and Griliches (1994), Autor, Katz, and Krueger (1998), Machin and Van Reenen (1998), and the survey by Acemoglu (2002).

The model is calibrated to match some key moments of the wage distribution before SBTC takes effect in 1970. We then systematically examine the implications of the model for the evolutions after 1970 of (i) the *second* moments (measures of inequality), (ii) the *first* moments (measures of average wages), and (iii) *lifetime* wage distribution (proportional to consumption in the model).

First, in the model overall wage inequality (measured by the cross-sectional variance of log wages) rises slowly in the 1970's, but substantially faster in the 1980's, as in the data. The slow rise early on is closely linked to the decline in *between*-group inequality (college premium) during the same time, which brings the wages of different education groups closer to each other, thereby compressing the overall distribution. However, counteracting this force is a rise in *within*-group wage inequality, which prevents overall inequality from falling in the 1970's. Therefore, the model is consistent with the behavior of both between-group and overall wage inequality during the 1970's.

The mechanism behind the non-monotonic behavior of the college premium can be explained as follows. Essentially, because human capital investment is a forward-looking decision, individuals increase their investment immediately after SBTC begins. Moreover, because college graduates have higher learning ability (by self-selection) than those with lower education, their investment increases more strongly, which increases their forgone earnings, and therefore reduces their relative wages in the short run. Over time, this higher investment begins to pay off, and the college premium starts to grow rapidly after the initial decline. As can be expected from this discussion, this mechanism is stronger for younger individuals who face a longer horizon, and thus, expect higher benefits from SBTC. Therefore, the college premium falls and rises substantially for young workers (but much less for the old), consistent with this behavior in the U.S. data (cf., Murphy and Welch (1992), and Card and Lemieux (2001)). Finally, the model is also consistent with the behavior of the relative wages of each education-experience group during this period (in particular, with the steepening (flattening) of cross-sectional profiles for high-school (college) graduates in the 1980's) documented by Katz and Murphy (1992) and Bound and Johnson (1992).

We then turn to the behavior of labor productivity growth (and median wage growth) in the model, which both fall sharply in the early 1970's in response to SBTC, and then recover very slowly. The slowdown is quantitatively large, with the median wage growth rate between 1970 and 1995 averaging about half its value before 1970. Labor productivity slowdown is sharper in the short run, but the recovery is also faster, averaging 80 percent of its value prior to SBTC in the period from 1970 to 1995.

Several authors have documented that consumption inequality in the United States has not increased nearly as much as wage inequality, especially at the upper tail (cf., Autor, Katz and Kearney (2004), and Krueger and Perri (2006)). This is also true in the present model where consumption inequality barely rises after 1970. The intuition for this result can be understood by noting that wage inequality in the model increases due to the fanning out of life-cycle wage profiles (see figure 10). As a result, those individuals whose wages rise the most later in life are exactly those whose wages remain lower early in life due to increased human capital investment, which then keeps the change in lifetime incomes small. This mechanism seems consistent with the observation that individuals in many high-skill occupations (doctors, engineers, professors, etc.) whose incomes have risen tremendously in the last decades, have had to go through longer periods of education, training, internship, certification, and so on, to be able to attain those high income levels.

While ideally one would also want to compare the model's implications for human capital accumulation to the data, this is not straightforward. This is because, as is well-understood since Becker (1964), Ben-Porath (1967) and others, on-the-job investment is a much broader concept than the relatively limited notion of on-the-job training directly provided by firms, which makes the total amount of such investment notoriously difficult to measure, let alone quantifying how much it has changed over time for different groups in the population. For example, Barron, Berger and Black (1997) piece together information from six different data sources and conclude that total investment on-the-job is at least seven times larger than the formal training provided by firms. The alternative approach of focusing on educational attainment is not likely to be very fruitful either. This is because, although formal education and on-the-job investment are likely to be positively related on average, this is not always the case, especially in response to major technological changes such as SBTC. For example, we provide an extension of the present framework in Guvenen and Kuruscu (2006, section 4) which retains many of the plausible implications for the wage distribution, but implies that college enrollment *falls*—whereas on-the-job investment *rises*—in the short run after SBTC. That extension, however, makes the model substantially more complicated, and is therefore beyond the scope of this paper. Thus, we keep the focus of this paper on the evolution of the wage distribution and do not venture into providing a model to study the details of educational attainment. See section 4.4 for further discussion of these points.

As noted above, in the baseline model we assume that individuals have perfect foresight after SBTC begins. While this assumption is common in the literature (among others, Greenwood and Yorukoglu (1997), Heckman et al. (1998)), there are good reasons to question the robustness of the results to such a stark assumption. In appendix A we relax this assumption and instead allow individuals to form some initial beliefs about the future growth rate of factor prices (SBTC), which they then update over time in a Bayesian fashion (observing only their own wage path).[4] We consider the case where these initial beliefs are unbiased as well as when they are very pessimistic (such that the average individual in 1970 forecasts no SBTC in the future and half of the population in fact expect the price of human capital to continually fall in the future.) Although the pessimistic priors case implies a smaller fall in the college premium than the unbiased case, the difference is not quantitatively large enough to overturn the main results of the paper (see figure 3).

In a companion paper, Guvenen and Kuruscu (2006), we theoretically study a slightly simplified version of the baseline model in this paper. Whereas in that paper we establish some results about the behavior of the wage distribution during SBTC theoretically, the main

[4] Of course for the exercise to make sense we assume that the prices of human capital and raw labor evolve stochastically.

contribution of the present paper is to provide a comprehensive assessment of the model's ability to explain the observed trends quantitatively. In addition, here we study progressively more general versions of that basic framework, by allowing for imperfect foresight and imperfect substitution in production (compared to the linear production in that paper).

An important precursor to our paper is Heckman et al. (1998). To our knowledge, their paper was the first one to emphasize that with human capital accumulation skill prices and observed wages differ, and that this could be important for understanding the recent rise in wage inequality. This observation also plays an important role in our model. However, the present paper also differs from theirs in several important respects. Two of these have already been mentioned above: one, the present paper extends the Ben-Porath model to allow for returns-to-skill; and two, a central thesis here is that individuals differ *significantly* in their ability to accumulate human capital, which is not the case in that paper. As a result, for example, the college premium falls significantly in our model even when the production function is linear, whereas imperfect substitution in the production function and changes in cohort sizes over time are essential for this result in that paper. In addition, whereas these authors focus on changes in wage inequality only, we also study the stagnation of average wage growth and the small rise in consumption inequality. Finally we also examine the robustness of our results to the lack of perfect foresight. This paper is also related to the macroeconomics literature that studies the slowdown in labor productivity. Important examples include Hornstein and Krusell (1996), Greenwood and Yorukoglu (1997), and Violante (2002). Greenwood and Yorukoglu emphasize the role of skill in facilitating the adoption of new technologies. They argue that the advent of computer technologies in the 1970s presented such a change, which increased the wages of skilled workers and resulted in a productivity slowdown due to the time it takes to utilize the new technologies effectively. Hornstein and Krusell also make a similar observation, but add that the acceleration in quality improvements during this period has exacerbated measurement problems, further reducing measured productivity growth. Violante (2002) develops a model of within-group inequality, in which vintage-specific skills, embodied technological acceleration, and labor market frictions combine to generate rising inequality.

In some interesting recent work, Heathcote, Storesletten and Violante (2005) and Krueger and Perri (2006) have constructed models that also generate a small increase in consumption inequality. In both of these papers individual wage processes are exogenous but feature idiosyncratic shocks, and improvements in the insurability of these shocks over time dampen the rise in consumption inequality. Compared to these papers, our model abstracts from several features that are likely to be important for a detailed study of consumption behavior. But our model highlights a different, and probably complementary, channel which suggests that even when the rise in wage inequality is entirely systematic (and substantial), life-time income inequality may not change much.

# 2 Baseline Model

The economy consists of overlapping generations of individuals who live for $S$ years. There is no population growth and we normalize the population size to 1, implying that the measure of $s$-year old individuals, $\mu(s)$, is equal to $1/S$.

## 2.1 Human Capital Accumulation: Extending the Ben-Porath Model

Individuals begin life with an endowment of "raw labor" (i.e., strength, health, etc.) which is constant over the life cycle, and are able to accumulate "human capital" (skills, knowledge, etc.) over the life cycle, which is the only skill that can be accumulated in this economy. There is a continuum of individuals in every cohort, indexed by $j \in [0,1]$, who differ in their ability to accumulate human capital, denoted by $A^j$ (also referred to as their "type"). Although, in some cases below, we will allow individuals to also differ in their raw labor endowment as well as in their beliefs, the heterogeneity in ability will be the crucial source of heterogeneity in the model.

Each individual has one unit of time endowment in each period that can be allocated between producing output and accumulating human capital. Let $l$ denote raw labor and $h^j_{s,t}$ denote the human capital in period $t$ of an $s$-year-old individual of type $j$. We assume that raw labor and human capital earn separate wages in the labor market, and each individual supplies both of these factors of production at competitively determined (potentially stochastic) wage rates, denoted by $P_{L,t}$ and $P_{H,t}$, respectively.[5]

Individuals begin their life with zero human capital, and accumulate human capital according to the following technology:

$$h^j_{s+1,t+1} = h^j_{s,t} + \underbrace{A^j((\theta_{L,t}l + \theta_{H,t}h^j_{s,t})i^j_{s,t})^\alpha}_{Q^j_{s,t}} \qquad (1)$$

where $i^j_{s,t}$ is the fraction of time devoted to human capital investment, henceforth referred to as "investment time"; and $Q^j_{s,t}$ is the newly produced human capital which will be referred to simply as "investment" in the rest of the paper. According to this formulation new human capital is produced by combining the existing stocks of raw labor and human capital with the available investment time.[6] A key parameter in this specification is $A^j$, which determines the

[5]The structure we have in mind is not one where an individual works at manual tasks (using raw labor only) some fraction of the time and at cognitive (or skill-intensive) tasks at other times. Instead the worker employs both factors of production simultaneously in producing output. For example, a college professor uses both his/her body and knowledge/skills at the same time when teaching, although probably at different proportions than a farmer, an auto mechanic, or a brain surgeon.

[6]The dependence of the weights in the human capital production function on $t$ is to stress that these could be time-varying.

productivity of learning. Due to the heterogeneity in $A^j$, individuals will differ systematically in the amount of investment they undertake, and consequently, in the growth rate of their wages over the life cycle. Another important parameter is $\alpha \in [0, 1]$, which determines the degree of diminishing marginal returns in the human capital production function. A low value of $\alpha$ implies higher diminishing returns, in which case it is optimal to spread investment over time. In contrast, when $\alpha$ is high the marginal return on investment does not fall quickly, and investment becomes bunched over time. In the extreme case when $\alpha = 1$, individuals either spend all their time on investment ($i^j_{s,t} = 1$) or none at all in a given period.

The main difference between the Ben-Porath (1967) model and the formulation in (1) is the introduction of raw labor as an additional factor into our model. When $l \equiv 0$ (and $\theta_{H,t}$ is normalized to 1), this model reduces to the standard Ben-Porath model. As will be clear in the analysis below, the reason for our deviation from the standard Ben-Porath model is because it is difficult to sensibly think about such notions as returns-to-skill and SBTC in that framework.

Investment in human capital takes place on-the-job as long as it does not exceed a fraction $\chi$ of an individual's time endowment in a period. If the individual wants to invest more, he enrolls in college and invests 100 percent of his time. Thus, the choice set for investment time is: $i^j_{s,t} \in [0, \chi] \cup \{1\}$, which is non-convex when $\chi < 1$. An upper bound less than 100 percent on on-the-job investment seems plausible as it could arise, for example, if the firm incurs fixed costs for employing each worker (administrative burden, cost of office space, etc.), or due to minimum wage laws.[7]

We assume that skills are general (i.e., not firm-specific) and labor markets are competitive. As a result, the cost of human capital investment is completely borne by workers, and firms adjust the hourly wage rate downward by the fraction of time invested on the job. Then, the observed wage income of an individual is given by

$$w^j_{s,t} = \underbrace{\left[P_{L,t}l + P_{H,t}h^j_{s,t}\right]}_{x^j_{s,t}}(1 - i^j_{s,t}) = \underbrace{x^j_{s,t}}_{\text{Potential earnings}} - \underbrace{x^j_{s,t}(t)\, i^j_{s,t}}_{\text{Cost of investment}}, \tag{2}$$

where $x^j_{s,t}$ is the "potential earnings" of an individual—that is, the income an individual would earn if he spent all his time producing for his employer. Therefore, wage income can be written as the potential earnings minus the "cost of investment," which is simply the forgone earnings while individuals are learning new skills. Since labor supply is inelastic, $w^j_{s,t}$ is also just a scaled version of the individual's observed "wage rate." (See Becker (1964) for a classic exposition of the basic human capital model.)

[7] In addition to its plausibility, such an upper bound is also important for a meaningful quantitative analysis. Otherwise, with a continuum of ability levels, there will be some individuals who invest slightly less than 100 percent of their time, appearing as employed while earning a wage income very close to zero. Because many of the statistics we analyze below involve the logarithm of wage rates as well as the variances of these logarithms, even a small number of such individuals can easily wreak havoc with the quantitative exercise.

## 2.2 A CES Aggregate Production Function

Let the aggregate factors used in production at a point in time be defined as

$$\begin{aligned} L_t^{net} &= \sum_{s=1}^{S} \mu(s) \int_j l\left(1 - i_{s,t}^j\right) dj, \quad \text{and} \\ H_t^{net} &= \sum_{s=1}^{S} \mu(s) \int_j h_{s,t}^j \left(1 - i_{s,t}^j\right) dj, \end{aligned} \tag{3}$$

where the integrals are thus taken over the distribution of individuals of all types and ages. The superscripts in $H_t^{net}$ and $L_t^{net}$ stress that these variables measure the actual amounts of each factor used in production (that is, *net* of the time allocated to human capital investment). The aggregate firm uses these two inputs to produce a single good, denoted by $Y$, according to the familiar CES production function:

$$Y = Z\left(\left[\theta_L L^{net}\right]^\rho + \left[\theta_H H^{net}\right]^\rho\right)^{1/\rho}, \tag{4}$$

where $\rho \leq 1$, and $Z$ is the total factor productivity (TFP). For simplicity we assume that capital is not used in production. Notice that human capital and raw labor enter the aggregate production function and human capital function with the same weights (compare equations (1) and (4)).[8]

The firm solves a static problem by hiring factors from households every period to maximize its profit: $Y - P_L L^{net} - P_H H^{net}$. The factor prices corresponding to human capital and raw labor are:

$$\begin{aligned} P_H &= \frac{\partial Y}{\partial H^{net}} = Z\theta_H^\rho \left(\theta_L^\rho \left[H^{net}/L^{net}\right]^{-\rho} + \theta_H^\rho\right)^{\frac{1-\rho}{\rho}}, \quad \text{and} \\ P_L &= \frac{\partial Y}{\partial L^{net}} = Z\theta_L^\rho \left(\theta_H^\rho \left[H^{net}/L^{net}\right]^{\rho} + \theta_L^\rho\right)^{\frac{1-\rho}{\rho}}. \end{aligned} \tag{5}$$

The price of human capital *relative* to raw labor has a simple expression:

$$\frac{P_H}{P_L} = \left(\frac{\theta_H}{\theta_L}\right)^\rho \left(\frac{H^{net}}{L^{net}}\right)^{\rho-1}. \tag{6}$$

[8] We have experimented with two other specifications of the human capital accumulation function that also seemed a priori plausible. In particular, we write the human capital function as $A^j((\lambda_{L,t} l + \lambda_{H,t} h_s^j) i_s^j)^\alpha$ and consider (i) weights that remain constant through SBTC: $\lambda_{L,t} = \overline{\lambda}_L$ and $\lambda_{H,t} = \overline{\lambda}_H$; and (ii) no role for raw labor in human capital production: $\lambda_{L,t} = 0$ and $\lambda_{H,t} = 1$. The first case had implications qualitatively similar to the baseline model described here, while the second displayed some implausible behavior even in steady state (that is, without SBTC). Overall, the simplest specification used here (originally proposed and studied in Guvenen and Kuruscu (2006)) also turned out to have the most plausible quantitative implications.

While the aggregate production function has the same CES form commonly used in the literature, its inputs are different than what is typically assumed. In most previous work $H^{net}$ and $L^{net}$ denote the labor supplied by workers with college and high school education respectively. Therefore, a change in the price of $H^{net}$ relative to $L^{net}$ has the same effect on all individuals within an education group. As a result, the college premium is simply equal to $P_H/P_L$ and satisfies the relationship in (6). A key implication of this equation is that a rise in the relative supply of high-skill workers will reduce the college premium. Several authors have emphasized this link to argue that the fall in the college premium during the 1970's resulted from the rapid increase in the supply of college-educated workers (cf., Katz and Murphy (1992), Juhn, Murphy and Pierce (1993), and Heckman et al. (1998)).

In contrast, in the present model, all workers have some endowment of human capital (which varies by ability and age) and $l$ (which is the same for all), and every worker contributes to both factors of production. Therefore, a change in the price of $H^{net}$ relative to $L^{net}$ affects all individuals differently depending on their ability level as well as their age. Moreover, as we show below, the college premium is now very different than $P_H/P_L$.

An important special case arises when $\rho = 1$. In this case, human capital and raw labor become perfectly substitutable, and the relative wage in equation (6) reduces to $P_H/P_L = \theta_H/\theta_L$. Therefore, this assumption eliminates the link between the relative supply of high-skill labor and the college premium mentioned above. To highlight the role of the mechanism proposed in this paper for the college premium, we make this assumption from this point on. Furthermore, in appendix B we show that $\rho = 1$ is an empirically plausible benchmark, but also analyze the more general case with $\rho < 1$.

**Skill-biased Technical Change.** Skill-*neutral* technological progress takes place at a constant rate: $Z_{t+1} = (1+g)\, Z_t$. The focus of this paper is on skill-*biased* changes in technology, which we model as follows. The productivities of human capital and raw labor evolve as:

$$\theta_{H,t+1} = \kappa + \theta_{H,t}, \quad \text{and} \quad \theta_{L,t+1} = -\kappa + \theta_{L,t}. \tag{7}$$

The growth rate of each factor's productivity is zero (i.e., $\kappa \equiv 0$) up to time 0. Skill-biased technical change is modeled as an unanticipated regime change in the growth rate of human capital's productivity relative to raw labor. Specifically, $\kappa = \kappa^* > 0$, for $t = 1, .., T$; and $\kappa$ reverts back to zero again for $t > T$. Therefore, SBTC is assumed to last for a finite period of time, which will be motivated below by empirical evidence.

In the baseline model, we assume that as soon as SBTC begins, individuals learn $\kappa^*$ and $T$, and therefore, have perfect foresight about the future paths of $\theta_{H,t}$ and $\theta_{L,t}$. This assumption, while probably too strong, has the advantage of providing a clear and simple benchmark. In appendix A, we relax this assumption, and study the case where individuals do not have perfect foresight, and learn about the growth rate of future skill prices in a Bayesian fashion.

## 2.3 Individuals' Lifetime Income Maximization Problem

Individuals are able to borrow and lend at a constant interest rate, $r$, which is sufficient for markets to be complete in this deterministic model. As is well-known, under complete markets, the consumption-savings decision can be disentangled from the lifetime income maximization problem. Therefore, to study the determination of wages it is sufficient to restrict our attention to the latter problem.

Let $V_s^j\left(h_{s,t}^j, \theta_{H,t}\right)$ denote the lifetime income of individual $j$ who is $s$-years-old at time $t$. The income maximization problem of the agent can be written recursively as:

$$V_s^j\left(h_{s,t}^j; \theta_{H,t}\right) = \underset{i_{s,t}^j \in [0,\chi]\cup\{1\}}{\text{Max}} \left[ Z_t(\theta_{L,t} l + \theta_{H,t} h_{s,t}^j)(1 - i_{s,t}^j) + \frac{1}{1+r} V_{s+1}^j\left(h_{s+1}^j; \theta_{H,t+1}\right)\right] \quad (8)$$

subject to

$$h_{s+1,t+1}^j = h_{s,t}^j + A^j((\theta_{L,t} l + \theta_{H,t} h_{s,t}^j) i_{s,t}^j)^\alpha, \qquad h_0 = 0, \quad (9)$$
$$V_{S+1}^j\left(h_{S+1}^j; \theta_{H,t}\right) \equiv 0,$$

where the evolution of factor prices are given by equation (7). It should be stressed again that our focus on the lifetime income maximization problem does not require the assumption of risk-neutrality; any concave utility function implies the same human capital investment behavior in this environment.

## 2.4 Optimal Investment Decision

When $\chi < 1$, the choice set for investment time in equation (8) is non-convex, making the problem difficult to study analytically. While this upper bound matters for the quantitative results, the main mechanisms for the results can be explained more clearly by deriving some analytical expressions. This is only possible when $\chi = 1$, which we therefore assume in the rest of this section.

Using equation (1) the opportunity cost of investing an amount $Q_{s,t}^j$ can be written as:

$$C^j(Q_{s,t}^j) \equiv (\theta_{L,t} l + \theta_{H,t} h_{s,t}^j) i_{s,t}^j = \left(\frac{Q_{s,t}^j}{A^j}\right)^{1/\alpha}. \quad (10)$$

When there is an interior solution for investment time $\left(i_{s,t}^j < 1\right)$, the optimal amount of investment satisfies the following first order condition:

$$C^j(Q_{s,t}^j)' = \left(\delta\theta_{H,t+1} + \delta^2\theta_{H,t+2} + ... + \delta^{S-s-1}\theta_{H,t+S-s-1}\right), \quad (11)$$

where $\delta \equiv (1+g)/(1+r)$ (see Guvenen and Kuruscu (2006) for derivation). The left hand side of this equation is the marginal cost, and the right hand side is the marginal benefit ($MB$) of investing in human capital. The latter is the present discounted value of the future stream of wages that is earned by an additional unit of human capital. An important implication of (11) is that an increase in *future* skill prices will immediately affect *current* investment behavior because of the forward-looking nature of this equation.

**Remark** Another important observation is that optimal investment, $Q^j_{s,t}$, only depends on the level of $\theta_H$—not on the levels of $\theta_L$ or $Z$. This is because the opportunity cost of investment depends on the prices of both raw labor and human capital (see equation (10)), whereas the marginal benefit is only proportional to the price of human capital. As a result, a higher level of $\theta_H$ (for example, due to SBTC) increases the marginal benefit more than the marginal cost, resulting in higher investment. This feature is an important difference between the current framework and the standard Ben-Porath model. In the latter, a higher price of human capital (which is the only factor of production since there is no raw labor) affects the cost and benefit of investment exactly the same way, leaving the trade-off—and therefore the investment decision—unaffected. It is precisely for this reason that it is difficult to think of the concept of returns-to-skill in that framework, because a higher price of human capital has no effect on the decision to invest. Instead in the present model $\theta_H/\theta_L$ is a measure of returns-to-skill, and affects investment in human capital without necessarily implying anything about aggregate productivity (which is captured by $Z$ and also has no effect on investment incentives for the same reason discussed for the Ben-Porath model).[9]

We now proceed to characterize the optimal investment choice, which can be solved for explicitly:

$$Q^j_{s,t} = \left(A^j\right)^{1/(1-\alpha)} \left[\alpha MB\right]^{\alpha/(1-\alpha)} . \tag{12}$$

This expression highlights the main sources of heterogeneity in this model: (i) individuals with higher learning ability invest more in human capital: $\partial Q^j_{s,t}/\partial A^j > 0$; (ii) more importantly, their investment responds more strongly to SBTC: $\partial^2 Q^j_{s,t}/\partial\theta_{H,t+k}\partial A^j > 0$ (for all $k > 0$); (iii) investment goes down over the life cycle: $\partial Q^j_{s,t}/\partial s < 0$; and finally, younger individuals respond more strongly to SBTC: $\partial^2 Q^j_{s,t}/\partial\theta_{H,t+k}\partial s < 0$.

**Price, Investment, and Quantity Effects.** Armed with this characterization of optimal investment behavior, we are now ready to discuss how the wage of a typical worker changes in response to SBTC. For $t > 0$, an individual's wage can be written as

$$w^j_{s,t} = Z_t \left[(\theta_{L,0} - t\kappa^*)\, l + (\theta_{H,0} + t\kappa^*)\, h^j_{s,t} - C^j(Q^j_{s,t})\right] .$$

The effects of SBTC on wages work through three separate channels, which can be seen

[9] Although an increase in the *growth rate* of $Z$ will increase investment rates, there is no evidence of increased TFP growth rate after the 1970s; in fact there is ample evidence to the contrary.

Figure 1: Cross-sectional Differences in Investment Time and Wages Over the Life-Cycle

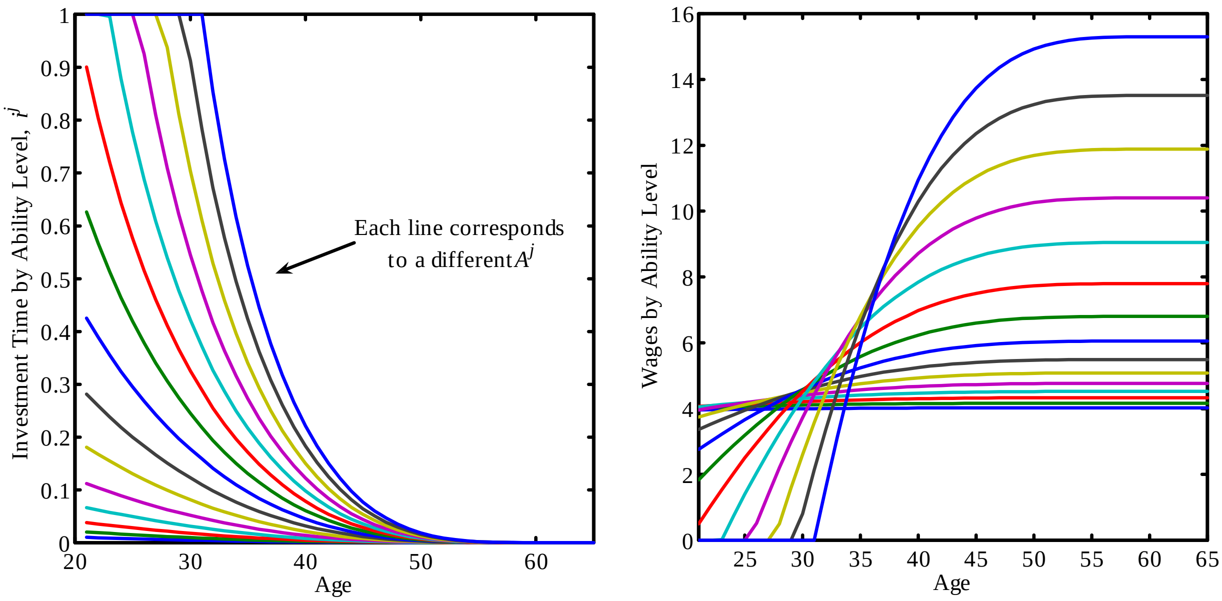


from this expression. First, individuals will increase their investment immediately in response to SBTC, resulting in a rise in the cost of investment, which will then reduce the wage rate. We call this the "investment effect." The discussion above makes clear that this effect is strongest (most negative) for those with high ability and/or those who are young.

Second, even in the absence of any response in the individuals' investment rate (that is, keeping $Q^j_{s,t}$ and $h^j_{s,t}$ fixed) the wage rate will change due to the fact that $\theta_{L,t}$ is falling and $\theta_{H,t}$ is rising over time. We call this the "price effect." The price effect will be positive (negative) for individuals who have a high (low) stock of human capital relative to their raw labor (that is, those with high (low) ability and/or labor market experience). Moreover, the price effect strengthens over time as $\theta_{H,t}/\theta_{L,t}$ rises during SBTC.

Third, and finally, the increased investment gradually raises the stock of human capital, $h^j_{s,t}$, which in turn gradually raises the wage rate. We call this the "quantity effect." Notice that the quantity effect is stronger for younger (and/or high-ability) individuals whose investment respond more strongly to SBTC than it is for older (and/or low-ability) individuals. These three effects are crucial for understanding the changes in the wage structure after SBTC, and we will refer to them throughout the paper.

### 2.5 Heterogeneity in Investment and Wages over the Life Cycle

Using the closed-form expressions derived above, we now construct the life-cycle profiles of human capital investment and the implied wage paths in steady state (figure 1). The figures are generated using the same parameter values as in the calibrated economy of the next section (except that $\chi = 1$ here) to highlight the aspects of the life-cycle investment behavior that will play a quantitatively important role in our results. The first point to note is the substantial cross-sectional heterogeneity in investment rates observed early in life (left panel), which ranges from virtually no investment for very low ability individuals all the way to full-time investment for those with highest ability (who enroll in college). This significant cross-sectional heterogeneity is at the heart of the mechanism that generates many of the results in this paper, including the decline in the college premium in the 1970's and the small rise in consumption inequality, among others. This feature is also one of the major differences between this paper and Heckman et al (1998), in which the cross-sectional heterogeneity in on-the-job investment rates is very small (compare, for example, figure 1 here to figure 3 in that paper).

Given the difficulty, noted above, of measuring on-the-job investment directly, one way to gauge the implications of a human capital model is to study the resulting life-cycle wages. In the right panel the wage profiles show significant fanning out over the life cycle, which is the only source of wage inequality in this model. The implication that systematic differences in growth rates are the major driving force behind the rise in wage inequality over the life cycle is supported by recent empirical studies that estimate wage and labor earnings processes from micro data sets (Baker (1997), Guvenen (2005, 2006), Huggett, Ventura and Yaron (2006a,b)). For example, the calibrated version of the present model in the next section implies that the cross-sectional wage inequality at age 55 is about 9 times the inequality at age 35. For comparison, Guvenen (2006, table 2) reports that the component of wage inequality that is due to *systematic* differences in growth rates (that is, net of the inequality due to idiosyncratic shocks) is about 10.5 times the inequality at age 35. Finally, although in figure 1 college graduates' wages are very low immediately after graduating, this is a consequence of setting $\chi = 1$; this does not happen when $\chi$ is calibrated appropriately (see figure 10), as we do in the next section.

# 3 Quantitative Analysis

We calibrate the baseline model to the U.S. data under the assumption that the U.S. economy was in steady state before SBTC took effect in 1970. The model is solved numerically and the results below are computed using simulated data. We then compare the evolution of the wage distribution implied by the model from 1970 to 2000 to the data. The U.S. wage data used in this paper are from the annual March Current Population Surveys (CPS) on full-time full-year male workers covering the period 1963 to 2003, and have been provided to us by David Autor; they are the same as the data used in Autor, Katz and Kearney (2005a,b). Our focus on male workers is to abstract from the significant changes during this period in females' labor market participation rates, fertility rates, family composition, and so on, which may potentially have affected the determination of females' wages (more so than males' wages). Having said that, it should be noted that many of the trends documented below have also been observed in females' wages, suggesting that the human capital channels emphasized in this paper could also have played an important role in the evolution of wages for that group as well.

## 3.1 Calibration

Individuals enter the economy at age 20 and retire at 65 ($S = 45$). The net interest rate, $r$, is set equal to 0.05, and the subjective time discount rate is set to $\beta = 1/(1+r)$, implying that individuals will choose a constant consumption path over their life cycle (given the absence of uncertainty and borrowing constraints).

**Aggregate Production Function.** The growth rate of neutral technology level, $Z$, is set equal to 1.5 percent per year. As will become clear below, *measured* TFP growth will be different than this number when the amount of investment on-the-job changes over time. In the baseline model, we take the curvature of the aggregate production technology, $\rho$, to be unity. In appendix B, we also explore the effect of imperfect substitution ($\rho < 1$) on our results. Notice that $\theta_L$ and $\theta_H$ always appear multiplicatively with raw labor and human capital, so the initial values of these parameters serve only as a normalization (given that $H$ and $L$ are also calibrated below). Therefore, we normalize $\theta_{L,t} + \theta_{H,t} = 1$ and set $\theta_{L,t} = \theta_{H,t} = 0.5$ for all $t < 1970$. We calibrate the change in the skill-bias of technology after 1970 below.

**Human Capital Accumulation.** The estimates of $\alpha$—the curvature of the human capital accumulation function—typically vary between 0.80 and 0.95 (see, for example, Heckman (1976), and the more recent estimates in Heckman et al. (1998) and Kuruscu (2006)). In Guvenen and Kuruscu (2006) we show, theoretically, that if $\alpha$ is higher than a certain threshold the college premium will fall and average wages will stagnate in the short run after SBTC (consistent with the data). Therefore, here we set $\alpha = 0.80$, a value close to the lower end of this empirically plausible range, to show that the plausible quantitative implications found in the next section do not require an extremely high value of $\alpha$. We have also experimented with values between 0.75 and 0.95, and found that they had a qualitatively small effect on our results.[10]

[10] These results are available upon request.

**Accounting for Idiosyncratic Shocks.** For a meaningful comparison of the model to the data, it is important to account for the fact that the model abstracts from idiosyncratic shocks, which are clearly present in the data. To this end, we assume that the logarithm of the observed wage in the data can be written as

$$\log \widetilde{w}_{s,t}^{j} = \log w_{s,t}^{j} + \upsilon_{s,t}^{j} + \xi_{s,t}^{j}, \tag{13}$$

where $w_{s,t}^{j}$ denotes the systematic (or life-cycle) component of wages, and is given by the baseline human capital model in this paper; $\upsilon_{s,t}^{j}$ represents a first-order autoregressive shock process, and $\xi_{s,t}^{j}$ is a transitory disturbance with variance $\sigma_{\xi}^{2}$. Both the innovation to the AR(1) process and $\xi_{s,t}^{j}$ are *i.i.d.* conditional on all individual characteristics (including $s$ and $A^{j}$). This specification is similar to the econometric processes for wages commonly used in the literature.[11] The key assumption we make is that the variances of these idiosyncratic shocks have been *stationary* during the period under study.[12] Under this assumption, and letting $var(\cdot)$ denote the cross-sectional variance of a variable, we have:

$$var\left(\log \widetilde{w}_{s,t}^{j}\right) = var\left(\log w_{s,t}^{j}\right) + \sigma_{\upsilon}^{2} + \sigma_{\xi}^{2}.$$

where $\sigma_{\upsilon}^{2}$ denotes the cross-sectional variance of the AR(1) process across all age and ability groups. Two points are easily noted from this expression. First, the *level* of the variance of wages in the model needs to be adjusted by $\left(\sigma_{\upsilon}^{2} + \sigma_{\xi}^{2}\right)$ before it can be compared to the data. Second, the *change* over time in the variance of observed wages will mirror that in the systematic component $\left(\Delta var\left(\widetilde{w}_{s,t}^{i}\right) = \Delta var\left(w_{s,t}^{i}\right)\right)$ which allows a direct comparison of the *trend* in the model variances to its empirical counterpart.

Similarly, the implications of the specification in (13) for the first moment of wages can also be seen easily: the average of observed log wages equals that of the systematic component, $E\left(\log \widetilde{w}_{s,t}^{j} | I\right) = E\left(\log w_{s,t}^{j} | I\right)$, where $I$ denotes a set of individuals—for example, those in the same age or education group. Therefore, both the level of, and the change in, the first moments of log wages in the model can be directly compared to the data.

[11] One caveat of this specification should be noted. Because the idiosyncratic shocks introduced here are multiplicative with $w_{s,t}^{j}$, it can be shown that if individuals take the existence of these shocks into account when making their investment decision, this would lead to a different optimal choice than the one that generated $w_{s,t}^{j}$. Although it is possible to modify the human capital problem and solve it in the presence of these idiosyncratic shocks, such an extension comes at considerable computational cost, so we do not tackle this potential complication here.

[12] While several studies have found the variances of idiosyncratic shocks to have increased during this period (Moffitt and Gottschalk (1994), Meghir and Pistaferri (2004), among others), it is important to note that these studies do *not* account for the possibility that the *dispersion* of wage growth rates could have increased during this time, which is the main thesis of the present paper. Therefore, that evidence is not informative about how the variances should be calibrated over time in our model.

### 3.1.1 Calibrating Model-Specific Parameters

We set $\chi$ equal to 0.50, which (together with the other parameters below) implies that in the initial steady state before SBTC, the lowest wage is 51 percent of the average (mean) wage in the economy. This choice is consistent with the minimum wage interpretation given for $\chi$ above: in 1969 the ratio of the bottom percentile of the wage distribution to the mean was—depending on the exact measure used—between 50 to 55 percent in the U.S. data (Hornstein, Krusell and Violante (2006, figure 1)).

**Distributions of Ability and Raw Labor.** Learning ability, $A^j$, is assumed to be uniformly distributed in the population with the same parameters for every cohort. As for the calibration of individuals' raw labor endowment, note that the present model is interpreted as applying to human capital accumulation after secondary school. But then, the assumption we made in the theoretical model—that individuals start out with the same human capital level—may be too restrictive because it seems likely that different individuals would have accumulated different amounts of human capital by the time they make the college enrollment decision. A simple way to model this heterogeneity is by assuming that the amount of raw labor, $l$, has a non-degenerate distribution in the population. We also assume $l$ to have a uniform distribution that is the same for all cohorts. Each distribution is fully characterized by two parameters, giving us four parameters to be calibrated.[13] The mean value of raw labor, $E\left[l^j\right]$, is a scaling parameter and is normalized to one, leaving three parameters: (i) the cross-sectional standard deviation of raw labor, $\sigma\left(l^j\right)$, (ii) the mean learning ability, $E\left[A^j\right]$, and (iii) the dispersion in the ability to learn, $\sigma\left(A^j\right)$. These are chosen to match the following three moments:

1. the average cross-sectional variance of log wages between 1965 and 1969,
2. the average level of the log college premium between 1965 and 1969,
3. the mean log wage growth over the *life cycle.*

As discussed above, we need an estimate of the variances $\sigma_v^2$ and $\sigma_\xi^2$ to obtain the target value for the cross-sectional wage inequality. Note that, for consistency, these estimates must be obtained from empirical studies that allow for heterogeneity in wage growth rates as implied by the human capital model in this paper.[14] Guvenen (2005) estimates such a specification

[13]Notice that we also need to calibrate the cross-sectional correlation of $l$ and $A$. Since we interpret the heterogeneity in $l$ as arising from investments made prior to college and high-ability individuals are likely to have invested more even before college, it seems reasonable to conjecture that $A$ and $l$ will be positively correlated. Indeed, Huggett et al (2006b) estimate the parameters of the standard Ben-Porath model from individual wage data allowing for heterogeneity in $A$ and $l$, and provide evidence that the two are strongly positively correlated (corr: 0.792). For simplicity we assume perfect correlation between the two. Furthermore, it will become clear later that the heterogeneity in $l$ does not play a significant role in this model, implying that the choice of perfect correlation is not likely to be critical.

[14]This requirement eliminates several well-known empirical papers, such as MaCurdy (1982), Abowd and Card (1989), and Meghir and Pistaferri (2004), among others, which restrict wage growth rates to be the same across the population.

Table 1: Baseline Parameterization

| Parameter | | Value |
|---|---|---|
| $R$ | Interest rate | 0.05 |
| $\beta$ | Time discount rate | $1/R$ |
| $\alpha$ | Curvature of human capital function | 0.80 |
| $S$ | Years spent in the labor market | 45 |
| $\rho$ | Curvature of aggregate prod function | 1.0 |
| $\chi$ | Maximum investment time on the job | 0.50 |
| $\Delta \log Z$ | Growth rate of neutral technology | .015 |
| $T$ | Duration of SBTC (years) | 26 |
| $E[l^j]$ | Average labor endowment (scaling) | 1.0 |
| *Parameters calibrated to match 1965-69 targets:* | | |
| $E[A^j]$ | Average ability | .071 |
| $\sigma(l^j)/E[l^j]$ | Coeff. of variation of labor endowment | .0503 |
| $\sigma[A^j]/E[A^j]$ | Coeff. of variation of ability | .245 |
| *Parameter calibrated to match 1995 wage inequality:* | | |
| $\Delta \log(\theta_H/\theta_L)$ | Annual change in skill-bias (1970-1995) | 2.21% |

and reports $\sigma_\xi^2$ to be 0.047. Similarly, $\sigma_\upsilon^2$ can be calculated to be 0.088 using the estimates in that paper (Table 1, row 2). The average cross-sectional variance of log wages in the U.S. data between 1965 and 1969 is 0.239, implying a target value for the first moment in the model ($var\left(w_{s,t}^j\right)$) of 0.104. Second, the log college premium in the U.S. data averaged 0.381 between 1965 and 1969 (and does not require any adjustments), which is the second empirical target we choose. Third, and finally, our target for mean log wage growth between ages 20 and 55 is 50 percent for a cohort of individuals who retire before 1970. This number is roughly the middle point of the figures found in studies that estimate life-cycle wage and income profiles from panel data sets such as the Panel Study of Income Dynamics (which typically report estimates between 40 and 65 percent; see, for example, Gourinchas and Parker (2002), Davis, Kubler and Willen (2002), Guvenen (2005)).[15]

Table 1 displays the implied values for the distributions of $A^j$ and $l^j$. Notice that the coefficient of variation of ability is more than four times that of raw labor. Overall, heterogeneity in $l$ has a much more modest effect on the quantitative results than does the heterogeneity in ability.

[15] Ideally we would like to use an estimate of average life-cycle wage growth during the period before 1970 (before SBTC), whereas PSID is only available starting 1968 on. However, we are not aware of any study that estimates the life-cycle (*not* cross-sectional) wage profiles using data from earlier periods. Our calibration implies a mean log wage growth of 62 percent for the cohort that enters the economy in 1968, consistent with the numbers found by these studies during the same period.

**Skill-Biased Technical Change.** The driving force behind the non-stationary changes in the model is a sustained increase in the relative productivity of human capital relative to raw labor, $\theta_H/\theta_L$. Specifically, $\theta_H$ grows and $\theta_L$ shrinks by $\kappa^*$ from 1970 until 1995. The ending year is chosen to be consistent with the observation that the rise in wage inequality seems to have slowed down, and productivity growth has started to pick up, by mid-1990's. This choice is also consistent with empirical evidence indicating a slowdown in the rate of skill-biased technical change in the second half of 90's (see, for example, Acemoglu (2002, p 28)).

The main quantitative experiment is the following: we choose $\kappa^*$ such that the model matches the total rise of 13 log points in the variance of log wages in the U.S. data between the first steady state (1965–69) and the ending year of SBTC, 1995. The resulting value of $\kappa^*$ is 0.0054, implying an average growth rate of 2.21 percent per year for $\theta_H/\theta_L$. As mentioned earlier, we specified the change in the productivities of each factor in their levels (equation (14)), rather than the more common specification in growth rates. Because $\kappa^*$ is very small relative to $\theta_H$ and $\theta_L$, the difference between the paths implied by the two specifications is very small (for the parameterization used in this paper), but the level specification simplifies computation substantially, especially in cases with Bayesian learning considered later.[16] Table 1 summarizes the baseline parameter choices.

# 4 Model Results

## 4.1 Evolution of Wage Inequality

In this section, we discuss the implications of the baseline model for the evolution of the wage distribution. To save space, each figure below also plots the results from the extensions of this model, which are examined later in appendix A. In the following discussion we only focus on the baseline model (thick solid line) and do not comment on the other graphs until later.

### 4.1.1 Overall wage inequality

We begin with the evolution of overall wage inequality during this period (figure 2). The baseline model is calibrated to match the total change in wage inequality between the first steady state and 1995, and not the evolution between these end points. Yet, the model seems to nicely capture the broad pattern during this period, with a slow increase in the 1970's that accelerates over time.

[16] Notice that, despite the fact that $\theta_L$ is falling during SBTC, the *absolute* productivity of raw labor, $Z\theta_L$, continues to grow (by 0.23 percent per year) due to the sustained growth in $Z$. Therefore, with this calibration SBTC results only in a *relative* fall in the productivity of raw labor relative to human capital.

Figure 2: The Evolution of Overall Wage Inequality: Model versus U.S. Data, 1965—2000.

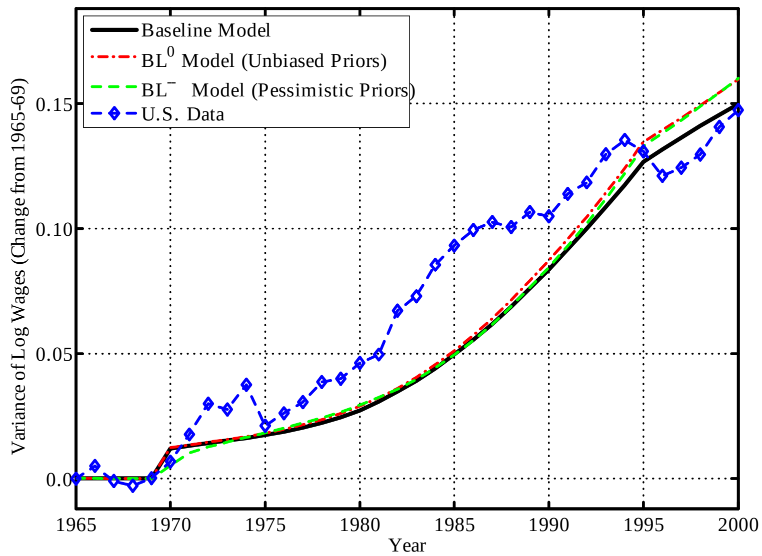


To understand this convex pattern two separate effects, which sometimes work in opposite directions, should be noted. First, $\theta_H/\theta_L$ increases at a roughly linear rate (as can be seen in the left panel of figure 4 below). If there was no change in investment rates in response to SBTC (and therefore, the distribution of human capital remained unchanged over this period) the price effect would increase wage inequality at the same constant rate as the relative price change. However, the investment rate *does* respond to SBTC, which is a key feature of this model. This effect works to offset the price effect early on, because individuals whose investment responds more strongly to SBTC are exactly those with higher ability, and thus who have relatively more human capital already. As a result, the rise in wage inequality is depressed early on. Over time, however, the differential investment response leads to an even larger dispersion in human capital levels, which reinforces the price effect, and leads to an accelerating rise in wage inequality. Overall, even though SBTC begins in 1970, most of the rise in overall wage inequality (11.5 out of the 13 log points) happens after 1980, consistent with the U.S. data.

One notable divergence occurs during the 1980's when inequality rises faster in the data compared to the model. Some authors have emphasized the role played by the erosion of the legal minimum wage due to high inflation in the late 70's, which resulted in the fall of wages in the lower tail of the distribution, thereby increasing inequality (cf., Card and Dinardo (2002)). This factor is not present in the model which might explain the divergence from the data during the 80's.

Figure 3: The Evolution of College Premium: Model versus U.S. Data, 1965-2000.

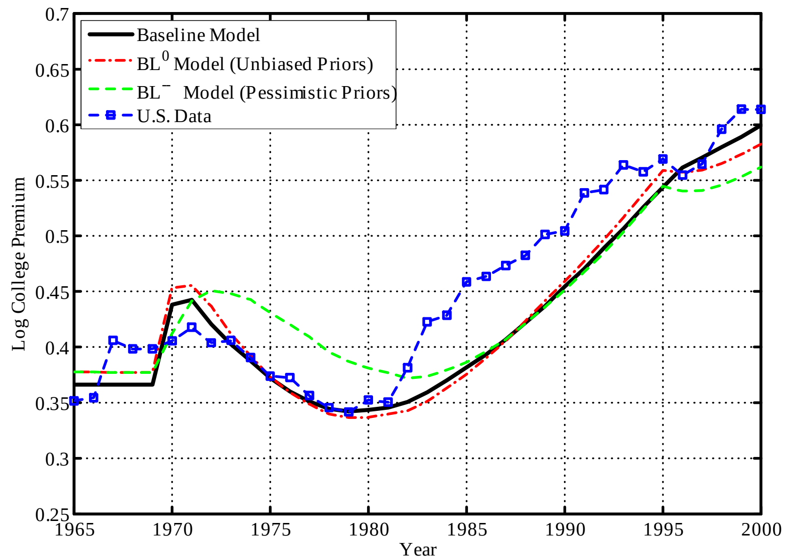


### 4.1.2 Between-group inequality (College premium)

Figure 3 plots the college premium in the model along with the empirical counterpart. The model is calibrated to match the average level of premium between 1965 and 1969.[17] In the model (thick solid line), the college premium falls throughout the 1970's followed by a robust increase in the next two decades, showing an overall pattern that is both qualitatively and quantitatively consistent with the data. To understand this non-monotonic behavior, we decompose the college premium:

$$\omega^* \equiv \frac{\overline{w}_c}{\overline{w}_n} = \frac{\left(\theta_L L_c^{net} + \theta_H H_c^{net}\right)/N_c}{\left(\theta_L L_n^{net} + \theta_H H_n^{net}\right)/N_n} = \frac{\left[\theta_L + \theta_H \left(H_c^{net}/L_c^{net}\right)\right]}{\left[\theta_L + \theta_H \left(H_n^{net}/L_n^{net}\right)\right]} \frac{\left(L_c^{net}/N_c\right)}{\left(L_n^{net}/N_n\right)}$$

[17] A college graduate is defined as an individual who has completed more than two years of full-time investment ($i = 1$). This is analogous to the definition adopted by Autor, Katz and Kearney (2005a) (i.e, those who enroll in college for more than two years) when constructing the empirical counterpart. One difference from Autor et al. (2005a) is that these authors adjust for compositional changes over time in constructing the college premium, whereas we do not. Earlier empirical studies that do not adjust for composition find a very similar shape for the college premium (c.f., Card and Lemieux (2001) and Acemoglu (2002) among others), so the difference is probably not crucial for our purposes.

Figure 4: Decomposing the College Premium

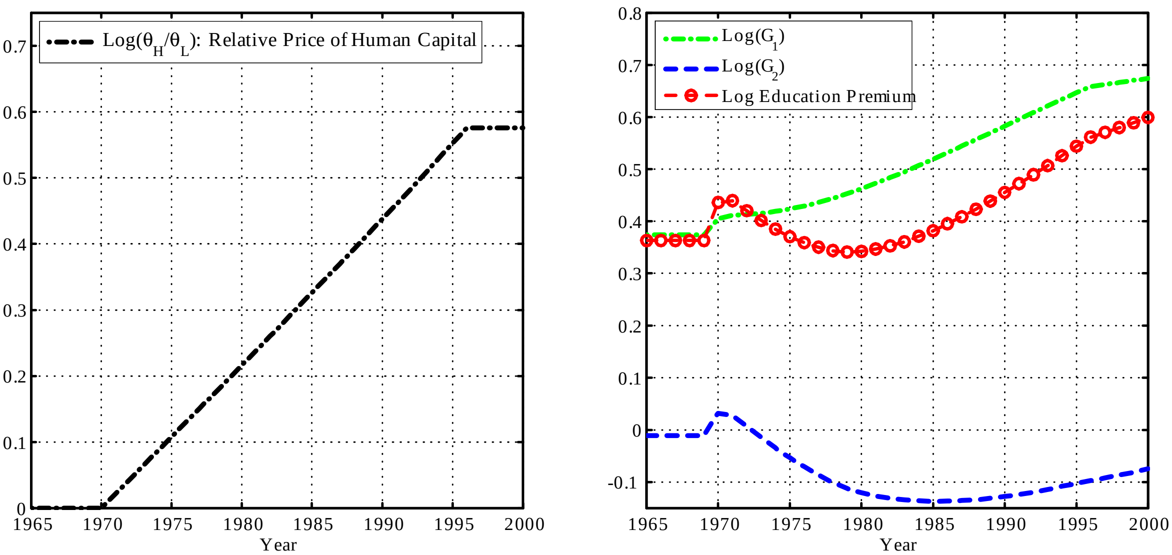


where

$$
\begin{aligned}
H_c^{net} &= \sum_{s=1}^{S} \mu(s) \int_{j \in C} h_s^j \left(1 - i_s^j\right) dj, \\
L_c^{net} &= \sum_{s=1}^{S} \mu(s) \int_{j \in C} l^j \left(1 - i_s^j\right) dj
\end{aligned}
$$

are, respectively, the human capital and raw labor supplied to the market by college graduates. Other aggregates are defined analogously, and the subscripts "$c$" and "$n$" denote college- and high-school-graduates, respectively. Note that we divide both the numerator and denominator by the measure of individuals in that group who are currently active in the labor market: $N_c = \sum_{s=1}^{S} \mu(s) \int_{j \in C} 1\left\{i_s^j \leq \chi\right\} dj$ to get average wages for each group. Note that $(L_c/N_c)$ is equal to the average hours devoted to the labor market (that is, average hours *not* spent on training) by college graduates. Finally, let $H_c^{net}/L_c^{net} = k_c$, and $H_n^{net}/L_n^{net} = k_n$. Divide and multiply the previous equation by $\theta_L$, and take logs to get

$$
\log \omega^* = \log \underbrace{\frac{1 + (\theta_H/\theta_L)k_c}{1 + (\theta_H/\theta_L)k_n}}_{G_1} \underbrace{\frac{(L_c^{net}/N_c)}{(L_n^{net}/N_n)}}_{G_2} = \log G_1 + \log G_2.
$$

The right panel of figure 4 plots the evolution of the logarithms of $G_{1t}$ and $G_{2t}$. The term

$G_{1t}$ depends on variables that adjust slowly (such as human capital stocks), and it grows monotonically over time. In contrast, there is a steep decline in $G_{2t}$, especially immediately after SBTC. The reason is that in response to SBTC workers with a college degree increase their investment time more than non-college workers (due to the difference in ability between the two groups). Because a rise in investment time reduces $L^{net}$, but has no effect on $N$ (conditional on working), $G_{2t}$ declines significantly in the short run after SBTC. Thus, the log education premium (line with circles) initially goes down together with $G_{2t}$, and over time it bounces back when the decline in $G_{2t}$ tapers off and the growth in $G_{1t}$ begins to dominate. Therefore, the differential investment response captured by $G_{2t}$ is crucial for the decline of college premium in the short run.

#### 4.1.3 College premium within experience groups

We now move one step further and examine the college premium *within* different experience groups. Several authors have documented that the fall and rise in the college premium in the U.S. data was largely due to this behavior among young workers, whereas the changes in the college premium among the old was very much muted (Katz and Murphy (1992), Murphy and Welch (1992)). Similarly, Card and Lemieux (2001) have found the same pattern to emerge in British and Canadian data (when individuals were grouped by age instead of potential experience).

In figure 5 we plot the college premium in the model for two different experience groups (1-15 years and 31-45 years). The college premium is higher among more experienced individuals before SBTC, which is consistent with the data (Card and Lemieux (2001) contains empirical counterparts of these graphs for several countries). After 1970, the initial decline and the strong rise in college premium is apparent among younger workers, but there is no fall among more experienced workers and the rise is slightly smaller as well. As a result, the gap between the older and younger workers widens initially (from 0.36 to 0.45) and then narrows to 0.22 in 1990, which is consistent with the empirical evidence documented in the studies mentioned above. The large initial fall among young workers is largely due to the fact that these individuals—who face a longer planning horizon, and hence have a larger marginal benefit from investing—respond to SBTC much more strongly than older individuals. In contrast, the slow but monotonic rise in the college premium among older workers is mainly driven by the price effect without a significant investment response. Therefore, the model studied in this paper offers a new, and in our view fairly plausible, explanation for the differences in the behavior of the college premium among different experience groups.

#### 4.1.4 Cross-sectional wage profiles by education and experience

In documenting the evolution of the relative wages of different groups, labor economists have gone one step further. Since the college premium is essentially the *ratio* of the wages of two groups, further insights can be gained by separately examining the behavior of its numerator

Figure 5: The College Premium By Experience Level in the Model

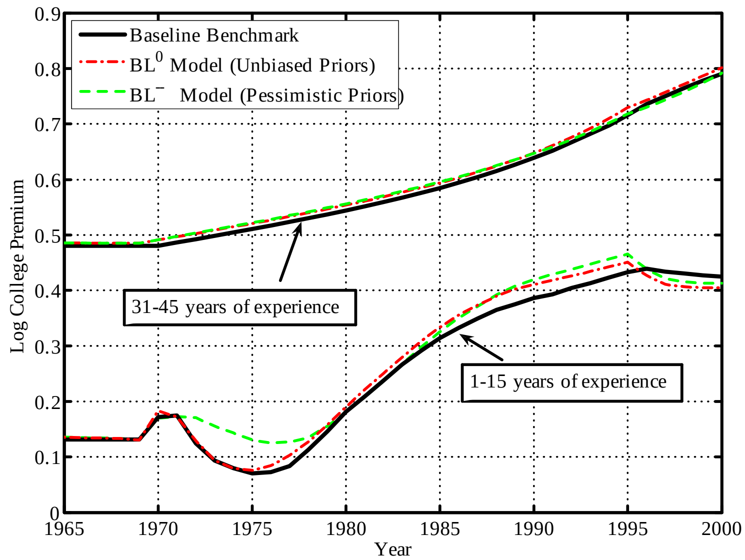


and denominator. For example, the aforementioned decline in the college premium among young workers in the 1970's is due to a fall in the numerator (college graduates' wages), whereas the substantial rise in the 1980's is mainly due to a fall in the denominator (high school graduates' wages). Two models that are consistent with the behavior of the ratio can differ greatly in their implications for the evolution of the wages of each education-experience group (i.e., the numerator and denominator), which makes this information useful for distinguishing between alternative explanations.

In Table 2, we report the changes over time in the wages of different education-experience groups in the U.S. data (reproduced from Katz and Murphy (1992, table 1)). The most striking fact that emerges from this table happens between 1979 and 1987 (last column). First, among high-school graduates, the average wage of workers with few years of experience plummet by 19.8 percent while older workers only see a small decline of 2.8 percent. As a result, the cross-sectional wage profile of high-school graduates significantly *steepens* during this period. Remarkably, the opposite happens among college graduates: young workers see a wage growth of 10.8 percent, whereas older ones only experience a small increase of 1.8 percent. Consequently, the cross-sectional wage profile *flattens* for this group.

We construct the model counterparts of the same statistics with one difference. As we discuss in Section 4.2, the model does not fully capture the magnitude of the slowdown in average wage growth. Given that our focus here is on the *relative* wage changes across education-experience groups, we normalize the data with the mean wage in a given year before calculating the statistics. This allows us to isolate the relative changes without being distracted by the overstated wage growth for all individuals. The model seems to capture

Table 2: Real Wage Changes By Education and Experience Groups, 1971-1987

| Group | | | Change in Log Average Real Wage (multiplied by 100) | |
|---|---|---|---|---|
| Education | Experience | Sample | 1971-79 | 1979-87 |
| 12 | Low | Data | **0.8** | **−19.8** |
| 12 | Low | Model | −1.9 | −9.7 |
| 12 | High | Data | **3.2** | **−2.8** |
| 12 | High | Model | −1.2 | −3.7 |
| 16+ | Low | Data | **−11.3** | **10.8** |
| 16+ | Low | Model | −7.8 | 13.3 |
| 16+ | High | Data | **−4.0** | **1.8** |
| 16+ | High | Model | 3.8 | 2.6 |

Notes: The empirical statistics reported are taken from Katz and Murphy (1992, Table 1). The low (high) experience group is defined as workers with 1 to 5 years of experience (26-35 years of experience) in Katz and Murphy (1992) and those with 1 to 15 years of experience (30-45 years of experience) in our model.

the changes for each education-experience group rather well, not only during the 1980's but also going back to the 1970's. For example, during the 1970's, both in the data and in the model, there is little difference in wage growth by experience levels among high-school graduates, whereas for college graduates there is a larger fall for younger individuals than for older ones. More importantly, the model is consistent with the signs and rough magnitudes of wage changes for three of the four education-experience groups from 1979 to 1987 noted above (see the last column of Table 2). For the fourth group—young high school graduates—the model implies a significant decline in their relative wages consistent with the data, but it does not capture the full magnitude (9.7 percent in the model versus 19.8 percent in the data). Overall, the average cross-sectional wage profile steepens for individuals with low education and flattens for those with a college degree during this period, consistent with the data.

There are three effects that drive the wage changes of high-school graduates of different ages in the 1980's. *First*, young high school graduates also respond to SBTC (even if it is not to the same extent as college graduates) by increasing their on-the-job investment, which reduces their measured wages. *Second*, there is selection: in response to SBTC the ability threshold for college enrollment falls, so the average ability pool of high school graduates—those who choose not to enroll in college—also falls, further reducing their wages. Neither one of these channels reduces the wages of *older* high school graduates: since they have a much shorter horizon they do not increase their on-the-job investment by much, nor do they decide to go back to college to create any compositional change. There is also a *third* effect: young workers have very little human capital, so the main factor they supply is raw labor. Therefore, they suffer from the lower returns to raw labor, but do not benefit from the higher returns to human capital. In contrast, older high-school graduates do have some human capital, so they are able to benefit from SBTC which partly offsets their loss on their raw labor endowment. Put differently, the price effect is negative for the young but close to zero for the old high school graduates. A combination of these three factors, which work in opposite directions for the young and old, explain why the former group experienced a large wage loss while the latter saw no significant change during the 1980's. Notice also that even though SBTC begins in 1970, the three mentioned effects strengthen gradually (as $\theta_H/\theta_L$ rises) over time, and only begin to make a noticeable impact on the wages of the young much later (1980's).[18]

The mechanism for the behavior of the wages of college graduates is similar, but the fact that $\chi < 1$ also plays a role. This is because high-ability individuals who want to increase their investment significantly in response to SBTC have to stay in college longer due to the upper limit on investment while working. As a result, college students accumulate significant amounts of human capital before entering the labor market. Since SBTC raises the value of human capital, the wages of young college graduates do not fall, unlike those of high school graduates (which can be seen in the right panel of figure 10 below).

#### 4.1.5 Within-group inequality

The analysis so far has focused on the evolution of some key moments of the wage distribution. However, a distribution typically contains much more information than what can be summarized by a few moments, and it is possible for a model to be consistent with some summary statistics, but generate patterns inconsistent with the data at a more disaggregated level. Juhn, Murphy and Pierce (1993) have documented an empirical regularity at a very disaggregated level which presents such a challenge. In figure 6 we report the same finding using our data set which covers a longer time span (solid line). The graph plots how each percentile of the wage distribution in 1963 (horizontal axis) has changed between 1963 and 2003 (vertical axis). The first point to note is that wage growth over this period has been systematically different for *every percentile* of the distribution. This shows that there is more to the rise in overall inequality than can be explained by differences in education alone.[19] Second, the relationship between a given percentile in 1963 and wage growth over the subsequent 40 years is almost linear, except at the very low end of the distribution. This implies that wage inequality has increased by a fanning out of the entire distribution, leaving the relative ranking of each percentile largely unchanged over time.

The model counterpart is also plotted in figure 6 (thick solid line). It shows the same general pattern of widening inequality that is spread quite evenly across the wage distribution as

[18] One popular explanation offered in the literature for the steepening of the cross-sectional wage profile of high-school graduates is that the quality of high school education fell significantly during the 70's leading to a fall in the wages of new high school graduates, whereas older workers were already well-vested and their wages were protected through union agreements, preventing their wages from falling. The present model generates the same result in a perfectly competitive model, and through a completely different mechanism that is also consistent with other aspects of the evolution of the wage distribution.

[19] Juhn, Murphy and Pierce (1993) also find the same pattern when they examine the wage distribution for each education group and each age group, making this point even stronger. We have generated the corresponding graphs from our model and they are also qualitatively consistent with the data. We do not discuss them for brevity; they are available upon request.

Figure 6: Log Real Wage Changes by Percentile: Model versus US Data, 1963—2003

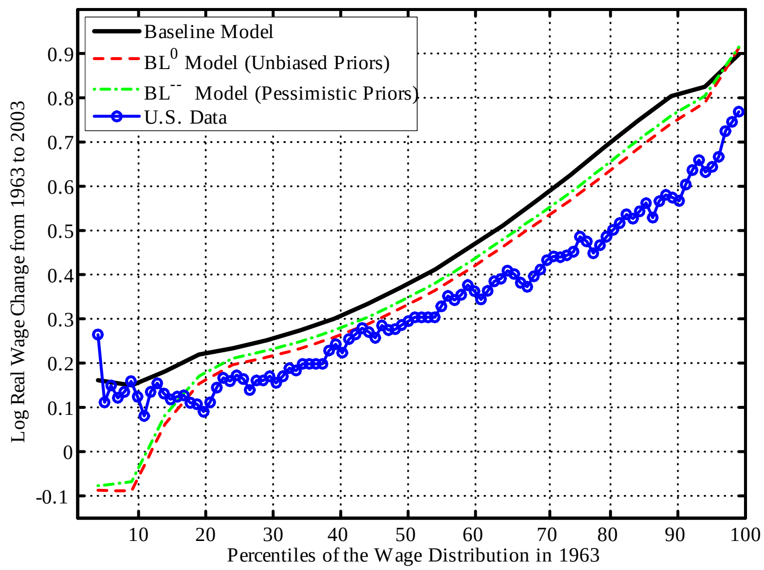


observed in the data. Therefore, despite the fact that the model displays significant nonlinearities in the relative wages in the short run, it displays an almost perfect linearity—that is, a stretching out of the entire wage distribution—in the long run. The mechanism behind this result should be clear from earlier discussions. Wage inequality arises entirely from differences in human capital accumulation rates, which in turn arises from differences in ability (for a given age). Because individuals' investment response to SBTC is monotonically increasing in their ability, those with high ability have both higher wages in 1963, and a higher wage growth in the subsequent 40 years; see figure 10 below. The existence of this same pattern in the data suggests that this mechanism appears to be an important channel behind the rise in within-group inequality.

## 4.2 Evolution of Average Wages

**Stagnation of Median Wages and the Productivity Slowdown.** We now turn from the second moments of the wage distribution to the first moment, that is, the changes in the average wages over this period. Macroeconomists and labor economists have documented two closely related trends: the slowdown in labor productivity and the stagnation of median wage growth, which both started with a sharp fall in 1973 and persisted until about 1995. For example, Juhn, Murphy and Pierce (1993) report that the median real wage has increased by 2.2 percent per year between 1963 and 1973, but actually *fell* by about 0.3 percent per year between 1973 and 1989. Similarly, labor productivity (measured as the non-farm business

Figure 7: The Growth Rate of Median Wages in the Model, 1965-2000

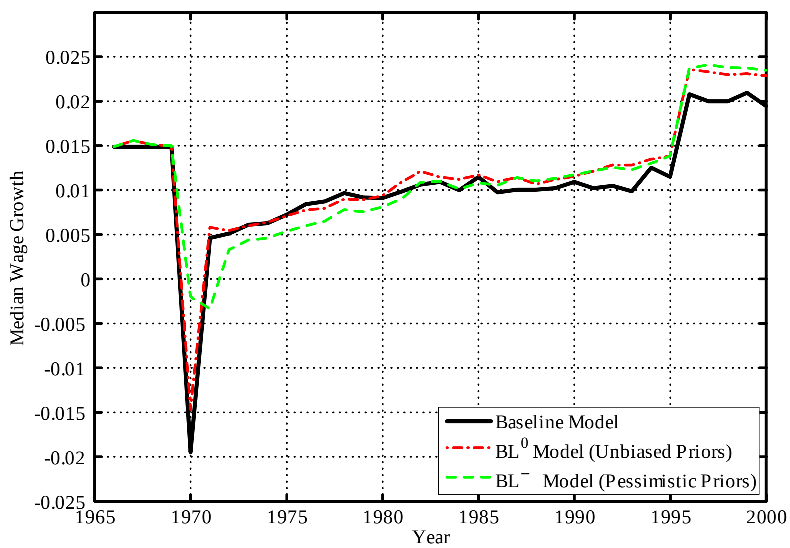


output per hour) has grown by 2.6 percent per year from 1955 to 1973, but only by 1.45 percent per year from 1973 to 1995.[20]

Figures 7 and 8 plot, respectively, the growth rates of median wages and labor productivity implied by the model.[21] Both series fall sharply immediately after SBTC starts in 1970. Thus, the model is able to generate the sharp initial slowdown, although this happens three years earlier than in the data. After the initial fall, the median wage continues to stagnate: it grows at 0.46 percent per year from 1970 to 1979, and averages 0.81 percent overall until 1995, representing a significant slowdown compared to the 1.5 percent growth during the period before 1970. Similarly, labor productivity in the model grows by only 0.6 percent per year during the 1970's, but recovers faster and averages 1.24 percent per year until 1995.

The are three reasons for the prolonged stagnation in wage growth. First, as noted earlier, workers respond to SBTC by increasing their on-the-job investment, which reduces average wages. The fraction of time invested before SBTC is 7.2 percent (or 2.9 hours in a 40-hour work week) and increases to reach 13.1 percent in 1995 (or 5.2 hours a week). Neither the initial investment level, nor the increase during SBTC appears implausibly large, especially considering that what matters for average wages is the change in $(1 - i)$, which goes from 93 percent down to 87 percent over 26 years. One reason for the relatively small change is that, as discussed above, the investment response is concentrated among young individuals, keeping the change in the population average of investment small. Another reason is that

[20] Authors' calculation from Bureau of Labor Statistics data.

[21] In the model labor productivity simply equals the mean wage rate since there is no capital.

Figure 8: The Growth Rate of Labor Productivity in the Model, 1965—2000

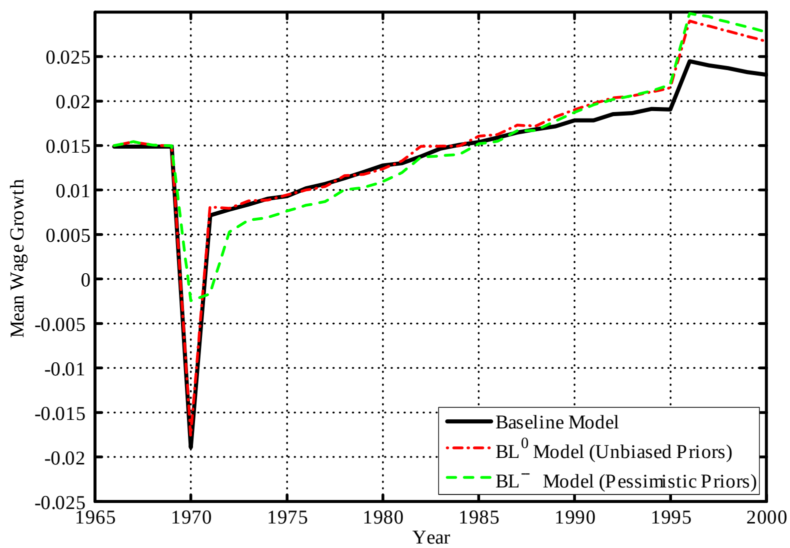


on-the-job investment is bounded from above by $\chi = 0.50$, limiting how much it can rise even for young individuals. Second, as mentioned earlier, SBTC also lengthens the duration of college education for those already planning to go to college. Consequently, the average ability of individuals who remain in the labor market continually falls during SBTC. Because, individuals with lower ability also have low human capital on average, this "selection effect" also depresses average wages and labor productivity after SBTC. Third, and finally, there is a pure price effect resulting from SBTC. Essentially, because the price of raw labor is falling as the the price of human capital is rising, and because the baseline calibration implies that the stock of raw labor is larger than the stock of human capital before SBTC, the change in relative prices puts further downward pressure on the average wage.

To sum up, during this period the labor market is composed of individuals who invest more on-the-job, but who also have lower ability than before, resulting in slow wage and productivity growth. Over time, the increase in the total human capital stock due to both types of investment begins to dominate, resulting in a recovery in both the median wage and labor productivity.

## 4.3 Evolution of Lifetime Wage (Consumption) Inequality

A rather surprising empirical finding from this period is that the rise in consumption inequality has been muted compared to the rise in wage inequality. Figure 9 (line with triangles) plots the variance of log consumption for several years between 1972 and 2000 calculated using the Consumer Expenditure Survey, which shows a very small rise of about 2 log points

Figure 9: The Evolution of Wage and Consumption Inequality: Model vesus U.S. Data, 1965—2000.

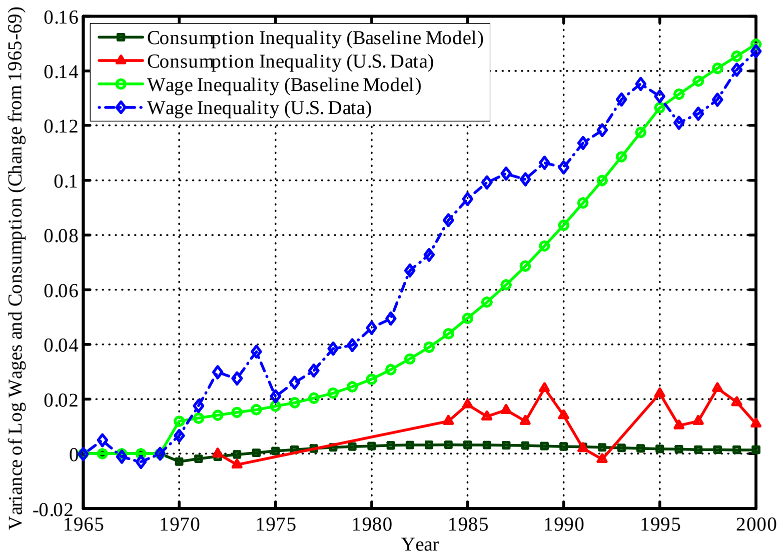


over this period (data taken directly from Heathcote, Storesletten, and Violante (2005)). Although there remains some disagreement about the exact magnitude of the rise in consumption inequality (mainly due to data problems), other studies have also documented findings broadly supporting this conclusion (cf., Krueger and Perri (2006) and Attanasio, Battistin and Ichimura (2004)). Moreover, the change between the 90th and 50th percentiles of the consumption distribution has not tracked the large rise in the 90-50 percentile wage inequality. Autor, Katz and Kearney (2004) document this fact and call it puzzling. The present model abstracts from many features that would be important for a detailed analysis of consumption inequality (such as incomplete markets, retirement savings, demographic changes, etc.). But the model can still address a simple but fundamental question: Has the substantial rise in cross-sectional wage inequality during this period resulted in a parallel rise in *lifetime* income inequality? The line marked with squares plots the evolution of lifetime income inequality in the model, which shows a very small increase of 0.2 log points during SBTC. Since, individuals consume a fixed fraction of their lifetime income in the present model, this is also the rise in consumption inequality.[22] (The figure also reproduces the graphs of wage inequality from

[22]There is a small complication that we sidestep in this discussion, but we fully acount for in the calculations. In particular, the statements in the text that consumption equals the annutized value of lifetime income is correct when individuals experience no shocks during their lifetime, since they begin with zero financial wealth. However, for those individuals who are in existence in 1970, consumption after that date is not simply equal to the annuity value of their remaining lifetime income, but also takes into account the fact that they over-saved or under-saved before the shock given the new path of prices and the implied lifetime wages. Therefore, the inequality in consumption calculated this way goes up more than it would have had we only looked at lifetime

figure 2 for comparison.)

At first blush, it seems surprising that wage inequality could rise in such a systematic fashion without a significant change in lifetime incomes. But note that, as mentioned earlier, wage inequality rises in this model because of a *fanning out* of wage profiles. This can be clearly seen in figure 10, which plots the life-cycle wage profiles in the steady state before SBTC (left panel), and after SBTC (right panel).[23] As can be seen in these figures, those individuals who experience a large increase in their wages later in life are exactly those who make larger investments and accept lower wages early on in response to SBTC. Because future gains are discounted compared to the early losses in calculating lifetime income, the rise in lifetime inequality remains small.

The small change in lifetime incomes we find here is partly anticipated from the findings in Kuruscu (2005), who quantifies the gain due to human capital investment in the standard Ben-Porath model, and finds it to be very small: the difference between following the optimal investment path versus doing no investment at all is less than 3 percent of lifetime income. Compared to that paper, the specification of the human capital model used here is different, as is our calibration: in particular, $\alpha$ is set equal to 0.80 here and 0.94 in Kuruscu (2005). As shown in that paper, a lower value of $\alpha$ generates a larger gain from human capital accumulation, so the total benefit of investment is larger in the present model: the population average of gains is 11.9 percent of lifetime income using the factor prices in 1995, and 7 percent using the factor price in 1983 (the mid point of the SBTC period). As can be expected, there is also a lot of heterogeneity across individuals, with the benefit ranging from a mere 0.57 percent for the lowest percentile of the ability distribution to 36.7 percent for the highest percentile. Nevertheless, because the *rise* in consumption inequality quantified in figure 9 is related to the *incremental* gains resulting SBTC, it remains much smaller than these *total* gains. To sum up, the present model offers a new mechanism that is consistent with a large increase in wage inequality but a small change in lifetime inequality. A fuller investigation of this model for consumption facts is left for future work.

## 4.4 The Rise in the Relative Supply of College Labor

While the main focus of this paper is on the evolution of the wage distribution, the model also makes predictions about the behavior of college enrollment, and consequently, about the change in the relative supply of college-educated labor during this period. Figure 11 plots (line with squares) the total hours worked by individuals with a college-equivalent degree or more relative to those with lower educational levels. This measure more than doubles from 1970 to 2000 in the U.S. data. The model counterpart (thick solid line) understates the level of relative supply before SBTC, which is perhaps not surprising, since no attempt was made to

income inequality calculated at birth.

[23]Haider (2001) estimates an econometric process for wages using micro data from 1968 to 1993 allowing for heterogeneity in growth rates. He finds that the rise in wage inequality during this period happened as an increase in the *dispersion* of wage growth rates, consistent with this figure.

Figure 10: Large Rise in *Cross-sectional* Wage Inequality — Small Rise in *Lifetime* Wage Inequality

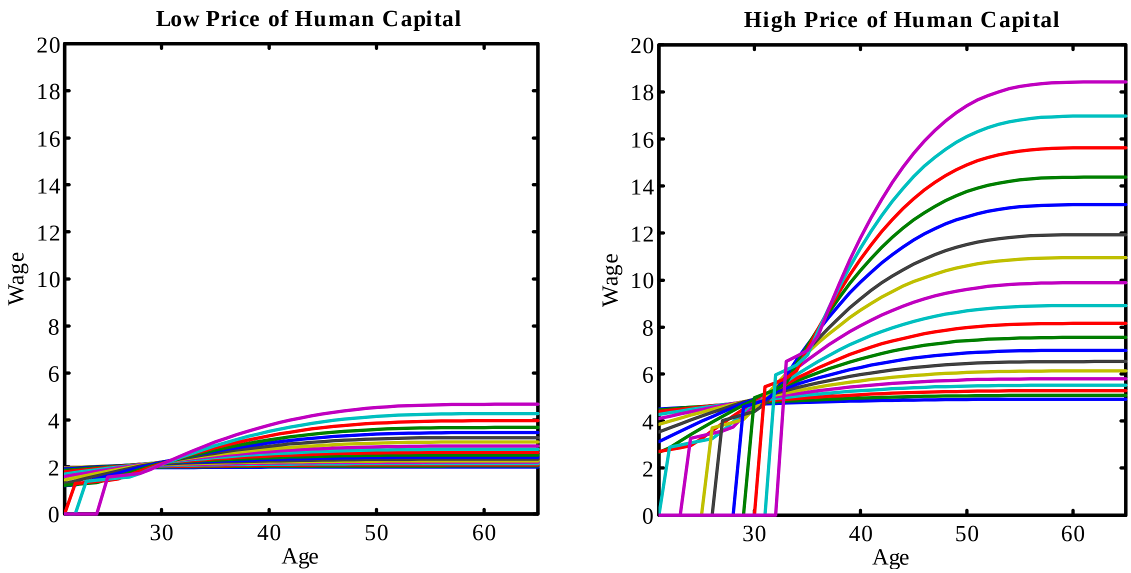


match any aspect of educational attainment in the calibration. However, the relative supply grows significantly in the model, by 0.36 over the entire period, compared to 0.33 in the data. This similarity seems surprising given that in the model college education is modeled merely as a by-product—depending on whether investment exceeds a certain threshold or not—and many potentially important features have been left out, such as tuition costs, changes in the availability of financial aid for college, changes in the quality of education, and so on. This result suggests that SBTC might have played a more important role than these factors in determining the overall rise in educational attainment during this period.[24]

However, despite these plausible implications for the long-run behavior, the model does not capture the behavior of college enrollment rates in the short run. In particular, college enrollment rates were stagnant in the U.S. data in the 1970's (cf., Card and Lemieux (2001)), whereas the model predicts an immediate rise after the onset of SBTC. As we show in Guvenen and Kuruscu (2006) this counterfactual implication is a direct consequence of the assumption we make in this paper that SBTC happens in a completely *disembodied* fashion: that is, the

[24] A related observation during this period has been made by Katz and Murphy (1992, p 52): "[F]or the 1963-87 period as a whole and most strongly for the 1980's, the groups with the largest increases in relative supplies tended to have the largest increases in relative wages." This observation is difficult to reconcile with a model where both the demand and supply are driven by exogenous factors. While one explanation for this fact has been provided by Acemoglu's (1998) model where demand endogenously responds to changes in the supply of different types of labor, the present model is also consistent with this observation since now both the supply of different types of workers and their wages (due to differential human capital accumulation) responds endogenously to changes in demand.

Figure 11: The Relative Supply of College Equivalent Labor: Model versus U.S. data, 1965–2000

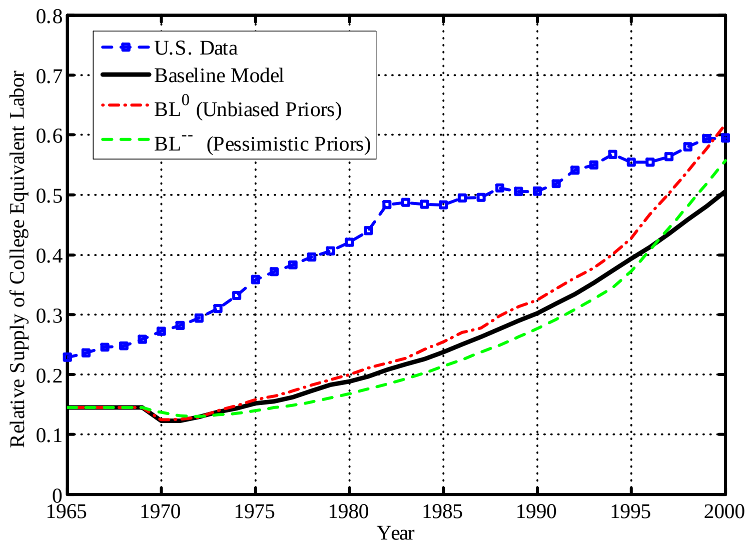


productivity of all human capital rises at the same rate (given by $\theta_H$) regardless of when it is acquired. Consequently, immediately after the beginning of SBTC individuals realize that they can gain immensely by investing today and capitalizing on all future improvements in technology, which causes enrollment rates to rise in the short run. In that paper, we also theoretically show that if part of SBTC comes embodied in new vintages of human capital, under certain parameter conditions college enrollment *falls* in the short run but rises in the long run. This is because, unlike with disembodied SBTC, now individuals do not gain from future improvements in technology by investing in the current vintage of human capital, which makes it optimal to spread investment more evenly over the life cycle. Therefore, investment early in life (i.e., college enrollment) falls, whereas on-the-job investment rises. Because of the latter (and in particular, because on-the-job investment rises *both* in the short run and long run) the present model's plausible implications for the short-run behavior of college premium and average wages are preserved. See Guvenen and Kuruscu (2006) for further details and proofs of these results. Although introducing disembodied SBTC would be a very valuable addition for a more detailed study of the trends in educational attainment, such an extension introduces several layers of complexity, which would distract from the main focus of this paper on the evolution of wages. We leave such an extension for future work.[25]

[25] There are also other reasons why college enrollment can respond differently to major changes in technology than on-the-job investment. For example, as Becker (1964, p.51) observes: "Training in new industrial skill is usually first given on the job since firms tend to be the first to be aware of its value, but as demand develops, some of the training shifts to schools." This is likely to be especially true in the 70's when many new technologies were developed and used in a decentralized fashion across different firms. It would take time for

# 5 Conclusion

In this paper we have studied a parsimonious overlapping generations model of human capital accumulation with some key ingredients: (i) significant heterogeneity in the ability to learn new skills, (ii) raw labor as a factor of production in addition to human capital, (iii) an aggregate production function that takes these two factors—rather than workers with different education levels—as inputs, and (iv) SBTC, modeled as a rise in the relative price of human capital. We have found that the resulting model makes considerable progress towards understanding seemingly disparate trends in the evolution of the wage distribution in a unifying framework.

The model has some other implications for the behavior of wages that have not been discussed in the paper for brevity. For example, the model is consistent with the flattening of, and the downward shift in, average life-cycle wage profiles for subsequent cohorts, documented by Kambourov and Manovskii (2005) from the 70's to 90's. Moreover, the model implies that the bulk of the rise in cross-sectional wage inequality happens at the top of the distribution: the rise in the 90-50 percentile wage differential is more than twice the rise in the 50-10 percentile differential, consistent with the finding in Autor, Katz and Kearney (2005b) (These results are available upon request). However, the model does not explain why college enrollment was stagnant in the 1970's, which suggests that this model, in its current form, may not be suitable for a detailed study of the educational trends; but attempting that as well may be too ambitious a goal for this paper. In Guvenen and Kuruscu (2006) we sketch an extension of the present framework where skill-biased technical changes come embodied in new vintages of human capital, which shows promise in that direction.

A contribution of this paper which could be of independent interest is the introduction of raw labor, which allows us to think about returns-to-skill in the Ben-Porath framework. The particular specifications for the human capital accumulation function and aggregate production function also make the model quite tractable, and thus suitable for potential extensions. Finally, in appendix A we also examine the robustness of our results to lack of perfect foresight about the future prices after SBTC. Although there are different channels that become operational with the introduction of uncertainty, Bayesian learning and heterogeneity in beliefs, overall these effects do not appear to be quantitatively large enough to overturn any of the substantive conclusions we reached in the baseline model.

a systematic body of knowledge to be distilled from the use of these technologies, which can then be taught at formal education institutions. Therefore, one could expect the initial response to SBTC to take place in the form of on-the-job training with a delayed rise in college enrollment.

# A Appendix: An Extended Model

## A.1 Relaxing Perfect Foresight: Bayesian Learning about SBTC

A fundamentally difficult question faced by researchers when studying periods of transition concerns the modeling of individuals' expectations of the future: what do individuals know and when do they know about it? For example, did individuals realize the advent of skill-biased technical change (or, call it the IT revolution, the information age, etc.) soon after it started in the 1970's as assumed in the baseline model? Or were they completely unaware of the fact that many companies were rapidly investing in new IT technologies,[26] and the demand for cognitive skills were rising during that time? Perhaps, as seems plausible, some workers were initially more aware than others—because of differences in education level, social networks, occupations, and so on—but over time all workers have learned about the new regime. But, if so, how fast was this learning process?

In the baseline model, we sidestepped these questions and followed earlier studies that assumed perfect foresight after such regime changes (among others, Greenwood and Yorukoglu (1997), Heckman et al. (1998), Greenwood and Jovanovic (1999)). But despite the obvious appeal of perfect foresight as a clear and simple benchmark, the questions raised above cannot be dismissed easily. Therefore, in this section we relax perfect foresight and consider several alternative assumptions about individuals' beliefs about the future after SBTC. In particular, we allow individuals to differ in their initial beliefs about the evolution of future skill prices—and as an extreme case, we also allow individuals to systematically underestimate the growth of future skill prices (i.e., have pessimistic prior beliefs)—but learn the truth over time in a Bayesian fashion. We now describe these cases in turn.

### A.1.1 Stochastic skill prices

A simple way to allow individuals' beliefs about skill prices to deviate from the truth in a rational manner is by making skill prices uncertain. Assume that the productivity of each factor (raw labor and human capital) follows a random walk with drift in levels:

$$\begin{aligned} \theta_{H,t+1} &= \kappa + \theta_{H,t} + \varepsilon_{t+1}, \quad \text{and} \\ \theta_{L,t+1} &= -\kappa + \theta_{L,t} - \varepsilon_{t+1}, \end{aligned} \tag{14}$$

where $\varepsilon \sim N\left(0, \sigma_{\varepsilon}^{2}\right)$. Note that the innovations to the productivity of each factor sum to zero, so these random shocks only affect the productivity of each factor *relative* to the other.[27] An important advantage of the random walk specification in the levels of $\theta_{H,t}$ and $\theta_{L,t}$ (instead of their logarithms) is that it makes the model tractable. In the special case when $\chi = 1$, this specification allows a closed-form solution for the optimal investment choice in the presence of uncertainty *and* Bayesian learning about future skill prices.

[26] See, for example, Greewood (1997) and Greenwood and Jovanovic (1999) for detailed descriptive evidence on the rapid pace of technology upgrading by firms during this period.

[27] With our parameter choices, the probability of $\theta_H$ or $\theta_L$ becoming negative will be virtually zero.

### A.1.2 Priors and the evolution of beliefs

We assume that before time 0, individuals' forecasts of $\kappa$ had converged to the true value ($\kappa \equiv 0$) after observing a sufficiently long history of skill prices. At the time of the shock (i.e., when $\kappa$ switches from 0 to $\kappa^*$), each individual receives a private signal about the new value of $\kappa$, creating differences in the initial beliefs about the future.[28] This initial heterogeneity seems plausible given that SBTC represents a structural shift (or a regime change), and it is unlikely that all individuals will initially agree on what it entails. In subsequent periods, each individual observes his own wage and updates his beliefs about the new value of $\kappa$. Since the source of uncertainty is aggregate, all individuals observe the same path of prices, and their beliefs eventually converge to each other (and hence to the truth).

We now describe the learning process more formally. First note that some cohorts of individuals are already in existence at time 0, whereas others will enter the economy after this date. Let $n$ be an index that uniquely identifies every individual in the history of this economy (which can be obtained by interacting ability type $j$ and the year of birth of each individual, $t_0$). Each individual $n$ who is alive at time 0 observes an initial private signal about the future growth rate, $\kappa_0^n = \kappa^* + \eta_0^n$, where $\eta_0^n \sim N\left(m_\eta, \sigma_\varepsilon^2/v\right)$. Therefore, individual $n$'s initial forecast is $\kappa_0^n$ with a precision of $v/\sigma_\varepsilon^2$. The population average of initial forecasts is $\kappa^* + m_\eta$. Prior beliefs are unbiased when $m_\eta = 0$, and pessimistic when $m_\eta < 0$. We consider both cases below.

Individuals who enter the economy at $t_0 > 0$ observe an initial private signal $\kappa_{t_0}^n = \overline{\kappa}_{t_0} + \eta_{t_0}^n$, where $\overline{\kappa}_{t_0}$ is the average forecast of *existing* individuals in period $t_0$, and $\eta_{t_0}^n \sim N\left(0, \sigma_{\kappa,t_0}^2\right)$ where $1/\sigma_{\kappa,t_0}^2$ is the precision of existing individuals at time $t_0$. This structure ensures that individuals who are born after the start of SBTC have the same average forecast ($\overline{\kappa}_{t_0}$), and the same precision $\left(1/\sigma_{\kappa,t_0}^2\right)$ as individuals already in existence in that year.

Every period an individual observes his own wage, which can be written as:

$$w_{s,t}^j = Z_t\left[\left(1-\theta_{H,t}\right)l + \theta_{H,t}h_{s,t}^j\right]\left(1 - i_{s,t}^j\right),$$

where we used the normalization $\theta_{H,t} + \theta_{L,t} = 1$ made above. Since individuals know the values of all variables except $\theta_{H,t}$, each wage realization reveals the price of human capital (and, consequently, raw labor) in that period. Two consecutive realizations of an individual's wage can then be used to identify $\kappa^* + \varepsilon_{t+1}$ ($= \theta_{H,t+1} - \theta_{H,t}$). Given individual $n$'s optimal forecast $\widehat{\kappa}_{t-1}^n$ at time $t-1$ and his wage realization at time $t$, his optimal forecast of $\kappa$ at time $t$ is recursively given by

$$\widehat{\kappa}_t^n = \left(\frac{v+t-1}{v+t}\right)\widehat{\kappa}_{t-1}^n + \left(\frac{1}{v+t}\right)\left(\theta_{H,t} - \theta_{H,t-1}\right), \tag{15}$$

with precision $\left(v+t\right)/\sigma_\varepsilon^2$.

[28] These shocks can be thought of as the "news shocks about future TFP" studied in recent work, for example, by Jaimovich and Rebelo (2006), and Christiano, Motto, and Rostagno (2006).

**Individuals' Income Maximization Problem.** Let $V_{s,t}^{j}\left(h^{j};\widehat{\kappa},\theta_{H}\right)$ denote the lifetime income of an individual with a human capital stock of $h$, and a current forecast, $\widehat{\kappa}$. Clearly, the dependence of an individual's lifetime income on $\widehat{\kappa}$ comes from the fact that his investment behavior depends on his perception of the sequence of future prices which depends on $\widehat{\kappa}$. Individuals solve:

$$V_{s,t}^{j}\left(h_{s}^{j};\widehat{\kappa},\theta_{H}\right)=\underset{i_{s}^{j}}{\text{Max}}\quad\left[Z_{t}(\theta_{L}l+\theta_{H}h_{s}^{j})(1-i_{s}^{j})+\left(\frac{1}{1+r}\right)E\left(V_{s+1,t+1}^{j}\left(h_{s+1}^{j};\widehat{\kappa}',\theta_{H}'\right)|\widehat{\kappa},\theta_{H}\right)\right]$$

subject to $V_{S+1,t-s+S+1}^{j}\left(h_{S+1}^{j};\widehat{\kappa},\theta_{H}\right)\equiv 0$, and equations (9, 14, 15). The optimality condition for investment choice at time $t$ is:

$$C^{j}(Q_{s,t}^{j})'=E_{t}\left(\delta\theta_{H,t+1}+\delta^{2}\theta_{H,t+2}+...+\delta^{S-s-1}\theta_{H,t+S-s-1}\right),\tag{16}$$

which only differs from (11) in the appearance of the expectations operator here. Substituting the future prices of human capital using equation (14) yields:

$$C^{j}(Q_{s,t}^{j})'=E_{t}\left\{\delta\left(\theta_{H,t}+\kappa+\varepsilon_{t+1}\right)+\delta^{2}\left(\theta_{H,t}+2\kappa+\varepsilon_{t+1}+\varepsilon_{t+2}\right)+...+\delta^{S-s-1}\left(\theta_{H,t}+\sum_{m=1}^{S-s-1}\left(\kappa+\varepsilon_{t+m}\right)\right)\right\},$$

Notice that the marginal benefit of investment is a linear function of both the shocks to the future prices of human capital (the $\varepsilon_{t+k}$'s) and the unknown parameter $\kappa$, which follows from the random walk structure in levels assumed above. As a result, the right hand side can be simplified to obtain:

$$C^{j}(Q_{s,t}^{j})'=b_{1,s}\theta_{H,t}+b_{2,s}\widehat{\kappa},\tag{17}$$

where $b_{1,s}$ and $b_{2,s}$ are some age-dependent positive constants.[29] The key point to observe here is that optimal investment only depends on the mean of the posterior beliefs, $\widehat{\kappa}$, but not on its variance, $\sigma_{\kappa,t}^{2}$. In other words, the *uncertainty* faced by an individual about future prices of human capital plays no role in the optimal investment decision. Nevertheless, this does not mean that individuals will make the same decisions as in the baseline model: even if $\widehat{\kappa}\equiv\kappa^{*}$ for all individuals in all periods, now $\theta_{H,t}$ is stochastic in equation (17), which will affect the average investment rate in the economy. We return to this point shortly.

**Belief Heterogeneity and the College Premium.** Despite the fact that any given individuals' uncertainty about skill prices does not affect investment behavior, the fact that each individual holds different mean beliefs does affect average investment, and therefore the behavior of the college premium. To see this, we compare the investment decisions of two otherwise identical individuals (i.e., same age and ability level) who only differ in their forecasts, given by $\kappa^{*}+\eta$ and $\kappa^{*}-\eta$, respectively. From equation (12), optimal investment is a convex function of marginal benefit (as long as $\alpha>0.5$), which implies that the average of these two individuals' investment will be higher

[29] $b_{1,s}=\frac{\left(1-(1+r)^{-S+s}\right)}{r}$ and $b_{2,s}=\left(\sum_{m=1}^{S-s-1}(1+r)^{-m}m\right)$

than if they both forecast $\kappa$ correctly.[30] Furthermore, it can also be shown that such a mean preserving spread will increase investment more for younger individuals (because $b_2/b_1$ is larger for these individuals, so the same mean preserving spread in forecasts of $\kappa^*$ will create a larger dispersion in the marginal benefit of young individuals), and that it will increase investment more among high-ability individuals. Putting these pieces together, it follows that even *without* SBTC, higher belief heterogeneity alone can result in higher investment among high-ability individuals, leading to a fall in the college premium. As individuals learn over time and beliefs converge to each other, this effect will weaken and the college premium will rise again. In the next section, we examine whether this is a quantitatively important channel for the behavior of (and especially, for the initial decline in) the college premium after SBTC.

### A.1.3 Quantitative Analysis of the Bayesian Model

**Model 1:** *Bayesian Learning with Unbiased Priors* **($\mathbf{BL^0}$)**

This model is obtained by assuming that the initial signal is unbiased: $E(\eta^n) = m_\eta = 0$. We refer to this version of this model as the $BL^0$ model.

**Model 2:** *Bayesian Learning with Pessimistic Priors* **($\mathbf{BL^-}$)**

We consider a second case where the initial signal is very pessimistic, so that the average forecast of the growth of skill prices is well below the truth. In particular, we assume $E(\eta^n) = -\kappa^*$ so that the average of the initial forecast in the population is $\widehat{\kappa}_0 = \kappa^* - \kappa^* = 0$. In this case, individuals on average do not realize the advent of SBTC, and because of the (symmetric) Normal distribution assumed for beliefs, half of the population initially forecast that skill prices will continually fall at all future dates. Moreover, in the calibration below, we will choose the parameters such that almost all individuals will initially underestimate the true growth rate of skill prices, and learning will be slow during SBTC. We refer to this pessimistic version as the $BL^-$ model.

**Calibration of Shocks and Priors.** The standard deviation of $\varepsilon$ is calibrated such that the model is consistent with the variability of the college premium observed in the data. In our data set, $\sigma(\Delta\omega_t^*)$ is 1.7 percent per year during the period 1963 to 2003. As could be anticipated (and will become clear below), when $\sigma_\varepsilon$ is small the stochastic versions of the model behave very much like the deterministic baseline model analyzed above. Therefore, we choose a slightly higher target for the volatility of the college premium, $\sigma(\Delta\omega_t^*) = 2.5$ percent, and set $\sigma_\varepsilon = 0.0025$ to match this target. The implied annual volatility of the relative skill prices, $\sigma(\Delta \log(\theta_H/\theta_L))$ is 1.4 percent.

[30] More precisely, let $Q^j_{s,t}(\widehat{\kappa})$ denote the investment level when current forecast is $\widehat{\kappa}$. Then we have:

$$\begin{aligned}\frac{Q^j_{s,t}(\kappa^*+\eta)+Q^j_{s,t}(\kappa^*-\eta)}{2Q^j_{s,t}(\kappa^*)} &= \frac{\left[(b_1\theta_{H,t}+b_2(\kappa^*+\eta))\right]^{\alpha/(1-\alpha)}}{2\left[(b_1\theta_{H,t}+b_2\kappa^*)\right]^{\alpha/(1-\alpha)}}+\frac{\left[(b_1\theta_{H,t}+b_2(\kappa^*-\eta))\right]^{\alpha/(1-\alpha)}}{2\left[(b_1\theta_{H,t}+b_2\kappa^*)\right]^{\alpha/(1-\alpha)}}\\ &= 0.5\left[1+\frac{b_2\eta}{(b_1\theta_{H,t}+b_2\kappa^*)}\right]^{\alpha/(1-\alpha)}+0.5\left[1-\frac{b_2\eta}{(b_1\theta_{H,t}+b_2\kappa^*)}\right]^{\alpha/(1-\alpha)} > 1\end{aligned}$$

Table 3: PARAMETERIZATION OF THE BAYESIAN MODELS

| Parameter | | | Value |
|---|---|---|---|
| $\sigma_\varepsilon$ | Standard deviation of SBTC shocks | | .0025 |
| $v$ | Measure of dispersion of initial beliefs | | 2.5 |
| Model: | | $BL^0$ | $BL^-$ |
| *Parameters calibrated to match 1965-69 targets:* | | | |
| $E\left[A^j\right]$ | Average ability | .069 | .069 |
| $\sigma\left(l^j\right)/E\left[l^j\right]$ | Coeff. of variation of labor endowment | .0503 | .0503 |
| $\sigma\left[A^j\right]/E\left[A^j\right]$ | Coeff. of variation of ability | .242 | .242 |
| *Parameter calibrated to match 1995 wage inequality:* | | | |
| $\Delta\log\left(\theta_H/\theta_L\right)$ | Annual change in skill-bias (1970-1995) | 2.13% | 2.24% |

Because beliefs and the speed of learning are inherently difficult to infer from the data (which is why the perfect foresight assumption is so popular!), we calibrate them to provide an upper bound on the effects of imperfect foresight. To this end, first note that the dispersion of the initial forecasts of $\kappa$ is given by $\sigma_\varepsilon^2/v$. A larger value of $v$ reduces the initial heterogeneity in beliefs *and* slows down the speed of learning, as can be seen from equation (15). We choose $v = 2.5$, which implies that in the $BL^-$ model 99.4 percent of individuals initially (i.e., in 1970) underestimate the true growth rate of skill prices. This choice also implies a slow rate of learning: 10 years after the start of SBTC, 86.7 percent of individuals still underestimate the true growth rate of skill prices. The choices of the remaining parameter values are the same as in the baseline model, except those that needed to be recalibrated to match the same targets described earlier in Section 3. The recalibrated values of these parameters are reported in Table 3.

### A.1.4 The College Premium (Bayesian model)

In the $BL^0$ model (dash-dot line in figure 3), the college premium displays a pattern that is qualitatively very similar to the baseline model, but the decline in the premium during the 70's is somewhat larger (11 log points) compared to the baseline (8 log points). Notice, however, that the baseline model differs from $BL^0$ in *two* dimensions: In addition to the fact that the latter model features imperfect foresight and Bayesian learning, it also features skill prices that are stochastic, which is not the case in the baseline model. In order to disentangle the two, we set $v = 10000$, which effectively eliminates all the heterogeneity in beliefs and Bayesian learning from the $BL^0$ model. We find that the college premium in this case is almost identical to that in the $BL^0$ model (not shown in the graph), which shows that the heterogeneity in beliefs in the parameterization of the $BL^0$ model is too small to have a quantitatively significant effect. Thus, the only quantitatively significant difference is that skill prices are stochastic in the $BL^0$ model, unlike in the deterministic baseline model.

We next turn to the $BL^-$ model. Despite the fact that priors are pessimistic and individuals learn slowly, the college premium still falls by 8 log points during the 70's (dashed line in figure 3). It then rises to reach 54 log points in 1995 (compared to 57 log points in the data). The decline in the college premium even in this case shows that the differential investment channel highlighted in the baseline model is strong enough even when the benefits of investment are underestimated on

Figure 12: The College Premium with High Uncertainty about Future Skill Prices

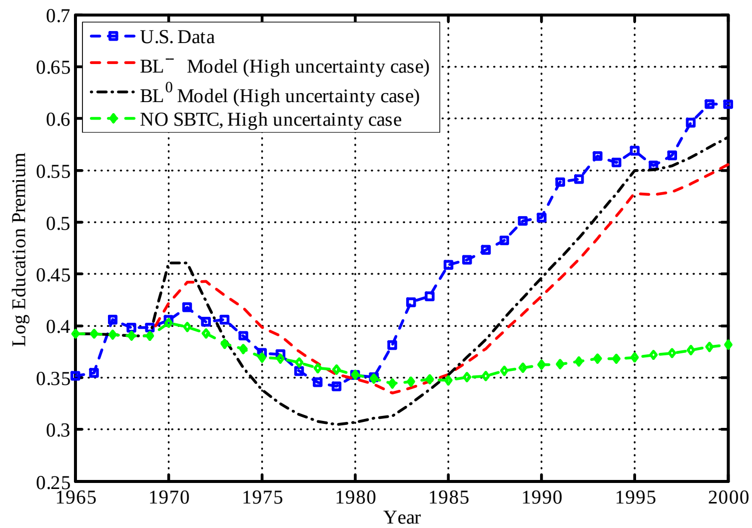


average. One effect of the pessimistic priors is that the college premium reaches its peak level in 1972 (compared to 1971 in the data), and its bottom in 1982 (instead of 1979 in the data). Given that the choice of 1970 as the beginning year of SBTC is somewhat arbitrary it is not clear how important these discrepancies are.

It is useful to discuss when, and how, imperfect foresight could affect the college premium in this model. To examine this, we further increase the heterogeneity in initial forecasts (set $\sigma_\varepsilon/\sqrt{v}$ to 50 percent of $\kappa^*$), and further slow down the speed of learning ($v = 3.5$). These choices imply $\sigma\left(\Delta\omega_t^*\right) = 3.4$ percent per year—double the value in the data. Hence, this case arguably serves as an upper bound on the effect of the lack of perfect foresight. Figure 12 plots the college premium with this new calibration. As can be seen here, the decline in the college premium is more significant now compared to the baseline calibration: it falls by 15.5 log points in the case with unbiased priors (compared to 11 log points before), and by 10.8 log points in the case with pessimistic priors (compared to 8 log points before). The behavior is not affected qualitatively, however.

Finally, to isolate the effect of belief heterogeneity, we set $\kappa^* \equiv 0$ in the last exercise so that there is no SBTC after 1970. Thus, the only thing that happens in 1970 is a pure "belief shock." Now, the college premium falls by about 5 log points during the 70's (line with diamonds), after which point it slowly recovers as individuals gradually learn the truth. This result could be anticipated from our previous discussion about the effect of belief heterogeneity on the evolution of the college premium. The conclusion we draw from this exercise is that if the increase in belief heterogeneity is sufficiently large, this alone will result in a decline in the college premium in the short run. However, for parameter values that we consider plausible, this channel does not seem strong enough to be quantitatively important.

To summarize, these results show that the decline in the college premium in the short run after SBTC does not critically depend on the assumption of perfect foresight. If individuals exhibit moderate heterogeneity in their beliefs immediately after SBTC begins, this has little effect on the college premium. Even when a substantial majority of the population underestimates the true growth rate of the price of human capital for a prolonged period of time, the college premium falls considerably in the short run. Furthermore, if the heterogeneity in beliefs is larger, the decline in the college premium only gets larger.

### A.1.5 Other implications

The remaining differences between the baseline model and the Bayesian counterparts are relatively modest. One difference is observed in the rise in within-group inequality (figure 6) where the baseline model matches the wage growth in the lower 20 percentiles quite well, whereas the models with Bayesian learning ($BL^0$ and $BL^-$) underestimate it quite significantly. Another difference is in the initial fall in average wages after SBTC (figures 7 and 8). The pessimistic priors model ($BL^-$) generates a smaller initial fall than both the baseline model and the $BL^0$ model, mainly because it features a weaker initial investment response to (an underestimated) SBTC. However, because of that weak investment response wages also recover more slowly in that model: the stagnation of median wage growth in the 1970's is 0.36 percent in the $BL^-$ model compared to 0.45 in the $BL^0$ model and 0.46 in the baseline model.

# B Appendix: Allowing for Imperfect Substitution in Production

## B.1 Estimating the Curvature of the Production Function

There are no existing estimates in the literature of $\rho$ that would guide our calibration. One difficulty with directly estimating $\rho$ from data is that our production function (4) features such inputs as human capital and investment time that are very difficult to measure directly in the data. There is, however, a large literature that has estimated a different elasticity, one that measures the degree of substitutability in a CES production function that takes the labor supplied by college and non-college workers as inputs (cf., Katz and Murphy (1992), Acemoglu (2002)). This elasticity, denote it with $\phi$, is obtained by running the following regression:

$$\log \omega_t^* = a_0 + a_1 t - \frac{1}{\phi} \log\left(N_{c,t}/N_{n,t}\right) + error, \tag{18}$$

where $N_c/N_n$ is the labor supply of college-educated workers relative to non-college workers, using the notation developed above.

Notice that all the variables that appear in this regression can be generated from our model as well, which suggests one way to calibrate $\rho$. Essentially, we can choose $\rho$ such that when we run the same regression above using simulated data from our model, we obtain the same estimate of $\widehat{\phi}$ as in earlier studies. This approach however presents its own challenges. The main difficulty is that when $\rho < 1$ and there are aggregate shocks (which are required to run the regression above), the equilibrium prices of raw labor and human capital will depend on $H$ and $L$, which in turn depend

Table 4: Estimating the Katz-Murphy (1992) Regression using Simulated Data

| | $\rho = 1$ | $\log \omega_t^* = a_0 + a_1 t - (1/\phi) \log (N_{c,t}/N_{n,t}) + error$ | | | |
|---|---|---|---|---|---|
| | Change from Baseline: | $\widehat{\phi}$ | $R^2$ | $corr\left(\omega^*, \frac{N_c}{N_n}\right)$ | $std\,(\Delta\omega_t^*)$ |
| (1) | Baseline | 2.47 | 0.88 | −0.87 | 0.025 |
| (2) | $\chi = 0.75$ | 2.49 | 0.84 | −0.82 | 0.029 |
| (3) | $N_c/N_n = 0.36$ | 2.79 | 0.85 | −0.86 | 0.023 |

Notes: T=30. The statistics are the medians of 100 simulations

on the distributions of human capital and investment time in the population (see equations (3) and (5)). One would then have to track the evolutions of these distributions over time, since they become state variables of the model, and use an algorithm such as the one developed by Krusell and Smith (1998) to make computations feasible. Given this additional complexity, and given that the model with $\rho = 1$ delivered quite plausible implications for wages in sections 3 and 4, it is useful to examine how the elasticity $\widehat{\phi}$ implied by this choice compares to empirical estimates.

Table 4 reports the value of $\widehat{\phi}$ implied by our model when $\rho = 1$. To obtain these results, we solved the stochastic version of the model introduced in appendix A, but without SBTC, and therefore, also without heterogeneity in beliefs (that is we set $\kappa = 0$ throughout in equation (14), and $\nu \equiv 0$). Then we simulated data for 30 years and estimated the regression in (18), repeating the exercise 100 times. The reported statistics are the median values from these estimations.

In the first row, the estimated elasticity is 2.47 and the regression has an $R^2$ of 0.88. The correlation of the college premium and the relative supply is also negative and very large. For comparison, Autor, Katz and Kearney (2005a) obtain $\widehat{\phi} = 2.05$ with an $R^2$ of 0.94, in the closest specification to ours (see column 3 of table 2 in their paper). Similarly, Hamermesh (1993) surveys the empirical estimates of $\widehat{\phi}$ that exist in the literature, which also concentrate around this value. In the second row, we relax the upper bound on on-the-job investment ($\chi = 0.75$), which may potentially alter the sensitivity of the college premium to relative supply of college graduates. This change makes little difference to the results. Finally, in the last row, the average relative supply is calibrated to twice its average value in row 1, which increases $\widehat{\phi}$ somewhat, to 2.79. Overall, however, these changes have little effect on the estimates of $\widehat{\phi}$. To sum up, the choice of $\rho = 1$, which implies perfect substitution between $H$ and $L$ in our model (in other words, *infinite* substitution elasticity!) generates, not only a finite elasticity between college and non-college workers, but also a relatively small value similar to that observed in the data.

To understand this result, consider the response of this economy to a positive innovation, $\varepsilon > 0$. From equation (14), this shock results in a permanent increase in $\theta_H/\theta_L$, and therefore, in a rise in human capital investment. At the extensive margin, this increases the relative supply of college educated workers, starting the year after the shock. At the intensive margin, on-the-job training rises differentially for high- and low-ability individuals, which causes the college premium to fall. Therefore, following a positive $\varepsilon$ shock, the supply of college workers rises while, at the same time, the college premium falls. This negative correlation makes it *appear as if* the high supply of college workers reduces the college premium, as would be the case in a CES production function. In other words, the disturbance term in the regression above is positively correlated with the relative supply of college workers, which when ignored biases the estimated coefficient $(1/\phi)$ upward (and the estimated elasticity $\phi$ downward). As a result, our model with $\rho = 1$ generates a finite substitution elasticity in the regression above. Instead, if in the model the supply of college labor were to change for completely exogenous reasons (and the workers to be be added were selected randomly from among each group) then the estimate of $(1/\phi)$ would be zero, correctly revealing the infinite supply elasticity. We conclude that if skill prices fluctuate in a persistent manner, the estimate of $\phi$ from the regression in (18) is likely to be downward biased due to an omitted variable bias. The true elasticity may be much higher—in the example presented, it is infinite.

## B.2 Model results

The results of the previous section suggests that $\rho = 1$ may be an empirically relevant benchmark. Nevertheless, it is of interest to know how the model would behave if there was some degree of imperfect substitution. In light of the results above (and given the described difficulty of using the indirect procedure above to exactly pinpoint what $\rho$ should be) we take a $\rho = 0.8$, which implies an elasticity between human capital and raw labor of $(1-\rho)^{-1} = 5$ (compared to an infinite elasticity in the models in the text).

Figure 13 illustrates the evolution of key variables.[31] The upper dash-dotted line in Panel A is the growth rate of relative productivity $\theta_H/\theta_L$ and the middle (bottom) line is the growth rate of relative prices, $P_H/P_L$, for $\sigma = 5$ ($\sigma = \infty$). Since we have targeted the change in total wage inequality, it is not surprising that it is very similar to the benchmark figure. As the figure illustrates the growth rate of relative price of human capital is larger during 1970-1995 when $\sigma = 5$ compared to the benchmark: the price of labor actually falls by 0.5% per year during 1970-1995 as opposed to the 0.23% increase in the benchmark. The reason can be explained as follows. We assume there is no change in $\theta_H/\theta_L$ after 1995. However $P_H/P_L$ declines after 1995 due to the increase in human capital stock until the economy reaches its steady state. Hence, a larger increase in $P_H/P_L$ is required during 1970-1995 to match the increase in total wage inequality (since individuals expect a decline in relative price of human capital after 1995). Panel C shows the evolution of the college premium. The larger increase in $P_H/P_L$ compared to the benchmark during 1970-1995 increases the price effect, which results in a smaller decline in the college premium(4.7 log points here as opposed to 8 log points in the baseline model).

Since the growth of $P_H/P_L$ is larger in this example compared to the benchmark, $P_L$ grows at a slower rate here. As panel D illustrates, this implies a smaller growth in average wages than in the benchmark since the stock of raw labor is larger than the stock of human capital. The median (mean) wage growth is 0.16 percent (0.62 percent) per year during 1970-1995 as opposed to 0.81 percent (1.26 percent) per year in the benchmark analysis. Although the average wage growth is smaller in this case compared to the benchmark, the decline in average wages in 1970 is smaller

[31] Note that in the baseline model we had $P_H/P_L = \theta_H/\theta_L$. Moreover, in that case we normalized $P_H/P_L$ to be one in 1969. To make the results here easy to compare to baseline case, we choose the level of $\theta_H/\theta_L$ in 1969 such that $P_H/P_L = 1$ in 1969 as before. This ensures that, in this exercise, the three parameters we choose to match the three targets in 1965-1969 period are the same as in the benchmark analysis. The only parameter that we need to recalibrate is $\kappa^*$, which is chosen, as before, to match the increase in total wage inequality.

Figure 13: Evolution of the US Wage Distribution ($\sigma = 5$)

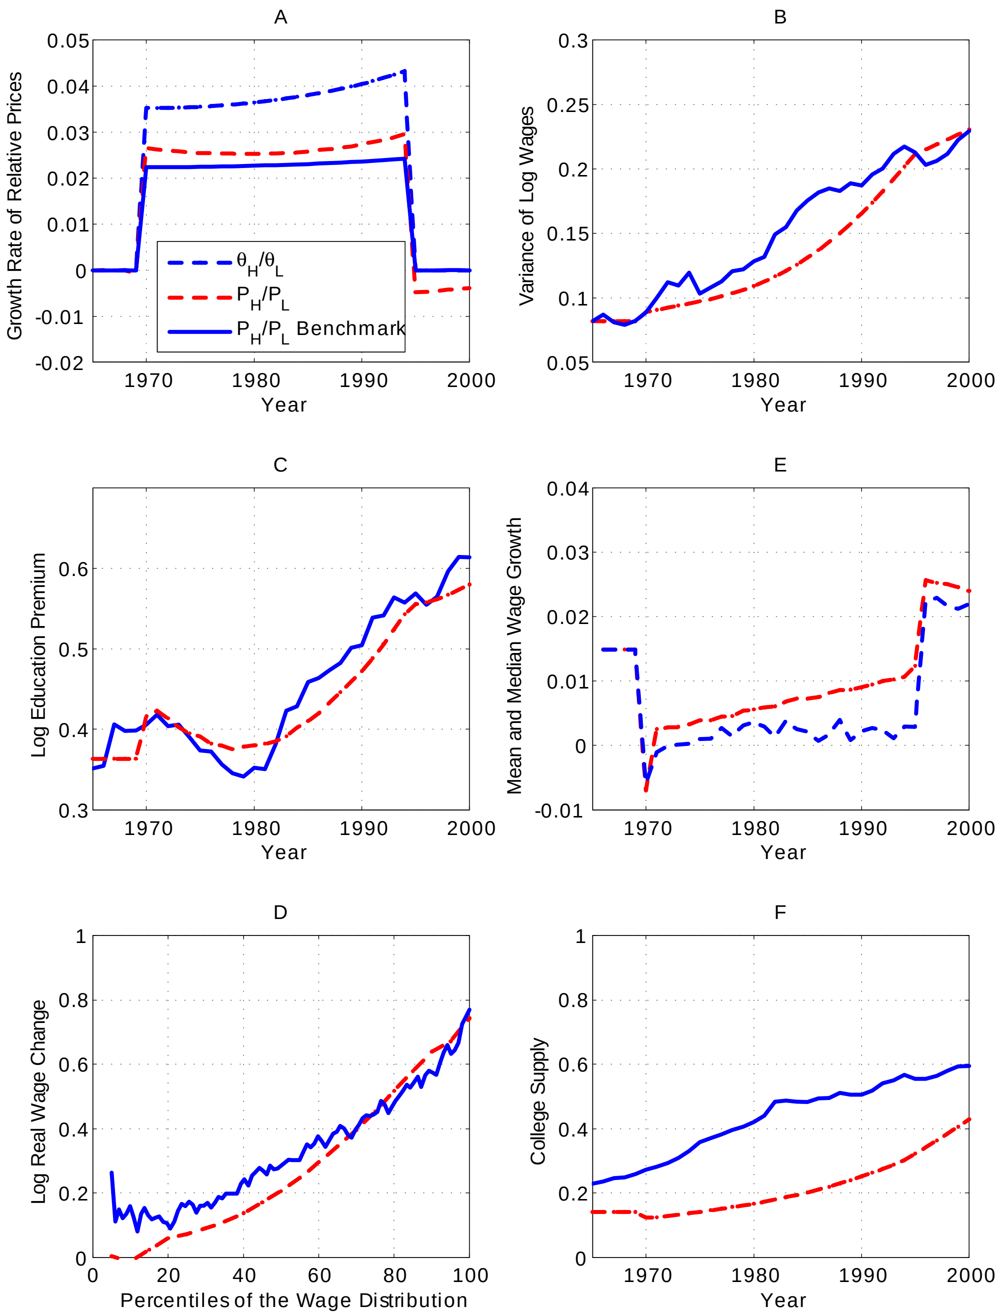

than the benchmark due to a smaller investment response. The smaller wage growth is also reflected in panel E which illustrates the wage growth across different percentiles. In this case, the model does a better job in matching the wage growth at upper percentiles of the wage distribution but it underestimates the wage growth at lower percentiles. Finally, panel F shows the relative supply of college graduates. Since the investment response is smaller in this case compared to the benchmark, the increase in college supply is also smaller.

To summarize, the investment response is smaller when $\sigma = 5$, which implies a smaller decline in college premium and average wages initially, and a smaller overall increase in college supply. By and large, the changes resulting from allowing for imperfect substitution appear to be rather small, and none of the substantive conclusions we reached in the baseline case are altered by this extension.

# References


[1] Abowd, John M., and David Card (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica,* Vol. 57, No. 2, pp. 411-445

[2] Acemoglu, Daron (1998): "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality," *Quarterly Journal of Economics,* Vol. 113, pp. 1055-1089

[3] Acemoglu, Daron (2002): "Technical Change, Inequality, and Labor Market", *Journal of Economic Literature*, March 2002, pp. 7-72.

[4] Attanasio, Orazio, Erich Battistin, and Hidehiko Ichimura (2004): "What really happened to consumption inequality in the US?" IFS Working paper.

[5] Autor, David, Lawrence F. Katz, Alan B. Krueger (1998): "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics.* 113, pp. 1169-1214.

[6] Autor David H., Lawrence F. Katz, and Melissa S. Kearney (2004): "Trends in U.S. Wage Inequality: Re-Assessing the Revisionists," Slides from earlier version of (2005a) paper.

[7] Autor David H., Lawrence F. Katz, and Melissa S. Kearney (2005a): "Trends in U.S. Wage Inequality: Re-Assessing the Revisionists," MIT working paper.

[8] Autor David H., Lawrence F. Katz, and Melissa S. Kearney (2005b): "Rising Wage Inequality: The Role of Composition and Prices," MIT working paper.

[9] Baker, Michael (1997): "Growth-rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings," *Journal of Labor Economics,* Vol. 15, pp. 338-375.

[10] Barron, John M., Mark C. Berger, Dan A. Black (1997): "How Well Do We Measure Training?", *Journal of Labor Economics*, Vol 15, pp. 507-529.

[11] Becker, Gary S. (1964): "Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education," reprinted 1993, University of Chicago Press.

[12] Ben-Porath, Yoram (1967): "The Production of Human Capital and the Life Cycle of Earnings." *Journal of Political Economy*, 75(4), pp. 352-365.

[13] Berman, Eli, John Bound and Zvi Griliches (1994): "Changes in the Demand for Skilled Labor within U. S. Manufacturing Industries: Evidence from the Annual Survey of Manufactures." *Quarterly Journal of Economics,* Vol 109, pp. 367-365.

[14] Bound, John and George Johnson (1992): "Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations," *American Economic Review,* Vol 82, pp. 371-92.

[15] Card, David, and John Dinardo (2002): "Skill Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles," *Journal of Labor Economics,* Vol 20, pp. 733-83.

[16] Card, David and Thomas Lemieux (2001): "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis," *Quarterly Journal of Economics*, 116, pp. 705-46.

[17] Caselli, Francesco (1999): "Technological Revolutions," *American Economic Review,* pp. 78–102.

[18] Christiano, Lawrence, Roberto Motto, and Massimo Rostagno (2006): "Monetary Policy and Stock Market Boom-Bust Cycles," Working Paper, Northwestern University.

[19] Davis, Steve J., Felix Kubler and Paul Willen (2003): "Borrowing Costs and the Demand for Equity Over the Life-cycle," Working Paper, University of Chicago.

[20] Gourinchas Pierre-Olivier, and Jonathan A. Parker (2002): "Consumption Over the Life Cycle," *Econometrica,* Vol. 70, pp. 47-89.

[21] Greenwood, Jeremy (1997): "The Third Industrial Revolution: Technology, Productivity and Income Inequality," AEI Studies on Understanding Economic Inequality, Washington, DC: The AEI Press, 1997.

[22] Greenwood, Jeremy, and Mehmet Yorukoglu (1997): "1974," *Carnegie Rochester Conference on Public Policy,* Vol 46, pp 49-94.

[23] Greenwood, Jeremy, and Boyan Jovanovic (1999): "The IT Revolution and the Stock Market," *American Economic Review,* Vol 89, No 2, pp 116-121.

[24] Guvenen, Fatih (2005): "An Empirical Investigation of Labor Income Processes," Working paper, University of Texas at Austin.

[25] Guvenen, Fatih (2006): "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?, *American Economic Review,* forthcoming.

[26] Guvenen, Fatih, and Burhanettin Kuruscu (2006): "Understanding the Evolution of the U.S. Wage Distribution: A Macroeconomic Analysis," Working Paper, University of Texas at Austin.

[27] Haider Steve J. (2001): "Earnings Instability and Earnings Inequality of Males in the United Stated: 1967-1991, *Journal of Labor Economics,* Vol. 19, pp. 799-836.

[28] Hamermesh, Daniel (1993): "Labor Demand," Princeton University Press.

[29] Heathcote, Jonathan, Kjetil Storesletten, Gianluca Violante (2005): "The Macroeconomic Implications of Rising Wage Inequality in the United States," Working Paper, NYU.

[30] Heckman, James J. (1976): "A Life-cycle Model of Earnings, Learning, and Consumption," *Journal of Political Economy* 84(4), 11-44.

[31] Heckman, James J.; Lochner, Lance J. and Taber, Christopher R. (1998): "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogenous Agents." *Review of Economic Dynamics*, 1(1), pp. 1-58.

[32] Hornstein, Andreas, and Per Krusell (1996): "Can Technology Improvements Cause Productivity Slowdowns?" *NBER Macroeconomics Annual,* Vol. 11, pp: 209-259.

[33] Hornstein, Andreas, Per Krusell, and Giovanni L. Violante (2006): "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," Working Paper, New York University.

[34] Huggett, Mark, Gustavo Ventura, and Amir Yaron (2006a): "Human Capital and Earnings Distribution Dynamics," *Journal of Monetary Economics,* Vol 53, pp. 265- 90.

[35] Huggett, Mark, Gustavo Ventura, and Amir Yaron (2006b): "Sources of Life-Cycle Inequality," Working Paper, Wharton School, University of Pennsylvania.

[36] Jaimovich, Nir, and Sergio T. Rebelo (2006): "Can News about the Future Drive the Business Cycle?" CEPR Discussion Paper No. 5877.

[37] Juhn, Chinchui, Kevin M. Murphy, and Brooks Pierce (1993): "Wage Inequality and the Rise in Returns to Skill," *Journal of Political Economy*, pp. 410-442.

[38] Kambourov, Gueorgui and Iourii Manovskii (2005): "Accounting for the Changing Life-Cycle Profile of Earnings," Working Paper, University of Pennsylvania.

[39] Katz, Lawrence F., and Kevin Murphy (1992): "Changes in Relative Wages: Supply and Demand Factors," *Quarterly Journal of Economics,* pp. 35-78.

[40] Katz, Lawrence F., and David H. Autor (1999): "Changes in the Wage Structure and Earnings Inequality," in Orley Ashenfelter and David Card, eds., Handbook of Labor Economics, Volume 3A, Amsterdam: Elsevier-North Holland, 1463 –1555.

[41] Krueger, Dirk and Fabrizio Perri (2006): "Does Income Inequality Lead to Consumption Inequality? Some Empirical Findings and a Theoretical Explanation," *Review of Economic Studies*, Vol. 73(1), 163-193.

[42] Krusell, Per and Anthony A. Smith, Jr. (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, Vol 106, No 5, pp. 867-896.

[43] Kuruscu, Burhanettin (2005): "Training and Lifetime Income," *American Economic Review*, Vol. 96, No 3, pp. 832-846.

[44] MaCurdy, Thomas E. (1982): "The use of time series processes to model the error structure of earnings in a longitudinal data analysis," *Journal of Econometrics*, Vol. 18, No 1, pp 83-114.

[45] Machin, Stephen and John Van Reenen (1998): "Technology and Changes in Skill Structure: Evidence from Seven OECD Countries." *Quarterly Journal of Economics.* 113:4, pp. 1215-44.

[46] Meghir, Costas, and Luigi Pistaferri (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica,* Vol. 72, pp. 1-32.

[47] Moffitt, Robert A., and Peter Gottschalk (1994): "Trends in the Autocovariance Structure of Earnings in the US, 1969-1987, Working Paper, Brown University.

[48] Murphy, Kevin M., and Finis Welch (1992): "The Structure of Wages," *Quarterly Journal of Economics,* Vol 107, pp. 285-326.

[49] Violante Giovanni L. (2002): "Technological Acceleration, Skill Transferability, and the Rise in Residual Inequality," *Quarterly Journal of Economics,* Vol. 117, pp 297-338.