Graduate Labor Economics

Notes to Accompany Lecture 3: The College Wage Premium and the "Canonical Model"

Brendan M. Price*
Federal Reserve Board

The study of inequality serves as something of an organizing principle for modern labor economics. The steady rise of inequality is one of the central social facts of our time, and many of the debates raging in our field are readily viewed through the lens of inequality: the distributional impacts of trade, technological change, and immigration; discrimination on the basis of race, sex, and other characteristics; the design of the safety net; the effects of labor-market institutions like minimum wages, unions, and legal regimes; intergenerational impacts of socioeconomic disadvantage; on and on and on. Our collective effort to understand and unpack the rise of inequality has also spurred the creation of many new methods that are now used to study a much wider set of questions. For all these reasons, inequality is the right place to start.

This lecture note will briefly summarize some overall trends in wage inequality, then zero in on systematic wage disparities between workers with different levels of education. We'll see that the rapid rise of the college wage premium can be understood as the result of long-term demand shifts favoring high-skill workers coupled with a slowdown in the relative growth of the supply of college-educated workers.

1 Inequality: some preliminaries

We'll start by considering different forms of inequality and how they are conventionally measured.

1.1 Notions of inequality

- As a threshold matter, we need to be clear about *what kind* of inequality we have in mind. There are several relevant constructs, each of them useful for answering different questions:
  - Wage inequality ($w_i e_i$): people receiving different wages for an hour of work. This may reflect either differences in productivity (or "efficiency", e_i) or differences in the price paid for each efficiency unit supplied (w_i). Empirically, the law of one price for (an efficiency unit of) labor doesn't always hold: many recent papers have shown that there is substantial cross-firm dispersion in wage-setting, with "good" firms paying above-average wages and "bad" ones paying below-average wages.
  - Earnings inequality ($w_i e_i h_i$): differences in earned income across workers (due to both hourly wages and hours worked). Earnings are measured at different frequencies in different datasets; typically we work with weekly, monthly, or annual earnings.

 Email: brendan.m.price@frb.gov. These notes are adapted from UC Davis course ECN 250A, which I taught in Winter 2018 and Spring 2019. Any views or opinions expressed here are my own and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff. This version: 2020/03/21.

- Compensation inequality ($w_i e_i h_i + b_i$): disparities in income inclusive of non-wage compensation b_i ("fringe benefits") like health insurance and pensions (Pierce, 2001, 2010).
- There's also *income inequality* (disparities in earned income + unearned income), *wealth inequality*, and *consumption inequality*. I won't have much to say about these, though the growing concentration of income and wealth among top-earners is a much-discussed phenomenon in its own right.

- For a given type of inequality, we face the additional question of whether to analyze *cross-sectional inequality* (across workers at a fixed point in time), *life-cycle inequality* (within workers over time), or *intergenerational inequality* (social mobility across generations).
- We'll be focusing mostly on cross-sectional wage inequality: our goal is to understand changes in how the labor market rewards different kinds of workers, and (hourly) wages are the best indication of that: in a competitive equilibrium—and in the absence of compensating wage differentials linked to differences in job (dis)amenities—the wage tells us the market price placed on a worker's time.

1.2 Measuring inequality

- Once we answer the question "inequality of *what*?", we have to decide how to measure it.
- Whichever measure we use, we typically work with *log* wages. Taking logs makes our wage measure "scale-free", facilitating comparisons over time and across countries, and lets us interpret regression coefficients as elasticities or semi-elasticities.
- Several measures are commonly used:
  - Variance of (log) wages. A nice property of the variance is that we can harness the law of total variance: $\text{Var}(w_i) = \text{Var}(\mathbb{E}(w_i \mid x_i)) + \mathbb{E}(\text{Var}(w_i \mid x_i))$. This lets us decompose wage variation into "between-group" and "within-group" components.[1]
  - Gini coefficient. This is often used for cross-country comparisons, but it's used less often in the labor literature.
  - Quantiles. Labor economists love quantiles because they let us characterize changes in inequality *at different points in the wage distribution*. One leading measure of wage inequality is the "log 90–10", the log ratio of the 90th percentile to the 10th.[2]
  - The log 90–10 is the sum of the log 90–50 and the log 50–10:

$$\log\left(\frac{w_{90}}{w_{10}}\right) = \log\left(\frac{w_{90}}{w_{50}}\right) + \log\left(\frac{w_{50}}{w_{10}}\right)$$

This lets us decompose total wage inequality into "lower-tail" and "upper-tail" components. Lower- and upper-tail inequality do not always move in lockstep. Where in the

[1] If we posit a linear model for the conditional expectation function, so that $w_i = x_i'\beta + \varepsilon_i$, then the within-group component is the variance of the residual, $\text{Var}(\varepsilon_i)$. For this reason, within-group inequality is often referred to as "residual inequality".

[2] One reason we often focus on the 10th and 90th percentiles—rather than, say, the 1st and 99th percentiles—is that the extreme tails of the income distribution are often not measured as well in the survey datasets that historically dominated the literature on wage inequality: the receipt of government transfers may be understated at the lowest percentile, while the top percentile is often topcoded in publicly available data. In recent years, the use of administrative records drawn from taxes and other databases has facilitated a hugely influential literature on the top 1 percent (see, e.g., Piketty and Saez, 2003).

wage distribution growth in inequality is concentrated in a given period may be informative about the underlying causes: for instance, changes in the minimum wage may influence the 50–10 but are unlikely to affect the 90–50.

- Recall that (small) changes in logarithms approximate percentage changes.
  - Mathematical basis: $\lim_{x\to 0}\log(1+x) = x$. For positive x, the log approximation understates the true percentage change. For negative x, it overstates the percentage change.
  - Example: real median US household income rose from \$57,230 in 2015 to \$59,039 in 2016, in 2016 dollars (FRED Economic Data). The log change is $100 \times \log\left(\frac{59{,}039}{57{,}230}\right) = 3.11$ log points, and the percent change is $100 \times \frac{59{,}039-57{,}230}{57{,}230} = 3.16$ percent.
  - The approximation is quite good for modest changes: for $x = .10$ (a 10 percent increase), the log approximation gives a 9.5 percent change. For large x it's way off: for $x = .50$, the log approximation is 40.5 percent, and for $x = 1.0$ the approximation is 69.3 percent.
  - We can always recover exact percent changes from log changes, using the transformation $\log(1+x) = \beta \implies x = \exp(\beta) - 1$.
- Whenever we talk about wage inequality, a key question is how we handle people who don't work. Treating non-workers as having a zero wage would paint a misleading picture of wage inequality since these people would certainly earn positive wages if they did work. Here are three commonly used approaches:
  - Exclude non-workers from the wage calculation. When doing so, we need to think hard about selection. If the lowest-paid workers exit from employment—say, in response to an adverse demand shock that reduces their earnings potential—then observed wage inequality might fall even if there were increased dispersion in the wages workers *could* command in the market.
  - Impute wages for non-workers on the basis of observable characteristics. That is, we can regress log wages on observables x_i among people who *do* work, then use the fitted regression model to predict wages for those who *don't*. Of course, these predictions may be biased if non-workers differ from workers on unobservable dimensions.
  - When working with quantiles, it sometimes makes sense to treat non-workers as zeroes—effectively assuming that, if they chose to work, they would still be at the left tail of the wage distribution. This, of course, is a very strong assumption.

2 The rise of wage inequality

Inequality is a supremely complex phenomenon, and there's no way I can do justice to it in the limited time we have here. So I'll just summarize a few key facts that motivate the papers we'll be studying in the next several lectures. Katz and Autor (1999), Goldin and Katz (2007), and Acemoglu and Autor (2011) offer good general overviews of the rise in US wage inequality. There is a large international literature as well.

- While recent decades have broadly been characterized by steadily increasing levels of wage inequality, the US wage structure actually *narrowed* significantly around mid-century.

- Goldin and Katz (2007): "The growth in wage inequality since the late 1970s was preceded by a substantial narrowing of the wage structure during the 'Great Compression' of the 1940s, when the male 90-10 log weekly wage gap decreased by 35 log points, and then by a period of little change in wage inequality during the 1950s and 1960s."

- After the broadly shared wage gains of the 1960s and the broadly suffered wage stagnation of the 1970s, wage inequality rose rapidly in the 1980s: between 1979 and 1987, the log 90–10 for full-time/full-year (FTFY) workers rose by 20 log points for men (from 1.25 to 1.45) and by 25 log points for women (from 1.05 to 1.30), according to Katz and Autor (1999).
  - Inequality rose *monotonically* during this decade: top-quantile wages rose relative to those in the middle quantiles, which in turn rose relative to the bottom ones (Juhn et al., 1993).
  - Inequality rose *between* education, experience, and occupation groups (though the gender wage gap narrowed significantly) and also *within* narrowly defined demographic groups.

- Whereas the 1980s were characterized by pervasive increases in wage inequality, the 1990s was a decade of labor market "polarization", with continued growth in upper-tail wage inequality coupled with stable or declining lower-tail wage inequality (Autor et al., 2006).[3] Polarization is especially apparent when one looks at the occupational composition of the workforce: recent decades have witnessed a steady "hollowing-out" of middle-class occupations.
  - This, at least, is the conventional wisdom. A recent working paper by Hunt and Nunn (2019) contends that, at least in the US context, the evidence for polarization is much weaker than generally believed, owing to a mix of methodological problems with widely used approaches and classification errors in occupational crosswalks.

- Real wages have fallen for some groups since 1980, e.g., men with less than a high school education (see Acemoglu and Autor, 2011). Many theories that predict *relative* wage declines for less-educated workers have difficulty explaining *absolute* declines, so this is an important part of the overall puzzle.

- These phenomena are not unique to the United States: many other developed countries have experienced rising wage inequality in recent decades (e.g., for Germany, see Dustmann et al., 2009), though generally not to the same degree as the United States. (I know less about wage structures in developing economies and will mainly be focusing on the OECD.)

- Researchers have advanced many possible causes for the rise of wage inequality in the US and other OECD countries, including (among others) skill-biased technological change, rising import competition from low-wage countries, deunionization, declines in the real minimum wage, and growing monopsony power in the labor market. We'll discuss many of these candidate explanations throughout the course.

[3] There is debate about the evolution of wage inequality during this period, especially in the 1990s, with Lemieux (2006) arguing that the apparent continued rise in residual wage inequality in the 1990s stems mostly from compositional shifts and from measurement error in the March Current Population Survey (CPS) rather than true changes in the pricing of unobserved skill. Autor et al. (2008) dispute these conclusions, emphasizing the continued rise of upper-tail inequality even when accounting carefully for compositional changes. We won't get into this debate.

3 Katz and Murphy (1992)

- A key metric for understanding wage inequality is the college wage premium, the ratio of wages paid to college graduates vs. high school graduates. The college wage premium is sometimes (a bit imprecisely) referred to simply as the "skill premium", reflecting the presumption that education confers productive skills rather than functioning merely as a signal of latent ability.
- The US college wage premium rose in the 1960s, fell in the 1970s, rose dramatically in the 1980s, and has risen more slowly since then—though the *post-graduate* premium has continued to rise robustly (Lindley and Machin, 2016).
- Katz and Murphy (1992, hereafter "KM") seek to explain the ups and downs of the college wage premium using a simple supply-demand framework. Their approach has shaped much of the ensuing literature, so much so that Acemoglu and Autor (2011) have dubbed it "the canonical model" of skill supply and demand.
  - KM conclude that "observed fluctuations in the rate of growth of the relative supply of college graduates combined with smooth trend demand growth in favor of more-educated workers can largely explain fluctuations in the college/high school differential over the 1963–1987 period."
  - They attribute the rapid growth in the skill premium in the 1980s not to an acceleration of demand growth for skilled labor, but to a *deceleration of growth in the supply of more-educated workers.*

3.1 The price and quantity of college vs. non-college workers

- The first big challenge that KM confront is how to measure both the skill premium and the relative supply of high-school-equivalent workers and college-equivalent workers. It's worth understanding the logic of how they construct their samples, as similar ideas are commonly used in labor economics.
  - KM use data from the 1964–1988 March Current Population Survey (CPS), which reports each respondent's total earnings and total weeks worked in the preceding calendar year. So, these CPS extracts pertain to earnings in years 1963–1987. What KM call the "wage" is in fact the weekly wage, computed as total annual earnings divided by total weeks worked. An hourly wage would be preferable, but the March CPS doesn't report annual hours worked prior to 1976.
  - KM construct two distinct samples: a "wage sample" for measuring changes in wages for workers of given skill, and a "count sample" for tallying up labor supply in different skill categories.
  - The wage sample consists of full-time wage/salary workers who participated in the labor force for at least 39 weeks in the past year and worked for at least one week. KM exclude the self-employed, those who worked part-year due to school, retirement, or military service, and those with very low weekly earnings. These restrictions aim to identify workers with high labor force attachment.
  - The count sample simply includes everyone who worked at least one week in the past year—so it adds self-employed workers, those enrolled in school for part of the year, etc.

    - Why do KM use two samples instead of one? Researchers should strive for simplicity, and using two samples here entails additional complexity. But KM's choice is well-justified. The wage sample tries to maximize comparability over time, so that measured changes in the skill premium reflect changes in skill prices rather than compositional changes in the effective labor inputs supplied by a worker in each group. The count sample tries to calculate aggregate labor supply within each skill group, so it (sensibly) includes a broader set of workers.

- Using these samples, KM next construct composition-adjusted measures of both wages and labor supply for each demographic group of interest (e.g., men as a whole, female college graduates, etc.).
    - To do so, they first assign each worker to one of 320 cells defined by the interactions of sex, four commonly used education categories—less than a high school degree (<HS), high school graduate (HSG), some college (SMC), and college-plus (CLG)—and 40 single-year potential experience bins.[4]
    - Rather than directly reporting wage changes for a particular group (say, female college graduates), KM report share-weighted averages of the wage changes experienced by the cells that constitute that group, with weights determined by a cell's average share of employment during the sample period. This adjustment removes any changes in wages that stems from shifts in the demographic makeup of a given group, in an effort to isolate changes in wages holding labor inputs constant.
    - When reporting changes in labor supply for a particular group (e.g., female college graduates), KM again account for compositional changes by incorporating within-group shifts in favor of higher-paid—and hence "more efficient"—subgroups (e.g., workers with more potential experience).

- Table 1 reports "fixed-weight" wage growth by sex, education, and experience groups over 1963–1987, broken out by subperiod. There are a lot of interesting patterns here, but we'll focus on education:
    - Over 1963–1987, wage growth is monotonically increasing across education levels: less than high school wages rose by 10.9 log points, whereas college graduate wages rose by 23.1 log points.
    - This overall pattern masks heterogeneity across time periods. In the 1960s, college graduates saw faster wage gains than other groups. The college wage premium then *fell* substantially in the 1970s before rising dramatically in the 1980s (see KM Figure 1c).
    - By construction, changes in the college/non-college wages reported here hold constant the sex and experience composition of each skill group. So it's plausible (though not uncontestable) to interpret these numbers as changes in the price of skill.

- Table 2 reports changes in the relative *supply* of each demographic group.

[4] Labor economists usually proxy for *actual* experience using *potential* experience, typically defined as (age − years of schooling − 6) or something similar. There are two reasons to do so. First, computing actual experience requires each worker's complete labor market history, which is seldom observed in survey data (though sometimes observed in administrative data). Second, actual experience is endogenous to labor market conditions encountered throughout a worker's career. If a treatment impacts accumulated experience, then actual experience may be a "bad control" (Angrist and Pischke, 2009, pp. 64–68).

- The share of workers with a college education rose dramatically between 1963 and 1987. But there are key differences in the speed of these changes across different subperiods: the supply of college graduates rose fastest in the 1970s and slowest in the 1980s.
- KM's results (discussed below) rely mainly on "time-series identification" related to fluctuations in the relative supply of college-educated workers. In empirical work, it's always wise to know "where the variation is coming from", both for assessing threats to identification and for interpreting the results. So, what drove the accelerating and then decelerating growth in the college share in the 1970s and 1980s? A major driver was the postwar baby boom: during the 1970s, large cohorts of highly educated young workers entered the labor market, rapidly raising average education levels in the workforce.[5] This process slowed in the 1980s as smaller post-boom cohorts came of age.

3.2 Demand shifts, supply shifts, or both?

- Before introducing their celebrated CES framework, KM use a simple supply-demand framework to argue that any explanation for the facts laid out above *must* incorporate shifts in the relative demand for certain types of workers—not merely shifts in relative supply.
- Treating each demographic group as a distinct factor of production, the economy-wide vector of (conditional) factor demands in year t is given by

$$X_t = D(W_t, Z_t)$$

  where W_t is a vector of factor prices and Z_t is a vector of demand shifters (e.g., technology).
- If the aggregate production function is concave, then the $K \times K$ matrix of cross-price effects on factor demands, D_w, is negative semidefinite. This implies that $dW_t'(dX_t - D_z dZ_t) \leq 0$: "changes in factor supplies (net of demand shifts) and changes in wages must negatively covary".
- KM test whether observed shifts in factor supplies and prices are consistent with a world of *stable factor demand* ($dZ_t = 0$). Under that polar assumption, a discrete analogue to the inequality above implies that $\Delta W_t' \Delta dX_t \leq 0$. Computing this dot product for different time intervals, KM find that it's mostly negative prior to the mid-1980s, but positive for periods inclusive of 1983–1987.
  - This implies that a "supply-only story" *could* explain the patterns in these data prior to the 1980s, but that demand shifts *have to be part of the story* over the full period.
  - This is clearest in Figure 3: during the full sample period 1963–1987, groups that exhibited increases in relative supply also saw increases in relative wages.
- When KM compute $\Delta W_t' \Delta X_t$ using *detrended* wages and supplies, the result is more consistently negative: that is, deviations of wages and supplies from trend are negatively correlated, suggesting they may be driven by supply factors. KM infer from this that a linear demand trend, coupled with higher-frequency supply fluctuations, might suffice to explain the evolution of relative factor prices.
- Section V of the paper looks at how skill demands have changed between and within industries. We'll skip this for now, but we'll see similar ideas when we look at Berman et al. (1994).

[5] The baby boomers achieved above-trend educational attainment in part due to the Vietnam Era, which incentivized men to attend college to avoid being drafted (Card and Lemieux, 2001b).

3.3 "The race between education and technology"

- Today, KM is best remembered for the regression model developed in Section VI. Their framework formalizes what Jan Tinbergen famously called a "race" between education and technology (a title later adopted by Goldin and Katz, 2010). KM's parsimonious model has surprising explanatory power, and it's a great example of using economic theory to motivate and interpret an empirical relationship.

- KM assume a CES production function in which college-equivalents (group 1) and high school-equivalents (group 2) are imperfect substitutes in production:

$$Y_t = \left[(A_{1t}x_{1t})^\rho + (A_{2t}x_{2t})^\rho\right]^{\frac{1}{\rho}}.$$

  where x_{gt} denotes group-g supply and A_{gt} is a g-augmenting technology shifter.

- Setting each group's wage equal to its marginal product, taking ratios to equate relative wages with the MRTS, and then taking logs yields KM's equation 17:

$$\log\left(\frac{w_{1t}}{w_{2t}}\right) = \frac{1}{\sigma}\left(D_t - \log\left(\frac{x_{1t}}{x_{2t}}\right)\right)$$

  where $D_t \equiv (\sigma - 1)\log\left(\frac{A_{1t}}{A_{2t}}\right)$ and $\sigma = \frac{1}{1-\rho}$ is the elasticity of substitution.

  - Note that σ is the key parameter that relates changes in skill supplies to changes in the skill premium: when σ is large, the groups are close substitutes, so that their relative prevalence doesn't affect relative wages much.

- To estimate this model empirically, KM need to compute the skill supplies x_{1t} and x_{2t}. They treat each high school graduate as one high school equivalent, each college graduate as one college equivalent, and each some-college or less-than-high-school worker as a linear combination of the high school and college groups, with weights determined by the degree to which "SMC" and "<HS" wages covary with those of "HSG" and "CLG" workers.

- KM assume that the demand term D_t can be modeled as a linear trend, implying the regression model

$$\log\left(\frac{w_{1t}}{w_{2t}}\right) = \alpha_0 + \alpha_1 t + \beta \log\left(\frac{x_{1t}}{x_{2t}}\right) + \varepsilon_t$$

  Estimating their model over 1963–1987, KM obtain an estimated elasticity of substitution of $\hat{\sigma} = -\frac{1}{\hat{\beta}} = 1.41$. Since $\hat{\sigma} > 1$, this implies, critically, that college and non-college workers are *gross substitutes*, which in turn implies that skill-biased technical change will increase the college wage premium. KM also find a 3.3 percent-per-year secular rise in the skill premium ($\hat{\alpha}_1 = .033$), suggesting steady shifts in labor demand towards college-educated workers.

4 Extensions and refinements

KM has stimulated lots of subsequent work that operates within the basic framework of the canonical model. I'll summarize some notable papers.

- Card and Lemieux (2001a) point out that, according to KM's model, the college wage premium should evolve proportionally over time for workers of different age groups. Empirically, however, the dramatic rise of the college wage premium in the US (and also in the UK and

Canada) has been largely concentrated among younger workers. Card and Lemieux explain this result using an augmented production function that allows for imperfect substitutability across workers of different experience levels within each education group: aggregate production is still a CES aggregate of high school (H) and college-educated (C) labor, but now the supply of each education group is itself a CES aggregate of the labor supplied by different education-experience cells (yielding a "nested" CES function):

$$Y_t = (\theta_{ht} H_t^\rho + \theta_{ct} C_t^\rho)^{\frac{1}{\rho}}, \quad H_t = \left(\sum_j (\alpha_j H_{jt}^\eta) \right)^{\frac{1}{\eta}}, \quad C_t = \left(\sum_j (\alpha_j C_{jt}^\eta) \right)^{\frac{1}{\eta}}$$

where $\sigma_A \equiv \frac{1}{1-\eta}$ is the cross-age elasticity of substitution within each education group. Card and Lemieux find $\hat{\sigma}_A \approx 4$–5, implying that younger and older workers are close but not perfect substitutes. With imperfect substitution between age groups, a slowdown in the trend towards rising collegiate attainment implies faster growth in the wage premium for younger workers than for older ones.

- Krusell et al. (2000) argue that the rising skill premium really reflects capital-skill complementarity coupled with the growth of the capital stock, rather than skill-biased technical change. Formally, Krusell et al. assume a nested CES production function that, in simplified notation, takes the form

$$Y_t = \left[u_t^\eta + (k_{et}^\rho + s_t^\rho)^{\eta/\rho} \right]^{1/\eta}$$

  where u_t is unskilled labor, s_t is skilled labor, and k_{et} is the stock of capital equipment. The elasticity of substitution between capital and skilled [respectively, unskilled] labor is $\sigma_s = \frac{1}{1-\rho}$ [$\sigma_u = \frac{1}{1-\eta}$]. Under capital-skill complementarity ($\sigma_s < \sigma_u$), an increase in the capital-labor ratio will increase the skill premium. Since falling capital prices have led to rapid growth in equipment capital, Krusell et al. find that capital deepening can explain most of what KM interpret as skill-biased shifts in demand.

Other papers that build on the canonical model include:

- Acemoglu (1998), who shows theoretically how shifts in the relative supply of skilled workers may endogenously lead to skill-biased technical change (as innovators seek to tap the growing market for skill-augmenting technology);

- Carneiro and Lee (2011), who provide evidence for a decline in the average "quality" of recent college graduation cohorts, which in turn implies that the quality-adjusted college premium has risen even faster than the raw premium; and

- Bowlus et al. (2017), who—using a methodology very different from that of Carneiro and Lee—find that correcting for unobserved changes in cohort quality substantially improves the explanatory power of KM's basic framework outside of the original 1963–1987 sample period.

Lastly, Acemoglu and Autor (2011) provide a very clear discussion of the canonical model, its empirical successes, and the facts that it struggles to explain.

References

Acemoglu, D. (1998). Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. *Quarterly Journal of Economics*, 113(4):1055–1089.

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, volume 4B, pages 1043–1171.

Angrist, J. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: an Empiricist's Companion.* Princeton University Press, Princeton, NJ.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2006). The Polarization of the US Labor Market. *American Economic Review Papers and Proceedings*, 96(2):189–194.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2008). Trends in US Wage Inequality: Revising the Revisionists. *Review of Economics and Statistics*, 90(2):300–323.

Berman, E., Bound, J., and Griliches, Z. (1994). Changes in the Demand for Skilled Labor within US Manufacturing: Evidence from the Annual Survey of Manufacturers. *Quarterly Journal of Economics*, 109(2):367–397.

Bowlus, A., Bozkurt, E., Lochner, L., and Robinson, C. (2017). Wages and Employment: The Canonical Model Revisited. NBER working paper 24069.

Card, D. and Lemieux, T. (2001a). Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis. *Quarterly Journal of Economics*, 116(2):705–746.

Card, D. and Lemieux, T. (2001b). Going to College to Avoid the Draft: The Unintended Legacy of the Vietnam War. *American Economic Review Papers and Proceedings*, 91(2):97–102.

Carneiro, P. and Lee, S. (2011). Trends in Quality-Adjusted Skill Premia in the United States, 1960–2000. *American Economic Review*, 101(6):2309–2349.

Dustmann, C., Ludsteck, J., and Schönberg, U. (2009). Revisiting the German Wage Structure. *Quarterly Journal of Economics*, 124(2):843–881.

Goldin, C. and Katz, L. F. (2007). Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing. *Brookings Papers on Economic Activity*, 2:135–165.

Goldin, C. and Katz, L. F. (2010). *The Race Between Education and Technology.* Belknap Press, Cambridge, MA.

Hunt, J. and Nunn, R. (2019). Is Employment Polarization Informative About Wage Inequality and Is Employment Really Polarizing? NBER working paper 26064.

Juhn, C., Murphy, K. M., and Pierce, B. (1993). Wage Inequality and the Rise in Returns to Skill. *Journal of Political Economy*, 101(3):410–442.

Katz, L. F. and Autor, D. H. (1999). Changes in the Wage Structure and Earnings Inequality. In *Handbook of Labor Economics*, volume 3, pages 1463–1555.

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

Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V., and Violante, G. L. (2000). Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica*, 68(5):1029–1053.

Lemieux, T. (2006). Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill? *American Economic Review*, 96(3):461–498.

Lindley, J. and Machin, S. (2016). The Rising Post-Graduate Wage Premium. *Economica*, 83(330):281–306.

Pierce, B. (2001). Compensation Inequality. *Quarterly Journal of Economics*, 116(4):1493–1525.

Pierce, B. (2010). Recent Trends in Compensation Inequality. In *Labor in the New Economy*, pages 63–98.

Piketty, T. and Saez, E. (2003). Income Inequality in the United States, 1913-1998. *Quarterly Journal of Economics*, 118(1):1–39.