MIT Open Access Articles

Comment

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

As Published: 10.1086/648298

Publisher: University of Chicago Press

Persistent URL: <https://hdl.handle.net/1721.1/134504>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.

Comment

David Autor, Massachusetts Institute of Technology and NBER

I. Introduction

Guvenen and Kuruscu's intensely ambitious paper sits astride the rugged border between labor economics and macroeconomics and from this vantage offers a comprehensive and remarkably parsimonious explanation for perhaps all of the key stylized facts describing the evolution of U.S. wage levels and wage inequality over the three decades between 1970 and 2000. To say that it does not succeed completely at this goal is akin to saying that not all of Columbus's ships arrived intact to America. Success in such a sweeping endeavor is rarely complete, but we must nevertheless admire the effort.

In these comments, I first lay out what I see as the five key facts about U.S. wage structure that the paper attempts to grapple with and that any unifying theory of wage inequality must accommodate. I then lay out the standard workhorse model often used by labor economists to interpret these facts. This framework, attributable to Jan Tinbergen, is sometimes referred to as the "education race" model, following the excellent book by Goldin and Katz, which takes its title from a quotation by Tinbergen. I then compare the Guvenen-Kuruscu model to the education race framework and consider how each model fares in explaining the key facts. With this comparison, I hope to highlight where both the Guvenen-Kuruscu and the education race models hit the mark or fall short. My goal is not to declare a winner—it is not clear that either model even receives a qualifying score—but to underscore which major facts are well understood and which remain elusive. I begin with the facts.

II. Key Facts on U.S. Wage Inequality

There are five main facts about U.S. wage and skill structure that the Guvenen-Kuruscu model grapples with.¹

A. Rising Wage Dispersion

The first fact is the rising unconditional dispersion of U.S. wages. Figure 1, using data from Autor, Katz, and Kearney (2008), plots the evolution of the 90/10 log wage gap for male workers for the period 1963–2005. The series labeled CPS March Full-Time refers to log weekly earnings of full-time, full-year workers, while the series labeled CPS May/ORG Hourly refers to log hourly earnings of all non-self-employed workers. The latter series, which is generally considered more accurate, unfortunately did not become available until 1973. Nevertheless, both series tell a comparable story for the period in which they overlap. Male wage dispersion is either flat or rising modestly in the 1970s, begins rising sharply in 1981, continues at a fairly torrid pace to 1988, and then increases at a slower (but nonzero) rate through the end of the sample period.² The net rise in 90/10 inequality in this time period is substantial. The May/ORG and March series imply rises of 36 and 47 log points, respectively, during 1973 through 2005, starting from a base of 121 log points.³ Residual wage inequality—that is, the inequality of log wage earnings remaining after netting out the estimated effects of education, sex, and potential experience using cross-sectional wage regressions estimated separately by year—followed a similar trajectory, though at reduced amplitude.⁴

Fig. 1. Evolution of male 90/10 log wage inequality, 1963–2005. Source: Autor et al. (2008, fig. 2).

B. The Rising, Falling, and Then Rising Return to Education

A second key fact is the dramatic fluctuation over four decades in the college/high school log earnings gap, which I will call the “college premium” for brevity. The evolution of this premium is depicted in figure 2. The composition-adjusted college premium rose from 40 log points to 47 log points between 1963 and 1971, then dropped back to 40 log points by 1981, before beginning a steep ascent, coincident with the rapid rise in 90/10 inequality. The premium rose by 14 log points between 1980 and 1989, by 10 additional log points between 1989 and 1999, and by an additional 2.5 log points between 1999 and 2005. Thus, similar to the 90/10, there was a sudden spectacular rise in the college premium commencing in the early 1980s. This rise then decelerated but did not reverse starting in the late 1980s. Unlike the 90/10, however, the college premium actually declined in the 1970s, whereas the 90/10 merely plateaued.

C. Age Group Differences in the Rising College Premium

Equally noteworthy is that the rise in the college premium was greater for young workers, those with 0–10 years of potential experience, than for prime-age workers, with 20–29 years of potential experience. This pattern, shown in figure 3 (Autor et al. 2008), is found in the United States,

Fig. 2. Evolution of the college/high school wage gap, 1963–2005. Source: Autor et al. (2008, fig. 4b).

Fig. 3. Evolution of the college/high school wage gap among inexperienced and experienced workers, 1963–2005. Source: Autor et al. (2008, fig. 6a).

the United Kingdom, and Canada and has been ingeniously interpreted by Card and Lemieux (2001b).

D. Rising Real Wages at the Top, Falling Real Wages at the Bottom

A fourth fact of considerable consequence for the evolution of the U.S. wage structure is that absolute wages fell for some groups—most notably, less educated males. Figure 4 shows trends in real log weekly wages for males within five mutually exclusive education groups, ranging from high school dropouts to those with postcollege education, during the years 1963–2005. Wages for each group are normalized at zero in 1963, with subsequent data points representing the difference in real log wage levels between the current year and the 1963 baseline.

Real wage levels rose in unison for all five education groups from 1963 to 1973 and then stagnated in unison from 1973 to 1981. From that point forward, the course of real earnings diverged, with college-educated workers seeing modest to sizable gains, workers with some college experiencing modest losses, and workers with high school or lower education experiencing substantial losses.

Autor et al. (2008) estimate that real, composition-adjusted wages of workers with less than a high school education fell by 15.7 log points

Fig. 4. Indexed changes in real male wages by education group, 1963–2005 (1963 = 100). Source: Autor et al. (2008, fig. 5a).

between 1979 and 1995. Real wages of high school graduates dropped by 9.8 log points. Drops of this magnitude are unlikely to be fully explained by any plausible indexation bias in the Personal Consumption Expenditure inflator (PCE) used by Autor et al.⁵

E. Boom and Bust in College Labor Supply

A final key fact that is at the heart of labor economists' understanding of the evolution of wage inequality is the sharp deceleration in the growth of the relative supply of college graduates to the U.S. labor market in the early 1980s. This deceleration is evident in figure 5, which plots the log relative supply of college versus noncollege workers between 1963 and 2005 (see Autor et al. [2008] for details of the construction of this index). College relative supply trends up at approximately 4.5 log points per year from 1963 to 1982, then discreetly and resolutely downshifts to approximately 2.0 log points per year from 1982 forward. This deceleration—which, tellingly, occurred just when wage inequality began its rapid ascent—was a direct consequence of events that had occurred years earlier. College-going rates declined in the United States in the mid-1970s in response both to the following college wage premium

Fig. 5. Log relative supply of college equivalent labor, 1963–2005. Source: Autor et al. (2008).

and the end of the Vietnam War (Card and Lemieux 2001a). This, in turn, implied a decline in the number of new college graduates 4–5 years later.

Also relevant was the end of the Baby Boom. For most of the twentieth century, new cohorts entering the labor market were significantly more educated than the cohorts that were retiring, and hence the average education of the working population rose steeply with each wave of entry and retirement. By 1982, the last of the Baby Boom cohorts had entered the labor market, and when smaller cohorts followed, the rate of increase in educational attainment of the workforce slowed. These two forces—declines in college-going and the end of the Baby Boom influx—led to a sharp deceleration in the supply of new college graduates.⁶

III. The Education Race Model

Before discussing the subtle and novel Guvenen-Kuruscü explanation for these five facts, it is useful to consider the benchmark model widely used in labor economics.⁷ This framework starts with a constant elasticity of substitution (CES) production function for aggregate output with two factors, college equivalents (c) and high school equivalents (h):

$$Q_t = [\alpha_t(\theta_{H,t}N_{ct})^\rho + (1 - \alpha_t)(\theta_{L,t}N_{ht})^\rho]^{1/\rho}, \quad (1)$$

where N_{ct} and N_{ht} are the quantities employed of college equivalents (skilled labor) and high school equivalents (unskilled labor) in period

t , $\theta_{H,t}$ and $\theta_{L,t}$ represent skilled and unskilled labor augmenting technological change, α_t is a time-varying technology parameter that can be interpreted as indexing the share of work activities allocated to skilled labor, and ρ is a time invariant production parameter. Skill-neutral technological improvements raise $\theta_{H,t}$ and $\theta_{L,t}$ by the same proportion. Skill-biased technological changes involve increases in $\theta_{H,t}/\theta_{L,t}$ or α_t . The aggregate elasticity of substitution between college and high school equivalents is given by $\sigma = 1/(1 - \rho)$.

Under the assumption that college and high school equivalents are paid their marginal products, we can use equation (1) to solve for the ratio of marginal products of the two labor types yielding a relationship between relative wages in year t , w_{ct}/w_{ht} , and relative supplies in year t , given by

$$\ln(w_{ct}/w_{ht}) = \ln[\alpha_t/(1 - \alpha_t)] + [(\sigma - 1)/\sigma] \ln(\theta_{H,t}/\theta_{L,t}) - 1/\sigma \ln(N_{ct}/N_{ht}). \quad (2)$$

This equation can be rewritten as

$$\ln(w_{ct}/w_{ht}) = (1/\sigma)[D_t - 1/\sigma \ln(N_{ct}/N_{ht})], \quad (3)$$

where D_t indexes relative demand shifts favoring college equivalents and is measured in log quantity units.

The impact of changes in relative skill supplies on relative wages depends inversely on the magnitude of aggregate elasticity of substitution between the two skill groups. The greater is σ , the smaller is the impact of shifts in relative supplies on relative wages and the greater must be fluctuations in demand shifts (D_t) to explain any given time series of relative wages for a given time series of relative quantities. Changes in D_t can arise from (disembodied) skill-biased technological change, nonneutral changes in the relative prices or quantities of nonlabor inputs, and shifts in product demand.

This model is sometimes called the “education race” model, after Jan Tinbergen, who envisaged the evolution of the education premium as reflecting a race between education and technology. In the Tinbergen view, technology is constantly exerting outward pressure on one blade of the Marshallian scissors (represented by D_t), while schooling is pushing outward on the other blade (represented by N_{ct}/N_{ht}). So long as these two forces move in parallel, that is, $\Delta D_t \approx \Delta[1/\sigma \ln(N_{ct}/N_{ht})]$, wage inequality (here, the return to education) is stable. When demand moves outward faster than supply, $\Delta D_t > \Delta[1/\sigma \ln(N_{ct}/N_{ht})]$, wage inequality rises, and vice versa when supply moves outward faster than demand.

This stylized and stripped down model is surprisingly powerful.⁸ Following the approach of Katz and Murphy (1992), one can directly estimate a version of equation (3) to explain the evolution from 1963 to 2005 of the overall log college/high school wage differential series shown above. The dependent variable in this model, $\ln(w_{ct}/w_{ht})$, is directly observed, as is one of the two explanatory variables, $\ln(N_{ct}/N_{ht})$. Rather than modeling unobserved demand shifts, D_t , one can proxy for them using a simple time trend (represent the outward thrust of demand in the spirit of Tinbergen's education race). Specifically, one may fit the equation

$$\ln(w_{ct}/w_{ht}) = \gamma_0 + \gamma_1 t - \gamma_2 \ln(N_{ct}/N_{ht}) + \varepsilon_t, \quad (4)$$

where γ_t provides an estimate of $1/\sigma$.

Fitting the 43 data points above for the college premium to the relative supply series using OLS, we obtain the following estimates for equation (4):

$$\ln(w_{ct}/w_{ht}) = 0.043 + 0.018 \times t - 0.411 \times \ln(N_{ct}/N_{ht}), \quad R^2 = 0.93.$$

(0.037) (0.001) (0.046)

Three noteworthy features of this estimate are (i) the fit is surprisingly good (note the R^2); (ii) secular relative demand shifts favoring college labor are quite important, as seen in the time trend; and (iii) shifts in relative supply play a critical role in explaining movements in the college wage premium. Note also that the our estimate of $\hat{\gamma} = 0.411$ implies that $\hat{\sigma} \approx 2.4$, which is within the bounds of standard estimates of the college/high school substitution elasticity.⁹

Figure 6 provides a glimpse of why this simple model works so well. This figure plots the detrended college/high school wage differential against the detrended college/high school relative supply series (where detrending simply means removing an OLS fitted linear time trend from each series). These series appear to be near mirror images of one another—and in particular, the inflexion point in the supply series appears “separated at birth” from the inflexion point in the wage premium series. It is difficult (at least for me) to look at this figure and conclude that relative supplies of educated labor do not play a central role in the determination of wage inequality (at least as measured by the college premium).

The model need not be this simple, however. One important extension developed by Card and Lemieux (2001a) considers a generalization where workers of different experience levels within the same educated

Fig. 6. Detrended measures of supply and wage differentials, 1963–2005. Source: Autor et al. (2008, fig. 4a).

group are imperfect substitutes for one another. This extension implies, with few additional assumptions, that when the supply of college graduates decelerates, as occurred in the United States in the 1980s, the college premium will rise more rapidly for young workers than for old workers—which is the pattern that we see in figure 3.

Before discussing the potential applicability of the education race model to the key facts above, I compare and contrast it with the Guvenen-Kuruscı model. As it turns out, they have almost nothing in common.

IV. The Guvenen-Kuruscı Model

Guvenen and Kuruscı start with a production function that looks familiarly like the Katz-Murphy model. In particular, the Guvenen-Kuruscı model’s aggregate production function can be written as

$$Q_t = Z \{ [(\theta_{H,t} N_{ct})^\rho + (\theta_{L,t} N_{lt})^\rho]^{1/\rho} \},$$

where Z is a neutral productivity shifter and the rest is equivalent to equation (1). The parameterization of this model that they study, however, is one where $\rho \rightarrow 1$, which implies that $\sigma \rightarrow \infty$. In other words, high- and low-skilled labor are perfect substitutes.

Due to this perfect substitutability, the earlier equation (2) for wage inequality becomes simply

$$\ln(w_{ct}/w_{ht}) = \ln(\theta_{H,t}/\theta_{L,t}). \quad (5)$$

Here, the demand for skills is completely elastic. The pure skilled wage premium is determined only by the state of technology, represented by $\theta_{H,t}/\theta_{L,t}$. While the Katz-Murphy model had three central ingredients—the relative supply of skills, the location of the relative demand curve, and the elasticity of substitution between skill groups in relative demand—the Guvenen-Kuruscu model features only one of these components, the relative demand curve, which might aptly be called a relative demand plateau since supply movements have no effect on the equilibrium price of skill.

Despite this utter parsimony, the model generates rich dynamics for, among other things, the measured college/high school wage premium, the inequality of earnings within and between skill groups, and even the relative supply of skill. How is this accomplished? The answer is heterogeneity and lots of it. There is heterogeneity in endowments, heterogeneity in learning capabilities, heterogeneity in investment windows, and heterogeneity in beliefs about the state of nature. These numerous sources of heterogeneity, combined with some strong assumptions on the evolution of $\theta_{H,t}/\theta_{L,t}$, give rise to a rich (almost bewildering) array of predictions, some of which fit the data quite well.

Guvenen and Kuruscu do an excellent job in explicating the model's many moving parts, and I will not attempt to improve on their work. I will instead name the main actors and discuss their roles. Following that, I compare how the model stacks up against the Tinbergen setup in explaining the key facts.

A. *The Big Bang*

The one and only forcing variable in the Guvenen-Kuruscu model is an unanticipated shock to the trajectory of the skill price, $\theta_{H,t}/\theta_{L,t}$. The growth rate of $\theta_{H,t}/\theta_{L,t}$ discontinuously rises in 1970 and then reverts to its prior trend in 1995. It is this shock, combined with the abundant sources of heterogeneity in the model, that leads to the dynamics. These dynamics work through three channels:

1. *Price effect.* When $\theta_{H,t}/\theta_{L,t}$ rises, this immediately raises inequality by amplifying preexisting differences in skill stocks.

2. *Investment effect.* A rise in skill prices stimulates further skill investments. In the Guvenen-Kuruscu model, workers partly or fully withdraw from the labor force to invest in human capital. In the short run, this lowers observed wages since the most able workers are those who most reduce labor supply to enable them to make the largest investments.

3. *Quantity effect.* Over the longer term, the investment effect leads to greater human capital stocks (denoted by l^i). Wages rise as workers accumulate skills and as they slow their rates of investment (i.e., as steady-state skill targets are achieved).

Of these three channels, the investment effect is the most critical. Following Becker (1964), workers acquire skills in the Guvenen-Kuruscu model in large part through investment in on-the-job skill acquisition (OJT for short). Workers may dedicate up to 50% of their work hours to OJT, reducing their effective labor supply and contemporaneous earnings accordingly. If a worker wants to invest more than 50% of his time in OJT, he “enrolls in college” and invests full-time for at least 2 years, at which point he is labeled a college graduate for purposes of the simulation.

Concretely, if worker a is twice as skilled as worker b , and a spends 50% of his work hours investing in human capital while b spends none, they earn the same net wage. Since the return to skill investments rises with the price of skill, an unanticipated rise in the skill price may initially reduce observed wages and then raise them. This occurs if the investment response—greater OJT in response to higher skill prices—dominates the price effect, that is, higher skill prices directly raising wages.

B. Heterogeneity in Investment Responses

If all workers were identical, shocks to skill price would have homogeneous effects on wages. In the Guvenen-Kuruscu model, wage effects are anything but homogenous. Heterogeneity enters through its interaction with each of the four channels below:¹⁰

1. *Heterogeneity in learning ability.* Perhaps the central source of heterogeneity in the model is cross-worker variation in learning ability, A^i . This parameter in large part determines a worker’s optimal investment path. Workers with lower A implicitly pay a higher price for acquiring skills, and they hence invest less in them. Most critically, given an unanticipated shock to the value of human capital (θ_H), workers with higher A optimally

increase their human capital investments by more, which greatly affects their earnings profile over the life cycle. The greater the extent of investment, the more that wages fall in the short run (due to lower effective labor supply) and the more they rise in the long run (due to higher prices and greater skills accumulation)—with both effects amplified by higher A .

2. *Heterogeneity in labor.* For reasons that elude me, Guvenen and Kuruscu also posit the presence of worker heterogeneity in raw labor endowment, l^i , with an assumed correlation of 1.0 between A and l . Thus, more able workers have an absolute advantage in both labor-intensive and human capital-intensive tasks. When high A workers differentially raise OJT in response to an increase in θ_H , average wages in the economy fall by even more than would occur if A and l were uncorrelated (i.e., since the workers who reduce their effective labor input the most are also the most productive).

3. *Heterogeneity in age.* Workers in the model spend a finite career of 45 years in the labor force, so older and younger workers face different investment horizons. An unanticipated rise in the skill premium induces a greater investment response among the young than the old, leading to relatively larger reductions in effective labor supply among young workers, who are then rewarded with higher earnings later in life.

4. *Heterogeneity in beliefs.* Given that the key exogenous force in the model is an unanticipated shock to skill prices, workers' investment responses to this shock will depend upon how rapidly and accurately they perceive the shock and how long they believe it will last. Guvenen and Kuruscu consider a few scenarios for the evolution of beliefs, but in practice these variations do not turn out to be empirically important.

In summary, the operative channels influencing wage inequality are in almost all senses different from the Tinbergen framework. In the Guvenen-Kuruscu model, the only macroeconomic factor is the skill price—all wage dynamics arise from microeconomic responses (observed and unobserved) to this price. In the Tinbergen framework, skill demand—though not its price—is exogenous, and the supply responses that follow are probably fairly viewed as endogenous (though Tinbergen offers no explicit model for this). The skill price, in turn, is determined by the intersection of the two blades of the Marshallian cross. Heterogeneity in microeconomic responses to skill demand is simply not part of the model, though it is not in any sense at odds with it.

V. Confronting the Evidence

I now return to the five key facts.

A. *Rising Wage Dispersion*

In the Tinbergen model, an increase in skill prices—reflecting a relative demand shift for college workers relative to noncollege workers (H vs. L) that outstrips the growth of supply—raises wage dispersion through two channels: an increase in the college wage gap and a rise in “within-group” inequality among H and L workers. This latter “residual channel” arises through the interaction between rising prices and heterogeneity in productivities among workers who have similar nominal credentials (college or noncollege).

While demand and supply play equal (and opposite) roles in the Tinbergen model, much analysis of U.S. wage inequality attributes most of the fluctuations in relative wages over recent decades to movements in relative supply. As Katz and Murphy (1992) first demonstrated, one can go a very long way toward explaining the evolution of the U.S. wage structure simply by overlaying the accelerations and decelerations in the growth rate of relative supply of college-educated labor on a smoothly trending outward shift in relative demand.

In the Guvenen-Kuruscu model, by contrast, there is no economic distinction between within-group and between-group inequality since all skills are perfect substitutes; although some workers are labeled college and others noncollege, this distinction is quantitative not qualitative, reflecting depth of skill rather than type of skill. Moreover, given perfect skill substitutability and perfectly elastic labor demand, fluctuations in the aggregate supply of skills have no impact on inequality.¹¹

But wage dispersion enters this model through so many other doorways that one hardly misses the supply channel. A positive shock to skill prices raises inequality through the price channel, lowers it initially through the investment channel, and then raises it again through the quantity channel. What is the net effect? Clearly, it will depend on model parameters, at least in the short run—though in the long run, the effect must be positive.

Which framework (Tinbergen’s or Guvenen and Kuruscu’s) does a better job with the facts? The key virtue of the Guvenen-Kuruscu model, in my view, is that it is capable of generating different time paths for “between-group” and within-group inequality—this despite the fact that

both are responding to the same causal force. (I place “between-group” in quotation marks here because the distinction between college and non-college groups is nominal.) In particular, within-group inequality—inequality among those labeled either college or noncollege—may rise initially following the price shock, reflecting heterogeneity in the investment response to changing skill prices (as well as the interaction between heterogeneity in skill endowments and changing skill prices). Between-group wage inequality—the mean wage gap between workers labeled college and noncollege—may initially fall due to greater investment responses among more skilled workers.

This additional richness comes with costs and benefits. One cost, as above, is that the definition of college versus noncollege workers has no economic counterpart in the model’s underpinnings; college and high school wages are set by the same pair of prices.¹² Thus, absent unmeasured heterogeneity in l , h , and effective labor supply, there would be no between-group inequality in the model.

A second cost is a complexity. As the discussion above underscores, with so many heterogeneity dynamics in the model—I count five parameters in Guvenen and Kuruscu’s table 1 that measure heterogeneity—one suspects that almost anything could happen under the right parameterization. This complexity allows heterogeneity in the Guvenen-Kuruscu model to do the work done by labor supply in the Tinbergen model, but to do it at considerable loss in transparency.

Perhaps the overriding consideration, however, is whether the rich dynamics of the model help it to better explain the data. The answer to this question turns on whether overall and residual inequality follow a different time trajectory than the college wage gap, which is a phenomenon that Guvenen and Kuruscu’s simulation predicts, or whether they do not, as the education race model predicts. As it turns out, this is a contested fact in the literature.

Using data from the March CPS, Katz and Murphy (1987), Juhn, Murphy, and Pierce (1993), Katz and Autor (1999), Acemoglu (2002), and Autor et al. (2005, 2008) find that there is a gap of one decade between the rise of the college premium, commencing in the 1980s, and the rise of overall and residual inequality commencing in the 1970s. This suggests that these are distinct phenomena. In contrast, Lemieux (2006), reiterating earlier conclusions from DiNardo, Fortin, and Lemieux (1996), finds that the path of the college gap and overall and residual inequality follow a similar trajectory in the May/ORG CPS (as distinct from the March CPS). This leads Lemieux to argue that these different dimensions of inequality are not truly distinct.¹³

To date, the difference in timing between the rise of between-group and within-group inequality in the March and May/ORG CPS has not been adequately explained (cf. Autor et al. [2005, 2008] and Lemieux [2006] for discussion). A dispassionate reader would certainly be justified in accepting Lemieux's unifying view of wage inequality, particularly because it has the appeal of parsimony. Indeed, despite my published conclusions to the contrary, I have been increasingly persuaded by Lemieux's view.

Surprisingly, strong countervailing evidence against the more parsimonious view is uncovered in a recent paper by Kopczuk, Saez, and Song (forthcoming). Using consistent social security earnings samples from 1973 through 2004 (arguably a more authoritative data source than the March or May/ORG CPS), they show that the Gini coefficient for male and pooled-gender (though not female) annual inequality started rising sharply in 1970, plateaued briefly from 1977 through 1980, rose even more steeply from 1980 through 1988, then grew more slowly through the end of their series in 2004. These patterns match the timing of overall and residual inequality in the March CPS, but they are at odds with the May/ORG CPS.

Based on Kopczuk et al.'s analysis of the social security data, I must conclude that the best evidence now favors the view that the timing of between-group versus overall inequality is distinct. And this in turn favors the Guvenen-Kuruscu model since it can do something that the education race model cannot, which is to explain this difference in timing. The fact that the Guvenen-Kuruscu model can rationalize this fact does not make the explanation correct—but it is a good start. For those keeping score at home, I award this round to the Guvenen-Kuruscu model.

B. The Rising, Falling, and Then Rising Return to Education

The crowning success of the education race model is its ability to explain the fluctuating college premium over numerous decades merely by positing a steady outward shift in relative demand, a stable (and plausible) elasticity of substitution between college and noncollege labor, and a set of sharp secular shifts in the growth rate of college supply. Figure 7 (from Autor et al. 2008) replicates the famous Katz and Murphy (1992) implementation of the education race model for the years 1963–87 (i.e., the years available to Katz and Murphy) and then extends the model's predictions out of sample through 2005. The model does an exemplary job of fitting the data through 1992, although it somewhat overpredicts the subsequent growth of the college premium. Goldin and Katz (2008) further demonstrate that this model does an impressive job with explaining skill

Fig. 7. Observed and predicted college/high school log wage gap based on Katz and Murphy (1992) fit of education race model for years 1963–87 and projected out of sample to 2005. Source: Autor et al. (2008, fig. 4b).

premiums over a much longer time period, 1915–2005. Notably, this model neither presumes nor requires a sudden unanticipated shock to skill demand. Its goodness of fit arises exclusively through movements in observable prices and quantities overlaid on a time trend.

Guvenen and Kuruscu’s explanation for the rising-falling-rising college wage premium is more subtle, but, to my eye, less credible. As a starting point, they invoke *ex cathedra* an unexpected acceleration in skill demands that starts in 1970 and ends in 1995—presumably because this appears to fit the data (notice that it is a few years prior to 1995 where the Katz-Murphy model “wants” more wage inequality than it gets). Under the baseline scenario, workers do not anticipate the acceleration in skill prices in 1970. But once the Big Bang hits, workers forecast its deceleration 25 years later.¹⁴ These assumptions do not make for a promising start.

The mechanism whereby the positive price shock leads to initially falling college wages is the investment channel as above: more skilled workers reduce their effective labor supply through OJT—or they exit the labor force to enroll in college—thereby depressing the observed wages of college workers versus noncollege workers. This idea has solid roots in Becker’s human capital model, but it is hard to reconcile with the facts. As the authors candidly admit (in their Sec. IV.D), the fall in the college

premium in the United States during the mid-1970s was met with a fall in college enrollments—opposite to the model's predictions. While the model can rationalize the fall in college wages (rather than enrollments) as reflecting an unmeasured rise in OJT, it is not particularly credible that workers unobservably increased on-the-job skills investments while observably deciding that school-based skill investments were not worthwhile.

C. Age Group Differences in the Rising College Premium

There is, however, another important empirical regularity that Guvenen and Kuruscu capture through the nuances of their model: the measured college premium rose by considerably more for the young than for the old during the 1980s, as shown in figure 3. Guvenen and Kuruscu's model offers an elegant explanation: facing the same skill price shock in the 1970s, young workers made larger skill investments than older workers since they had more remaining work years in which to garner the returns. Accordingly, by 1980, they had greater human capital than older workers possessing similar nominal education levels. Hence, the college wage gap rose by more for the young than the old in the 1980s as the investment effect abated and the quantity effect took hold.

While this is ingenious, a simple extension to the education race model offered by Card and Lemieux (2001b) does at least as well with these facts, and it taxes credulity less. A deceleration in college production will induce a greater scarcity of young college workers than old college workers. If, plausibly, workers of different age groups within an education group are imperfect substitutes, this will cause the college wage premium to rise by more for young workers than old workers when the college crunch hits.

This slowdown in experience-group relative supplies is shown in figure 8. The relative supply of young college workers (0–9 years of experience) falls off sharply in 1975, while the relative supply of older college workers (20–29 years of experience) does not decelerate until 1995. Comparing figure 3 (the college gap by experience group) with figure 8 shows a clear correspondence: the experience group specific college premium rises when experience group specific relative supply decelerates. This parsimonious explanation fits the data well, so this round goes to Tinbergen.¹⁵

D. Rising Real Wages at the Top, Falling Real Wages at the Bottom

One of the least well understood facts about U.S. wage inequality is the rapid fall in real wages for less educated workers during the 1980s. What explains this fall?

Fig. 8. Evolution of college/high school relative supply among inexperienced and experienced workers, 1963–2005. Source: Autor et al. (2008, fig. 6b).

The education race model has trouble with this fact. Given that there was no increase in the relative supply of low-skilled workers during this period depressing the low-skilled wage—just the opposite—the education race model leaves only one remaining channel for low-skill wages to fall while high-skill wages are rising: an absolute decline in θ_L , implying a form of “skill-destroying” technical change.¹⁶

Is this plausible within the confines of the simple two-factor model? Not especially. Technical change that is massively biased in favor of high-skilled workers and against low-skilled workers would have to offer substantial TFP gains for it to dominate the existing technology and hence be adopted (Krugman 2000). But U.S. productivity growth was torpid at the time that wage inequality was soaring. This makes it hard to tell a credible story within this simple model that can rationalize these wage declines—though many have tried.¹⁷

In the Guvenen-Kuruscu model, the fall in the real wage of less-skilled workers works through three channels. The first channel is skill-destroying technical change. Guvenen and Kuruscu posit that, during the Big Bang (1970–95), θ_L trends downward at about 0.5% a year. Offsetting this, Guvenen and Kuruscu assume that factor-neutral TFP grows at 1.5% per year—so raw labor productivity is rising (in net). The observed fall

in real wages is therefore not purely mechanical, though the decline in θ_L helps it along.

The second channel operating against low-skill wages is the investment effect. An increase in OJT means that low-skill wages (which are already afflicted with slowing productivity growth) are further depressed by a rise in OJT and a concomitant reduction in effective labor supply.

Finally, the assumed perfect correlation between raw labor endowments and learning ability means that the low-skill workers who reduce effective labor supply the most are those with the highest ability. This magnifies the wage losses stemming from rising OJT. In net, real wages of college workers stagnate and noncollege wages fall during the initial years of the Big Bang.

This explanation for falling real wages is no more or less satisfying than the earlier explanation for a falling then rising college premium in the 1970s and 1980s, respectively. If one accepts the implication that the 1970s was a period of dramatically increased investment in human capital on the job (but not in school), then the explanation for falling wages at the bottom and stagnating wages at the top is coherent and ingenious. If one rejects this implication, then the explanation is less appealing.

It appears hard to deny that the real earnings power of low-skill workers fell during the 1970s and 1980s—not just their effective labor supply. High school dropout rates, nonemployment, criminal participation and incarceration rose among low-skilled workers and low-income families during these decades. This is hard to square with the view that the 1970s were a renaissance for skills investment among the less educated.

Neither the education race model nor the Guvenen-Kuruscuk model can, in my view, be declared a winner on this score. They are coequals in defeat.

E. Boom and Bust in College Labor Supply

The boom and then bust in the supply of new college graduates depicted in figure 5 can arguably be rationalized by either model. In the education race model, the fluctuations in college supply seem largely to follow the trajectory of the college premium, though of course with a built-in delay of 4 or more years (i.e., the time it takes to earn a college degree).¹⁸

In the Guvenen-Kuruscuk model, the rise in college attainment is also a response to changes in skill prices. But here the fit is less good. Guvenen and Kuruscuk predict a brief dip and then accelerating rise in college supply from 1970 to 2000. The crucial deceleration of college supply in the early 1980s is not predicted by their model, as is seen in figure 13 of their paper. Of course, this is not a serious problem for the Guvenen-Kuruscuk

model since the realized college premium is independent aggregate supplies. For the education race model, however, the supply deceleration is both a logical response to the falling wage premium and a leading causal force in the subsequent rise of wage inequality.

VI. Conclusion

If I have done my job well, I have made a solid case that we do not have a perfectly satisfactory explanation of the key facts surrounding the rise of U.S. wage inequality. The most successful models explain between-group inequality well, within-group inequality less well, and real wage levels least of all. In particular, the canonical education race model is a major success on the first front, a partial success on the other, and a failure on the third. The Guvenen-Kuruscu model is at its best where the canonical model has weak footing—in particular, in explaining the disparate timing of rising between versus within-group inequality. Is the Guvenen-Kuruscu model a stand-alone explanation for all of the major facts? Clearly not. Nor is it probably intended to be. For example, I doubt the authors would wish to argue that skill supplies have no effect on prices—only that one can shut down this channel and still get rich results. The question that only time can answer is whether the Guvenen-Kuruscu model provides enough additional points of contact between theory and data to drive analysis and understanding forward. If yes, the paper will clearly have moved the frontier. If not, this work will remain admirable for its ambition and its craft.

Endnotes

1. These are not the only facts that they confront. But, in my view, they are the most robust and important.

2. The rate of growth after 1988 differs between the two series, and there is not a definitive explanation for their divergence. See Autor et al. (2005, 2008) and Lemieux (2006) for discussion.

3. Autor et al. (2008) show that the evolution of the 90/10 ratio is driven by two different trends. From 1979 to 1988, both 90/50 and 50/10 wage inequality increase rapidly. After 1988, 90/50 wage inequality continues its near-linear expansion path through 2005 (the end of their sample), whereas the 50/10 ratio plateaus in 1988 and then reverses somewhat thereafter. Thus, the deceleration in the growth of the 90/10 ratio is driven by a reversal of lower-tail inequality in combination with ongoing expansion of upper-tail inequality.

4. Indeed, the comparable trajectory of residual and overall wage inequality in the CPS May/ORG series has led some researchers to conclude that these are really one phenomenon (see Lemieux 2006). I say more about this below.

5. Note that the PCE is thought to overstate inflation by less than the more commonly used Consumer Price Index.

6. Another possibility is that young Americans were influenced by Richard Freeman's classic 1976 book, *The Overeducated American*, which documented that the college premium

had fallen and estimated that the net social return to further rises in college enrollment were probably negative. Thus, Richard Freeman may be personally responsible for the epochal rise in U.S. wage inequality—for which a generation of researchers (myself among them) owe him thanks.

7. The lineage of this explanation runs through Katz and Murphy (1992), Autor, Katz, and Krueger (1998), Katz and Autor (1999), Card and Lemieux (2001b), Acemoglu (2002), Autor et al. (2008), and Goldin and Katz (2008).

8. Although it should be stressed that sophisticated versions of this model are successfully applied by, among others, Card and Lemieux (2001a), Acemoglu, Autor, and Lyle (2004), Carneiro and Lee (2008), and Goldin and Katz (2008).

9. Standard estimates range from 1.0 to 2.5 (Katz and Autor 1999). The estimate here is on the high side because we impose linearity of the time trend in log relative demand. If we allow a bit more flexibility, the estimates of σ fall toward approximately 1.7, which is closer to the estimates of Katz and Murphy (1992), as well as those of many subsequent authors.

10. There is also initial heterogeneity in human capital stocks even prior to the “Big Bang,” and this is amplified by the “price effect” as detailed below.

11. Labor demand is perfectly elastic because this is a one good economy with constant returns to scale production that uses two perfectly substitutable factors, l and h . Thus, there is neither diminishing marginal rate of substitution in production nor diminishing marginal utility in consumption.

12. In the education race framework, there is a distinct college wage and a distinct non-college wage. Of course, there may also be within-group heterogeneity in the skills of college and noncollege workers. But college and noncollege skills are not linear syntheses of one another.

13. Lemieux also argues that part of the rise in residual inequality in the 1990s can be explained by the changing composition of the U.S. labor force. See Autor et al. (2005) and Lemieux (2006) for discussion.

14. Although Guvenen and Kuruscu consider more plausible learning processes, all of them have the feature that workers forecast a deceleration at some future point, which seems a bit odd. (Why wouldn’t priors give further accelerations and decelerations equal weight?)

15. Perhaps most persuasively, Card and Lemieux (2001b) document that the rise in the college premium among the young commenced 5 years later in the United Kingdom than in the United States, which is exactly as predicted because the slowdown in relative supply of young college workers also started 5 years later in the United Kingdom.

16. Given gross complementarity between H and L workers ($\sigma > 1$), as is suggested by all available evidence, an increase in θ_H holding constant θ_L raises wage inequality and raises the wages of L workers (due to q -complementarity). Thus, L wages can only fall if θ_L falls. And θ_L must fall farther still if either H or θ_H is rising since both raise earnings of L workers. Note that, under the parameterization in eq. (1), one could also interpret the decline in L wages to indicate that α_t fell, so a larger share of tasks was allocated to H than L workers. This is a distinction without a difference, however, since $\theta_{L,t}$ and $\theta_{H,t}$ can always be redefined to include α_t .

17. Of course, many sophisticated (and more credible) alternative explanations have been offered. Greenwood and Yorukoglu (1997) is a brave and early effort. Acemoglu (1999) offers a model in which rising skill supplies cause firms to endogenously adopt a strongly skill-complementary technology in lieu of a comparatively skill-neutral technology that was appropriate when skills were scarce. Beaudry and Green (2003) study a setting in which a rise in skill supplies causes capital starvation in the low-skill intensive sector, causing low-skill wages to fall. Unfortunately, the rapidly falling U.S. minimum wage during the 1980s does not offer a credible explanation for falling low-skill wages, since the minimum was set too low to be relevant for almost all male workers in this period (Autor, Manning, and Smith 2009). Levy and Temin (2007) offer a richer institutional explanation for the evolution of low-skilled and high-skilled wages.

18. Indeed, Freeman’s (1976) book, which arguably caused the college skill shortage also predicted its reversal. Using a “cobweb” model, Freeman predicted that the supply response to the falling wage premium would overshoot the mark, leading the wage premium to rebound.

References

- Acemoglu, Daron. 1999. "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence." *American Economic Review* 89, no. 5:1259–78.
- . 2002. "Technical Change, Inequality and the Labor Market." *Journal of Economic Literature* 40, no. 1:7–72.
- Acemoglu, Daron, David Autor, and David Lyle. 2004. "Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury." *Journal of Political Economy* 112, no. 3:497–551.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2005. "Residual Wage Inequality: The Role of Composition and Prices." Working Paper no. 11628, NBER, Cambridge, MA.
- . 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics* 90, no. 2:300–323.
- Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113, no. 4:1169–1213.
- Autor, David H., Alan Manning, and Christopher L. Smith. 2009. "The Minimum Wage's Role in the Evolution of U.S. Wage Inequality over Three Decades: A Modest Reassessment." Working paper, Massachusetts Institute of Technology.
- Beaudry, Paul, and David A. Green. 2003. "Wages and Employment in the United States and Germany: What Explains the Differences?" *American Economic Review* 93, no. 3:573–602.
- Becker, Gary S. 1964. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. Chicago: University of Chicago Press.
- Card, David, and Thomas Lemieux. 2001a. "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis." *Quarterly Journal of Economics* 116, no. 2:705–46.
- . 2001b. "Dropout and Enrollment Trends in the Postwar Period: What Went Wrong in the 1970s?" In *Risky Behavior among Youths: An Economic Analysis*, ed. Jonathan Gruber. Chicago: University of Chicago Press.
- Carneiro, Pedro, and Sokbae Lee. 2008. "Trends in Quality-Adjusted Skill Premia in the United States, 1960–2000." Working Paper no. CWP02/09 (September), Centre for Microdata Methods and Practice (CEMMAP), Institute for Fiscal Studies, London.
- DiNardo, John, Nicole Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach." *Econometrica* 64, no. 5:1001–44.
- Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce. 1993. "Wage Inequality and the Rise in Returns to Skill." *Journal of Political Economy* 101, no. 3:410–42.
- Freeman, Richard. 1976. *The Overeducated American*. New York: Academic Press.
- Goldin, Claudia, and Lawrence Katz. 2008. *The Race between Education and Technology*. Cambridge, MA: Harvard University Press.
- Greenwood, Jeremy, and Mehmet Yorukoglu. 1997. "1974." *Carnegie-Rochester Conference Series on Public Policy* 46:49–95.
- Katz, Lawrence F., and David H. Autor. 1999. "Changes in the Wage Structure and Earnings Inequality." In *Handbook of Labor Economics*, vol. 3A, ed. Orley Ashenfelter and David Card. Amsterdam: Elsevier.

- Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107, no. 1:35–78.
- Kopczuk, Wojciech, Emmanuel Saez, and Jae Song. Forthcoming. "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data since 1937." *Quarterly Journal of Economics*.
- Krugman, Paul. 2000. "Technology, Trade and Factor Prices." *Journal of International Economics* 50, no. 1:51–71.
- Lemieux, Thomas. 2006. "Increased Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?" *American Economic Review* 96, no. 3:461–48.
- Levy, Frank, and Peter Temin. 2007. "Inequality and Institutions in 20th Century America." Working Paper no. 13106, NBER, Cambridge, MA.