

Discussion of Guvenen and Kuruscu  
by Francesco Caselli

Between 1970 and 2000 the wage structure in the US underwent significant changes (several other countries experienced broadly similar patterns). Some of the key facts are as follows.

- The college premium (the log-difference in average wages between college graduates and other workers) for younger cohorts of workers declined in the 1970s, and then increased very steeply in the 1980s and 1990s.
- The college premium for older workers declined less (or not at all) initially, and the subsequent rise was also less pronounced.
- *Within* wage inequality (the variance of the residual in a log-wage regression controlling for education, experience, gender, etc.) increased substantially throughout the period.
- Mean and median wages have stagnated (at least in the 1970s and 1980s)

On these four facts the evidence seems overwhelming and there is little serious controversy that I know of. Some authors, including the authors of the present paper, complete the list with a fifth, more controversial, fact:

- Available data show an increase in consumption inequality that is an order of magnitude smaller than the increase in wage inequality.

The goal of this extremely ambitious paper is to provide a *unified* as well as *quantitative* explanation for all five facts. It does so by developing a theory of costly on-the-job learning, with heterogeneity in learning abilities, in a dynamic environment where the rate of growth of the demand of learned skills accelerates. This results in a most elegant and stimulating exercise that also matches the quantitative features of the data quite well.

In this comment I begin by reviewing the mechanisms at work in the model in greater detail. I then discuss informally the empirical plausibility of the proposed explanation, and conclude that the authors fail to make the empirical case for the quantitative importance of their key mechanism, namely costly on the job learning. I then compare the explanation in this

paper to existing explanations in the literature. One superficial advantage of the current paper is that it uses a (mostly) unique mechanism to fit all five facts above. In contrast, the existing literature takes a more piecemeal approach, using different stories for different facts. Since the various bits and pieces in the literature seem all rather plausible, however, and collectively more plausible than the single one proposed here, I am not inclined to be unduly swayed in this case by appeals to Occam's razor. This paper *is*, however, more ambitious than most of the existing literature in that it sets a higher bar for explaining the facts, as it seeks to provide a quantitative account. Future work will have to respond to the challenge thrown by Guvenen and Kuruscu.

## 1 Paper Overview

The paper is built on three key ingredients. The first is *costly human-capital accumulation* throughout a worker's lifetime, in the form of unpaid time spent acquiring productive skills while formally on the job. Begin with a completely stationary environment (i.e. no change in the rate of change of technology). Then a typical worker's lifetime wage profile is upward sloping. The worker spends the early part of his life using all of his productive time acquiring human capital, i.e. in education. He then joins the labor force, but continues to invest some of his time further building up human capital. The cost of this time is the foregone earnings. Time investments in human capital decline over time, as the retirement horizon approaches. Throughout the working years the wage increases for two reasons. First, human capital becomes larger and larger, which means that the worker earns more income for unit of time spent working. Second, the fraction of time devoted to working (as opposed to accumulating human capital) also increases over time.

The costly human-capital accumulation model has rich implications for wage inequality. To show this, Guvenen and Kuruscu introduce a second ingredient, *heterogeneity in ability to learn marketable skills*. Specifically, the addition to a worker's skill capital per unit of time spent learning is higher for some workers (called high ability) than for others (low ability). Hence, high learning-ability workers will spend longer in school, and then will also spend a larger fraction of their formal work time further accumulating skills. They will therefore have greater educational achievement, higher wages on average, and steeper life-time wage profiles. Figure 1 draws the implications

of this for the relative wage profiles of college and high-school graduates in the initial steady state, say circa 1965. Because of the steeper wage profile of college graduates the college premium is larger for older cohorts.

Figure 1 here

It is useful to note here the importance of the assumption that ability only affects the production function for human-capital accumulation. If ability also affected productivity directly, high-ability individuals would have a correspondingly higher opportunity cost of time, and their age-wage profile would be flatter (and possibly they would achieve less schooling).<sup>1</sup>

The third key ingredient of the model is an *increase in the rate of growth of the market price of learned skills*, or an increase in skilled biased technical change (SBTC), which is assumed to have occurred between 1970 and 2000. This change has two orders of effects on wages. First, there is a *price effect*, which is simply the (exogenous) increase in the price of learned-skills. The price effect tends to increase the relative wages of high-ability workers. Second, there is an *investment effect*. Agents respond to the increased rate of growth of the price for skills by spending more time accumulating skills. Crucially, high-ability agents' investments respond disproportionately, so the investment effect *reduces* the relative wages of high-ability workers. In general, therefore, the acceleration in SBTC has ambiguous effects on the wage structure.

Most of the interesting action in the model comes out of this ambiguity, and particularly from the fact that the relative strength of the price effect and the investment effect vary over time, between ability groups, and between age groups. Regarding the latter, notice that price effects tend to be larger for older workers (who have accumulated larger stocks of skills), but the investment effect is more relevant for younger workers (who will reap the benefits for a longer time span).

Consider the initial impact of the acceleration on the cross-sectional wage profile of college graduates. In Figure 2 the top solid line shows the initial profile (what we called the 1965 profile) and the dashed line the profile shortly after the onset of faster SBTC - what one might see in 1975, say. Since the shock is an increase in the rate of change of prices for skills, the price effect is

---

<sup>1</sup>High-ability workers already have a higher opportunity cost of time in the model as it is, but the setup is such that the net effect always goes in the direction of making high-ability workers invest more.

positive, particularly for older workers who have a larger accumulated stock of skills. But since prices have had very little time to increase, the effect is very small. On the other hand, there is a potentially sizable negative investment effect for younger workers, as they foresee considerably high future prices for learned skills. Since the high prices are far in the future, on the other hand, the investment effect will be very small for the old. As a result, the only significant impulse on wages is the negative investment effect for younger workers.

The qualitative effects on the wage profile of high-school graduates, showed by the two lower curves in Figure 2, is very similar. However, as discussed, low-ability workers respond less to the higher future returns to skills because their investments in accumulating knowledge are less productive. Hence, the wage profile for high-school graduates changes less. The net effect, as shown in the Figure, is an initial decline in the college premium for younger workers and very little change for older workers.

Figure 2 here

The subsequent rise in college premia is largely explained by the fact that, eventually, the price effect seriously kicks in for all groups. In the 1980s, young college graduates continue to invest much more than young high-school graduates, but in the meanwhile the price of human capital has increased sufficiently that the net result is an increase in the college premium. In the 1980s, the higher college premium for older workers reflects both the higher price of human capital and the increased on-the-job investment in human capital by this generation in the 1970s. The subsequent increase for both age groups is showed in Figure 3. This mechanism therefore explains the first two facts of the initial list.

Figure 3 here

A potential difficulty in which this mechanism runs is that one of its robust predictions seems to be that the cumulative growth of the college premium should be larger for older workers than for younger workers (as is implicit in Figure 3). Relative to the 1960s, at any subsequent point in time the older should enjoy a larger price effect and a smaller investment effect. This is indeed the case even in the main quantitative exercise of the paper: between 1970 and 2000 the old college premium is predicted to rise 30 log

points for the old and 25 for the young (Figure 7). In the same figure, the data show a 20 log-point increase for the young and 30 for the old. Given that the interplay between price effect and investment effect are at the heart of the paper I am not sure this point is entirely minor.

I now turn to the implications of unpaid learning time, heterogeneity, and SBTC for *within* inequality. As far as I can tell the key to the success of the paper in predicting a relatively smooth increase in within inequality over time is that there are pervasive crossing of wage-age profiles within narrowly defined education groups. The crossings are due to the fact that higher-ability agents (conditional on educational achievement) sacrifice more pay early on in order to accumulate skills, so their take-home pay is less than lower-ability agents with similar educational background. Later on relative wages switch as the high-ability reap the benefits of greater life-time skills.

The key effects of SBTC are still the price effect and the investment effect, but with crossings the directions of these effects on inequality *for the young* is reversed. In particular, for the young, the investment effect now increases inequality, as the high-ability respond disproportionately, leading to a further decline of their wages relative to the low-ability ones. The price effect instead reduces inequality among the young. For the old, instead, the investment effect reduces inequality, and the price effect increases it, as was the case for the college premium. So if the investment and price effects on inequality go in opposite directions for the young and the old, how can inequality go up in all groups? Because the investment effect dominates for the young (leading to more inequality) and the price effect dominates for the old (also leading to more inequality)! This is all shown in Figure 4.

Figure 4 here

A slight complication on this discussion of within inequality is that crossings, while pervasive, are not universal. There may be educational groups in which the wage profile of high ability agents is entirely above the wage profile of low-ability ones. This can happen because ability also affects the amount of skills learned in school, so it is possible that in some educational groups the high ability have such a disproportionate stock of knowledge that their greater on-the-job investments are insufficient to bring their earnings below those of the lower-ability members of their age-education group. Applying the now-familiar reasoning to these cases we should have that inequality within the young of these groups should have declined in the 1970s, not increased.

The model therefore suggests the following test: identify education groups where the model predicts no crossings, and check that inequality within these groups declines in the 1970s. It would be interesting to know the answer to this question.

The last of the well-established empirical facts that the paper matches is the stagnation of the median wage. This stagnation is due to three factors. The first two are entirely within the economic logic of the paper's central argument. Namely, an increased (and increasing) fraction of time on the job spent in unpaid learning - rather than producing; and a deterioration of the average quality of the labor force at any point in time, as the most talented spend more time in formal schooling. The third contributing factor is a by-product of a technical assumption that is entirely outside the economic argument of the paper (this is that agents are also endowed with a stock of "raw labor," whose price has declined for exogenous reasons). It was not entirely clear what was the relative contribution of the "economic mechanisms," v. ad hoc assumptions in allowing the model to match the facts on median wages.

## **2 How important is unpaid learning effort on the job?**

The most striking lacuna in the paper is that there is virtually no effort whatsoever to convince the reader that unpaid time spent learning while formally on the job is a quantitatively important phenomenon, much less that it has increased substantially since the 1960s, which is the crucial channel giving rise to all the (interesting) wage effects in the model. It seems inconceivable that if significant numbers of workers were engaged in this type of activity the literature on labor relations would not have produced a description of the phenomenon as well a number of case studies. Similarly, had there been significant changes in the nature and use of time at work it is hard to imagine that no anecdotal accounts would have surfaced in the popular press. The authors' case would have been much more persuasive had they been able to point to at least some case studies/anecdotes. It is hard for me not to interpret their apparent inability to do so as circumstantial evidence against their contention.

In thinking about the plausibility of the story it is helpful to give more de-

tailed consideration to what exactly the on-the-job investments in the model may represent in the real world. At one end, we can give them a *literal interpretation*. In the literal interpretation, workers spend a certain number of each work-day, *while at the workplace*, reading books and manuals, experimenting with software, etc., rather than contributing (directly or indirectly) to production. Their wages are curtailed proportionately to the fraction of time spent learning.

Essentially, the literal interpretation amounts to a *mismeasurement theory* of changes in the wage structure. For example, the college premium declines in the 1970s because the decline in effective labor hours, which is much larger for college graduates, is underestimated. Conversely, if time spent learning was not (erroneously) reported in employer and employee surveys (e.g. the CPS) as regular work time we would have observed no decline in the college premium. Even setting aside the issue of the extent to which common sense is compatible with an explanation of declining college premia primarily based on greater time set aside by college graduates to learn new skills, one needs to wonder why firms and workers would misreport labor hours.

Potentially more promising seems to be the following *lose interpretation*. In this interpretation all effort is productive, but while working workers learn useful skills, most of which are portable (i.e. are worker specific but not employer- or job-specific). Because workers are learning skills that will pay off in future employment, their current employers charge them for the implicit teaching by paying them less than their marginal productivity in the current job. Furthermore, since high-ability workers benefit more from this on-the-job learning, they are also charged at a higher rate. It may then be possible to obtain the result that with SBTC the “tax on wages” charged by firms will increase disproportionately for young, high-ability workers. The lose interpretation amounts therefore to a theory of changes in wage inequality based on a (time varying) wedge between wages and marginal products.

Whether a model along these (somewhat more plausible) lines would be able to quantitatively match the opening facts on the wage structure is of course a moot question. Clearly some special feature would be required to insure that the charge for learning survives competition among firms.

### 3 This paper v the literature

The developments addressed in this paper are the focus of a very large literature, and many who have followed this literature feel that it has been relatively successful at explaining the main facts. Most of this literature emphasizes the role of technical change. In this respect, Guvenen and Kuruscu are entirely in line with widespread thinking. But they argue that the literature misses one of the fundamental mechanisms through which SBTC leads to the wage facts. In this section I outline how the Guvenen and Kuruscu explanation differs from received wisdom (which in the process, I re-evaluate).

#### 3.1 College premium

The “majority view” in the profession is that changes in college premia during the 1970s, 1980s, and 1990s result from the interaction of SBTC, which continuously increases the relative demand for skills learned in college, and the relative supply of college graduates.<sup>2</sup> Specifically, the standard explanation for the decline in the college premium in the 1970s is that in that decade there was a large increase in the rate of increase in the relative supply of college-educated workers, which then subsided in the 1980s and 1990s. This pattern is actually reproduced in Figure 13 in the paper. The view is that the fast increase in relative supply more than compensated for the increase in relative demand leading to a temporary compression in the premium. In the 80s and 90s the growth in the supply for college graduates slowed down, which allowed the continued increase in relative demand for skills learned in college to show up as a steep rise in the premium. As a result, the college premium was higher by the end of this period.

Guvenen and Kuruscu propose a completely different interpretation of the 1970s. College graduates’ earnings fell because college graduates worked less - as they shifted to more unpaid human-capital time investments while employed. It then recovered, despite continued greater investments by college graduates, due to the continued increase in the relative price per unit of human capital actually supplied to the employer.

As to why the changes in the college premium (both down in the 1970s and up in the subsequent decades) were more pronounced for younger workers the “standard view” adds a simple (though eminently plausible) twist.

---

<sup>2</sup>Whether explaining the facts requires an acceleration in SBTC, or simply that SBTC continued at the same rates as before the 1970s is more contentious.

Namely, young college graduates are imperfect substitutes for old college graduates. By definition, the increase in the relative supply of college graduates in the 1970s must have been driven by the young, so the “supply glut” hypothesis is consistent with a decline in the “young college premium” with little change in the “old college premium”. I think this can explain also the subsequent pattern in the 1980s and 1990s: the glut of young college graduates of the 70s became a glut of middle-aged college graduates in the subsequent period, preventing the relative wages of this group from raising very much then. Meanwhile, the fall off in the relative supply of young college graduates allowed the relative wages of this group to soar in the 80s and 90s. The latter trend was almost certainly amplified by changes in college curricula in the 1980s that made the skills of young college graduates even more useful in the new technological environment.

In contrast to this (well-rehearsed) argument Guvenen and Kuruscu argue that older college graduates in the 1970s did not respond by significantly increasing human-capital investments, so their relative earnings did not decline as much as those of young college graduates, who shifted into learning mode much more aggressively.

In terms of fitting the college-premium facts, I already mentioned one dimension where Guvenen and Kuruscu’s story falls somewhat short. By the end of the SBTC episode, say by the year 2000, the overall increase in the college premium relative to 1970 is counterfactually larger for the older cohorts, both because the gap in human-capital between old college graduates and high-school graduates has grown more than the equivalent gap for young workers, and because the gap in human-capital investments has grown more for young workers. The standard view, in contrast, should have no problem matching a larger increase for the young.

## **3.2 Within inequality**

In my reading of the literature, SBTC seems to command a significant majority of supporters (without admittedly even coming close to the status of a consensus view) as an explanation for the behavior of the college premium – especially when combined with the labor-supply patterns. The explanation for the rise in within-group inequality is much more a matter of controversy, and candidates that are not based on changes in technology enjoy greater support. Nevertheless, SBTC is clearly a viable option, and I would venture the guess that it would still command the support of a plurality, if not a

majority, of labor and macroeconomists.

The fundamental SBTC explanation for increased within inequality is simply that technical change has not only increased the relative demand for skills learned in college, but also the relative demand for innate skills, or ability (essentially, IQ). Since innate ability obviously varies widely even within narrowly-defined education and demographic groups, this naturally leads to a rise in within-group inequality.

It is sometimes argued that the innate-skill mechanism is inconsistent with the fact that the college premium declines in the 1970s, while within-group inequality increased throughout the SBTC era. Clearly no such inconsistency exists. For SBTC to predict that within-group inequality and the college premium move in lockstep it would have to be the case that innate skills and skills learned in college are perfect substitutes. There is no reason to believe this should be the case. Since the innate-skill distribution is presumably roughly constant, when innate skills and learned skills are imperfect substitutes, the increase in the relative supply of college graduates in the 1970s can depress the college premium without necessarily preventing the “innate skill” premium to rise.<sup>3</sup>

The Guvenen and Kuruscu story for the increase in within inequality is very different. It rests on the “crossing” property of the wage profiles within educational groups, together with disproportionate increases in time devoted to learning by young people with high ability conditional on education (and correspondingly large price effects for old people with high ability, always

---

<sup>3</sup>A simple example of a view of production and of SBTC that is consistent with the patterns of the college premium and of within-group inequality is the following. Output is given by  $Y = [(A_h L_h)^\rho + (A_c L_c)^\rho]^{1/\rho}$ ,  $\rho < 1$ .  $L_u$  and  $L_s$  are the flows of productive services from high-school graduates and college graduates, assumed to be imperfect substitutes, and defined, respectively, as  $L_h = \int a^\sigma \mu_h(a) da$  and  $L_c = \int a^\sigma \mu_c(a) da$ , where  $a$  is innate skill and the  $\mu$ s are the densities of innate skills among high-school and college graduates, respectively. Assume for simplicity that the  $\mu$ s are uniform on  $[0, \bar{a}_h]$  and  $[0, \bar{a}_c]$ , and define  $N_h = \int_0^{\bar{a}_h} \mu_h(a) da$  and  $N_c = \int_0^{\bar{a}_c} \mu_c(a) da$  the masses of high-school and college graduates, respectively. Suppose now that SBTC results in both an increase in the relative productivity of skills learned in college, i.e. an increase in  $A_c/A_h$ , and an increase in the elasticity of output to innate skill, i.e. an increase in  $\sigma$ . It is immediate to show that within-wage inequality (i.e. the ratio of the wages of two workers with the same schooling but different  $a$ ) does not depend on either  $A_c/A_h$  or  $N_s/N_h$ , but it will increase with an increase in  $\sigma$ . On the other hand the college premium, i.e. the average wage earned by college graduates divided by the average wage of high-school graduates increases with  $A_c/A_h$  and decreases with  $N_s/N_h$  (and is independent of  $\sigma$  if  $\bar{a}_c = \bar{a}_h$ . If  $\bar{a}_c > \bar{a}_h$  then the college premium also increases in  $\sigma$ ).

conditional on education). As mentioned above, however, there should be educational groups where the crossing property fails, and in such groups we should observe falling within inequality among the young in the 1970s. A test along these lines would buttress the Guvenen and Kuruscu story against the standard story, that predicts increased inequality within all groups.

### **3.3 Stagnation of the median and mean wage**

The stagnation of median and mean wages through most of the last quarter of the last century remains contentious. However, one popular explanation is closely related to the previous theories of SBTC. This explanation emphasizes the fact that the period in question has seen the advent of the information-technology (IT). IT is a General Purpose Technology, and the arrivals of such technologies are widely thought to entail transitional adjustment and learning costs that show up as slow downs in TFP growth, labor productivity, and average wages. A substantial literature models these effects.

In part, the Guvenen and Kuruscu explanation is entirely in line with the GPT view. One of the three contributing factors to the slowdown in wages is the increase in time taken off from production by workers to get acquainted with the technology. In this sense, the Guvenen and Kuruscu paper may be interpreted as an attempt to model a particular case and quantify the importance of the "adjustment to GPT" argument. As mentioned, however, the relative importance of the learning cost effect on wages and other, more ad hoc mechanism in the paper is not clear.

### **3.4 Consumption inequality**

The potential ability of the model to explain why the rise in wage inequality has seemingly not been matched by a similar increase in consumption inequality is perhaps the most important "victory" of Guvenen and Kuruscu over conventional accounts of SBTC. The latter, as far as I can tell, are unable to avoid predicting a significant rise in consumption inequality, because they can't help predicting an increase in the variance of permanent income.

The problem here is that a quick glance at the evidence that consumption inequality did not rise is enough to destroy any confidence in this claim. The most credible evidence that consumption inequality has increased considerably less than income inequality comes from the Consumer Expenditure

Survey (CEX). Heathcote, Perri, and Violante (2008), among others, document a large and growing discrepancy between aggregate consumption as inferred from CEX and in the national Accounts (see Figure 3 in their paper). Clearly this points to a significant and increasing fraction of missing consumption in the CEX. At a minimum, this should make us very cautious in accepting quantitative assessments of the change in consumption inequality based on these data.

Even accepting that the increase in consumption inequality has not matched the increase in wage inequality, we do not necessarily have to conclude that the variance in permanent income has not changed. We now painfully know that over the last few decades the financial sector has allocated credit based on wildly unrealistic assumptions on the creditworthiness of lower-income households. This has fuelled a consumption boom in the lower and middle segments of the income distribution. Because the consumption increase was excessive relative to the true increase in permanent income these households are now beginning to default.<sup>4</sup>

## 4 College premia and labor supplies: another look

## 5 Conclusions

The paper by Guvenen and Kuruscu is an ambitious, elegant, original, and well-executed attempt to qualitatively and quantitatively explain a variety of patterns in the wage and consumption data over a 30-year period, where such data shows momentous changes. The main shortcoming of the paper is that it does not even attempt to defend the empirical plausibility of its main mechanism, much less establish it. In the absence of any evidence, even merely anecdotal, that unpaid learning time at work is an important phenomenon, and one that has undergone significant changes in recent decades, one has to fall back on one's priors on what is plausible. From this perspective, the view that changes in the wage structure are driven by changes in the percentage of time formally spent at work for which the worker does not get paid may appear somewhat far fetched.

---

<sup>4</sup>Incidentally, given the default option, the behavior of households needs not be irrational. It's the lenders that clearly miscalculated.

In contrast, existing alternatives based on SBTC, imperfect substitutability between different groups of workers, and the inherent costs of adapting to GPTs continue to appear both plausible and to have the potential to adequately account for the facts quantitatively. Admittedly, however, there have been few attempts of embedding them in a unified quantitative model. Perhaps one fruitful way of reading the Guvenen and Kuruscu paper is as a challenge to do just that.

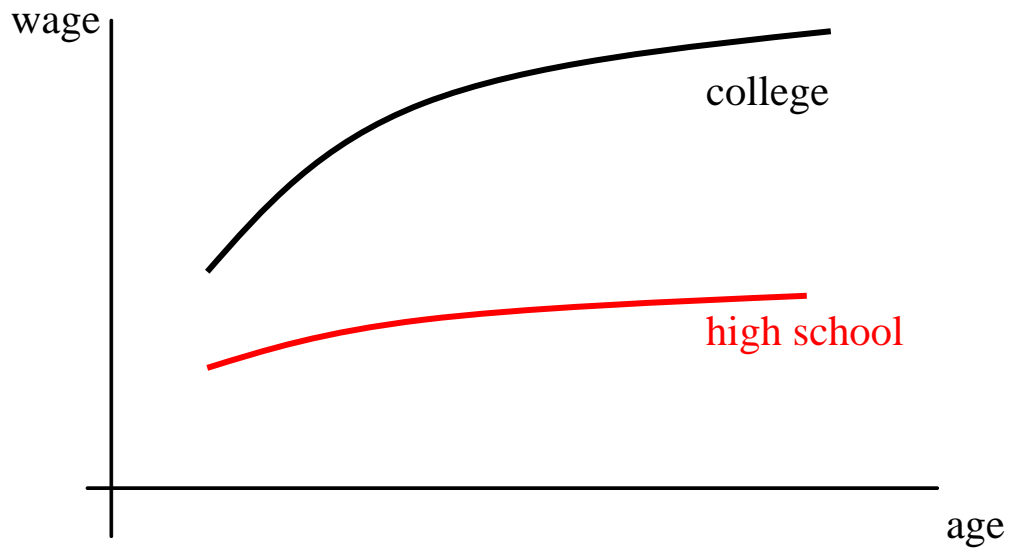


Figure 1: Cross-sectional relation between age and earnings in 1960s




Figure 2: Cross-sectional relation between age and earnings: **1960s v. 1970s**




Figure 3: Cross-sectional relation between age and earnings: **1960s v. 1990s**




Figure 4: SBTC and within inequality