

NBER WORKING PAPER SERIES

MULTIDIMENSIONAL HUMAN CAPITAL AND THE WAGE STRUCTURE

David J. Deming

Working Paper 31001

<http://www.nber.org/papers/w31001>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

March 2023

I have nothing to disclose. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by David J. Deming. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Multidimensional Human Capital and the Wage Structure
David J. Deming
NBER Working Paper No. 31001
March 2023
JEL No. J24

ABSTRACT

This paper reviews and synthesizes the literature on the macroeconomic implications of human capital theory. I begin with a review of the canonical model of education and the wage structure pioneered by Tinbergen (1975) and developed more fully by Goldin and Katz (2007). I also review innovations such as the task framework developed by Acemoglu and Autor (2011). The canonical model does a surprisingly good job of predicting changes in the wage structure in the U.S. and other developed countries over the last half-century. Relative to the canonical model, the task framework adopts a more flexible view of technology and does a better job of fitting non-monotonic changes in the wage structure. Yet the task framework does not fully explain why educated workers have done so well since 1980, nor does it explain other recent facts such as flattening returns to cognitive skills and growing returns to non-cognitive, “higher-order” skills such as teamwork. To understand these recent trends, we must move beyond a single index view of human capital, toward richer, multi-dimensional frameworks. I conclude with a discussion of the nascent literature on the implications of multi-dimensional human capital for the wage structure, which raises more questions than it answers.

David J. Deming
Harvard Kennedy School
Malcolm Wiener Center for Social Policy
79 JFK St
Cambridge, MA 02138
and Harvard Kennedy School
and also NBER
david_deming@harvard.edu

1 Introduction

According to human capital theory, education and training are investments where the costs are paid up front and the benefits are earned later in the form of increased productivity and earnings. At least since Becker (1962), economists have repeatedly found that differences in the quantity and quality of schooling are surprisingly powerful predictors of earnings differences between people and across countries.

In terms of both prediction and policy impact, human capital theory is one of the economics profession's biggest success stories. Over the last seventy years, the share of the world's population with at least some secondary school education increased from 13 percent to 51 percent, and the share with some tertiary education increased nearly sevenfold, from 2.2 percent to 14.6 percent (Lee and Lee 2016). Education spending in the U.S. has more than doubled as a share of GDP over this same period, and growth rates are even faster in developing countries (Deming 2022). Global economic growth exploded over this same period, especially in countries that invested more in education (e.g. Hanushek and Woessmann 2010).

Not surprisingly, scholarly study of human capital has also blossomed over the last half-century. Beginning with the Mincer (1958) model, hundreds of studies estimate the economic return to schooling investments. The search for instrumental variables that affect schooling but are unrelated to unobserved determinants of earnings was a central focus of labor economists for decades and spawned the “credibility revolution” in economics (Angrist and Pischke 2010).

To date, education and labor economists have focused mostly on the microeconomic implications of human capital theory. Education increases earnings by augmenting the marginal product of labor. People invest more in their own education when the economic return is higher, and so on. This literature has been a spectacular success. The causal impact of education on earnings is one of the most well-established and important findings in social science.

By their nature, causal estimates of the return to education are *ceteris paribus* tests of human capital theory. Yet the value of skills learned in school is contextual and depends on the technology of workplace production and – more broadly – the economic environment. Numeracy and the ability to understand and analyze data are more valuable now than before computers were invented. Social skills are more important now that workers are assigned to flexible problem-focused teams, rather than taking spots on a factory assembly line. Education conveys a mix of general and specific skills, and these skills have an economic return that is determined by the intersection of supply and demand. Skills can be valuable in one context but worthless in another; skills can be scarce in some economies and plentiful in others.

Macroeconomic tests of human capital theory go beyond the individual return to education and focus instead on the wage structure of an economy. They suggest, for example, that causal estimates of the return to a year of education are determined by the equilibrium between the supply of and the demand for skills. Estimates of the return to education are likely to be higher when skill demand outraces skill supply – not when educated workers are scarce in an absolute sense, but scarce relative to demand. This explains why relative wages for U.S. college graduates rose rapidly in the 1980s and have stayed high through the present day, and why the return to education is higher in India than in South Korea, where the shares of young people with a college degree are 15 percent vs. 50 percent respectively. Clearly, micro and macro tests of human capital theory are strong complements, and so education economists should be experts in both literatures.

This chapter reviews the evidence on the macroeconomic implications of human capital theory. I start with a simple competitive, closed-economy model with two types of workers (high and low skill) who are imperfect substitutes and are both paid their marginal products. This basic framework is sometimes called the canonical model, or the supply-demand-institutions (SDI) framework (Tinbergen 1975, Katz and Murphy 1992, Goldin and Katz 2007, Acemoglu and Autor 2011). In these models, technology is factor-augmenting,

meaning it complements either low- or high-skilled labor. In the “education race” model of Tinbergen (1975) and Goldin and Katz (2007), technology is always skill-biased, and so the return to skills and/or education is determined by a “race” between the increasing supply of skills and increasing relative demand.

I then discuss the task framework developed by Acemoglu and Autor (2011), a prominent extension of the canonical model which adds a mapping between skills and job tasks. In the canonical model, technology operates only by shifting the relative prices of high and low-skilled labor. The task framework allows technology to directly replace labor in some tasks, which permits real wage declines when machines substitute for workers as well as non-monotone changes in the wage structure. Overall, the canonical model and the task framework do a good job of explaining changes in the returns to education and broad trends in the wage structure, both within and across countries and over time (Katz and Murphy 1992, Autor et al. 1998, Card and Lemieux 2001, Acemoglu and Autor 2011).

Despite their strengths, these existing models of education and technology are inadequate for understanding recent trends in the wage structure. Notably, they do not explain why the demand for college-educated labor continues to increase even as the return to cognitive skill has flattened and perhaps even declined since 2000 (Castex and Kogan Dechter 2014, Autor et al. 2020, Edin et al. 2022). I discuss evidence that the economic return to “non-cognitive”, higher-order skills has increased over this same period (Deming 2022). This pattern of evidence requires us to move beyond a single index view of human capital, toward richer, multi-dimensional frameworks.

I discuss three important limitations of the single index model of human capital. First, I review evidence of the growing importance of teamwork and social interaction, and I show that social skills are best viewed as a different variety of human capital, distinct from and complementary to cognitive skills. I develop the model of team production in Deming (2017) and show how it helps explain recent trends in the wage structure and yields testable predictions such as cognitive skill and social skill being complements in a Mincerian earnings

function. Second, I show that strong vintage effects for technical skills can help explain the variation in returns to human capital for young vs. old workers over time. Third, I discuss how the value of human capital may vary in imperfectly competitive labor markets, when workers with scarce combinations of skills receive higher rents within the firm (e.g. Kline et al. 2019, Jäger and Heining 2022).

In each case, there are many threads left hanging by existing work. In economics we are only beginning to confront the multidimensional and context-specific nature of human capital. I hope and expect that future reviews of this important literature will have much more to say about how and why skills “matter” for thinking about inequality and the wage structure. The paper proceeds as follows. Section 2 develops the canonical model and discuss related empirical work. Section 3 develops a task framework in the spirit of Acemoglu and Autor (2011) and also discusses related evidence. Sections 4 through 6 review and synthesize the evidence for vintages, varieties, and diverging valuations of human capital. Section 7 concludes.

2 The Canonical Model

The canonical model begins with an aggregate production function for a single composite good Y . Y is produced using two factors L and H , representing low- and high-skilled workers respectively. The model includes several simplifying assumptions, each of which can be relaxed. They include:

1. Only labor – not capital – as factors of production;
2. A closed economy, so no trade or offshoring;
3. Workers vary in their productivity but supply labor inelastically, e.g. each worker i has l_i or h_i efficiency units of labor (depending on their type) and $L = \int_{i \in L} l_i di$ and $H = \int_{i \in H} h_i di$; and

4. Perfect competition, meaning workers of each skill type are always paid their marginal product.

A key feature of the canonical model is that low- and high-skilled workers can be imperfect substitutes in production.¹ One way to think about this is that each skill group works in different occupations or types of jobs, which are combined in some fashion to produce total output. If low- and high-skilled workers are perfect substitutes, then there is effectively only one job, with wages in proportion to marginal products, and the relative supply of each worker type doesn't matter. If the production function is perfect complements (Leontief), workers of each type must be combined in fixed proportions.

The most general specification is thus a constant elasticity of substitution (CES) production function for the aggregate economy:

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

where $\sigma \in [0, \infty)$ is the elasticity of substitution between H and L and A_H and A_L are technology terms that augment the productivity of each skill group.² As $\sigma \rightarrow \infty$, equation (1) converges to the perfect substitutes case $Y = A_L L + A_H H$, so worker skill types can be combined in arbitrary proportions to produce output and relative supplies don't affect relative wages.

The elasticity of substitution σ determines the extent to which a policy that increases the supply of high-skilled workers can “scale up”. While most estimates of the return to education come from small-scale interventions, there is also evidence from large-scale educational expansions in countries like India and Taiwan (Chen and Chen 2021, Khanna 2023).

¹The notion of imperfect substitution between worker “types” in the canonical model can be generalized beyond low- and high-skill groups. We could have technical or production workers and managerial workers as the two types, for example. The CES production function in equation (1) can also be generalized to incorporate multiple goods (not just an aggregate), sectors, and factors of production.

²Equation (1) is sometimes expressed more simply as $Y = [(A_L L)^\rho + (A_H H)^\rho]^{\frac{1}{\rho}}$ with $\sigma = \frac{1}{1-\rho}$, $\rho \leq 1$. Alternatively, a share parameter is sometimes added, e.g. $Y = \left[\alpha (A_L L)^{\frac{\sigma-1}{\sigma}} + (1-\alpha) (A_H H)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$ with α governing the relative weight in production that is placed on low-skilled relative to high-skilled workers.

If skill groups are perfect substitutes, micro estimates of the return to education will scale up infinitely, because type H and L workers are interchangeable even if they have different skill levels. To see this more directly, note that under perfect competition workers are paid their marginal products ($\frac{\partial Y}{\partial L}$ and $\frac{\partial Y}{\partial H}$ respectively) and so we can derive the “skill premium” ω as the relative wage for high-skilled workers, which in log terms is:³

$$\ln \omega = \ln \left(\frac{w_H}{w_L} \right) = \ln \left(\frac{\frac{\partial Y}{\partial H}}{\frac{\partial Y}{\partial L}} \right) = \frac{\sigma - 1}{\sigma} \ln \left(\frac{A_H}{A_L} \right) - \frac{1}{\sigma} \left(\frac{H}{L} \right) \quad (2)$$

Equation (2) shows that as the elasticity of substitution between low-skilled and high-skilled workers increases (e.g. as $\sigma \rightarrow \infty$), the relative supply term $\frac{H}{L}$ is a less important determinant of the skill premium. In the limiting case of perfect substitution, the log skill premium is just equal to the skill bias of technology, represented here as the ratio $\frac{A_H}{A_L}$.

More generally, an increase in the relative supply of skills $\frac{H}{L}$ reduces the demand curve more for lower values of σ . We can show this by differentiating the log skill premium with respect to relative supply:

$$\frac{\partial \omega}{\partial \ln \left(\frac{H}{L} \right)} = -\frac{1}{\sigma} \quad (3)$$

Equation (3) is always negative, with a slope equal to $\frac{1}{\sigma}$. It is worth thinking through the mechanism through which increases in relative supply affect the skill premium. In equation (1), an increase in H without any corresponding increase in L would lower the wages of high skilled workers because the labor demand curve is downward sloping, e.g. $\frac{\partial w_H}{\partial \left(\frac{H}{L} \right)} < 0$. Conceptually, high-skilled labor is being used more intensively but less productively on the margin. One natural interpretation is that high-skilled workers substituting for the job functions – or tasks – previously performed by low-skilled workers. This substitution is implicit in the canonical model, but the task framework models it explicitly, a point we will

³ $w_L = \frac{\partial Y}{\partial L} = A_L^{\frac{\sigma-1}{\sigma}} \left[A_L^{\frac{\sigma-1}{\sigma}} + A_H^{\frac{\sigma-1}{\sigma}} \left(\frac{H}{L} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{1}{\sigma-1}}$ and $w_H = \frac{\partial Y}{\partial H} = A_H^{\frac{\sigma-1}{\sigma}} \left[A_H^{\frac{\sigma-1}{\sigma}} + A_L^{\frac{\sigma-1}{\sigma}} \left(\frac{H}{L} \right)^{-\frac{\sigma-1}{\sigma}} \right]^{\frac{1}{\sigma-1}}$, so the ratio is just $\frac{w_H}{w_L} = \left(\frac{A_H}{A_L} \right)^{\frac{\sigma-1}{\sigma}} \left(\frac{H}{L} \right)^{-\frac{1}{\sigma}}$.

return to in Section 3.

The relationship in equation (3) holds for a given skill bias of technology $\frac{A_H}{A_L}$. Strikingly, it suggests that without technological progress, growth in the supply of skills will always *reduce* the skill premium. Given the huge increases in educational attainment that have occurred over the last half-century, both in the U.S. and around the world, we would expect to see rapidly falling skill prices. If anything, the opposite has occurred, suggesting that technology has become more skill biased over time (e.g. Goldin and Katz 2007). To see this more precisely, differentiate equation (2) with respect to the skill bias term:

$$\frac{\partial \ln \omega}{\partial \ln \left(\frac{A_H}{A_L} \right)} = \frac{\sigma - 1}{\sigma} \quad (4)$$

Equation (4) shows the relationship between skill-biased technological change (SBTC) and the elasticity of substitution between high- and low-skilled labor. SBTC increases the skill premium as long as $\sigma > 1$.⁴ All of the empirical studies discussed below estimate an elasticity of substitution that is substantially greater than one, especially in the long-run, so I focus on the $\sigma > 1$ case going forward (Bils et al. 2022).⁵

Putting equations (3) and (4) together, we can see that the skill premium ω is the equilibrium between two countervailing forces – downward pressure from increases in the supply of skills and upward pressure from the increasing skill bias of technology. This is the essence of Tinbergen’s “education race” framework, which posits that the labor market return to skill is the outcome of a race between education and technology (Goldin and Katz 2007).

⁴If $\sigma < 1$, SBTC reduces the skill premium. Consider the case where worker skill groups are perfect complements. If L and H are combined in fixed proportions, then a technological improvement that makes high-skilled workers more productive will generate an even larger increase in the unit demand for low-skilled workers, because you need relatively more of them to leverage the benefits of high-skilled workers becoming more productive. While $\sigma < 1$ might hold in the very short-run or for a small firm, this is an unlikely case in practice, and nearly all empirical estimates find an elasticity substantially greater than one (Katz and Murphy 1992, Bils et al. 2022, Hendricks and Schoellman 2023).

⁵The aggregate production function is an abstraction which combines substitution possibilities across industries, firms, and factors of production. For this reason, the aggregate elasticity of substitution is almost certainly greater than for individual firms or sectors, and the long-run elasticity is probably larger than the short-run elasticity.

Empirical tests of the canonical education race model start with time series data on skill supplies and skill premia ($\frac{H}{L}$ and ω respectively), leaving the elasticity of substitution and the skill bias of technology (σ and $\frac{A_H}{A_L}$) as parameters to be estimated. Tinbergen’s hypothesis is that technological progress is always skill-biased (e.g. $\frac{\partial \ln(\frac{A_H}{A_L})}{\partial t} > 0$).

Katz and Murphy (1992) operationalize the model by treating college graduates and high school graduates as type H and L workers respectively, estimating the skill premium $\frac{\partial \ln \omega}{\partial \ln(\frac{H}{L})}$ as the relative return to a college degree and fitting a linear time trend through the data to represent steady skill bias in technological progress. These are strong assumptions, but nonetheless are a good place to start because they allow us to obtain estimates of σ for various time periods. We can also assess the fit of the model to the data.

Katz and Murphy (1992) fit the education race model to Current Population Survey (CPS) data for the period 1963-1987 and obtain an estimate of $\hat{\sigma} = 1.4$. The model fits the data surprisingly well, during a period where both relative demand and relative supply moved upward and downward over different sub-periods (see Figure 1, reproduced from Figure IV of their paper, for details). The model predicts both the downward trend in the college wage premium in the 1970s and the rapid growth in relative wages for college graduates beginning in the early 1980s.

The upward climb in the college wage premium in the early 1980s is an undisputed fact, but the interpretation is unclear even within the framework of the canonical model. While the model obtains an estimate of σ by imposing a linear trend in $\frac{A_H}{A_L}$, one could also assume a value of σ and back out the trend in $\frac{A_H}{A_L}$ (see Panel D of Figure 1 for an example). In general, fluctuations in the skill premium are more tightly linked to fluctuations in skill bias for higher values of σ , because it is easier to substitute between worker skill groups. In the limiting case of perfect substitutes, changes in the skill premium perfectly reflect changes in $\frac{A_H}{A_L}$ because relative skill supplies don’t matter. This would imply a very sudden and dramatic shift in relative demand for high-skilled workers beginning in the early 1980s.

Goldin and Katz (2007) estimate a version of the canonical model using U.S. data back to

1915. Figure 2 reproduces Figure 2 from Goldin and Katz (2007), which shows that the model fits the data surprisingly well over the last century. The two exceptions are the 1940s and the 1970s, where in both cases the skill premium was much lower than the model predicts. Goldin and Katz (2007) attribute this to wage setting policies, unions and other institutional factors, and call it the supply-demand-institutions (SDI) framework.

Other work provides estimates of the canonical model for more recent years (Autor et al. 1998, Acemoglu and Autor 2011, Autor et al. 2020). Acemoglu and Autor (2011) and Autor (2017) extend the Katz and Murphy (1992) model forward to the late 2000s, which yields much higher estimates of the elasticity of substitution σ . This is because the skill premium has grown much less in recent years than the model would have predicted.

Autor (2017) shows this by fitting the parameter estimates from Katz and Murphy (1992) to the longer period. The model fits well out of sample through the early 1990s, but then starts to diverge. From 1992 to 2012, the skill premium grew by 11.6 log points, whereas the estimates from Katz and Murphy (1992) suggest it should have grown by 30.4 log points (Autor 2017). Some authors deal with this by augmenting the Katz and Murphy (1992) specification with a separate trend post-1992 or a quadratic time trend, both of which allow the growth in relative demand to decelerate. This fits the data better but doesn't explain the economic phenomenon.

What explains slower growth of the skill premium over the last several decades? Beaudry et al. (2016) present evidence of a “great reversal” in demand for skills around the year 2000 and argue that this can be explained by the boom-and-bust cycle of capital investment following the information technology revolution in the 1980s. However, this only considers the demand side. Autor (2017) shows that growth in the relative supply of college educated labor in the 2000s explains more of the flattening skill premium than demand deceleration. Estimates of the canonical model suggest a smooth deceleration in the skill premium beginning in the early 1990s, not a sharp decline around 2000 (Autor 2017).

Another possibility is that the elasticity of substitution between high- and low-skilled

labor is higher in the long-run than in the short-run. Perhaps Katz and Murphy (1992) estimated the parameters of the canonical model during a period of rapid technological change, and the elasticity was lower because firms had less time to adjust. Bils et al. (2022) use evidence on returns to schooling across countries and over time and estimate a long-run elasticity of substitution of around 4. Hendricks and Schoellman (2018) use the wage gains from cross-country migration to calibrate estimates of the canonical model and find values of σ between 5 and 8. A technology “shock” in the late 1970s or early 1980s coupled with a gradual increase in σ as firms adjust could explain the flattening of the skill premium starting in the early 1990s.

A final possibility is that college graduates have become less productive over time. As a larger share of young people attend college, the marginal student may be less academically talented. Alternatively, college quality may suffer when many students attend, a phenomenon known as “cohort crowding” (Bound and Turner 2007).

Carneiro and Lee (2011) find that workers who grew up in regions with higher college-going rates earn less than comparable workers in the same labor market, suggesting that marginal college graduates are less productive. Their estimates imply that declines in the average productivity of college graduates reduced the college premium by about 6 percentage points between 1960 and 2000. The more general lesson is that the labor market return to college quality may vary by cohort for any number of reasons. Card and Lemieux (2001) and Bowlus et al. (2021) add the possibility of imperfect substitution across cohorts to the canonical model. I discuss this evidence further in Section 5.

Overall, the canonical model delivers an intuitive framework that does a surprisingly good job of summarizing long-run trends in educational wage premia in the U.S. over the last century. The model’s predictions also broadly hold in cross-country evidence. The median earnings premium for a tertiary education across OECD countries is around 50 percent, with lower values in countries such as Australia, Israel, Norway, and South Korea where educational attainment is especially high (OECD 2021). Broecke et al. (2018) use cross-

country data on measures of adult literacy and numeracy to test the canonical model’s prediction that the supply of skills and the skill premium are negative related. They find that higher relative skill supplies depress earnings at the top of the distribution (e.g. the 90/50 earnings ratio) but not the bottom (50/10), suggesting perhaps that institutional factors matter more for low-earning workers. Berman et al. (1998) and Machin and Van Reenen (1998) show that the predictions of the canonical model hold across many different countries.

The canonical model also delivers important predictions for the macroeconomic literature on growth and development. A common approach is development accounting, where one asks how much of the cross-country variation in income can be statistically explained by differences in human capital.

Development accounting works best when workers of different skill levels are perfect substitutes in production (e.g. when $\sigma \rightarrow \infty$ in equation 1). This is because in the perfect substitutes case, the aggregate production function is just $Y = A_L L + A_H H$, and relative supplies don’t affect relative wages. Thus we can simply “add up” human capital across countries. For example, if college graduates earn 50 percent more than high school graduates, they must be 50 percent more productive. The more these ratios differ across countries, the more human capital “matters” in a development accounting framework.

Because skill premia are remarkably similar across less- and more-developed countries, development accounting exercises that assume perfect substitution across skill groups typically find that education is not a very important contributor to cross-country income differences (Bils and Klenow 2000, Caselli 2005).⁶ However, Jones (2014) shows that relaxing the perfect substitutes assumption greatly increases the importance of human capital in explaining cross-country income differences. Table 2 in Jones (2014) shows the impact of imperfect substitution (e.g. lower values of σ) on development accounting. The Katz and Murphy (1992) estimate of $\sigma = 1.4$ suggests that human capital explains essentially all of cross-country in-

⁶Another limitation of development accounting is that it relies on good measurement of human capital. Hanushek and Woessmann (2012) show that including international achievement test scores into a development accounting framework substantially improves the explanatory power of human capital, and Rossi (2020) finds a similar result when adding test scores and measures of labor market experience.

come differences, whereas higher values of σ imply a smaller role. Hendricks and Schoellman (2023) use the wage gains from migration to calibrate a development accounting exercise that allows for larger long-run values of σ . They estimate that human capital accounts for roughly 60 percent of cross-country income differences. Thus the canonical model’s prediction that relative skill supplies affect a country’s wage structure is critical for understanding the contribution of human capital to economic growth.

The canonical model’s core insight is that the economic return to human capital is an equilibrium between supply and demand. Policies that expand access to education will have diminishing returns for the marginal enrollee, suggesting that small-scale causal estimates of the return to education will overstate returns in general equilibrium. One example is Khanna (2023), who studies a program in India that expanded public schooling in only half of the country’s regions. He combines micro evidence from the discontinuous change in funding across regions with macro evidence across areas and cohorts to estimate both reduced form and general equilibrium returns to education. He finds that educational expansion reduced the skill premium by about 35 percent relative to small-scale estimates from a regression discontinuity (RD) design.

A second insight from the canonical model is that technological progress can increase inequality even if it makes everyone better off, by benefiting the highly skilled more than others. Broad macroeconomic trends and case studies of specific firms all strongly suggest that the impact of computers and information technology has been skill-biased (that is, they have increased $\frac{A_H}{A_L}$). An unusually clear demonstration comes from Akerman et al. (2015), who study the staggered adoption of broadband internet in Norway and find that it increased the relative productivity of skilled workers, particularly those that work in the kinds of abstract, non-routine occupations that are complemented by digital technology. Lindner et al. (2022) estimate the impact of firm-level technology shocks on worker wages and find evidence for skill-biased technological change.

Despite its strengths, the canonical model has some important limitations. One is the

limited, “black box” view of technology. Technology augments units of labor in the canonical model (through the A_H and A_L terms) but does not otherwise appear in the production process at all. Yet common sense suggests that technology often changes work by directly replacing job tasks that were previously done by people.

The factor-augmenting form of technological progress yields empirical predictions that are sometimes at odds with the data. As noted by Acemoglu and Autor (2011), real earnings have declined for low-skilled workers in the U.S. over the last few decades. Yet in the realistic case where $\sigma > 1$, the canonical model implies that increases in A_L or A_H will always increase wages for both skill groups, even if one group benefits more than the other. A different view of technology is required to generate real wage declines. For example, technology such as computers and information technology have targeted impacts on jobs related to information processing, which may yield subtle and non-monotonic impacts on the wage structure (Autor et al. 2003).

Another limitation of the canonical model is its “black box” view of how workers apply their skills to produce output. Would type L and H workers have similar earnings if they worked in the same jobs? Do college-educated engineers have the same skills as college-educated nurses, even though they perform very different sets of job tasks? Skills do not produce output – rather, workers deploy their skills to specific job tasks to produce output, and these job tasks may be changing over time along with technology. The canonical model is silent on these intermediate steps.

3 The Task Framework

Acemoglu and Autor (2011) develop a task-based framework to address the shortcomings of the canonical model. They start with a production function for the aggregate economy, but it takes tasks as its main inputs rather than low- or high-skilled labor:

$$Y = \exp \left[\int_0^1 \ln y(i) di \right] \quad (5)$$

The production function combines a continuum of tasks i at production level $y(i)$ to generate aggregate output Y . Each task has its own production function:

$$y(i) = A_L \alpha_L(i) l(i) + A_M \alpha_M(i) m(i) + A_H \alpha_H(i) h(i) \quad (6)$$

There are now three types of workers – L , M , and H – with the A terms as factor-augmenting technologies as in the canonical model, the α terms as each worker type’s productivity in task i , and l , m , and h as the total labor supplied by each worker type assigned to task i . It is also possible to add capital as k/K or to add more worker skill groups. The key departure from the canonical model is that tasks can be performed by any worker skill group or by capital.

However, to say more about the allocation of workers to tasks and relative wages, we need to take a stand on whether tasks differ and whether workers are better at some tasks than others. Acemoglu and Autor (2011) impose a critical assumption to make the task framework tractable. They assume that the continuum of tasks i can be arrayed on a single dimension of “complexity” and that high-skilled workers are relatively more productive in the most complex tasks, followed by medium-skilled and low-skilled workers. Mathematically, they assume that $(\frac{\alpha_L(i)}{\alpha_M(i)})$ and $(\frac{\alpha_M(i)}{\alpha_H(i)})$ are both strictly decreasing along the task continuum i and continuously differentiable, which yields a smooth comparative advantage structure.

The equilibrium is defined by two task thresholds I_L and I_H , where type L workers perform the tasks from complexity levels 0 to I_L , type M workers perform the tasks between I_L and I_H , and type H workers perform the tasks between I_H and 1. The “boundary” tasks I_L and I_H are endogenous and respond to changes in skill supply and demand. This structure can be generalized to multiple vertically differentiated types and is similar to Ricardian trade models with a continuum of goods and skill groups as countries (Dornbusch et al. 1977, Eaton

and Kortum 2002).

Another critical assumption that makes the model work is a “law of one price for skills”. In a competitive labor market, all tasks employing workers of a given skill group must pay the same wage, and thus the marginal product of all workers in a skill group must be the same for all the tasks they perform.

To see this, let $p(i)$ be the unit price of task i and normalize the price of the final good Y to 1 so that $\exp\left[\int_0^1 \ln p(i) di\right] = 1$. Since workers of each skill level are paid the same wage for all tasks they perform, we can define a common task price $P_L = p(i)\alpha_L(i)$ for all $i < I_L$ and likewise for P_M and P_H within their respective task thresholds. The production function in (6) combined with the comparative advantage structure of the task productivity schedule yields the following expressions for wages:

$$w_L = P_L A_L \quad i < I_L$$

$$w_M = P_M A_M \quad I_L < i < I_H$$

$$w_H = P_H A_H \quad i > I_H$$

The Cobb-Douglas production function in (5) implies equal output shares for each task (e.g. $p(i)y(i) = p(i')y(i') = Y$ for any $i \in [0, 1]$). This yields a simple expression for the labor shares of each skill group in each task, which are just equal to total labor supply divided by the fraction of total tasks performed by each group:

$$l(i) = \frac{L}{I_L} \quad i < I_L$$

$$m(i) = \frac{M}{I_H - I_L} \quad I_L < i < I_H$$

$$h(i) = \frac{H}{1 - I_H} \quad i > I_H$$

A final implication is that the threshold tasks must have the same marginal product for worker types on either side of each boundary, and thus in equilibrium they could be supplied at the same level by either type of worker. Specifically, for threshold task I_L we have the “no arbitrage” condition $\frac{P_L A_L L}{I_L} = \frac{P_M A_M M}{I_H - I_L}$, and for threshold task I_H the condition is $\frac{P_H A_H H}{1 - I_H} = \frac{P_M A_M M}{I_H - I_L}$.

Since wages equal marginal products, we can use the “no arbitrage” conditions above to express relative wages in terms of relative supplies and the endogenous allocation of tasks to skill groups:

$$\begin{aligned} \frac{w_M}{w_L} &= \frac{P_M A_M}{P_L A_L} = \left(\frac{I_H - I_L}{I_L} \right) \left(\frac{L}{M} \right) \\ \frac{w_H}{w_M} &= \frac{P_H A_H}{P_M A_M} = \left(\frac{1 - I_H}{I_H - I_L} \right) \left(\frac{M}{H} \right) \end{aligned} \tag{7}$$

The expressions in (7) show that relative wages depend on relative supplies just as in the canonical model, but also on the share of the task continuum occupied by each worker group. Task shares are determined by the threshold tasks I_L and I_H , which respond to changes in the structure of comparative advantage across skill groups.

An increase in A_H holding all else constant would make type H workers more productive in all tasks. This would in turn shift the I_H task threshold downward, expanding the set of tasks that these workers perform and thus raising their relative wages.

There are also indirect effects on other skill groups. For example, if I_L stays constant and I_H shifts downward, there would be excess supply of type M workers. This equilibrium adjustment moves the I_L threshold down as well, shifting some of the tasks previously

performed by type L workers over to type M workers. However, the direct effect is always larger than the indirect effect, meaning an increase in A_H reduces $\frac{w_M}{w_L}$. An increase in A_L would yield results in the opposite direction. An increase in A_M would increase the wages of type M workers relative to the other two groups, but the impact on $\frac{w_H}{w_L}$ is ambiguous. See Acemoglu and Autor (2011) for formal derivations of these comparative statics.

The Acemoglu and Autor (2011) task framework is significantly more flexible than the canonical model. In fact, the two models are identical if we replace the continuum of tasks in equation (5) with only two tasks, allow for type L and H workers to each have comparative advantage in one of them, and fix $\sigma = 1$ to yield a Cobb-Douglas production function.

The empirical predictions of the two models are also identical if the task framework includes only two skill groups L and H . There would be only one task threshold (call it I^*), and so multiplying the two relative wage terms in equation (7) would yield $\frac{w_H}{w_L} = \left(\frac{1-I^*}{I^*}\right) \left(\frac{L}{H}\right)$. By comparison, the expression for relative wages in the canonical model is $\frac{w_H}{w_L} = \left(\frac{A_H}{A_L}\right)^{\frac{\sigma-1}{\sigma}} \left(\frac{H}{L}\right)^{-\frac{1}{\sigma}}$.

In both cases, relative wages depend on the skill bias of technological change and on relative supplies. In the task framework, skill bias endogenously determines the task boundary I^* , and changes in skill bias affect relative wages with a magnitude that depends on the structure of comparative advantage. This is substantively similar to a model where the responsiveness of relative wages to technology depends on the elasticity of substitution. The task framework nests the canonical model as a special case while allowing for more flexibility to fit the facts.

The endogenous assignment of worker skill groups to tasks helps explain sorting patterns and relative wage changes across occupations in the U.S. and other countries over the last several decades. A large literature in economics documents employment polarization in the U.S. and other OECD countries, with growth in both low- and high-skilled occupations relative to middle-skilled occupations (e.g. Acemoglu and Autor 2011, Autor and Dorn 2013, Goos et al. 2014, Akerman et al. 2015). Employment polarization has been linked to the

declining quality-adjusted price of computer capital, which directly replaces workers in “routine” physical and information processing tasks that are disproportionately represented in the middle of the wage distribution (Autor et al. 2003). At the same time, the evidence suggests that computer and information technologies (IT) complement high-skilled workers by increasing access to information (e.g. Autor et al. 2003, Beaudry et al. 2010, Michaels et al. 2014, Akerman et al. 2015).

The task framework handles these empirical trends more naturally than the canonical model by allowing for non-monotone changes in the wage structure. The computer and IT revolution can be modeled as technological progress that gives machines comparative advantage over type M workers in some tasks.⁷ This reallocates tasks of medium complexity to machines and away from type M workers, shifting their equilibrium task thresholds I_H and I_L inward and lowering their relative wages (e.g. $\frac{w_M}{w_L}$ declines and $\frac{w_H}{w_M}$ increases). In principle the impact on $\frac{w_H}{w_L}$ is ambiguous because it depends on how the “excess supply” of type M workers is reallocated. However, since computer capital also complements type H workers, an increase in the high-skill bias of technology A_H means that type M workers mostly shift into the low-skill sector, increasing $\frac{w_H}{w_L}$.

The specific channel in the model is occupational downgrading of type M workers, who now perform lower-skilled tasks. Several authors present evidence that workers with high school degrees and/or two-year college degrees have shifted out of higher-paying clerical and production positions and into the service sector, which is consistent with occupational downgrading (Acemoglu and Autor 2011, Autor and Dorn 2013, Abel et al. 2014, Altonji, Kahn and Speer 2016, Beaudry et al. 2016).

How does the task framework and the related literature on automation fit into our understanding of education, skills and the wage structure? In Acemoglu and Autor (2011), workers are divided into low, medium, and high-skilled along a single index of human capital. While

⁷Formally, add capital K as a factor to the task production function in (6) and specify a range of tasks $[I', I'']$ within $[I_L, I_H]$ for which α_K increases enough to make it more profitable to perform the tasks with capital rather than type M labor.

it is possible to add many more skill groups, the single index assumption is necessary to make the model analytically tractable (Costinot and Vogel 2010, Lise and Postel-Vinay 2020).

Acemoglu and Autor (2011) do not explicitly model technological replacement of routine tasks. Instead, changes in the wage structure since 1980 are explained as technology replacing middle-skilled tasks, which are assumed to be routine. This is still an improvement on the canonical model, which cannot easily accommodate predictions about which types of jobs are affected or how workers shift up and down the occupational ladder.

Acemoglu and Restrepo (2018) develop a task-based model where technology can be factor-augmenting but can also directly replace workers in job tasks. They measure the impact of automation technologies as a decline in the industry-level labor share and show that such labor share declines are correlated with adoption of robots and specialized software during the 1980-2016 period. They also find reductions in hourly wages for industries and demographic groups most exposed to automation.

They allow for different impacts of technology by education and other demographics, but do not model the impact of education directly. Rather, they assume that only routine tasks can be automated, and that the impact of education is fully mediated by the routineness of a worker's industry and occupation. Since highly educated workers are employed in industries and occupations that are much less routine, their model delivers the result that most of the rising return to education and the change in the wage structure between 1980 and 2016 is explained by task displacement. This leaves open the question of why educated workers sort into non-routine jobs in the first place.

Autor et al. (2022) link data on new job titles to patent applications and estimate the impact of new technologies on both margins – task replacement, but also the creation of “new work”. They find that technological innovation since 1980 has had a polarizing effect, creating new job titles both in personal services with low educational requirements, but also in well-educated and highly-paid professional occupations. New technologies clearly affect returns to education and the wage structure, but the exact mechanisms remain unclear.

4 Varieties of Human Capital

A growing body of work emphasizes the importance of “non-cognitive” or “soft” skills like patience, self-control, conscientiousness, teamwork, and critical thinking. Many studies find positive labor market returns to measures of non-cognitive skills or personality traits (e.g. Heckman et al. 2006, Lindqvist and Vestman 2011). These findings are consistent even though the measurements of skills vary so widely. Indeed, the very terms “non-cognitive” and “soft” reveal our lack of understanding of what these skills are and how to develop them. Deming (2022) calls them higher-order skills following the Bloom et al. (1956) taxonomy of educational objectives. We know that higher-order skills matter for wage determination, but we know much less about how and why these skills matter.

One approach treats human capital as a vector of skills with weights that vary across occupations and/or industries. Sanders and Taber (2012) review the literature on heterogeneous human capital and its implications for life-cycle wage growth, including the extent to which skills are fully transferable or specific to occupations, sectors and/or job tasks. Several papers use industry-specific demand shocks or earnings losses of displaced workers to understand the contribution of specific human capital to earnings (Shaw 1989, Neal 1995, Kambourov and Manovskii 2009, Pavan 2011). These papers typically find larger earnings losses upon reemployment for displaced workers who switch industries or occupations. Lachowska et al. (2020) estimate that over half of the earnings losses for displaced workers can be explained by the loss of match-specific human capital between workers and employers.

Other papers study sorting patterns and earnings changes for job switchers based on the similarity of job tasks and show that task-specific human capital is important (Gathmann and Schönberg 2010, Yamaguchi 2012, Taber and Vejlin 2020). These studies show evidence for varieties of human capital based on sorting patterns across jobs, but do not measure them directly. Other work goes further by imposing assumptions such as different skills for blue vs. white collar occupations (e.g. Willis and Rosen 1979, Keane and Wolpin 1997, Lindenlaub 2017).

Two recent papers incorporate direct measures of multiple skills into models of occupational sorting and human capital accumulation – Guvenen et al. (2020) and Lise and Postel-Vinay (2020). Lise and Postel-Vinay (2020) estimate a structural search model where workers sort themselves into jobs based on their cognitive, interpersonal, and manual skills, all of which also accumulate on-the-job. They find different patterns of sorting and dynamic returns to each type of skill. Interpersonal skills have moderate returns that are roughly fixed over a worker’s career, manual skills have lower returns but adjust quickly, while cognitive skills have the highest returns but take the longest time to accumulate (Lise and Postel-Vinay 2020). This implies an especially high cost of job mismatch for cognitive skills. They also show how a single index model of skills would deliver incorrect predictions about the sources of life-cycle wage growth. Guvenen et al. (2020) estimate a dynamic model of occupational choice and to quantify the impact of job mismatch for career earnings. They find that skill mismatch depresses current wages but also future wage growth by stunting skill accumulation, and that moving from the bottom to the top decile of match quality would increase annual earnings by 11 percent.

These papers provide important evidence that skills are multidimensional and that the match between worker skills and jobs is quantitatively important for wage determination. What are the macroeconomic implications of human capital varieties for the wage structure? Even if human capital is heterogeneous at the worker or occupation level, this variation may “average up” so that a single index model of human capital is a good enough approximation for the aggregate economy. But if workers skills differ in important ways, the single index model will be unable to explain patterns in the wage structure.

4.1 Declining Returns to Cognitive Skills and the College Wage Premium

The U.S. college wage premium grew from 40 percent to 60 percent between 1980 and 2000, and then from 60 percent to 68 percent between 2000 and 2017 (Autor et al. 2020). Between

2000 and 2020, the share of young people obtaining a four-year college degree grew from 29 percent to 39 percent, with most of the growth occurring in the last decade (Digest of Education Statistics, 2021). In the framework of the canonical model, these facts are explained by rapid relative growth in skill supplies and steady growth in skill demand. In the U.K., the share of workers with a university degree has tripled since 1993, yet the college wage premium has remained constant (Blundell et al. 2022). Overall, the evidence suggests steady growth in relative demand for college-educated labor.

A single index model of human capital would suggest that returns to college attainment and returns to cognitive skills follow the same pattern. Instead, the return to cognitive skills has declined since 2000. Castex and Kogan Dechter (2014) estimate labor market returns to both education and cognitive skill in the National Longitudinal Survey of Youth (NLSY) 1979 and 1997 samples, which allows them to compare estimates from the 1980s and 1990s to the post-2000 period. They find that a one standard deviation increase in the Armed Forces Qualifying Test (AFQT) score – a widely-used measure of cognitive skill – was associated with about 10 percent higher hourly wages in the 1980s and early 1990s but only 4.5 percent in the 2000s and early 2010s.⁸ In contrast, the economic return to a bachelor’s degree increased by 6 percentage points unconditionally and by nearly 15 percentage points after controlling directly for cognitive skills in both waves (Castex and Kogan Dechter (2014)). Their results hold for all demographic groups and are robust to measurement error, test time and other details. Using test scores and administrative earnings records for roughly half of the Swedish male population, Edin et al. (2022) show that the return to cognitive skills declined by about 25 percent between 2000 to 2013. Beaudry et al. (2016) show that the demand for cognitive skill-intensive jobs began to decline sharply around 2000.

While the return to cognitive skills has declined, the return to “non-cognitive” or “soft” skills has increased. In the same paper, Edin et al. (2022) estimate that the economic return to interviewer-rated skills like social maturity, energy, intensity and emotional stability roughly

⁸Castex and Kogan Dechter (2014) estimate separate models by gender and find that the returns decrease from 9.6% to 3.3% for males and from 10.8% to 6.2% for females.

doubled between 1992 and 2013. Using the same data as Castex and Kogan Dechter (2014), Deming (2017) finds that the labor market return to noncognitive skills increased significantly across NLSY waves.⁹ Falk et al. (2021) show that children from higher SES families score higher on measures of non-cognitive skills, and Attanasio et al. (2020) find growing inequality in these skills across two British cohorts born 30 years apart. Hermo et al. (2022) find evidence of growing demand for reasoning skills relative to crystallized knowledge.

Overall, the evidence shows growing demand for college-educated labor over a period where the return to cognitive skills and cognitive skill-intensive jobs declined. This is impossible to reconcile with a single index model of human capital. Yet an agnostic approach that infers skill demands from occupational sorting also doesn't fully address this puzzle. To make progress we need a theoretical framework that illuminates why and where certain skills matter, and that makes sharp empirical predictions that fit the data.

4.2 Team Production and Social Skills

Deming (2017) shows that the relative decline in cognitive skill-intensive employment identified by Beaudry et al. (2016) is driven by Science, Technology, Engineering and Math (STEM) jobs. STEM occupations shrank as a share of all jobs in the U.S. labor force between 2000 and 2012, after growing during the previous two decades. In contrast, non-STEM jobs that typically require college degrees such as managers, teachers, nurses, physicians, lawyers and economists all grew faster in the 2000s than in previous decades. All of these jobs require college degrees, and all of them require relatively high amounts of social interaction and teamwork.

Teamwork has become more common in part because of advances in IT and computerization, which have shifted jobs toward flexible, team-based settings that facilitate adaptation and group problem-solving (e.g. Lindbeck and Snower 2000, Caroli and Van Reenen 2001, Bartel et al. 2007). Team production becomes more desirable as the complexity of work

⁹Deming (2017) separates noncognitive skills into intrapersonal and interpersonal skills, where the former is measured by self-esteem and self-control, and the latter is measured by prosociality and extraversion.

increases, because a well-functioning team can operate more efficiently than individuals by exploiting comparative advantage between team members (Becker and Murphy 1992, Bolton and Dewatripont 1994, Garicano and Rossi-Hansberg 2006).

If we take seriously the idea that workers engage in team production, we must also consider the value of being able to work in a team. Deming (2017) develops a theoretical model of team production where social skills reduce the coordination cost of “trading tasks” between workers, allowing them to better exploit comparative advantage. The model starts by specifying the worker’s production function using the task framework of Acemoglu and Autor (2011):

$$y_j(i) = A_j \alpha_j(i) l_j(i) \tag{8}$$

Worker j ’s production function for task i is equal to the worker’s cognitive skill A_j times a task-specific productivity parameter $\alpha_j(i)$ times labor supplied to task i , $l_j(i)$. Workers supply a single unit of labor inelastically to the production of a continuum of tasks indexed over the unit interval with a Cobb-Douglas technology:

$$Y_j = \exp \left[\int_0^1 \ln y_j(i) di \right]; \int_0^1 l_j(i) di = L_j = 1 \tag{9}$$

Equation (8) shows that workers with the same cognitive skill A_j can vary in their productivity over individual tasks, which allows for the possibility of comparative advantage. Workers can increase their total output Y_j by producing tasks in which they have a comparative advantage and trading them with other workers for mutual benefit, just as countries trade goods. This grounds the value of social skills in economic theory.

Deming (2017) models social skills as inverse “iceberg” trade costs as in Dornbusch et al. (1977) and Eaton and Kortum (2002). Let $S_{j,k} \in (0, 1)$ be a depreciation factor that is applied to any task trade between workers - $S_{j,k} = S_j * S_k$ for $j \neq k$ and let $S_{j,j} = 1$ so that self-trade is frictionless. Workers with higher social skills pay a lower coordination cost to trade with

others, allowing them to earn higher wages by specializing in their most productive tasks and “trading” for the rest. Workers with high cognitive skill but low social skill have high average productivity but will perform too many tasks themselves. In this sense, we can think of social skills as lowering social “gravity” and expanding opportunities to work productively with others.

To see this, consider the simple case of bilateral task trade in a competitive market where labor is the only factor of production. Identical firms hire pairs of workers and pay market wages equal to output Y_j times an exogenous output price P^* . Workers maximize output subject to their labor supply constraint, and firms maximize total revenue $P^* * (Y_1 + Y_2)$ minus wages $(w_1 + w_2)$.

Define the comparative advantage schedule over tasks as:

$$\gamma(i) = \frac{A_1 \alpha_1(i)}{A_2 \alpha_2(i)} \tag{10}$$

and index the continuum of tasks in (9) in order of decreasing comparative advantage for worker 1, so that $\gamma'(i) < 0$ by construction. This is very similar to the comparative advantage structure in Acemoglu and Autor (2011). Deming (2017) assumes that $\gamma(i) = \frac{A_1}{A_2} \exp[\theta(1 - 2i)]$, although this specific functional form is not necessary to derive the empirical predictions of the model.

The parameter θ indexes the variance of task productivity and thus the steepness of the comparative advantage schedule $\gamma(i)$. $\theta = 0$ represents the limiting case where workers with higher cognitive skill are more productive in all tasks, as in the standard human capital model.

Deming (2017) shows that the costless equilibrium in this model (e.g. where $S_{j,k} = 1, \forall j, k$) is similar to Acemoglu and Autor (2011) in the case of two worker types. Workers trade tasks with each other at “prices” defined by efficiency units of labor and a budget constraint equal to total labor supply. The price of task i is determined by the equilibrium wage paid to worker j for a unit of labor divided by the worker’s productivity in that task:

$$p_j(i) = \frac{w_j}{A_j \alpha_j(i)} \quad (11)$$

Since tasks will be performed by the worker with the lowest “price” in equilibrium and since $\gamma^i(i) < 0$, there will be a marginal task i^* in which the workers are equally productive. Worker 1 performs the tasks in the interval $(0, i^*)$ and worker 2 performs the tasks in the interval $(i^*, 1)$. Relative wages ω are determined by each worker’s task share, $\omega = \frac{i^*}{1-i^*}$.¹⁰

Equilibrium wages are increasing in the worker’s own cognitive skill, the cognitive skill of her co-worker, and the variance of task productivity θ .

The equilibrium with social skills involves two task thresholds i^L and i^H , with an “un-traded” zone of tasks where coordination costs outweighs the benefits of comparative advantage. Define $S^* = S_1 * S_2$ as the symmetric cost of trading tasks between workers 1 and 2, with self-trade normalized to 1 as above. Worker 1 will produce her own tasks rather than trading if the “price” of doing so is lower, or from equation (11):

$$\frac{w_1}{A_1 \alpha_1(i)} < \frac{w_2}{S^* A_2 \alpha_2(i)} \quad (12)$$

Equation (12) expressed in terms of relative wages and the comparative advantage schedule is just $\omega < \frac{\gamma(i)}{S^*}$. By similar logic, worker 2 will produce her own tasks if $\omega > S^* \gamma(i)$.

In equilibrium there will be two task thresholds, $\gamma(i^L) = \frac{\omega}{S^*}$ and $\gamma(i^H) = S^* \omega$, and because $\gamma'(i) < 0$, $i^H > i^* > i^L$ as long as $S^* < 1$. Tasks in the interval $(0, i^L)$ will be produced by worker 1, tasks in the interval $(i^H, 1)$ will be produced by worker 2, and tasks in the interval (i^L, i^H) will be “nontraded”. As $S^* \rightarrow 1$, the two task thresholds converge to a single i^* . Solving the two task thresholds together for ω yields an expression for the range of the nontraded task interval:

$$i^H - i^L = -\frac{\ln S^*}{\theta} \quad (13)$$

¹⁰The marginal task is $i^* = \frac{A_1}{A_1 + A_2 \exp[\theta(2i^* - 1)]}$.

Equation (13) shows that the size of the nontraded zone $i^H - i^L$ is decreasing in the variance of task productivities θ and (inversely) scales the gains from trade. Equation (13) also shows that there are many values of S^* and θ for which $i^H - i^L > 1$, meaning it is optimal for all workers to produce their own tasks (i.e. autarky). Figure 3 reproduces Figure II of Deming (2017), which shows how the equilibrium task thresholds and the nontraded zone vary with θ .

There are two interpretations of the variance parameter θ . First, it could measure the routineness of an occupation. Autor et al. (2003) define a task as routine if it can be accomplished by following explicitly programmed rules, which implies that there is an objectively correct approach. Routine jobs have lower task variance and thus less scope for comparative advantage. Second, θ could be an economy-wide technology parameter that has increased over time, following the literature on declining demand for routine tasks (Autor et al. 2003, Acemoglu and Autor 2011). In either case, the empirical prediction is that the return to social skills is increasing in θ .

Using panel data from the NLSY79 and NLSY97, Deming (2017) shows that workers with higher social skill sort into nonroutine occupations and that the wage gains from switching are increasing in social skills. I also find that the economic return to social skills – conditional on cognitive skill and other covariates - is higher in the 2000s than in the 1980s.¹¹

Another key prediction from the model is that cognitive skill and social skill are complements in a Mincerian earnings regression. In the model this is because reducing coordination costs is more valuable when workers have more of value to “trade”. This prediction contrasts with many other multidimensional assignment models, where the separability of skills is assumed for tractability (e.g. Lindenlaub 2017, Lise and Postel-Vinay 2020). Deming (2017) finds a positive and statistically significant interaction between cognitive skills and social skills in both NLSY waves, and notably does not find the same complementarity between

¹¹Deming (2017) shows that the return to intrapersonal non-cognitive skills is higher in the NLSY97 than the NLSY79, which is consistent with the increasing return to non-cognitive skills showed by Edin et al. (2022) and others.

cognitive skill and other widely used measures of non-cognitive skills such as those used in Heckman et al. (2006).

What are the implications of the growing importance of social skills for the wage structure? Social skill-intensive occupations grew by nearly 12 percentage points as a share of all jobs in the U.S. economy between 1980 and 2012, and real hourly wages for these jobs grew around 25 percent compared to less than 10 percent for other occupations (Deming 2017). This suggests growing relative demand for social skill, and flat or declining demand for cognitive skill. However, because these skills are complements, the jobs with the most employment and earnings growth are those where both types of skill are required.

A few other papers consider the economics of team production. Jarosch et al. (2021) and Herkenhoff et al. (2018) develop models where teams of workers learn from each other, and they show that having highly paid coworkers increases future wage growth because of human capital spillovers. Weidmann and Deming (2021) develop an experimental method to identify individual contributions to group performance, while Bonhomme (2021) proposes an econometric framework for identifying individual contributions to teams in observational data. Finally, Jäger and Heining (2022) show that the unexpected death of workers in high-skilled occupations has a negative impact on other workers in the same firm, suggesting that workers are complements in production.

Taken together, this evidence suggests that cognitive skills and social skills are conceptually distinct and that they work together in non-obvious ways to explain an important recent trend in the wage structure – rising returns to education and social skills, but declining returns to cognitive skills.

5 Human Capital Vintages

Another deviation from the single index view of human capital is the possibility of age- or cohort-based variation in skill premia. Does a college degree earned two decades ago have

the same value as a degree earned last year, or are there vintage effects in human capital?

Card and Lemieux (2001) show that between 1959 and 1995 the college wage premium was relatively constant for older men but started to grow rapidly for younger men in the 1980s. They find this same pattern in the U.S., the U.K. and Canada. Card and Lemieux (2001) argue that relative wage growth among young college graduates is driven by slower growth in the supply of skills combined with imperfect substitutability between younger and older workers. In other words, deceleration in college attainment created relative scarcity in the supply of young college graduates. Thus the “twisting” of the college premium by age could be evidence of vintage effects in human capital.

Card and Lemieux (2001) develop an extension of the canonical model with imperfect substitution by age group. They write down a multi-level CES production function where the upper level is expressed in terms of skill supplies for type L and type H workers at time t :

$$\begin{aligned}
 L_t &= \left[\sum_{j=1}^J (\alpha_j L_{jt}^\eta) \right]^{\frac{1}{\eta}} \\
 H_t &= \left[\sum_{j=1}^J (\beta_j H_{jt}^\eta) \right]^{\frac{1}{\eta}}
 \end{aligned} \tag{14}$$

where $\sigma_A = \frac{1}{1-\eta}$ is the elasticity of substitution across j age groups and α_j and β_j are technology parameters that are constant by age group. If $\eta = 1$, age groups are perfect substitutes and the setup is isomorphic to the canonical model.¹²

If the aggregate production function does not depend on age groups once they are “added up” in equation (14), we can employ the usual assumptions that define a competitive equilibrium and write cohort-specific relative wages as:

¹²The only difference is the cohort-specific technology parameters α_j and β_j but in the empirical work these get absorbed as part of the L_j 's and H_j 's.

$$\ln \omega_{jt} = \ln \left(\frac{A_t^H}{A_t^L} \right) - \frac{1}{\sigma} \ln \left(\frac{H_t}{L_t} \right) + \ln \left(\frac{\alpha_j}{\beta_j} \right) - \frac{1}{\sigma_A} \left[\ln \left(\frac{H_{jt}}{L_{jt}} \right) - \ln \left(\frac{H_t}{L_t} \right) \right] \quad (15)$$

with $\sigma = \frac{1}{1-\eta}$ as the aggregate elasticity of substitution.

The first two terms after the equal sign are familiar from equation (2) and the canonical model, where relative wages depend on skill bias ($\frac{A_t^H}{A_t^L}$) and aggregate relative skill supplies ($\frac{H_t}{L_t}$). The third term adds age-specific technology parameters, and the last term is the gap between cohort-specific and aggregate relative supply. Cohort effects will be empirically unimportant if cohorts are perfect substitutes ($\eta = 1$) or if the effects are the same size for all cohorts, in which case the two terms in the last expression - $\ln \left(\frac{H_{jt}}{L_{jt}} \right)$ and $\ln \left(\frac{H_t}{L_t} \right)$ exactly cancel out.

Card and Lemieux (2001) estimate the model by restricting cohort effects to be the same for the ten oldest cohorts in the data, which solves the age-time-cohort identification problem. This is perhaps justified by the data, which show a flat college wage premia for older cohorts. They find an elasticity of substitution between age groups σ_A of between 4 and 6 and an aggregate elasticity σ of between 1.1 and 1.6, which is very similar to Katz and Murphy (1992).

Card and Lemieux (2001) interpret “twisting” in the college wage premium by age as evidence of cohort supply constraints. However, these patterns might also reflect cohort variation in the quality of education. Without an instrument or additional assumptions, it is not possible to distinguish these two explanations. Carneiro and Lee (2011) show that the marginal college attendee is less academically prepared than the average attendee and then “quality adjust” measures of skill supply for the 1960-2000. Their interpretation of the rising college premium for young college graduates is based on shifting quantities rather than prices – because fewer young people graduated from college in the early 1980s, the composition of college graduates is relatively more skilled, which drives up the price of their skill.

Heckman et al. (1998) derive from the Ben-Porath model of human capital investment an age range near retirement where workers optimally supply the same number of efficiency

units of labor, which has become known as the “flat spot” approach. In this range, human capital supplies are assumed to be constant, and so cohort variation in wages represents human capital prices rather than quantities. Bowlus and Robinson (2012) use the flat spot approach to identify price series for different levels of human capital. They find substantial variation over time in absolute prices for all levels of schooling, but little variation in relative prices since the series mostly move together. This suggests that the increase in the college wage premium starting in the 1980s was driven by growth in the “quality” of labor supplied by college graduates, similar to Carneiro and Lee (2011).

Bowlus et al. (2021) measure skill prices using the flat spot method and then estimate the canonical model in more recent data, which allows them to adjust for cohort differences in the relative quality of college graduates. They find substantial improvement over time in the relative quality of college-educated workers. Applying this adjustment to the Katz and Murphy (1992) framework for recent data yields much larger estimates of the elasticity of substitution σ and a lesser role for skill bias $\frac{A_H}{A_L}$.

However, just like Katz and Murphy (1992) and others, Bowlus et al. (2021) cannot separately identify σ from the time trend in $\frac{A_H}{A_L}$ without additional assumptions. To make progress, they model the time trend in $\frac{A_H}{A_L}$ as a function of private investment in IT, effectively using direct measures of expenditure on technology to model skill-biased technological change (Beaudry et al. 2016). They also allow for permanent trend changes in skill bias after recessions, following the literature on technology upgrading during economic downturns (e.g. Hershbein and Kahn 2018).

Bowlus et al. (2021) show that using direct measures of technological change on the demand side can improve the predictions of the canonical model. What about the supply side? If fluctuations in the college wage premium are driven by investments in IT and other technology, might we expect bigger “vintage” effects in human capital for workers with technical skills?

5.1 Vintage Human Capital and Technical Skills

College teaches a mix of general and specific skills. The average college wage premium conceals large variation in returns by field of study (e.g., Arcidiacono 2004, Altonji et al. 2012, Altonji, Arcidiacono and Maurel 2016, Kirkeboen et al. 2016). Heterogeneous returns to field of study are quantitatively important. Lemieux (2014) estimates that field of study and occupation matching can explain about half of the total return to a college degree. Altonji, Kahn and Speer (2016) show that growing earnings differences by college major are explained by changing returns to job tasks, and that high-paying majors were relatively more insulated from earnings losses following the Great Recession.

College majors that teach technical skills have particularly high returns, mostly because they are gateways to high-paying jobs. Deming and Noray (2020) show that the wage premium for engineering and computer science majors is very high when they work in computer science and engineering jobs, but modest otherwise. Kinsler and Pavan (2015) develop a structural model with major-specific human capital and show that science majors working in science-related jobs earn about 30% more than science majors in unrelated jobs, even after controlling for SAT scores, high school GPA, and worker fixed effects. Leighton and Speer (2017) develop a measure of human capital specificity based on the concentration of returns to college majors across occupations. They find that specific majors such as engineering and nursing have high early career payoffs relative to more general majors such as Psychology and Philosophy.

In other words, the wage premium for technical fields of study derives mostly from working in technologically intensive high-paying fields, rather than selection on ability or some other explanation.

Technical skills are in high demand because they are always changing, and thus always scarce. Deming and Noray (2020) measure changes in the skill content of occupations using job vacancy data and show that technology-intensive jobs change especially rapidly. In their model, returns to work experience are a race between gains from on-the-job learning and

losses from skill obsolescence. They then show that applied majors such as computer science and engineering have initially very high returns but slower earnings growth due to obsolescence. This is intuitive - computer science majors today learn a very different curriculum than computer science majors twenty years ago, but economics or history curricula have changed much less. It suggests that vintage effects in human capital are likely to be much stronger in technology-intensive jobs and fields of study.

Deming and Noray (2020) provide empirical evidence for an older line of work in economics on vintage human capital and skill obsolescence (Rosen 1975). Weiss and Lillard (1978) compare scientists with similar levels of work experience and find greater earnings growth for graduates of more recent vintages. Neuman and Weiss (1995) infer skill obsolescence from the shape of wage profiles in “high-tech” fields, and Thompson (2003) studies how new technology changed age-earnings profiles in the Canadian Merchant Marine in the late nineteenth century. MacDonald and Weisbach (2004) use the example of earnings losses among older architects after the development of drafting software to develop a model of “has beens”, where skill obsolescence among older workers is increasing in the pace of technological change. In Chari and Hopenhayn (1991) and Kredler (2014), new technologies require vintage-specific skills, and an increase in the rate of technological change raises the returns for newer vintages and flattens the age-earnings profile.

One implication of vintage capital for technical skills is that there are individual and societal tradeoffs to investing in general versus specific skills. In Gould et al. (2001), workers make precautionary investments in general education to insure against obsolescence of technology-specific skills. Krueger and Kumar (2004) show that an increase in the rate of technological change increases the optimal subsidy for general versus vocational education, because general education facilitates the learning of new technologies. Hanushek et al. (2017) find that workers with vocational education have lower youth unemployment rates at labor market entry, but higher rates later in life.

Vintage human capital may also help explain the puzzle of rapid relative wage growth for

young college graduates in the early 1980s identified by Card and Lemieux (2001). Deming and Noray (2020) use vacancy data collected from classified ads to show that job skill change was particularly rapid for technology-intensive jobs in the late 1970s and early 1980s. They also show that the wage growth for young college graduates in Card and Lemieux (2001) was concentrated among those working in STEM jobs. Several related papers show that delayed or imperfect computer adoption by older workers can explain slower wage growth and early retirement during this period (Friedberg 2003, Weinberg 2004, Hudomiet and Willis 2022). Barth et al. (2022) estimate the impact of firm-level software adoption on wages and find that it has a relatively greater impact on younger workers.

Overall, there is strong evidence that imperfect substitution among human capital vintages can explain important features of the wage structure. Yet few studies make this connection directly. We need more research that directly measures how technology affects work and changes the returns to specific skills.

6 Valuation of Human Capital – Rent-Sharing and the Law of One Price

The canonical model and the task framework assume that labor markets are perfectly competitive. Since firms take market wages as given, workers are paid only based on their skills and there are no firm-specific pay premia. Yet we know that productivity varies tremendously across firms (Syverson 2011). Do firms have wage-setting power, or are productive firms simply those that hire productive workers? Do labor markets follow the law of one price for skill?

Card et al. (2018) review the literature on imperfect competition and the impact of firm-specific pay premia on inequality and the wage structure. They present evidence that firm productivity shocks do not fully pass through into worker wages, which suggests that firms have some wage-setting power. They also show using matched employer-employee data

that time-invariant firm fixed effects can explain about 20 percent of the variance of wages. To understand these facts, they develop a model of imperfect competition where workers have heterogeneous preferences for non-wage amenities such as location, commute time, and workplace culture. Amenity variation gives firms some power to set wages, because they are imperfect substitutes from the worker’s perspective.

Several studies decompose the sources of earnings inequality using large administrative datasets with longitudinal matched employer-employee data. Card et al. (2013) study rising wage inequality in West Germany from 1985 to 2009 and find that rising worker and firm-specific pay premia play an important role, as does the growing sorting and segregation of high-paying workers to high-paying firms. Song et al. (2019) find that about two-thirds of the growth in U.S. earnings inequality between 1978 and 2013 is explained by rising worker effects, which could plausibly be interpreted as rising returns to skill. The other third is explained by an increasing covariance between worker effects and firm effects, and almost none is explained by growth in firm effects.

Lamadon et al. (2022) build an equilibrium model of imperfect competition using U.S. tax data for the 2001-2015 period and use it to estimate worker and firm rents. They find that the average worker is willing to pay 13 percent of their wages to stay in their current job and that rents are shared roughly equally between firms and workers. However, despite substantial rent-sharing, firm effects are small because productive firms also have good nonwage amenities, which diminishes the impact of firm-specific pay premia on earnings inequality. They find that worker effects explain about 72 percent of earnings variation, compared to only 4 percent for firm effects. They also find strong evidence that high-skilled workers sort into high-paying firms due to production complementarities, and they estimate that worker-firm sorting explains about three times as much of earnings variation as firm effects alone.

The high degree of labor market sorting and worker-firm rent-sharing in the U.S. economy suggests that some firms may have substantial wage-setting power, and thus that the “price”

of human capital may differ substantially across workers and firms.

How does imperfect labor market competition affect returns to skill and the wage structure? With a single-index model of human capital, the distinction is mostly a matter of degree. The college wage premium could in theory be decomposed into 1) higher marginal product, or the “true” return to skill; and 2) gains from sorting into higher-wage firms. Haanwinckel and Soares (2021) estimate a task-based equilibrium model of the Brazilian labor market that allows for imperfect competition and rent-sharing and maintains the single index comparative advantage assumption in Acemoglu and Autor (2011). They find that increased educational attainment reduces the college wage premium as predicted by the canonical model, but also reduces the amount of labor market sorting. Skill-biased technological change has the opposite effect.

In essence, imperfect competition magnifies equilibrium wage differentials between workers with different levels of skill through the sorting channel. But it does not qualitatively change the predictions about skill premia and the wage structure made by standard models.

However, the conclusions can be quite different when human capital is multidimensional. Imperfect labor market competition may be especially important for workers with firm-specific human capital or hard-to-replace skills. An important piece of evidence comes from Kline et al. (2019), who study what happens to wages when firms are granted allowances for commercially valuable patents. They show that being granted a valuable patent increases firm productivity, and they estimate the pass-through of this positive productivity shock to worker wages. They find that workers capture roughly 30 cents of every dollar of patent-induced surplus in the form of higher wages. Interestingly, they find much greater pass-through for high-earning and longer-tenured workers, and almost none for new hires or for workers who subsequently leave the firm. Using data on unexpected worker deaths, Jäger and Heining (2022) find that workers within a firm are much closer substitutes for each other than new hires, and that workers with longer firm tenure and in specialized occupations are hardest to replace.

The evidence suggests that firms share rents with workers who have specific human capital to prevent them from quitting. Bloesch et al. (2022) extend this logic to occupations, arguing that the structure of firm production gives hold-up power to workers in occupations that are critical to the production process. They develop a generalization of the Kremer (1993) O-ring production function where the combination of task complementarities and specific skills gives workers hold-up power over inframarginal rents. Following Jäger and Heining (2022), they show that productivity losses from worker deaths are particularly high for managerial, professional and technical occupations. They then show using job vacancy data that these same occupations are more differentiated in terms of skill requirements from other jobs within a firm, suggesting higher rates of hold-up power.

Bloesch et al. (2022) provide interesting new evidence on how imperfect competition in labor markets may have differential impacts across occupations. This is potentially important, because imperfect competition may matter for wage inequality if some labor markets are much less competitive than others. Deming (2021) presents evidence that labor market returns to experience vary greatly by occupation, which is consistent with the idea that some occupations allow for the development of specific human capital and thus increase worker hold-up power.

Since occupations are ultimately just bundles of job tasks, we can extend this logic further to think about the importance of skill bundling for labor market competition. In a single index model of human capital, there is only one skill and thus no bundling. However, workers may have more market power if they possess unusual and important bundles of skills. Edmond and Mongey (2019) show that over the last two decades, low-skill occupations have become more alike while high-skill occupations have become more differentiated. They use these motivating facts to develop a model where workers supply individual bundles of skills and the rents earned by workers decrease as occupations become more similar. In the model, technological replacement of tasks causes occupations to unbundle, dissipating rents and decreasing within-occupation inequality. Choné et al. (2021) develop a similar model to

study skill bundling and wage inequality.

While this work is in its infancy, it is likely to yield important insights about the role of human capital in driving wage inequality when labor markets are imperfectly competitive. Systematic variation in rent-sharing elasticities across firms and labor markets is hard to justify with a single index view of human capital, but easier to explain if workers possess indivisible skill “bundles” and firm-specific work experience.

7 Conclusion

This paper reviews and synthesizes research on the macroeconomic implications of human capital theory. Macroeconomic test of human capital theory go beyond individual returns to education, focusing instead on the wage structure of an economy. I start with the “canonical” model of wage determination, also called the supply-demand-institutions (SDI) framework (Tinbergen 1975, Katz and Murphy 1992, Goldin and Katz 2007).

The canonical model has two key features. First, low-skilled and high-skilled labor are imperfect substitutes in production, and thus a shortage or a surplus of one labor type will affect the college wage premium even if technology is held constant. Second, technology takes a factor-augmenting form, meaning it always makes workers more productive. Skill-biased technological change (SBTC) occurs when technological change complements high-skilled labor relatively more than low-skilled labor.

The canonical model does a remarkably good job of explaining trends in the U.S. wage structure going back more than a century. It also does a good job of explaining the evolution of upper-tail (90/50) wage inequality across developed countries. The college wage premium rose rapidly in the U.S. and many other countries beginning in the early 1980s. This growth continued but decelerated significantly in the 1990s and early 2000s. The canonical model provides a simple and appealing framework within which to test different hypotheses about the slowdown in the college wage premium. For example, growth in the supply of educated

labor accelerated in the 2000s, which can explain the flattening of the college wage premium even if demand was constant (Autor 2017). Another possibility is that the long-run elasticity of substitution between skill groups is higher than the short-run elasticity because firms adjusted over time to the digital technology “shock” of the computer age (Bils et al. 2022).

One limitation of the canonical model is that technology can be skill-biased but cannot replace workers directly, nor does it allow for some types of jobs and tasks to be more affected by technological change than others. A recent innovation in this literature is the task framework, which takes tasks (rather than workers) as the main input into production and allows for workers of different skill levels to perform different tasks depending on the economic environment (Acemoglu and Autor 2011). The task framework also allows for an analysis of automation technologies that replace labor directly. Acemoglu and Restrepo (2018) find that reduction in employment and wages due to automation technology can explain a significant share of the rise in U.S. wage inequality between 1980 and 2016. As they show, this is because routine tasks are more exposed to automation risk, and highly educated workers are much more likely to be employed in non-routine industries and occupations.

Compared to the canonical model, the task framework is a more realistic and flexible framework for understanding changes in the wage structure over the last half-century. Nonetheless, it leaves some key questions unanswered. Why does education so strongly determine a worker’s industry and occupation, leaving them less exposed to automation risk? In Acemoglu and Autor (2011), workers are divided into low, medium, and high-skilled along a single index of human capital, and higher skilled workers have a comparative advantage in tasks of greater complexity. Is this single index model of human capital supported by the data?

I argue that existing models are inadequate for understanding recent trends in the wage structure. In particular, I discuss evidence that labor market returns to cognitive skills have been flat or even declining since 2000, while the returns to various measures of “non-cognitive” skills have been increasing. In my view, this evidence requires economists who

want to understand the importance of skills for the wage structure to move beyond a single index view of human capital, toward richer, multi-dimensional frameworks.

At least three features of the wage structure are best understood as arising from the multidimensionality of human capital. First, I discuss growing evidence of the importance of teamwork in the labor market, and I show that social skills are a different variety of human capital, distinct from and complementary to cognitive skills (Deming 2017). Second, I discuss evidence for vintage effects in technical skills, which help explain life-cycle returns to college majors and cohort effects more generally (Card and Lemieux 2001, Deming and Noray 2020). Third, I discuss recent evidence that the valuation of human capital depends on the extent of labor market competition. In particular, there is an emerging body of evidence that highly-skilled workers and workers with unusual combinations of skills may be able to extract greater rents from firms when labor markets are imperfectly competitive (Edmond and Mongey 2019, Kline et al. 2019, Bloesch et al. 2022).

We know that education and skills are a primary determinant of labor earnings, and that variation in the supply and demand for skills play an important role in determining the wage structure. Yet as an education economist, I am ashamed to admit that we mostly still don't know why. Yet I am not discouraged. Rather, I am hopeful that this last part of the paper will suffer from technological obsolescence. With any luck, scholars of younger vintage armed with new techniques and ideas will answer the important outstanding questions raised in this paper, and future scholars will have a much richer understanding of the impact of education and skills on the wage structure.

References

- Abel, J. R., Deitz, R. and Su, Y.: 2014, Are recent college graduates finding good jobs?, *Current issues in economics and finance* **20**(1). 3
- Acemoglu, D. and Autor, D.: 2011, Skills, tasks and technologies: Implications for employ-

- ment and earnings, *Handbook of labor economics*, Vol. 4, Elsevier, pp. 1043–1171. (document), 1, 2, 3, 3, 3, 4.2, 4.2, 4.2, 6, 7
- Acemoglu, D. and Restrepo, P.: 2018, The race between man and machine: Implications of technology for growth, factor shares, and employment, *American Economic Review* **108**(6), 1488–1542. 3, 7
- Akerman, A., Gaarder, I. and Mogstad, M.: 2015, The skill complementarity of broadband internet, *The Quarterly Journal of Economics* **130**(4), 1781–1824. 2, 3
- Altonji, J. G., Arcidiacono, P. and Maurel, A.: 2016, The analysis of field choice in college and graduate school: Determinants and wage effects, *Handbook of the Economics of Education*, Vol. 5, Elsevier, pp. 305–396. 5.1
- Altonji, J. G., Blom, E. and Meghir, C.: 2012, Heterogeneity in human capital investments: High school curriculum, college major, and careers, *Annu. Rev. Econ.* **4**(1), 185–223. 5.1
- Altonji, J. G., Kahn, L. B. and Speer, J. D.: 2016, Cashier or consultant? entry labor market conditions, field of study, and career success, *Journal of Labor Economics* **34**(S1), S361–S401. 3, 5.1
- Angrist, J. D. and Pischke, J.-S.: 2010, The credibility revolution in empirical economics: How better research design is taking the con out of econometrics, *Journal of economic perspectives* **24**(2), 3–30. 1
- Arcidiacono, P.: 2004, Ability sorting and the returns to college major, *Journal of Econometrics* **121**(1-2), 343–375. 5.1
- Attanasio, O., Blundell, R., Conti, G. and Mason, G.: 2020, Inequality in socio-emotional skills: A cross-cohort comparison, *Journal of Public Economics* **191**, 104171. 4.1
- Autor, D.: 2017, How long has this been going on? a discussion of ‘recent flattening in the higher/education wage premium: Polarization, skill downgrading, or both?’ by robert g.

- valletta, *NBER Conference on Research and Income in Wealth “Measuring and Accounting for Innovation in the 21st Century,”* Washington, DC (March 10–11, 2017). 2, 7
- Autor, D., Chin, C., Salomons, A. M. and Seegmiller, B.: 2022, New frontiers: The origins and content of new work, 1940–2018, *Technical report*, National Bureau of Economic Research. 3
- Autor, D. and Dorn, D.: 2013, The growth of low-skill service jobs and the polarization of the us labor market, *American Economic Review* **103**(5), 1553–97. 3
- Autor, D., Goldin, C. and Katz, L. F.: 2020, Extending the race between education and technology, *AEA Papers and Proceedings*, Vol. 110, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 347–351. 1, 2, 4.1
- Autor, D. H., Katz, L. F. and Krueger, A. B.: 1998, Computing inequality: have computers changed the labor market?, *The Quarterly Journal of Economics* **113**(4), 1169–1213. 1, 2
- Autor, D. H., Levy, F. and Murnane, R. J.: 2003, The skill content of recent technological change: An empirical exploration, *The Quarterly journal of economics* **118**(4), 1279–1333. 2, 3, 4.2
- Bartel, A., Ichniowski, C. and Shaw, K.: 2007, How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills, *The quarterly journal of Economics* **122**(4), 1721–1758. 4.2
- Barth, E., Davis, J. C., Freeman, R. B. and McElheran, K.: 2022, Twisting the demand curve: Digitalization and the older workforce, *Journal of Econometrics* . 5.1
- Beaudry, P., Doms, M. and Lewis, E.: 2010, Should the personal computer be considered a technological revolution? evidence from us metropolitan areas, *Journal of political Economy* **118**(5), 988–1036. 3

- Beaudry, P., Green, D. A. and Sand, B. M.: 2016, The great reversal in the demand for skill and cognitive tasks, *Journal of Labor Economics* **34**(S1), S199–S247. 2, 3, 4.1, 4.2, 5
- Becker, G. S.: 1962, Investment in human capital: A theoretical analysis, *Journal of political economy* **70**(5, Part 2), 9–49. 1
- Becker, G. S. and Murphy, K. M.: 1992, The division of labor, coordination costs, and knowledge, *The Quarterly journal of economics* **107**(4), 1137–1160. 4.2
- Berman, E., Bound, J. and Machin, S.: 1998, Implications of skill-biased technological change: international evidence, *The quarterly journal of economics* **113**(4), 1245–1279. 2
- Bils, M., Kaymak, B. and Wu, K.-J.: 2022, Labor substitutability among schooling groups, *Technical report*, National Bureau of Economic Research. 2, 4, 7
- Bils, M. and Klenow, P. J.: 2000, Does schooling cause growth?, *American economic review* **90**(5), 1160–1183. 2
- Bloesch, J., Larsen, B. and Taska, B.: 2022, Which workers earn more at productive firms? position specific skills and individual worker hold-up power, *Position Specific Skills and Individual Worker Hold-up Power (March 1, 2022)*. 6, 7
- Bloom, B. S., Engelhart, M. D., Furst, E. J., Hill, W. H., Krathwohl, D. R. et al.: 1956, Taxonomy of educational objectives: the classification of educational goals: handbook i: cognitive domain, *Technical report*, New York, US: D. Mckay. 4
- Blundell, R., Green, D. A. and Jin, W.: 2022, The uk as a technological follower: Higher education expansion and the college wage premium, *The Review of Economic Studies* **89**(1), 142–180. 4.1
- Bolton, P. and Dewatripont, M.: 1994, The firm as a communication network, *The Quarterly Journal of Economics* **109**(4), 809–839. 4.2

- Bonhomme, S.: 2021, Teams: Heterogeneity, sorting, and complementarity, *arXiv preprint arXiv:2102.01802* . 4.2
- Bound, J. and Turner, S.: 2007, Cohort crowding: How resources affect collegiate attainment, *Journal of public Economics* **91**(5-6), 877–899. 2
- Bowlus, A. J. and Robinson, C.: 2012, Human capital prices, productivity, and growth, *American Economic Review* **102**(7), 3483–3515. 5
- Bowlus, A., Lochner, L., Robinson, C. and Suleymanoglu, E.: 2021, Wages, skills, and skill-biased technical change: The canonical model revisited, *Journal of Human Resources* pp. 0617–8889R1. 2, 5
- Broecke, S., Quintini, G. and Vandeweyer, M.: 2018, Wage inequality and cognitive skills: Reopening the debate, *Education, skills, and technical change: Implications for future US GDP growth*, University of Chicago Press, pp. 251–286. 2
- Card, D., Cardoso, A. R., Heining, J. and Kline, P.: 2018, Firms and labor market inequality: Evidence and some theory, *Journal of Labor Economics* **36**(S1), S13–S70. 6
- Card, D., Heining, J. and Kline, P.: 2013, Workplace heterogeneity and the rise of west german wage inequality, *The Quarterly journal of economics* **128**(3), 967–1015. 6
- Card, D. and Lemieux, T.: 2001, Can falling supply explain the rising return to college for younger men? a cohort-based analysis, *The Quarterly Journal of Economics* **116**(2), 705–746. 1, 2, 5, 5, 5.1, 7
- 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. 2, 5
- Caroli, E. and Van Reenen, J.: 2001, Skill-biased organizational change? evidence from a panel of british and french establishments, *The Quarterly Journal of Economics* **116**(4), 1449–1492. 4.2

- Caselli, F.: 2005, Accounting for cross-country income differences, *Handbook of economic growth* **1**, 679–741. 2
- Castex, G. and Kogan Dechter, E.: 2014, The changing roles of education and ability in wage determination, *Journal of Labor Economics* **32**(4), 685–710. 1, 4.1, 8
- Chari, V. V. and Hopenhayn, H.: 1991, Vintage human capital, growth, and the diffusion of new technology, *Journal of political Economy* **99**(6), 1142–1165. 5.1
- Chen, C.-L. and Chen, L.-C.: 2021, The impact of higher education expansion on the educational wage premium in taiwan: 1985 to 2015, *International Journal of Financial Studies* **9**(3), 38. 2
- Choné, P., Kramarz, F. et al.: 2021, Matching workers’ skills and firms’ technologies: From bundling to unbundling, *Technical report*, Center for Research in Economics and Statistics. 6
- Costinot, A. and Vogel, J.: 2010, Matching and inequality in the world economy, *Journal of Political Economy* **118**(4), 747–786. 3
- Deming, D. J.: 2017, The growing importance of social skills in the labor market, *The Quarterly Journal of Economics* **132**(4), 1593–1640. 1, 4.1, 4.2, 9, 4.2, 4.2, 4.2, 11, 7
- Deming, D. J.: 2021, The growing importance of decision-making on the job, *Technical report*, National Bureau of Economic Research. 6
- Deming, D. J.: 2022, Four facts about human capital, *Journal of Economic Perspectives* **36**(3), 75–102. 1, 4
- Deming, D. J. and Noray, K.: 2020, Earnings dynamics, changing job skills, and stem careers, *The Quarterly Journal of Economics* **135**(4), 1965–2005. 5.1, 7

- Dornbusch, R., Fischer, S. and Samuelson, P. A.: 1977, Comparative advantage, trade, and payments in a ricardian model with a continuum of goods, *The American Economic Review* **67**(5), 823–839. 3, 4.2
- Eaton, J. and Kortum, S.: 2002, Technology, geography, and trade, *Econometrica* **70**(5), 1741–1779. 3, 4.2
- Edin, P.-A., Fredriksson, P., Nybom, M. and Öckert, B.: 2022, The rising return to noncognitive skill, *American Economic Journal: Applied Economics* **14**(2), 78–100. 1, 4.1, 11
- Edmond, C. and Mongey, S.: 2019, Unbundling labor, *Technical report*, Technical report, Working Paper. 6, 7
- Falk, A., Kosse, F., Pinger, P., Schildberg-Hörisch, H. and Deckers, T.: 2021, Socioeconomic status and inequalities in children’s iq and economic preferences, *Journal of Political Economy* **129**(9), 2504–2545. 4.1
- Friedberg, L.: 2003, The impact of technological change on older workers: Evidence from data on computer use, *ILR Review* **56**(3), 511–529. 5.1
- Garicano, L. and Rossi-Hansberg, E.: 2006, Organization and inequality in a knowledge economy, *The Quarterly journal of economics* **121**(4), 1383–1435. 4.2
- Gathmann, C. and Schönberg, U.: 2010, How general is human capital? a task-based approach, *Journal of Labor Economics* **28**(1), 1–49. 4
- Goldin, C. and Katz, L. F.: 2007, The race between education and technology: the evolution of us educational wage differentials, 1890 to 2005, *Technical report*, National Bureau of Economic Research. (document), 1, 2, 2, 7
- Goos, M., Manning, A. and Salomons, A.: 2014, Explaining job polarization: Routine-biased technological change and offshoring, *American economic review* **104**(8), 2509–2526. 3

- Gould, E. D., Moav, O. and Weinberg, B. A.: 2001, Precautionary demand for education, inequality, and technological progress, *Journal of Economic Growth* **6**(4), 285–315. 5.1
- Guvenen, F., Kuruscu, B., Tanaka, S. and Wiczer, D.: 2020, Multidimensional skill mismatch, *American Economic Journal: Macroeconomics* **12**(1), 210–244. 4
- Haanwinckel, D. and Soares, R. R.: 2021, Workforce composition, productivity, and labour regulations in a compensating differentials theory of informality, *The Review of Economic Studies* **88**(6), 2970–3010. 6
- Hanushek, E. A., Schwerdt, G., Woessmann, L. and Zhang, L.: 2017, General education, vocational education, and labor-market outcomes over the lifecycle, *Journal of Human Resources* **52**(1), 48–87. 5.1
- Hanushek, E. A. and Woessmann, L.: 2010, Education and economic growth, *Economics of education* **60**, 67. 1
- Hanushek, E. A. and Woessmann, L.: 2012, Do better schools lead to more growth? cognitive skills, economic outcomes, and causation, *Journal of economic growth* **17**, 267–321. 6
- Heckman, J. J., Lochner, L. and Taber, C.: 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. 5
- Heckman, J. J., Stixrud, J. and Urzua, S.: 2006, The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior, *Journal of Labor economics* **24**(3), 411–482. 4, 4.2
- Hendricks, L. and Schoellman, T.: 2018, Human capital and development accounting: New evidence from wage gains at migration, *The Quarterly Journal of Economics* **133**(2), 665–700. 2

- Hendricks, L. and Schoellman, T.: 2023, Skilled labor productivity and cross-country income differences, *American Economic Journal: Macroeconomics* **15**(1), 240–68. 4, 2
- Herkenhoff, K., Lise, J., Menzio, G. and Phillips, G. M.: 2018, Production and learning in teams, *Technical report*, National Bureau of Economic Research. 4.2
- Hermo, S., Päälyssaho, M., Seim, D. and Shapiro, J. M.: 2022, Labor market returns and the evolution of cognitive skills: Theory and evidence, *The Quarterly Journal of Economics* **137**(4), 2309–2361. 4.1
- Hershbein, B. and Kahn, L. B.: 2018, Do recessions accelerate routine-biased technological change? evidence from vacancy postings, *American Economic Review* **108**(7), 1737–72. 5
- Hudomiet, P. and Willis, R. J.: 2022, Computerization, obsolescence and the length of working life, *Labour Economics* **77**, 102005. 5.1
- Jäger, S. and Heining, J.: 2022, How substitutable are workers? evidence from worker deaths, *Technical report*, National Bureau of Economic Research. 1, 4.2, 6
- Jarosch, G., Oberfield, E. and Rossi-Hansberg, E.: 2021, Learning from coworkers, *Econometrica* **89**(2), 647–676. 4.2
- Jones, B. F.: 2014, The human capital stock: a generalized approach, *American Economic Review* **104**(11), 3752–3777. 2
- Kambourov, G. and Manovskii, I.: 2009, Occupational specificity of human capital, *International Economic Review* **50**(1), 63–115. 4
- Katz, L. F. and Murphy, K. M.: 1992, Changes in relative wages, 1963–1987: supply and demand factors, *The quarterly journal of economics* **107**(1), 35–78. 1, 4, 2, 5, 7
- Keane, M. P. and Wolpin, K. I.: 1997, The career decisions of young men, *Journal of political Economy* **105**(3), 473–522. 4

- Khanna, G.: 2023, Large-scale education reform in general equilibrium: Regression discontinuity evidence from india, *Journal of Political Economy* **131**(2), 000–000. 2, 2
- Kinsler, J. and Pavan, R.: 2015, The specificity of general human capital: Evidence from college major choice, *Journal of Labor Economics* **33**(4), 933–972. 5.1
- Kirkeboen, L. J., Leuven, E. and Mogstad, M.: 2016, Field of study, earnings, and self-selection, *The Quarterly Journal of Economics* **131**(3), 1057–1111. 5.1
- Kline, P., Petkova, N., Williams, H. and Zidar, O.: 2019, Who profits from patents? rent-sharing at innovative firms, *The quarterly journal of economics* **134**(3), 1343–1404. 1, 6, 7
- Kredler, M.: 2014, Experience vs. obsolescence: A vintage-human-capital model, *Journal of Economic Theory* **150**, 709–739. 5.1
- Kremer, M.: 1993, The o-ring theory of economic development, *The quarterly journal of economics* **108**(3), 551–575. 6
- Krueger, D. and Kumar, K. B.: 2004, Skill-specific rather than general education: A reason for us–europe growth differences?, *Journal of economic growth* **9**(2), 167–207. 5.1
- Lachowska, M., Mas, A. and Woodbury, S. A.: 2020, Sources of displaced workers’ long-term earnings losses, *American Economic Review* **110**(10), 3231–66. 4
- Lamadon, T., Mogstad, M. and Setzler, B.: 2022, Imperfect competition, compensating differentials, and rent sharing in the us labor market, *American Economic Review* **112**(1), 169–212. 6
- Lee, J.-W. and Lee, H.: 2016, Human capital in the long run, *Journal of development economics* **122**, 147–169. 1
- Leighton, M. and Speer, J.: 2017, Labor market returns to college major specificity. 5.1

- Lemieux, T.: 2014, Occupations, fields of study and returns to education, *Canadian Journal of Economics/Revue canadienne d'économique* **47**(4), 1047–1077. 5.1
- Lindbeck, A. and Snower, D. J.: 2000, Multitask learning and the reorganization of work: From tayloristic to holistic organization, *Journal of labor economics* **18**(3), 353–376. 4.2
- Lindenlaub, I.: 2017, Sorting multidimensional types: Theory and application, *The Review of Economic Studies* **84**(2), 718–789. 4, 4.2
- Lindner, A., Muraközy, B., Reizer, B. and Schreiner, R.: 2022, Firm-level technological change and skill demand. 2
- Lindqvist, E. and Vestman, R.: 2011, The labor market returns to cognitive and noncognitive ability: Evidence from the swedish enlistment, *American Economic Journal: Applied Economics* **3**(1), 101–128. 4
- Lise, J. and Postel-Vinay, F.: 2020, Multidimensional skills, sorting, and human capital accumulation, *American Economic Review* **110**(8), 2328–76. 3, 4, 4.2
- MacDonald, G. and Weisbach, M. S.: 2004, The economics of has-beens, *Journal of Political Economy* **112**(S1), S289–S310. 5.1
- Machin, S. and Van Reenen, J.: 1998, Technology and changes in skill structure: evidence from seven oecd countries, *The quarterly journal of economics* **113**(4), 1215–1244. 2
- Michaels, G., Natraj, A. and Van Reenen, J.: 2014, Has ict polarized skill demand? evidence from eleven countries over twenty-five years, *Review of Economics and Statistics* **96**(1), 60–77. 3
- Mincer, J.: 1958, Investment in human capital and personal income distribution, *Journal of political economy* **66**(4), 281–302. 1
- Neal, D.: 1995, Industry-specific human capital: Evidence from displaced workers, *Journal of labor Economics* **13**(4), 653–677. 4

- Neuman, S. and Weiss, A.: 1995, On the effects of schooling vintage on experience-earnings profiles: Theory and evidence, *European economic review* **39**, 943–955. 5.1
- Pavan, R.: 2011, Career choice and wage growth, *Journal of Labor Economics* **29**(3), 549–587. 4
- Rosen, S.: 1975, Measuring the obsolescence of knowledge, *Education, income, and human behavior*, NBER, pp. 199–232. 5.1
- Rossi, F.: 2020, Human capital and macroeconomic development: a review of the evidence, *The World Bank Research Observer* **35**(2), 227–262. 6
- Sanders, C. and Taber, C.: 2012, Life-cycle wage growth and heterogeneous human capital, *Annu. Rev. Econ.* **4**(1), 399–425. 4
- Shaw, K. L.: 1989, Life-cycle labor supply with human capital accumulation, *International Economic Review* pp. 431–456. 4
- Song, J., Price, D. J., Guvenen, F., Bloom, N. and Von Wachter, T.: 2019, Firming up inequality, *The Quarterly journal of economics* **134**(1), 1–50. 6
- Syverson, C.: 2011, What determines productivity?, *Journal of Economic literature* **49**(2), 326–365. 6
- Taber, C. and Vejlin, R.: 2020, Estimation of a roy/search/compensating differential model of the labor market, *Econometrica* **88**(3), 1031–1069. 4
- Thompson, P.: 2003, Technological change and the age–earnings profile: Evidence from the international merchant marine, 1861–1912, *Review of Economic Dynamics* **6**(3), 578–601. 5.1
- Tinbergen, J.: 1975, *Income distribution: Analysis and policies*. (document), 1, 7

- Weidmann, B. and Deming, D. J.: 2021, Team players: How social skills improve team performance, *Econometrica* **89**(6), 2637–2657. 4.2
- Weinberg, B. A.: 2004, Experience and technology adoption. 5.1
- Weiss, Y. and Lillard, L. A.: 1978, Experience, vintage, and time effects in the growth of earnings: American scientists, 1960-1970, *Journal of Political Economy* **86**(3), 427–447.
5.1
- Willis, R. J. and Rosen, S.: 1979, Education and self-selection, *Journal of political Economy* **87**(5, Part 2), S7–S36. 4
- Yamaguchi, S.: 2012, Tasks and heterogeneous human capital, *Journal of Labor Economics* **30**(1), 1–53. 4

Figure 1

Reproduced from Figure IV of Katz and Murphy (1992)

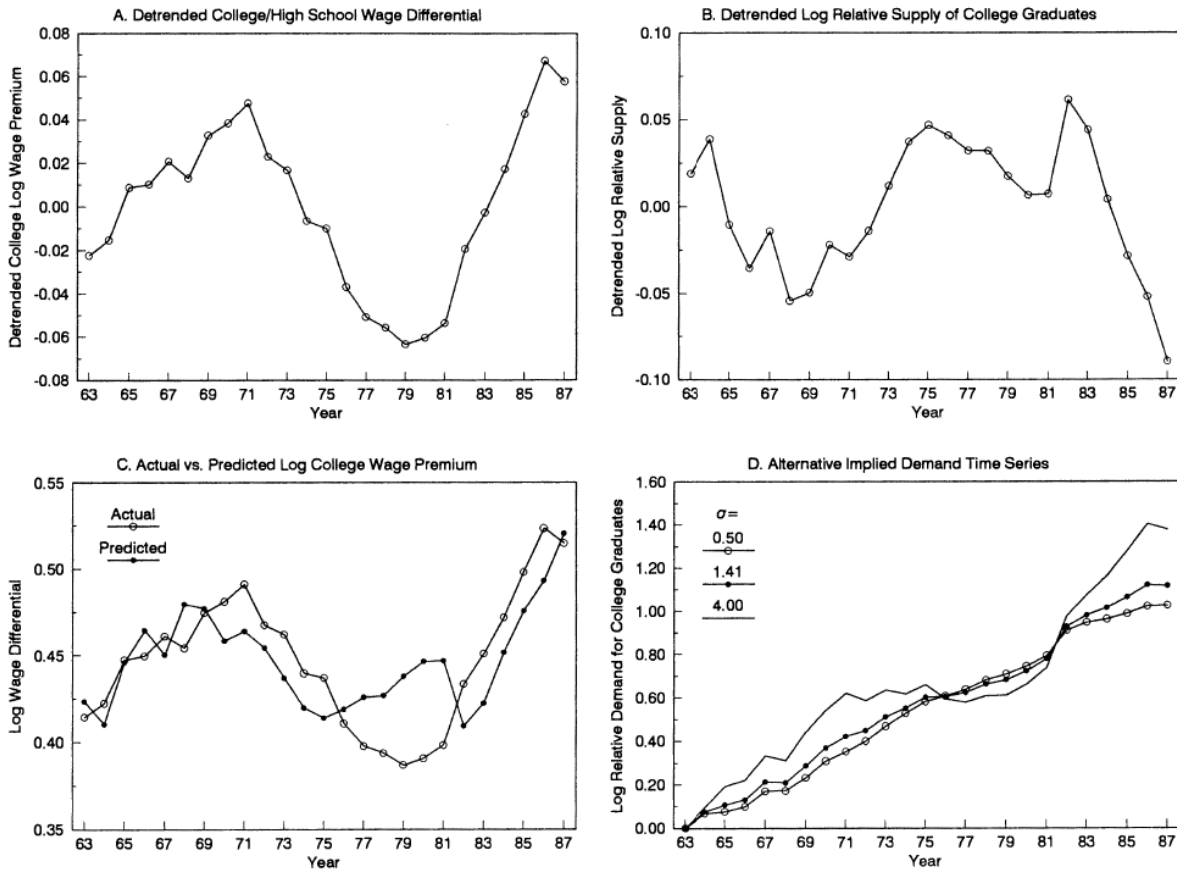


FIGURE IV
Education Differentials: Supply and Demand Analysis

Figure 2

Reproduced from Figure 2 of Goldin and Katz (2007)

Figure 2
Actual versus Predicted College Wage Premium: 1915 to 2005

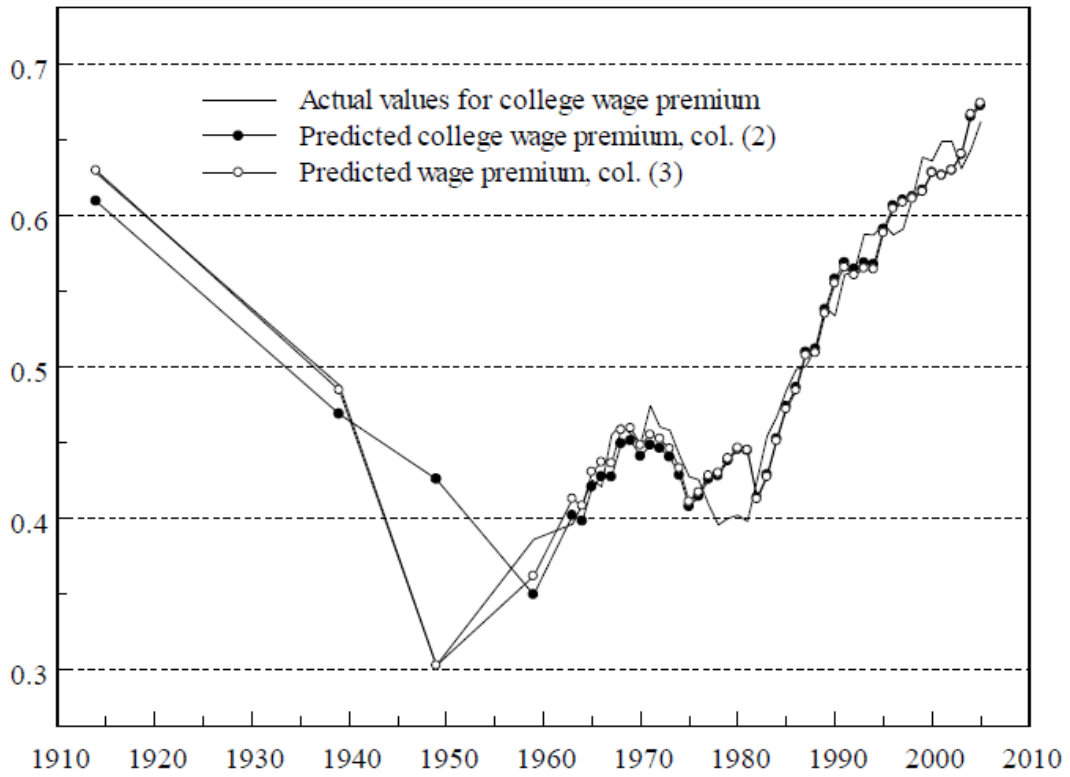


Figure 3

Reproduced from Figure II of Deming (2017)

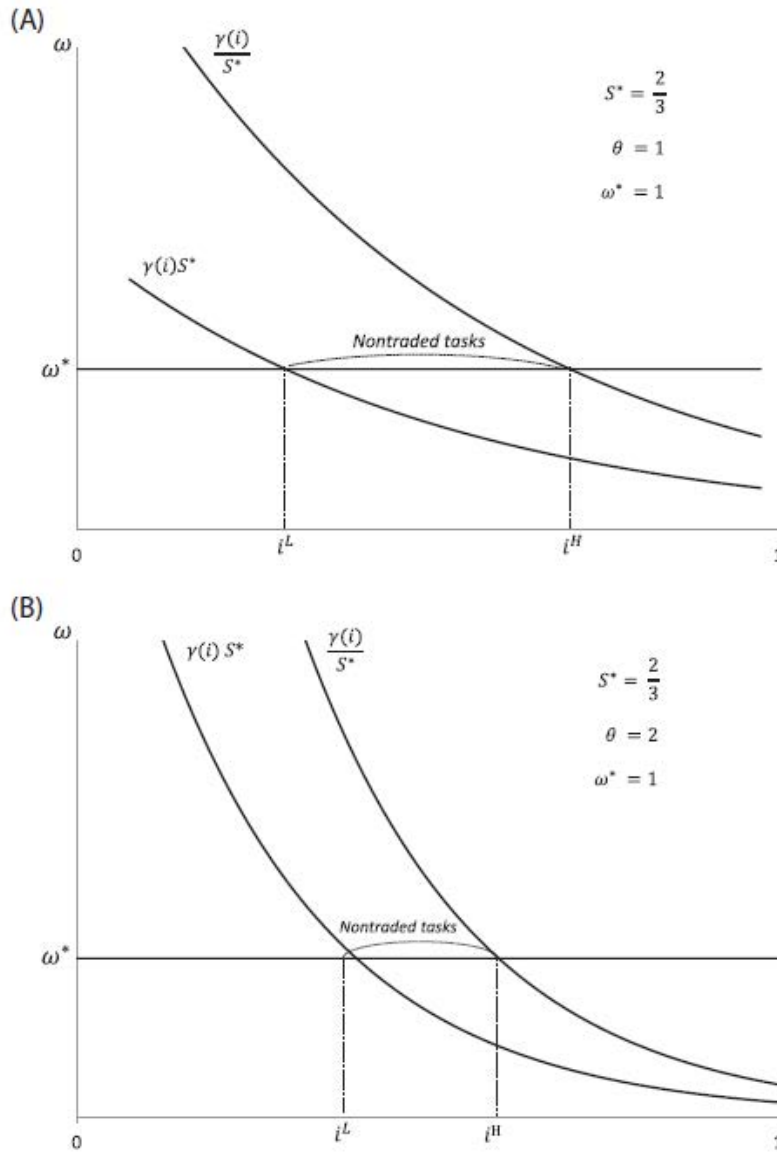


FIGURE II

Equilibrium Task Thresholds with Different Values of Theta

Panel A illustrates the equilibrium task thresholds i^L and i^H from the model in Section II when $S^* = \frac{2}{3}$, $\theta = 1$, and $\omega^* = 1$. Panel B illustrates the equilibrium task thresholds i^L and i^H from the model in Section III when $S^* = \frac{2}{3}$, $\theta = 2$, and $\omega^* = 1$ (see the text for details).

Table 1*Reproduced from Table 2 of Jones (2014)*

TABLE 2—HUMAN CAPITAL AND INCOME: GENERALIZED ACCOUNTING

	Elasticity of substitution between unskilled labor, H_1 , and skilled aggregate, $Z(H_2, \dots, H_N)$						
	1	1.2	1.4	1.6	1.8	2.0	∞
<i>Panel A. Relaxing Assumption 3</i>							
H^R/H^P	∞	358	21.9	8.6	5.4	4.1	1.3
Λ	∞	269	16.4	6.5	4.0	3.1	1.0
Success (percent)	∞	1,050	163	88	64	54	25
<i>Panel B. Relaxing Assumptions 3 and 4</i>							
H^R/H^P	∞	306	18.6	7.3	4.6	3.5	1.1
Λ	∞	269	16.4	6.5	4.0	3.1	1.0
Success (percent)	∞	950	147	79	58	48	23

Notes: This table compares Israel and Kenya, which represent the 85th and 15th percentile countries respectively ranked by income per worker. H^R/H^P is the ratio of human capital stocks. Λ is the ratio of H^R/H^P at the indicated elasticity of substitution to H^R/H^P for the infinite elasticity of substitution case. Success is the consequent percentage of the income variation that is explained by variation in capital inputs. Figure 2 summarizes accounting for a broader set of rich and poor countries and shows that Israel and Kenya provide a useful benchmark, as discussed in the text.