

Discussion of “Recent Flattening in the Higher Education Wage Premium:
Polarization, Skill Downgrading, or Both?” by Robert G. Valletta


David Autor
MIT and NBER

January 2017

Robert Valletta’s paper “Recent Flattening in the Higher Education Wage Premium...” illuminates one of the leading puzzles for contemporary U.S. labor economics: the unexpected ‘flattening’ of the premium to higher education in the United States in the 2000’s. This single metric—the college/high-school wage premium—has been the North Star guiding neoclassical analysis of the evolution of wage inequality during a period of rapidly shifting wage structures. Two impactful papers by Beaudry, Green and Sand (2014, 2016, BGS hereafter) argue that since approximately the year 2000, this North Star has become an increasingly dubious point of navigation. Specifically, BGS highlight the failure of the college premium to rise in the 2000s following two decades of steep increases. They interpret this deceleration as reflecting the maturation of the information technology revolution, which in turn has spurred a slackening in the pace of workplace IT investments and a consequent slowdown in the trend of rising demand for highly educated labor. A key piece of evidence favoring BGS’s narrative is the precipitous fall in U.S. investment information processing equipment and software in the U.S. after 1999 (Figure 1), which seems to have precisely the right timing to explain a falloff in IT-augmentation of skilled labor demand.

Valletta’s careful analysis extends and probes the BGS findings, verifies their robustness, and considers their interpretation in the light of both the BGS conceptual framework and an alternative framing offered by Acemoglu and Autor (2011). There are many things to admire about Valletta’s paper: it is empirically rigorous, intellectually ecumenical, and commendably ambitious in synthesizing and adjudicating between two conceptual models that are not, to a first approximation, speaking the same language. My remarks focus exclusively on one question that is core to both Valletta’s BGS’s work: when did rising demand for college-educated labor decelerate? I argue below that (1) the recent flattening of the skill premium in the 2000s is *not* surprising in light of the canonical supply-demand model; and (2) what *is* surprising is that the underlying demand for college labor decelerated sharply and (to date) inexplicably almost a decade beforehand. These observations render the phenomenon that Valletta tackles no less consequential. But they may suggest a different set of explanations for the slowdown than those focusing on discontinuous changes in economic trends in the 2000s.

Figure 1: Private Fixed Investment in Information Processing Equipment and Software as a Percentage of Gross Domestic Product, 1949–2014



Source: FRED, Federal Bank of St. Louis. <http://research.stlouisfed.org/fred2/graph/?g=GXC> (accessed 8/3/2014). This graphic originally appeared in Autor (2015).

Modeling school

Following an extraordinarily influential series of papers that includes Goldin and Margo (1992), Katz and Murphy (1992), Murphy and Welch (1992), Card and Lemieux (2001), and Goldin and Katz's magisterial 2008 volume *The Race Between Education and Technology*, labor economists have applied a remarkably simple and surprisingly powerful calibrated supply-demand model (the 'canonical model') to rationalize the over-time fluctuations in the skill premium and the accompanying evolution of wage inequality. This so-called canonical model takes its inspiration from the observation by Nobel Laureate Jan Tinbergen in 1974 that there appears to be an ongoing 'race' between technology and schooling, with technological advancements progressively raising the demand for educated labor and the school system simultaneously secularly raising its supply. When technological advancement surges faster than educational production, the relative scarcity of educated labor rises, and the skill premium rises with it—that is, technology pulls ahead of education in this two-person race. Conversely, when educational production surges ahead of technologically induced demand shifts, the skill premium falls.

While many elements of this description seem far too simple (e.g., history provides many examples of technologies that replace rather than complement skills), this framework provides a surprisingly good high-level description of what we see in the data. The canonical model provides a benchmark for interpreting the evolution of the skill premium. I apply this model here to address the question of whether we should be surprised—and if so, how much—by the slowdown in the skill premium after 2000. Before applying the model, I review its rudiments, and I refer readers to Acemoglu and Autor (2011) for a fuller development.

The canonical model posits two skill groups, high and low. It draws no distinction between skills and occupations (tasks), so that high skill workers effectively work in separate occupations (perform different tasks) from low skill workers. In most empirical applications of the canonical model, it is natural to identify high skill workers with college graduates (or in different eras, with other high education groups), and low skill workers with high school graduates (or in different eras, those with less than high school). Critical to the two-factor model is that high and low skill workers are imperfect substitutes in production. The elasticity of substitution between these two skill types is central to understanding how changes in relative supplies affect skill premia.

Suppose that the total supply of low skill labor is L and the total supply of high skill labor is H . Naturally not all low (or high) skill workers are alike in terms of their marketable skills. As a simple way of introducing this into the canonical model, suppose that each worker is endowed with either high or low skill, but there is a distribution across workers in terms of efficiency units of these skill types. In particular, let \mathcal{L} denote the set of low skill workers and \mathcal{H} denote the set of high skill workers. Each low skill worker $i \in \mathcal{L}$ has l_i efficiency units of low skill labor and each high skill worker $i \in \mathcal{H}$ has h_i units of high skill labor. All workers supply their efficiency units inelastically. Thus the total supply of high skill and low skill labor in the economy can be written as:

$$L = \int_{i \in \mathcal{L}} l_i di \text{ and } H = \int_{i \in \mathcal{H}} h_i di.$$

The production function for the aggregate economy takes the following constant elasticity of substitution form

$$Y = \left[(A_L L)^{\frac{\sigma-1}{\sigma}} + (A_H H)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \quad (1)$$

where $\sigma \in [0, \infty)$ is the elasticity of substitution between high skill and low skill labor, and A_L and A_H are factor-augmenting technology terms.¹ The elasticity of substitution between high and low skill workers plays a pivotal role in interpreting the effects of different types of technological changes in this canonical model. We refer to high and low skill workers as *gross substitutes* when the elasticity of substitution $\sigma > 1$, and *gross complements* when $\sigma < 1$.

In this framework, technologies are *factor-augmenting*, meaning that technological change serves to either increase the productivity of high or low skill workers (or both). This implies that there are no explicitly skill-replacing technologies. Depending on the value of the elasticity of substitution, however, an increase in A_H or A_L can act either to complement or (effectively) substitute for high or low skill workers (see below).

Assuming that the labor market is competitive, the low skill unit wage is simply given by the value of marginal product of low skill labor, which is obtained by differentiating (1) as

$$w_L = \frac{\partial Y}{\partial L} = A_L^{\frac{1}{\sigma}} \left[A_L^{\frac{\sigma-1}{\sigma}} + A_H^{\frac{\sigma-1}{\sigma}} (H/L)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{1}{\sigma-1}}. \quad (2)$$

¹This production function is typically written as $Y = \left[\gamma (A_L L)^{\frac{\sigma-1}{\sigma}} + (1-\gamma) (A_H H)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$, where A_L and A_H are factor-augmenting technology terms and γ is the distribution parameter. I suppress γ (i.e., set it equal to 1/2) to simplify notation.

Similarly, the high skill unit wage is

$$w_H = \frac{\partial Y}{\partial H} = A_H^{\frac{\sigma-1}{\sigma}} \left[A_L^{\frac{\sigma-1}{\sigma}} (H/L)^{-\frac{\sigma-1}{\sigma}} + A_H^{\frac{\sigma-1}{\sigma}} \right]^{\frac{1}{\sigma-1}}. \quad (3)$$

Combining (2) and (3), the skill premium—the high skill unit wage divided by the low skill unit wage—is

$$\omega = \frac{w_H}{w_L} = \left(\frac{A_H}{A_L} \right)^{\frac{\sigma-1}{\sigma}} \left(\frac{H}{L} \right)^{-\frac{1}{\sigma}}. \quad (4)$$

Equation (4) can be rewritten in a more convenient form by taking logs,

$$\ln \omega = \frac{\sigma-1}{\sigma} \ln \left(\frac{A_H}{A_L} \right) - \frac{1}{\sigma} \ln \left(\frac{H}{L} \right). \quad (5)$$

The log skill premium, $\ln \omega$, has been a central object of study in the empirical literature on the changes in the earnings distribution. Equation (5) shows that there is a simple log linear relationship between the skill premium and the relative supply of skills as measured by H/L . Equivalently, equation (5) implies:

$$\frac{\partial \ln \omega}{\partial \ln H/L} = -\frac{1}{\sigma} < 0. \quad (6)$$

This relationship corresponds to the second of the two forces in Tinbergen's race (the first being technology, the second being the supply of skills): for a *given skill bias of technology*, captured here by A_H/A_L , an increase in the relative supply of skills reduces the skill premium with an elasticity of $1/\sigma$. Intuitively, when high and low skill workers are producing the same good but performing different functions, an increase in the number of high skill workers will necessitate a substitution of high skill workers for the functions previously performed by low skill workers.² The downward sloping relationship between relative supply and the skill premium implies that if technology, in particular A_H/A_L , had remained roughly constant over recent decades, the remarkable increase in the supply of skills (seen, for example, in Table 1 of Valletta's paper) would have led to a significant decline in the skill premium. The lack of such a decline is a key reason why economists believe that the first force in Tinbergen's race—changes in technology increasing the demand for skills—must have also been important throughout the 20th century (cf. Goldin and Katz, 2008).

More formally, differentiating (5) with respect to A_H/A_L yields:

$$\frac{\partial \ln \omega}{\partial \ln(A_H/A_L)} = \frac{\sigma-1}{\sigma}. \quad (7)$$

Equation (7) implies that if $\sigma > 1$, then relative improvements in the high skill augmenting technology (i.e., in A_H/A_L) increase the skill premium. This can be seen as a shift out of the

²In this interpretation, we can think of some of the “tasks” previously performed by high skill workers now being performed by low skill workers. Nevertheless, this is simply an interpretation, since in this model, there are no tasks and no endogenous assignment of tasks to workers. One could alternatively say that the H and L tasks are imperfect substitutes, and hence an increase in the relative supply of H labor means that the H task is used more intensively but less productively at the margin.

relative demand curve for skills. The converse is obtained when $\sigma < 1$: that is, when $\sigma < 1$, an improvement in the productivity of high skill workers, A_H , relative to the productivity of low skill workers, A_L , shifts the relative demand curve inward and reduces the skill premium. Nevertheless, the conventional wisdom is that the skill premium increases when high skill workers become relatively more—not relatively less—productive, which is consistent with $\sigma > 1$. Most estimates put σ in this context to be somewhere between 1.4 and 2 (Johnson, 1970; Freeman, 1986; Heckman, Lochner and Taber, 1998).

The key equation of the canonical model links the skill premium to the relative supply of skills, H/L , and to the relative technology term, A_H/A_L . This last term is not directly observed. Nevertheless, the literature has made considerable empirical progress by taking a specific form of Tinbergen's hypothesis, and assuming that there is a log linear increase in the demand for skills over time coming from technology, captured in the following equation:

$$\ln \left(\frac{A_{H,t}}{A_{L,t}} \right) = \gamma_0 + \gamma_1 t, \quad (8)$$

where t is calendar time and variables written with t subscript refer to these variables at time t . Substituting this equation into (8), we obtain:

$$\ln \omega_t = \frac{\sigma - 1}{\sigma} \gamma_0 + \frac{\sigma - 1}{\sigma} \gamma_1 t - \frac{1}{\sigma} \ln \left(\frac{H_t}{L_t} \right). \quad (9)$$


Equation (9) implies that “technological developments” take place at a constant rate, while the supply of skilled workers may grow at varying rates at different points in time. Therefore, changes in the skill premium will occur when the growth rate of the supply of skills differs from the pace of technological progress. In particular, when H/L grows faster than the rate of skill biased technical change, $(\sigma - 1) \gamma_1$, the skill premium will fall. And when the supply growth falls short of this rate, the skill premium will increase. Surprisingly, this simple equation provides considerable explanatory power for the evolution of the skill premium—though its limitations are also immediately evident.

Doing the Katz-Murphy

Using data from Autor (2014), I fit this simple model to fifty years of U.S. data for 1963-2012. Figure 1 provides the key input into this estimation: the observed log relative supply of U.S. college vs. non-college labor for years 1963-2012, measured in efficiency units and normalized to zero in the base year.³ Figure (2) highlights the steep rise in production of college-educated labor in the United States in the post-War period—specifically until the late 1970s—followed by a sharp deceleration after 1980. This deceleration is frequently interpreted as the key driver of the rapid rise in the skill premium after 1980 (Katz and Murphy, 1992). Notably, there is also a steep *acceleration* of supply after 2004. All else equal, one would expect this supply acceleration to depress the skill premium

³Extensive details on the calculation of these series are provided in Acemoglu and Autor (2011).

Figure 2: Efficiency Units of College vs. Non-College Labor Supply, 1963-2012



Notes: Source data from Autor (2014).

absent any slowdown of the secular trend rise in relative demand after 2004. This observation highlights that the evolution of the skill premium is *not* a sufficient statistic for fluctuations in demand for skilled labor; one must also account for supply.

Using the data series in Figure 2, I fit equation (9) to obtain the following estimate:


$$\ln \omega_t = \text{constant} + 0.0151 \times t - 0.302 \cdot \ln \left(\frac{H_t}{L_t} \right). \quad (10)$$

(0.0013) (0.0429)

This simple OLS model implies that: (1) the relative demand curve for college versus non-college labor is shifting outward by approximately 1.5 log points per year; and (2) that increases in the relative supply of skilled labor buffer the impact of shifting demand on wage inequality. Specifically, the point estimate of -0.30 on the relative supply term implies an elasticity of substitution of $\hat{\sigma} = 1/3.31$. While the explanatory power of this time-series model is high ($R^2 = 0.94$), the point estimate for the elasticity of substitution is more than *twice* as high as Katz-Murphy's 1992 estimate of 1.41. This implies that either the elasticity of substitution is changing over time or that the linear time trend is not doing an adequate job of capturing trends in relative demand.

Figure 3 explores these possibilities. The series plotted in blue (with circular markers) corresponds to the measured (i.e., observed) skill premium in each year. This series depicts the now familiar rise in the skill premium from the early 1960s (start of the series) to the early 70s, the sharp fall between 1971 and 1981, the steep and continuous rise from 1982 to 1999, and then the much shallower rise from 2000 to 2012 (end of the series). The red series (diamond markers) performs a within-series extrapolation by re-estimating equation (10) using only data from 1962 to 1992 (the period of best fit), and recovering estimates of the time-trend and the elasticity of substitution

Figure 3: Observed, Predicted, and Fitted Evolution of the Log College/Non-College Hourly Earnings Gap, 1963-2012



Notes: Source data from Autor (2014).

$(\hat{\gamma} = 0.028, \hat{\sigma} = 1/ - 0.631 = -1.59)$. The plotted series then projects this estimate forward to 2012 using the estimated parameters from the 1962-1992 fit in combination with the *observed* evolution of aggregate skill supplies ($\ln H_t/L_t$). Notably, the time trend and elasticity recovered from this procedure are extremely similar to those obtained by Katz-Murphy's in 1992, and using data for 1963 through 1987. The similarity of the current estimates implies that Katz-Murphy's within-sample point estimates continue to closely track the observed data for an additional five years *out of sample*.

As the figure reveals, however, this projection badly misses the mark after 1992. Adjusting for the evolution of aggregate skill supplies, the growth in the skill premium is far more modest after 1992 than the extrapolation projects. Between 1992 and 2012, the observed college/non-college log earnings gap rises by 11.6 log points. But the projection based on data to 1992—applying the observed evolution of skill supplies to 2012—predicts an increase of 30.4 log points, nearly three times as large as what occurred. A summary judgment is that the evolution of the skill premium has been *surprising* since 1992.

The element of surprise

Economic literature noted this surprise some time ago. Card and DiNardo (2002) first pointed out this discrepancy in their broad critique of the Skill Biased Technical Change literature. Autor, Katz and Kearney (2008) proposed an ad-hoc workaround, which was to allow for a trend deceleration in the evolution of skill demands after 1992. Goldin and Katz (2008) and Autor (2014) pursue a

related approach by applying a quadratic time trend in the time-series model, thereby allowing a smooth deceleration of the trend demand shift. The series in Figure 3 labeled “Fitted gap: quadratic trend” (green series, triangular marker) shows just how well this works. Conditional on the quadratic trend the fit is impressively close. But of course, this is simple reverse engineering. This flexibility was added to the model because the data demanded it, not because the theory suggested it.

These various exercises raise an urgent question: after accounting for fluctuations in the supply of skilled labor, when did the ‘flattening’ of demand for skill commence? Here, I draw a distinction between flattening in the *skill premium* and flattening (or deceleration) in the movement of the underlying demand schedule. As noted above, it would be entirely possible for the skill premium to decline even as demand was accelerating—if skill supplies rose fast enough. Figure 1 makes clear that skill supplies *accelerated* after 2004. Was this *supply-side* change an important contributor to the observed ‘flattening’ of the skill premium?

The series plotted in gold (square markers) in Figure 3 addresses this question. The log-relative supply of college workers (Figure 2) rose at an annual rate of 4.31 log points between 1963 and 1982, by 1.79 log points between 1982 and 2004, and by 2.61 log points between 2004 and 2012 (i.e., a 45 percent *increase* after 2004). The gold series in Figure 3 (labelled “Supply trend 1984-2004 continues post 2004) replaces the observed values of $\ln(H_t/L_t)$ with a counterfactual series in which log relative supply rises at the 1963-1982 of 1.79 log points per annum. Surprisingly (at least to me), this substitution makes a *substantial* difference for inference. The estimated college premium rose by only 1.65 log points between 2004 and 2012. This exercise implies that had the relative supply of college-educated labor *not* accelerated after 2004, the skill premium would have risen by 5.47 log points rather than a measly 1.65 log points. I submit based on this evidence that had there been no *supply acceleration* after 2004, Beaudry et al. would have had a more difficult time making the argument that there was a *demand deceleration* in the 2000s.


How long has this been going on?

The evidence in Figure 3 in no way obviates the claim that demand for college workers ‘flattened’ by the lights of the canonical model. It instead underscores that the raw skill premium, not purged of the impact of supply forces, could generate misleading inferences about the trajectory of the demand for skilled labor.

To address this shortcoming, Figure 4 plots the implied log relative demand shift favoring college vs. non-college labor for 1962-2012, again using the estimated value of $\sigma = 1.59$ based on fitting equation equation (9) to data for 1962-1992. The plotted (scatter) points in Figure 4 are not regression estimates. They correspond instead to the calculated values of $\gamma_t = \omega_t - (1/\hat{\sigma}) \ln(H_t/L_t)$ in each year, where we treat σ as known.⁴ To guide interpretation of these data points, the figure also contains three regression lines. The red (solid) line depicts a pure linear extrapolation, fitted and projected using data for 1962-1992. This corresponds to the implied path of relative demand

⁴Equivalently, they are the time dummies from a saturated regression (no error term) of $\omega_t - (1/\hat{\sigma}) \ln(H_t/L_t)$ on a full set of year indicators.

Figure 4: Implied Evolution of the Demand for College vs. Non-College Labor using $\sigma = 1.59$



Notes: Source data from Autor (2014).

from 1992 through 2012 had there been no deviation after 1992. The green (short-dash) series is the quadratic fit to this set of scatter points. The orange (long-dash) series is a linear spline that allows for a discreet slope shift in 1992 (and otherwise fully overlays the initial trend from 1963-1992).

This plot highlights three key patterns. A first is that the trajectory of (implied) relative demand for educated labor is astonishingly linear for the initial thirty years of the series, 1963-1992. This linearity is in no sense mechanical: the relative demand shift estimates plotted in Figure 4 are extracted from a college wage premium series that fluctuates dramatically over three decades, rising for the first ten years of the time interval, falling for the next nine, and then increasing with remarkable rapidity thereafter. The linearity of the (implied) underlying demand trend therefore reflects the uncanny success of the relative supply term $\ln H_t/L_t$ in explaining the fluctuations in the premium, leaving little behind but a smooth secular underlying demand shift favoring college-educated labor.

The second pattern immediately visible in Figure (4) is that the steady secular demand shift favoring college-educated labor *decelerates* after 1992, and does so abruptly. Estimates of (9) fit using a linear spline (orange long-dash series) imply that the relative for college labor rose by 2.80 log points per year between 1963 and 1992 and then *decelerated* to 1.84 points thereafter (a fall of one-third). This pattern, while occasionally noted in the literature (cf. Acemoglu and Autor 2011), has not been rigorously explained by any formal model—though of course there are many informal explanations.

The third takeaway from Figure 4 is that it is hard to see any evidence of a discontinuous deceleration in the demand for educated labor in the 2000s. Whether fit using the linear spline

(allowing all the post-1992 points to cluster along one axis) or a quadratic trend, which allows the deceleration to cumulate over the full sample, there's almost nothing in this figure that suggests a trend break in demand in the 2000s.⁵ Rather, this evidence suggests that the trend movements in relative demand in the 2000s were a continuation of those commencing circa 1992.

Conclusion: Timing is everything

These fact patterns lead me to draw a distinct inference from BGS: we *should not* be surprised by the evolution of the skill premium—or even the weaker job prospects of college-educated workers—in the 2000s. These outcomes are consistent with steadily rising demand for college-educated labor and a surprising surge in new college entrants in the U.S. labor market after 2003, which depressed the skill premium as it would be predicted to do. We *should* however be deeply puzzled by the sudden trend shift in demand after 1992, which ushered in (at least) 20 years of slower (though still non-negligible) growth in demand for skilled labor.

This development is not altogether bad news, however. Had demand for skilled labor continued to rise after 1992 at its pre-1992 pace, the estimates in Figure 3 suggest that the U.S. would have seen *substantially* more growth of between-group inequality—specifically, a meteoric 30 log point rise in the college premium between 1992 and 2012, nearly three times as large as the economically significant rise of 11 log points that actually occurred. This ‘good news’ is at best partial, however. In the canonical model, relative demand shifts intrinsically convey good economic news because they imply ongoing factor-augmenting technological progress.⁶ Thus, this slowdown may be read to support BGS’ view that as Information Technology has matured, the pace of accompanying labor augmentation has slackened. If so, however, we would want to caveat their conclusion to note that this slowdown started about ten years prior to the date that BGS pinpoint, and that it occurred during a period in which aggregate productivity growth was robust and U.S. IT investment was rising extraordinarily rapidly (Figure 1).

⁵If one squints, it's possible to see that some of the points immediately after 2000 fall slightly below the regression line whereas those immediately before fall slightly above it—implying a possible further deceleration after 2000. But then the last three points in the series (2010-12) again lie slightly above the regression line, suggesting a tiny re-acceleration. This is pretty thin evidence.

⁶This is also true for technological progress that raises A_L or both A_H and A_L . Presuming as the model does that technological progress is always factor augmenting—there are no factor-retarding technological regresses—any movement of A_H or A_L must be correspond to an *increase* of either or both term and hence rising productivity.

References

Acemoglu, Daron, and David H. Autor. 2011. "Skills, Task and Technologies: Implications for Employment and Earnings." *Handbook of Labor Economics*, 4: 1043–1171.

Autor, David H. 2014. "Skills, Education, and the Rise of Earnings Inequality Among the "Other 99 Percent"." *Science*, 344(6186): 843–851.

Autor, David H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives*, 29(3): 3–30.

Autor, David H, Lawrence F Katz, and Melissa S Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics*, 90(2): 300–323.

Beaudry, Paul, David a. Green, and Benjamin M. Sand. 2014. "The Declining Fortunes of the Young Since 2000." *American Economic Review*, 104(5): 381–386.

Card, David, and John E. DiNardo. 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics*, 20(4): 733–783.

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(2): 705–746.

Freeman, Richard B. 1986. "Chapter 6 Demand for education." *Handbook of Labor Economics*, 1: 357–386.

Goldin, Claudia, and Lawrence F. Katz. 2008. *The Race Between Education and Technology*. Cambridge:Harvard University Press.

Goldin, Claudia, and Robert A. Margo. 1992. "The Great Compression: The Wage Structure in the United States at Mid-Century." *The Quarterly Journal of Economics*, 107(1): 1–34.

Heckman, James J, Lance Lochner, and Christopher Taber. 1998. "Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents." *Review of economic dynamics*, 1(1): 1–58.

Johnson, George E. 1970. "The demand for labor by educational category." *Southern Economic Journal*, 190–204.

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

Murphy, Kevin M., and Finis Welch. 1992. "The Structure of Wages." *Quarterly Journal of Economics*, 107(1): 285–326.

Tinbergen, Jan. 1974. "Substitution of Graduate Labor by Other." *Kyklos*, 27(2): 217–226.