Life-Cycle Wage Growth and Heterogeneous Human Capital

Carl Sanders[1] and Christopher Taber[2]

[1]Department of Economics, Washington University in Saint Louis, Saint Louis, Missouri 63130; email: carlsanders@wustl.edu

[2]Department of Economics, University of Wisconsin-Madison, Madison, Wisconsin 53706; email: ctaber@ssc.wisc.edu

Annu. Rev. Econ. 2012. 4:399–425

First published online as a Review in Advance on April 5, 2012

The *Annual Review of Economics* is online at economics.annualreviews.org

This article's doi:
10.1146/annurev-economics-080511-111011

JEL codes: J24, J3, J62

1941-1383/12/0904-0399$20.00

Erratum

Keywords

task-specific human capital, occupation, industry

Abstract

Wages grow rapidly for young workers, and the human capital investment model is the classic framework to explain this growth. While estimation and the theory of human capital have traditionally focused on general human capital, both have evolved toward models of heterogeneous human capital. In this article, we review and evaluate the current state of this literature. We exposit the classic model of general human capital investment and extend it to show how a model of heterogeneous human capital can nest previous models. We then summarize the empirical literature on firm-specific human capital, industry- and occupation-specific human capital, and task-specific human capital and discuss how these concepts can explain a wide variety of labor market phenomena that traditional models cannot.

1. INTRODUCTION

Workers' wages grow substantially during the early part of their careers. For example, in the National Longitudinal Survey of Youth 1979 (NLSY79), the typical wage growth for a male with exactly 12 years of schooling is 47% from ages 18 to 28, and wages approximately double between ages 18 and 45. This aspect of the data is well known, and any model of wages must account for life-cycle wage growth. The canonical model of wage growth is the on-the-job model of human capital investment (e.g., Becker 1964, Ben-Porath 1967). The standard model considers human capital to be a one-dimensional skill that is useful at all firms. However, Becker (1964) considers a second type of human capital, firm-specific human capital, which is useful only at the firm at which it was acquired.

Although treating skills as either completely general or completely firm specific is pedagogically useful, it is not realistic. Most skills can be easily transferred to some other job with little loss in efficiency, but they cannot be transferred to all jobs. We can see this in the economics profession. Someone who teaches introductory economics at the University of Wisconsin-Madison can easily transfer these skills to another economics department, such as Washington University in St. Louis. However, this skill would be virtually useless if the economist became a ballroom dancer. Intermediate cases also exist: These skills are also presumably useful if the economist became a high school history teacher, although they will not transfer perfectly. Although scholars have recognized this issue for a long time, it is only recently that empirical work has incorporated this partially transferable type of heterogeneous skill.

The full literature on life-cycle wage growth is much too large to summarize here. We instead present a selected survey. We see the literature moving forward toward more general models of heterogeneous human capital. Our main goal is to discuss different pieces of this literature that have evolved somewhat independently. We hope this proves useful for future work that attempts to combine the elements.

With this goal in mind, we start with the theory and then move to the empirical work. We begin with a model of general human capital based on the Ben-Porath (1967) model and then extend it to incorporate a more general type of heterogeneous human capital in the presence of search frictions. In presenting the empirical work, we begin with the general framework and the work on returns to seniority, which introduces the key selection problem. We then discuss occupation- and industry-specific human capital and finally survey the relatively recently developed task-based approach. We argue that the usefulness of the occupation and industry categories in thinking about specific human capital results from them serving as measurable proxies for the underlying task requirements of a job and that future research in human capital should focus on a deeper analysis of task-specific human capital at both the theoretical and empirical levels.

2. HUMAN CAPITAL THEORY

2.1. General Human Capital

We start with a generalized discrete-time version of the Ben-Porath (1967) model. Although the exposition is cleaner with continuous time, we use discrete time to be consistent with the next section. We assume that workers work for T years and that agents are risk neutral so they maximize their expected lifetime earnings.[1] Workers borrow and lend freely with return R.[2]

Human capital has two inputs: previous human capital and time. Workers divide their time into two activities, investing in human capital (s_t) and producing output ($1 - s_t$). We write the human capital production function as

$$H_{t+1} = \mathcal{H}(s_t, H_t) \tag{1}$$

and assume that it is strictly increasing and concave in s_t. The rental rate on human capital is normalized to unity so that worker productivity during a period is

$$(1 - s_t)H_t.$$

That is, if the worker did not invest in human capital at all ($s_t = 0$), he would produce H_t. However, the time spent investing in human capital does not augment output. We are assuming that markets are competitive, so for the classic reasons in Becker (1964), workers finance human capital investment through foregone earnings. Thus the worker's wage is his productivity $(1 - s_t)H_t$.

We write the model using Bellman's equation with V_t as the value function with one state variable, human capital (H_t). For periods prior to the final one, we write the Bellman equation

$$V_t(H_t) = \max_{s_t} (1 - s_t)H_t + \frac{1}{R} V_{t+1}(H_{t+1}) \tag{2}$$

subject to

$$H_{t+1} = \mathcal{H}(s_t, H_t), \tag{3}$$

$$0 \leq s_t \leq 1. \tag{4}$$

For the final period, the problem simplifies to

$$V_T(H_T) = (1 - s_T)H_T. \tag{5}$$

When s_t is interior, the first-order condition for investment is

$$H_t = \frac{1}{R} \frac{\partial V_{t+1}(H_{t+1})}{\partial H_{t+1}} \frac{\partial \mathcal{H}(s_t, H_t)}{\partial s_t}. \tag{6}$$

Equation 6 states that the marginal cost of human capital investment equals the marginal return. If we measure human capital investment as s_t, the marginal cost is H_t. The marginal return to this investment is the discounted value of the increase in human capital.

Because $\mathcal{H}$ is concave, as R increases, human capital investment (s_t) decreases. Moreover, as the future value of human capital $\left[\frac{\partial V_{t+1}(H_{t+1})}{\partial H_{t+1}}\right]$ increases, so does human capital investment.

The model is completed with the envelope condition

$$\frac{\partial V_t(H_t)}{\partial H_t} = (1 - s_t) + \frac{1}{R} \frac{\partial V_{t+1}(H_{t+1})}{\partial H_{t+1}} \frac{\partial \mathcal{H}(s_t, H_t)}{\partial H_t}. \tag{7}$$

[1]For this simple model, it will not matter as consumption and human capital are separable. When we add uncertainty in job offers in Section 2.2, they will no longer be equivalent.

[2]That is, $R = 1 + r$, where r is the interest rate.

Although empirical work suggests a more general form (see Browning et al. 1999 for discussion), to get clearer predictions we use the Ben-Porath functional form:

$$\mathcal{H}(s_t, H_t) \equiv A(s_t H_t)^{\alpha} + (1-\sigma)H_t, \tag{8}$$

where (A, α, σ) is a parameterization of the model in which σ represents human capital depreciation, and A and α are similar to the scale and share parameters in the Cobb-Douglas production function.

This simplifies the model considerably because one can get a closed-form solution for

$$\frac{\partial V_t(H_t)}{\partial H_t} = \frac{1-\left(\frac{1-\sigma}{R}\right)^{T+1-t}}{1-\left(\frac{1-\sigma}{R}\right)}. \tag{9}$$

To get some intuition for the expression in Equation 9, notice that when in the final period of work the future value of human capital $\frac{\partial V_{T+1}(H_{T+1})}{\partial H_{T+1}}$ would be zero. This makes perfect sense as after the final year of life, human capital is worthless. Also this object is strictly decreasing in t. This is intuitive for the same reason. As one gets closer to retirement, there are fewer years to utilize human capital, so its value falls.

One can also show that when investment is interior,

$$H_t s_t = \left[\frac{1}{R}\frac{\partial V_{t+1}(H_{t+1})}{\partial H_{t+1}} A\right]^{\frac{1}{1-\alpha}}. \tag{10}$$

The monetary cost of investment is foregone earnings $H_t s_t$, so the left-hand side of Equation 10 is a measure of investment. We see from this expression that investment must fall as workers age. Furthermore, the value of $H_t s_t$ depends only on α, σ, A, and R. In particular, it does not depend on initial human capital, so two otherwise identical individuals with different initial human capital will still have exactly the same foregone earnings $H_t s_t$. This property is often referred to as Ben-Porath neutrality.

One constraint of the model is that $s_t \leq 1$. Because $H_t s_t$ decreases with age, early in life this constraint may bind, in which case $s_t = 1$ and earnings are zero. Ben-Porath (1967) interprets this as schooling.[3]

To give a basic sense of the model, we simulate a version of it in **Figures 1** and **2**.[4] **Figure 1** presents human capital investment. There is a period of human capital specialization in the beginning of the life cycle in which $s_t = 1$ (schooling). After that, investment decreases rapidly and then reaches zero at retirement. This relatively simple model can explain schooling, the rapid increase in earnings we see shortly after labor market entry, and the decrease in earnings near retirement.

Figure 2 presents human capital over ages as well as earnings, $(1 - s_t)H_t$. Because the rental rate on human capital is one, these are in the same units, and H_t is also potential earnings. One can see the importance of the distinction between potential earnings and wages. Until very late in the working life, potential earnings are higher because human

[3]We do not distinguish between parents making decisions for their children and children making them on their own. As long as parents make decisions to maximize the net present value of the children's earnings, the distinction is irrelevant.

[4]These parameters were chosen loosely to match an average data profile from a fixed-effect model with high school graduates from the Current Population Survey matched outgoing rotation group. In this simulation, $\alpha = 0.80$, $A = 0.121$, $\sigma = 0.01$, $R = 1.05$, and we parameterize the terminal level of human capital as 16.00.

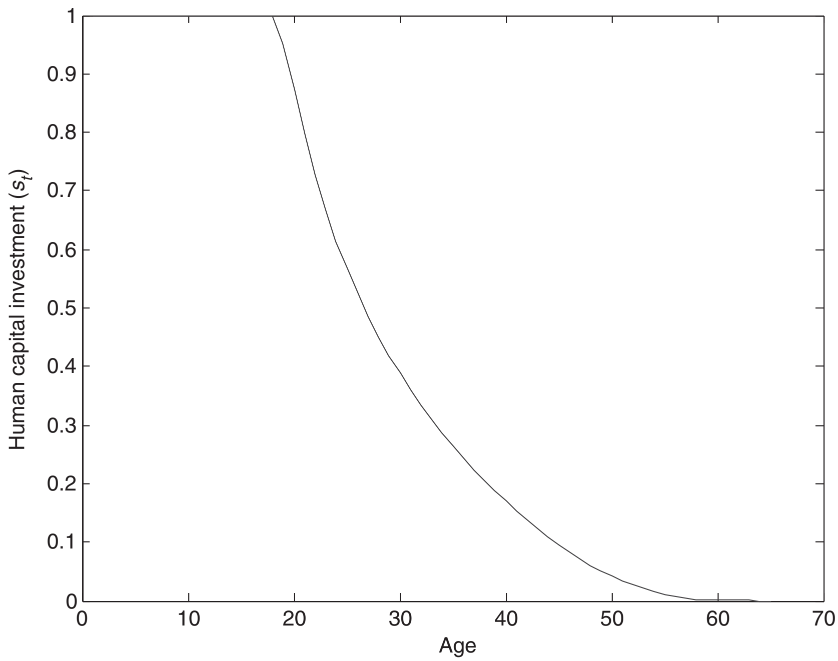

Figure 1

Human capital investment.

capital investment (s_t) is positive. Whereas our model is fundamentally a model of human capital accumulation, H_t, the empirical work below does not look at human capital, but rather wages, $(1 - s_t)H_t$. This is important to keep in mind when interpreting the empirical work below, and Heckman et al. (1998a) show that this distinction can be important.

2.2. Multidimensional Skills

We now extend the model to a more general framework with multidimensional skills and search frictions. Although our framework here has not been written down precisely in this form before, the pieces can be found in many previous models. Key papers on either heterogeneous human capital or human capital with search include Rosen (1983), Heckman & Sedlacek (1985), Stevens (1994), Acemoglu & Pischke (1998), Felli & Harris (1996), Neal (1998), and Rubinstein & Weiss (2006). The closest to our framework is Lazear (2009).

A key aspect of the Ben-Porath (1967) model is that human capital investment declines with age. This is also true in our more general model in much the same way. To increase clarity, we focus on the differences between the models rather than the similarities. These can be illustrated clearly in a two-period model, so in the rest of this section we assume that $T = 2$. Extending the model to more than two periods is straightforward. We also treat period 1 as post-schooling, so it could be the first period the worker enters the labor market.

The primary difference between this model and the classic model above is that human capital H_t is now a multidimensional vector. Let $H_t^{(m)}$ denote the m-th component of human

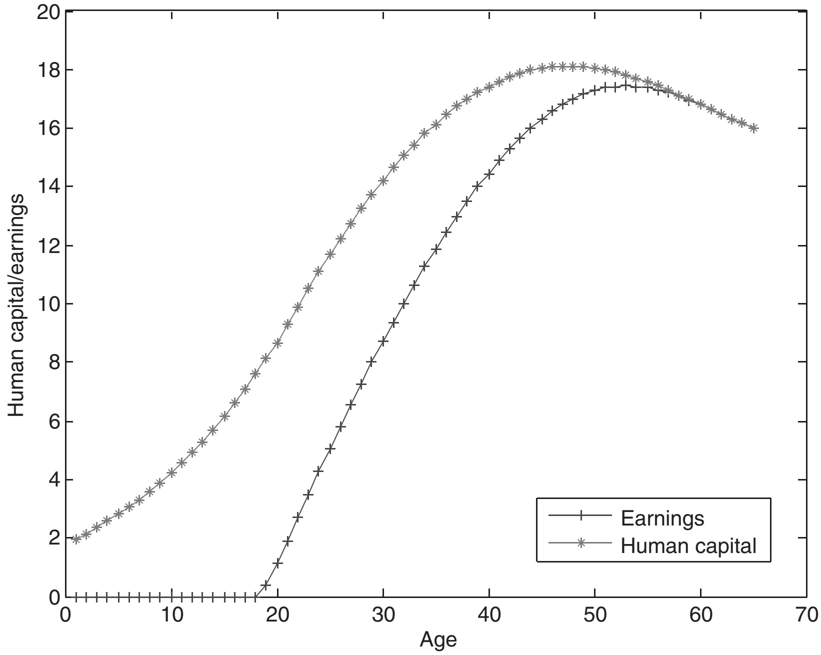

Figure 2

Human capital and earnings over the life cycle.

capital. Because initial H_1 human capital is exogenous, we do not incorporate it into the human capital production function; this would be a simple extension. We also assume that the only input into the m-th skill is m-specific time so that

$$H_2^{(m)} = \mathcal{H}^{(m)}\left(s_1^{(m)}\right), \tag{11}$$

where $s_1^{(m)}$ is the time devoted to m, and s_1 is the corresponding vector of investments.[5]

A second feature of this model is that different types of jobs will pay different wages. We let π be the vector of skill prices that a job pays. Thus an individual with human capital H_t is paid $\pi' H_t$ per efficiency unit at a π job. That is, during the first period a worker employed at a type π job earns

$$\pi_1' H_1 \left[1 - \sum_{m=1}^{M} s_1^{(m)}\right]. \tag{12}$$

For the reasons mentioned above, in the second period there will be no investment, so a worker at a π job will be paid $\pi' H_2$.

We intentionally use broad notation so as not to be specific about what this vector includes. General human capital is a special case in which H_t is one dimensional and $\pi = 1$ at all firms. We show in the **Supplemental Appendix** that firm-specific human capital can

[5]This production function may vary across first-period firms, although in this model the first-period firm is taken as given, so that is not relevant here.

also be included by assuming that each firm corresponds to a dimension of H_t and that the payoff to that skill is zero at any other firm (follow the **Supplemental Material** link from the Annual Reviews home page at **http://www.annualreviews.org**). We also show in the **Supplemental Appendix** that this framework can easily be generalized so that different dimensions correspond to different occupations or industries. Then occupation- or industry-specific human capital is modeled to be valuable only in the occupation or the industry in which it was acquired. We intend for the notation to be even more general than this. It can include broad categories of skills such as cognitive skills, or various noncognitive skills such as the Big Five personality traits. It could also include narrow task-specific skills such as typing, welding, or public speaking. The question of what exactly is in this vector is a primary subject of debate in the empirical work we survey in later sections, although usually those papers do not use this explicit theoretical framework.

Once we include firms in the model, an issue arises as to why workers would ever work at more than one firm. The most natural way to account for movement is to include search frictions in the labor market. There are many ways to model the wage-determination process in an on-the-job search model such as Burdett & Mortensen (1998) and Cahuc et al. (2006).[6] These different models require different assumptions on what types of contracts and outside options are available to workers and firms. As the goal of this model is to help interpret the empirical work discussed below, we abstract from this issue by assuming, as described above, that wages are linear with prices π.

However, ignoring the wage-setting process abstracts from a couple of interesting aspects of human capital that arise in these models. In particular, in a bargaining model there are two potential inefficiencies in human capital investment. The first is the familiar holdup problem that results in underinvestment in human capital. If a worker switches to a different firm in the second period and if that firm collects some of the rents from first-period investment, then during the first period the worker will underinvest in skill. The second, more novel, potential inefficiency is specific to this heterogeneous human capital framework. In some wage-determination models, firms lose rents if workers switch to a slightly better firm. Even when this switching is efficient, firms can avoid it by overinvesting in skills that are relatively valuable at the current firm and underinvesting in skills that are more general. We discuss these cases in the **Supplemental Appendix**. Although these issues have not been extensively studied in the empirical literature, they are potentially important future areas of interest.

We start by solving the model for period 2 first. In that period, the worker receives one offer from another firm. We could easily allow more than one offer at a time, but that adds complexity to the notation with little extra economic insight. To keep the notation simple, one can allow the offer distribution of π to have a point mass at zero to denote the state of the world in which the worker receives no offers. The worker always has the option to stay at the first-period firm.[7]

We work backward starting with the period 2 wage. The wage at time 2 for a worker at a π firm is

$$w_2(H_2, \pi) = \pi' H_2. \tag{13}$$

[6]Burdett & Coles (2010) and Fu (2011) provide equilibrium versions of human capital models with search frictions.

[7]Relaxing this assumption to allow for exogenous separations is straightforward and is also done in the **Supplemental Appendix**.

The value function at time 1 for a worker employed during period 1 at a productivity π_1 firm is

$$V_1(H_1, \pi_1) = \left(1 - \sum_{m=1}^{M} s_1^{(m)}\right)\pi_1' \, H_1 + \frac{1}{R} E \max\{\pi_1' H_2, \pi' H_2\}. \tag{14}$$

The Emax term in Equation 14 represents that the worker chooses the firm that offers the highest wage in the second period.

For any particular skill m, we can write the first-order condition for human capital as

$$\pi_1' H_1 = \frac{1}{R}\left[Pr(\pi_1' H_2 \leq \pi' H_2) E\left(\pi^{(m)} | \pi_1' H_2 \leq \pi' H_2\right) + Pr(\pi_1' H_2 > \pi' H_2)\pi_1^{(m)}\right] \frac{\partial \mathcal{H}^{(m)}\left(s_1^{(m)}\right)}{\partial s_1^{(m)}}. \tag{15}$$

The first terms in the brackets in Equation 15 represent the worker's expected return to human capital in cases in which the worker switches to a different firm in the second period. The second term represents the expected return in the case in which the worker stays at the current firm.

An interesting aspect of the problem in Equation 15 is that one can isolate two different strategies: investment in specific skills and investment in general skills. One can see this as two different local optima in the objective function. One local optimum represents specialization in firm 1, and the other represents workers who specialize little. To see this, consider a simple version of the model in which there are only two types of skills. Take the first skill to be immutable [i.e., $\mathcal{H}^{(1)}(s_1^{(1)}) = H_1$ so that ${s_1}^{(1)} = 0$]. Skill 2 is firm specific in that it is valueless at any outside firm [$\pi^{(2)} = 0$]. Using a Ben-Porath-style production function for specific human capital but ignoring the human capital component, we write

$$H_2^{(2)} = A\left[s_1^{(2)}\right]^{\alpha}. \tag{16}$$

We take $\alpha = 0.4$, $A = 1.5$, $1/R = 0.95$, $H_1 = (1, 1)$, and $\pi_1 = (1, 1)$.We first simulate a version of the model that is concave. In this first model, the outside offer $\pi^{(1)}$ is distributed as a standard log normal. In **Figure 3*a*** we graph the value function as a function of time spent investing [$s_1^{(2)}$ in this case]. One can see that it is a nicely behaved model with a clear optimum. We contrast this with an alternative case in which the outside offer distribution is a single point $\pi = 1.7$ and the probability of an offer is 80% (**Figure 3*b***). The value function is not concave with two distinct local optima. The nonconcavity comes from the degree of specialization. The low-investment local optimum comes from the case in which the worker invests little and would take the outside job if it were offered. The second corresponds to the case in which the worker would stay at the current firm if offered the outside job. This suggests that in more complicated specifications, one might find two or more different types of workers: some that specialize in specific skills and plan to stay at this type of firm and others that invest in more general skills and plan to leave.

We have emphasized the role of human capital investment on wage growth, but the addition of search frictions also yields an additional source of wage growth over the life cycle. As workers age, they will have had more time to have received outside offers with higher values of π. Thus even if human capital did not change over time and regardless of

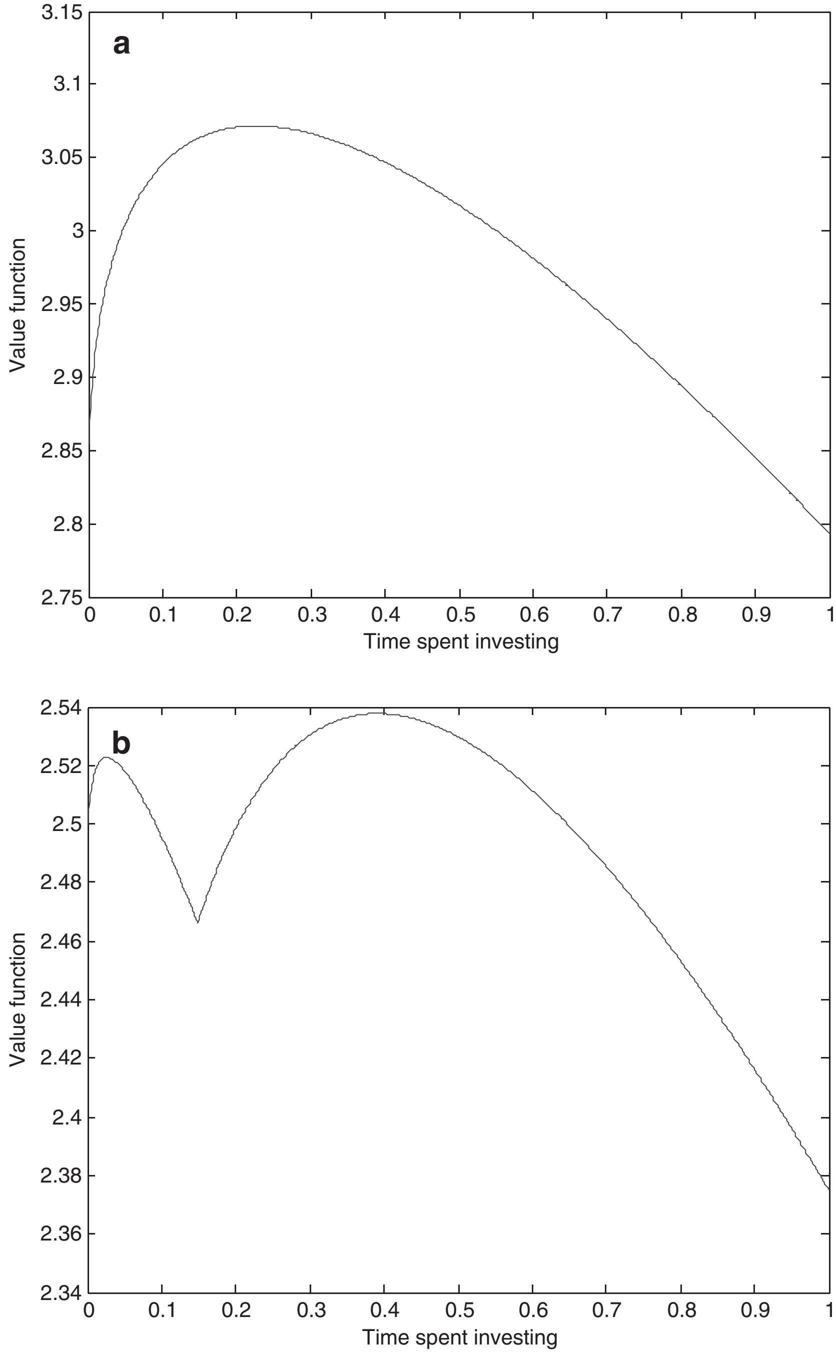

Figure 3

Value function by (*a*) investment level case 1 and (*b*) investment level case 2.

its dimension, we would expect average wages to grow as the average value of π would increase between periods 1 and 2.[8]

A related additional source of wage growth could arise from heterogeneity in human capital production across firms. In the model above, firms differ in output technology but not their efficacy in helping the worker produce additional human capital. Even in the two-period model, if we allow for heterogeneity in human capital investment production across firms, workers will tend to choose firms with higher human capital productivity in the first period and higher π in the second. This idea can be found in Rosen (1972) and is built on by Heckman et al. (2003). We could incorporate this into our model by assuming the human capital production function is

$$\mathcal{H}(s;\theta),$$

where θ varies across firms. Then a firm could be characterized by the pair (θ, π). Because they have a longer time to reap the benefits of human capital accumulation, young workers will put relatively more weight on firms with a better human capital production function, whereas older workers will put relatively more weight on the rental rate π.

3. EMPIRICAL WORK

In Section 2, we intentionally leave the interpretation of the components of the multi-dimensional human capital vector vague. For the past 25 years, there has been a debate over what appropriate empirical analog of this theoretical vector is best supported by the data. For example, researchers have examined models of firm-specific human capital, occupation-specific levels, industry-specific levels, and even more general interpretations in terms of "career" matches and related concepts. In this section, we discuss how at times all these have been found to have explanatory power for human capital investment models. In our view, these approaches succeed or fail to explain the data to the extent that they serve as proxies for similarities in the underlying tasks that workers have to perform. In addition, thinking in terms of investments in skills is a better way to explain wage growth than occupation tenure, firm tenure, or other discrete categories that have been used. We take an approximately chronological approach in this discussion, first discussing empirical implementation of the classic Ben-Porath model; then the debate over the returns to firm-specific tenure versus general experience; and then occupational tenure, industry tenure, and extensions of these concepts, all of which are special cases of the model above. Finally, we summarize the empirical evidence that these older measures are best understood as proxies for task-specific human capital.

3.1. Empirical Implementations with General Human Capital

The oldest and most closely linked literature to our theoretical framework estimates structural versions of the Ben-Porath model. The set of papers that explicitly estimate a Ben-Porath model is quite small. Early work on this topic includes Heckman (1976), Haley (1976), and Rosen (1976). Heckman et al. (1998a) represent a relatively recent version in

[8]In our model, on-the-job search leads to wage gains as individuals move across jobs. Topel & Ward (1992) show that at least one-third of wage growth for young workers occurs at job switches. Although we do think this is a key aspect of life-cycle wage growth, owing to space constraints we abstract from it to focus on human capital.

LEARNING ABOUT MATCH QUALITY

Another form of wage growth from which we completely abstract is learning. This literature concerns information shocks, in which firms and workers are uncertain about their match quality but learn it over time. This leads to wage growth as workers and firms find the best matches over time. As in the human capital literature, there is a similar trend in the learning literature of moving from a simple worker/firm match to more heterogeneous types of matches. The seminal paper is Jovanovic (1979), who analyzes the case of a firm and worker learning about their match and looking at the implications for turnover rates. Miller (1984) analyzes a model in which individuals learn about their match with occupations, whereas Neal (1999) develops a model in which individuals have both a match with a "career" (which can be common to a number of firms) and a firm. The most recent work (Antonovics & Golan 2012) allows for heterogeneous skills. In their model, worker must experiment across different tasks to learn their true underlying skill levels in a framework analogous to the task-specific human capital models discussed here. Sanders (2012) estimates a model with both skill-based human capital accumulation and skill uncertainty and finds that skill uncertainty plays a primary role in job mobility but only a minor role in wage growth.

which they estimate the model separately for different schooling groups and different ability groups. That human capital is endogenously accumulated is crucial for their explanation of the changing wage structure. These results, as well as the other work up to that point, are presented in Browning et al. (1999). Subsequent work includes Taber (2002), who incorporates progressive income taxes into the estimation, and Kuruscu (2006), who estimates the model nonparametrically.

Despite the widespread use of the Ben-Porath (1967) model as a human capital theory model, by far the most common empirical approach to account for life-cycle wage profiles is the Mincer (1958, 1974) model. We can write the specification as

$$\log(w_i) = \beta_0 + \beta_1 S_i + \beta_2 E_i + \beta_3 E_i^2 + u_i, \tag{17}$$

where, for individual i, w_i is wages, S_i is years of schooling, E_i is experience (usually Age $- S_i - 6$), and u_i is the error term. Mincer (1974) formally derives the specification as an approximation of the Ben-Porath (1967) model. However, most authors that use it do not explicitly give this interpretation. The very large literature on Mincer models is surveyed and discussed in Heckman et al. (2006a). Murphy & Welch (1990) argue that the data are best approximated by a quartic in experience rather than a quadratic.[9]

Although human capital investment is endogenous in our model, most of the empirical work we discuss below either explicitly or implicitly uses a learning-by-doing model. The distinction between learning by doing and endogenous human capital is important in this literature. A learning-by-doing model is a special case of our model above in which human capital increases without the input s_t. However, this does not mean that human capital is completely exogenous. In many models, such as Keane & Wolpin (1997) and Imai & Keane (2004), individuals must work for their human capital to increase. In fact, the main point of Imai & Keane (2004) is that this human capital motivation leads younger workers to supply more labor to the market than they otherwise would. Heckman et al.

[9]However, Heckman et al. (2006a) show that this is not important in computing the internal rate of return.

(2003) discuss the issues of distinguishing between these two models. Rosen (1972) points out that even if there is learning by doing at a particular firm, if that amount differs by firm and workers can choose their firm, then this type of framework can look very much like the standard Ben-Porath model.

3.2. The Returns to Experience Versus the Returns to Tenure: Firm Tenure and Selection Bias

Firm-specific human capital is a special case of our general model that can be modeled by treating different elements of our vector of human capital as different firms (as discussed in the **Supplemental Appendix**). To our knowledge, this type of model has never been directly estimated in the empirical literature. Virtually the whole literature on returns to seniority, as well as the literature discussed below on more general types of human capital, uses a regression framework to analyze human capital accumulation. In this section, we focus on the fundamental econometric problem faced when estimating specific human capital models: how to disentangle growth in wages due to accumulating specific human capital from selection on match quality. If we observe workers who have one year of tenure at a firm receiving significantly lower wages than those with 10 years of tenure, it could be that either (*a*) the workers became that much more productive at that firm or (*b*) the workers who stayed for 10 years have done so because they are good matches with the firm. Although we highlight this issue in this section, the exact same problem arises in the more general models that follow. A good deal of the human capital literature that attempts to measure the true causal returns to tenure—specific to the firm, industry, occupation, or task—must deal with this fundamental identification problem. For this reason we dedicate a fair amount of space to the problem.

We write a simplified version of the model as

$$\log(w_{it}) = \beta E_{it} + \alpha T_{it} + \theta_i + \eta_{ij(i,t)} + \varepsilon_{it}, \tag{18}$$

where w_{it} is the wage that person i receives at time t. T_{it} represents tenure on the current job, and E_{it} is experience as above. The distinction is that E_{it} increases regardless of the firm at which the worker is employed, but T_{it} increases only for job stayers. We let $j(i, t)$ denote the job held by person i at time t. If the worker moves to a new firm $[j(i, t) \neq j(i, t-1)]$, she starts with $T_{it} = 0$ and her tenure accumulates as long as she stays on that job [i.e., $T_{it} = T_{it-1} + 1$ when $j(i, t) = j(i, t-1)$]. The error terms are person specific, θ_i; match or firm specific, $\eta_{ij(i,t)}$; and idiosyncratic, ε_{it}. We have written the model as linear for expositional purposes, but empiricists typically include higher-order terms as well as other regressors. This literature does not allow tenure at one firm to affect wages at another, something that the subsequent literature relaxes.

The major issue in this literature is that T_{it} is likely positively correlated with $\eta_{ij(i,t)}$. Following the logic of our model above, the reason is that the larger $\eta_{ij(i,t)}$ is, the less likely it will be that an outside firm will be able to hire the worker away. Thus people with better matches will tend to have higher tenure. We would also generally expect $\eta_{ij(i,t)}$ to be positively correlated with E_{it} because people who have been in the labor market for longer would have gotten more chances to get a better match draw. Although the theory does not clearly sign the effect, researchers worry about the correlation of E_{it} and T_{it} with θ_i as well.

A number of different papers have estimated versions of Equation 18. We focus on two of them: Altonji & Shakotko (1987) and Topel (1991).[10] Altonji & Shakotko (1987) address the problem by instrumenting T_{it} with

$$\tilde{T}_{it} \equiv T_{it} - \overline{T}_{ij(i,\,t)}, \tag{19}$$

where

$$\overline{T}_{ij(i,\,t)} \equiv \frac{1}{N_{ij(i,\,t)}} \sum_{t \in \tau_{ij(i,\,t)}} T_{it} \tag{20}$$

is the average value of tenure that one would see for an individual i at job j. τ_{ij} is the set of time periods in the data at which worker i worked for firm j, and N_{ij} is the number of observations in t_{ij}.

This instrument has the nice feature of being uncorrelated with θ_i and $\eta_{ij(i,t)}$ by construction. However, the problem remains that E_{it} is still likely to be positively correlated with $\eta_{ij(i,t)}$. This is problematic. For example, one cannot simply construct an analogous instrument $\tilde{E}_{it} = E_{it} - \overline{E}_{ij(i,\,t)}$ because it would be equal to $\tilde{T}_{it}$. That is, within a job spell, experience and tenure are perfectly collinear as they each increase linearly with time. The positive correlation between E_{it} and $\eta_{ij(i,t)}$ leads to a positive bias on β. Because E_{it} and T_{it} are positively correlated, this leads to a negative bias on α.

Altonji & Shakotko (1987) estimate the model in the Panel Study of Income Dynamics (PSID) and include higher-order terms in the model as well as a dummy for the first year on the job. They provide a range of estimates that suggest that the effect of 10 years of seniority on log wages is less than 0.10 (which is less than half the ordinary-least-squares result). They also model the bias from the positive correlation between E_{it} and $\eta_{ij(i,t)}$ and show that under reasonable assumptions on the magnitude of this correlation, this should have a small overall impact on the results.

Topel (1991) takes a different approach and gets a different answer. He notices that for people who do not change jobs,

$$\begin{aligned} \log(w_{it}) - \log(w_{it-1}) &= \beta(E_{it} - E_{it-1}) + \alpha(T_{it} - T_{it-1}) + \varepsilon_{it} - \varepsilon_{it-1} \\ &= \beta + \alpha + \varepsilon_{it} - \varepsilon_{ijt-1}. \end{aligned} \tag{21}$$

Thus one can get a consistent estimate of $\beta + \alpha$ by first differencing log wages and taking means for job stayers. Therefore, if we knew β, we could construct a consistent estimate of α.

To estimate β, Topel (1991) considers new hires for which by definition $T_{it} = 0$. Thus for new hires, the wage equation is

$$log(w_{it}) = \beta E_{it} + \theta_i + \eta_{ij(i,\,t)} + \varepsilon_{ijt}. \tag{22}$$

He then suggests running a regression of $\log(w_{it})$ on E_{it} for new hires, and then one can estimate α as

$$\hat{\alpha} = \widehat{\beta + \alpha} - \hat{\beta}. \tag{23}$$

[10]Another important paper is Abraham & Farber (1987).

As in Altonji & Shakotko (1987), if $\hat{\beta}$ is biased upward, $\hat{\alpha}$ will be biased downward. Thus Topel also interprets his estimate as a lower bound of the effect.[11] He implements the second stage by using this approach in the PSID and actually finds quite different results with returns to seniority after 10 years somewhere in the 0.25 range.

Topel (1991) attempts to explain why his results differ from Altonji & Shakotko (1987). Altonji & Williams (2005) also try to reconcile the results. They show that the differences arise from three factors. The first is how the overall economy-wide trends in wages are modeled. The second is precisely how tenure is measured—an issue because the PSID contains annual data, so tenure and wages are measured only at an annual level. The third is that, as Topel recognizes, there is an upward bias in his estimator from the second stage because θ_i is likely to be negatively correlated with E_{it} for low-wage workers. Altonji & Williams (2005) argue that there is no perfect way to address all these problems but that the best approach leads to an intermediate value that 10 years of experience increases log wages by approximately 0.11.

3.3. Occupation- and Industry-Specific Human Capital

Although incorporating firm-specific human capital into the model is an improvement over purely general human capital, it is not enough. Few skills are either purely general or purely firm specific. Fortunately, occupational and industry codes are commonly available in data. For this reason, a large literature has expanded the empirical specification (Equation 18) to allow human capital to be either occupation or industry specific. These include essentially three different types of studies. The first recognizes that people have different skills in different occupations. This is the essence of the classic Roy (1951) model in which villagers choose whether to be fisherman or hunters. Heterogeneity in skills leads to a self-selection problem. The second type either explicitly models the accumulation of occupation/industry skill or measures the loss in wages when workers switch industries or occupations. The third focuses on the demand side and allows industries and occupations to enter the aggregate production function as different inputs (either as two inputs into one production function or as inputs into different production functions). We focus mostly on the last two types and discuss them separately, even though there is obviously some overlap.

Much of the empirical literature attempting to estimate models with occupation- and industry-specific capital suffers from the same econometric problems as the returns to firm seniority literature we discuss in Section 3.2. To distinguish between increases in productivity by tenure and self-selection on match quality, they commonly adapt the econometric methods we summarize there.

3.3.1. Evidence on supply: industry specific. One of the earliest extensions of the general/specific human capital distinctions was industry-specific human capital; that is, workers have some endowment of human capital that is only productive within a certain industry. A significant number of early papers used this framework to think about interindustry wage differentials, such as Carrington (1993), Carrington & Zaman (1994), Kim (1998), Neal (1998), and Kletzer (1989). Here we instead focus on papers that explore the importance of industry-specific capital for wage dynamics.

[11]He actually does something related, but a bit different because new workers will not have exactly zero years of tenure and he wants to use data on more than just new workers. Instead he uses the model $w_{it} - \widehat{\alpha + \beta} T_{it} \approx \beta E^0_{ij(i,t)} + \theta_i + \eta_{ij(i,t)} + \varepsilon_{it}$, where E^0_{ij} is experience at the start of the job that is held at time t.

Neal (1995) finds that much of the wage losses for displaced workers can be attributed to interindustry moves instead of firm-specific human capital. He considers the following thought experiment: Take a group of workers and randomly reassign them from their current jobs to other jobs. Some will be assigned to the same firm and thus remain in the same industry, some will be assigned to different firms within the same industry, and others will be assigned to firms in different industries. Using notation similar to Equation 18, suppose wages are determined according to

$$\log(w_{it}) = \beta E_{it} + \alpha_f T_{it}^f + \alpha_I T_{it}^I + u_{it}, \tag{24}$$

where i and t continue to index individuals and time, respectively. The error term u_{it} presumably includes firm and industry match components, but we do not explicitly define them. Assuming that β, α_f, and α_I are positive, workers who are randomly reassigned to the same firm will experience smaller wage losses than those who are assigned to a different firm in the same industry, who in turn will have less losses than those who switch both industry and firm, and the differences in average wages among these groups will be consistent estimates of β, α_f, and α_I. Neal uses a variety of econometric methodologies to correct for the selection of workers into firms and industries to attempt to recover estimates of the relative returns to industry tenure versus firm tenure. Using the Displaced Worker Survey to focus on individuals who lost their jobs through plant closings, he finds that the wage losses among industry switchers are significantly higher than those for stayers, and that the differences are increasing in predisplacement job tenure and experience. He interprets these findings as evidence that industry-specific human capital plays an important role in wage growth.

Parent (2000) considers the same question with the PSID and the NLSY79, which allow him to directly measure industry tenure. He adopts an econometric framework analogous to Equation 24 and uses Altonji & Shakotko's (1987) approach by instrumenting both types of tenure with the deviation between it and its spell-specific mean. His results are consistent with Neal in that they show a large return to industry-specific tenure and a small return to firm-specific tenure. Dustmann & Meghir (2005) use administrative data from Germany and make use of workers who lose their jobs through firm closures. They develop a control function approach that makes use of these data. Although results vary among skill types, they find relatively large returns to general experience and firm tenure, and generally small returns to sector-specific experience.

Weinberg (2001) examines the same basic problem in a very different way by using industry-specific demand shocks. He finds that younger workers respond more strongly to these negative shocks. He also finds that the wages of displaced workers fall more when the demand in their predisplacement industry is lower. Both these facts are consistent with industry-specific human capital.

3.3.2. Evidence on supply: occupation and career specific. The idea that human capital is occupation specific is not new as it is very much at the heart of Adam Smith's (1776) famous example of the pin factory and also a key aspect of Friedman & Kuznets (1954). However, it is only relatively recently that economists tried to quantify this effect.

Kambourov & Manovskii (2009) extend Parent (2000) by accounting for occupation-specific tenure as well. Their specification is analogous to

$$log(w_{it}) = \beta E_{it} + \alpha_f T_{it}^f + \alpha_I T_{it}^I + \alpha_o T_{it}^o + u_{it}, \tag{25}$$

where everything is defined as in Equation 24, but now T_{it}^{o} is occupational tenure. Following Altonji & Shakotko (1987), they use the PSID and instrument each type of tenure with the deviation between it and its spell-specific mean. Their results show a large return to occupation-specific tenure and a small return to firm-specific tenure. For example, at the two-digit level, their instrumental variable results lead to an eight-year return to occupation-specific human capital of 0.14 log point, but the return to employer-specific skills is 0.0060. The industry results are somewhere in the middle, varying across specifications. For example, at the two-digit level, the effect is 0.02. They conclude that the bulk of specific human capital is best viewed as occupation specific. Sullivan (2010a) estimates a similar specification using the NLSY. His results are less robust. In particular, he finds that how occupational changes within a job are accounted for makes a large difference. If he ignores them, he finds results similar to Kambourov & Manovskii (2009). However, he shows that many of these within-firm occupation changes appear to be real occupational changes rather than just measurement error. More importantly, when he treats them as true occupational changes, he finds that the magnitude of the returns to occupational tenure is similar to the return to industry tenure.

Pavan (2011) is critical of Altonji & Shakotko's (1987) approach. He points out that Sullivan (2010a) actually finds a negative return to firm-specific human capital, which he replicates in the NLSY. He develops a structural model based on Neal's (1999) concept of career-specific human capital in which a career is defined as a (three-digit) industry/occupation pair. Thus if one switches occupations but not industries, his career does not change. Pavan estimates a model with general human capital, firm-specific human capital, and career-specific human capital. He also allows for career learning as workers are unsure of their optimal career match. He finds that general skills, career-specific skills, firm-specific skills, and job shopping are all important components of wage growth. He also simulates the data from his model and applies the generalized Altonji & Shakotko (1987) approach. Most troubling for the above method, he finds a negative coefficient on firm-specific tenure despite the fact that this return is positive in his model. He concludes that one needs a structural approach to deal with the selection problem.

The first such approach was pioneered by Keane & Wolpin (1997). They use the NLSY to estimate a dynamic discrete choice model in which workers choose whether to provide their labor to either blue- or white-collar occupations. With only two occupations, they also allow human capital in one occupation to affect skills in the other, so this differs slightly from the standard interpretation of occupational-specific human capital, but it is similar in spirit. Their model allows them to look at the interaction between schooling and occupation choice. In their baseline model, they find that an additional year of schooling increases white-collar skill by 9.4% and blue-collar skill by 1.9%. They also find that an additional year of white-collar experience increases white-collar skill by 11.7%, and the first year of blue-collar experience augments blue-collar skill by 14%, with some attenuation as workers get older. Sullivan (2010b) extends their model to five occupations, includes search, and also includes firm-specific human capital. He finds that eliminating firm-specific and occupation-specific capital on wages would lead them to decline by 2.8%. This is small when compared with the wage changes that would arise if workers were not allowed to shop for better firms and occupations.

Occupation-specific human capital is typically modeled under the assumption that the specific human capital stocks cannot be transferred across occupations. Shaw (1984) hypothesizes that, rather than occupational tenure, the amount of skill that is transferable across occupations is the relevant way to think about human capital across occupations. Her paper provides a bridge between the occupation-specific human capital literature and the task-specific case we discuss in the next section. She attempts to proxy for the task similarity across occupations, even though her paper was written many years before the rest of the literature developed. She augments the traditional Mincer equation with an "occupational investment" term

$$log(w_{it}) = \beta E_{it} + \alpha T_{it}^{f} + \lambda O_{it} + \varepsilon_{it}, \tag{26}$$

implying that worker i at time t will earn wages depending on her firm tenure T_{it}^{f}, overall experience E_i, and also occupational investment O_{it}. The key innovation is this final variable, which is not defined simply as occupational tenure, but as a weighted sum of all past investment, in which the weights depend on how substitutable skills are between the two occupations,

$$O_t \equiv \gamma_{t-1} O_{t-1} + \gamma_{t-2} O_{t-2} + \ldots, \tag{27}$$

and the γ are weights that determine how much skills are substitutable between the occupation chosen in that period and the current occupation. If the worker chooses occupations that share no skills, all the weights would be zero. Using proxies for the γ, Shaw (1984) shows that her measure of occupational investment is a stronger predictor of wages than both overall labor market experience and current occupational tenure. As her paper was written before the data sets on tasks became more widely used and the econometric methods were developed, there is no explicit allowance for tasks in this empirical model and selection is an issue. However, the motivation behind the concept of occupational investment is nearly identical to the more recent focus on multidimensional skills and predates it considerably.

3.3.3. Evidence on labor demand. The studies from the supply side emphasize within-worker wage growth over their careers as a function of their firms or occupational tenures. On the demand side, changes in the relative prices of skills combined with information on the skill composition of firms or industries can provide information about the importance of occupation- or industry-specific skills. The seminal paper in this literature is Heckman & Sedlacek (1985). They implement a general equilibrium version of the Roy (1951) model in which they use a task-specific framework to model occupational and tenure skills. They specify and estimate different supply and demand conditions for the manufacturing and nonmanufacturing sectors. Heckman et al. (1998a) extend a similar framework into a dynamic environment in which workers explicitly invest in skills and workers are not perfect substitutes. Heckman et al. (1998b,c) use the estimates of this model to show the importance of incorporating heterogeneous human capital into labor demand. For example, Heckman et al. (1998b) show that one can overstate the impact of a tuition subsidy on college by an order of magnitude if one treats college- and high school–educated individuals as perfect substitutes.

Heckman et al. (1998a) do not distinguish human capital by industry or occupation, but rather by education. Lee (2005) and Lee & Wolpin (2006) use a related approach but model heterogeneity by occupation. Lee (2005) distinguishes between white- and blue-collar occupations. Lee & Wolpin (2006) add pink collar as well and allow for different industries by estimating different production functions in the goods and service sectors.

This approach is extended even further by Keane & Johnson (2012), who allow separate entries by occupation, age, gender, and education.

Hu & Taber (2011) use a related model, but a different approach. They do not explicitly model industries and occupations but consider a model with multiple sectors and multiple types of human capital within each sector. One can think of the sectors as industries and the human capital types as occupations. Although all sectors hire all skill types, they differ in their composition of skill intensity. For example, both Facebook and a law firm likely have computer programmers and lawyers, but Facebook will have a higher concentration of programmers and a law firm will have a higher concentration of lawyers. This framework in some sense combines the model of Heckman & Sedlacek (1985) with that of Heckman et al. (1998a) and its predecessors, in the sense that Heckman & Sedlacek have multiple sectors with one type of skill in each sector, whereas Heckman et al. have essentially one sector with multiple worker types within it. By modeling these separately, Hu & Taber (2011) can distinguish between occupation or skill-type productivity shocks and industry or sector shocks. They simulate these different types of shocks in their model and show that one would expect sector-specific shocks to have smaller effects on displaced workers than shocks substitutable with skill type. Using the Current Population Survey Displaced Workers supplement, they find evidence supporting this model.

3.4. Task-Specific Human Capital

As mentioned above, we view work on occupation and industry human capital as being motivated by data availability rather than as the natural way to classify human capital. In this section, we discuss the most natural extension of the specific human capital concept, task-specific human capital. Although the explicit formalization goes back at least to Heckman & Sedlacek (1985), it is only recently that it has been widely adopted.[12] Studies of occupation-specific human capital often note that the classifications of occupation and industry may be arbitrary, but they are proxies for the similarity of underlying tasks performed and skills required across different firms. In this framework, a worker's multidimensional skills can be exploited differently across different firms, occupations, and industries. If a particular skill is used significantly across all firms, occupations, and industries, it has the characteristics of general human capital. If the skill is used only by one firm, it is firm-specific human capital. Multidimensional skills allow for a more general framework that can incorporate other substitution patterns as well. One of the major factors driving interest in task-specific human capital is the development of data sets on worker tasks and skills; we summarize those data sets in Section 3.4.1. In Section 3.4.2, we discuss the evidence that a multidimensional skill framework can fit the data on worker choices and outcomes better than the traditional firm-specific/general human distinction, or even the occupation- and industry-specific human capital literature. In Section 3.4.3, we discuss direct estimates of skill accumulation rates and their implications for life-cycle wage dynamics.

Although the focus of this review is the role of multidimensional skills in wage growth, the concept of task-specific human capital has been applied in a broad literature attempting

[12]Gibbons & Waldman (2004, 2006) analyze a model with what they call task-specific human capital. However, this is not the same task-specific capital vector as we discuss here, as theirs is related to a specific task a firm assigns a worker and is destroyed when the worker is moved to a new position within the firm.

to explain time trends in the relative wages and labor supply choices of different skill groups. For example, what are the effects of the increasing importance of computerization in the labor market for workers whose comparative advantages lay in manual skills? How have changes in the prices of different skills affected the wage gaps between workers of different education levels? Many papers have focused on the effects of what is called skill-biased technical change to explain these changes in the structure of the demand for labor. Although this literature shares the same focus on the heterogeneous nature of human capital as this article, a survey of it is beyond our scope. Some important papers that address the changes in returns to specific skill categories are Autor et al. (2003), Spitz-Oener (2006), and Ingram & Neumann (2006) (for a broader review of the literature, see Acemoglu & Autor 2011). For the rest of this section, we focus on the literature studying life-cycle wage growth.

3.4.1. Data. Without data on tasks performed or individual skills, it is difficult to get any direct estimates of the returns to individual skills. This is perhaps the primary reason research had focused on categories such as occupation or industry because those measures are readily available in standard data sets. Over the past few years, researchers using a series of data sets of specific worker tasks moved the literature away from focusing on the occupation/industry classification and instead toward modeling and estimating the underlying task-specific human capital. It is useful to understand the strengths and weaknesses of the publicly available data on worker tasks to see how the human capital literature has developed. The primary data set that contains occupation-level data from the United States is the Dictionary of Occupational Titles. Outside the United States, there are major occupation-level-task data sets from Germany (the German Qualification and Career Survey) and the United Kingdom (the British Skills Survey). Because all the surveys contain similar information across different countries, in the interest of brevity we summarize the Dictionary of Occupational Titles (for a paper that uses all three data sets to look at task-specific human capital, see Borghans et al. 2006).

The Dictionary of Occupational Titles and its follow-up, O*NET (henceforth we refer to both as the DOT), are administered by the US Bureau of Labor Statistics. The DOT provides cross-sectional measures from the United States of the types of tasks performed and skills required at the occupation level. The surveys are administered to a representative sample of workers. The DOT has experienced a number of iterations, with the fourth (and final) revised edition published in 1991. After its publication, the Bureau of Labor Statistics scrapped the original DOT framework and began a program to build its successor O*NET from the ground up. This was in part because of problems with the original DOT design (for a detailed discussion of the design and issues with the DOT, see Miller et al. 1980).

The data contents of the DOT involve occupation codes linked with descriptions of the tasks performed in the occupation. Composite scores of the occupation on a number of different criteria are formed. The primary classification for occupations is the requirements for working with “data,” “people,” and “things.” These categories necessarily have some element of subjectivity, but in general the “data” category relates to the necessity of processing and using information, “people” involves the necessity of relating to others in the course of work, and “things” relates to the ability to use/manipulate physical objects.

O*NET largely revamped the structure of the DOT while maintaining its core purpose of tracking the skill usage of US occupations. In particular, the data-gathering and -summarizing procedure is explicitly stated, which makes the data easier to use and interpret. Moreover, there has been an attempt to move away from expert scoring of occupations toward self-reported survey data about tasks and skill usage by workers. With O*NET there is also a good deal more data released to the public: Instead of a score for "data," the scores for many of the individual component questions are released (e.g., "How much Mathematics do you use on your job?" "Does your job require manipulation of precise equipment?"). This allows the researcher to decide how to reduce the data into the scores he is interested in.

Both the DOT and O*NET have weaknesses that other data sets can potentially address. Even though they have gone through a number of revisions, they are only cross-sectional. The DOT also suffers from some issues of data reliability and representativeness. Also, as in most of these data sets, it is difficult to separate occupational requirements from worker skills, which are difficult concepts to distinguish even theoretically. Most papers using these data sets attempt to use them as occupational requirements by choosing the appropriate questions as data.

3.4.2. Evidence for task-specific human capital. We believe that the best (currently available) empirical analog to multidimensional human capital is task-specific human capital. There is a large literature on worker skills touching on disparate questions that provide evidence for this claim. In this section, we focus on a few key pieces of evidence that tasks performed and skills required by jobs can explain residual wage gaps, occupation and industry choices, and job mobility better than models that focus only on discrete occupation, industry, and career categories. Then in Section 3.4.3 we address the relationship among skill investment, job transitions, and wage growth that was the focus of the theory developed above.

The residual wage gap is an important topic in labor economics; conditioning on a set of "standard" observable characteristics (age, experience, education, cognitive test scores) cannot explain significantly more than 30% of variation in individual wages [see Mortensen (2005), who focuses on search frictions explaining this rather than unobserved human capital]. Before the use of data on noncognitive skills, typically the maximum allowed amount of across-worker heterogeneity in skills was a one-dimensional "ability" vector. However, a number of studies have argued that distinguishing between cognitive and noncognitive skills (particularly interpersonal skills) is useful for explaining these wage gaps. For example, Heckman & Rubinstein (2001) provide evidence that high school dropouts who complete their high school equivalency (GED) in the United States earn higher wages than those who simply drop out but that the difference is explained by differences in cognitive ability. Conditional on Armed Forces Qualification Test scores (which proxy for cognitive skills), those with GEDs actually earn lower wages than other dropouts. Heckman & Rubinstein show that differences in noncognitive skills can explain this residual wage gap.

Reinforcing this hypothesis, Heckman et al. (2006b) posit a two-dimensional cognitive/noncognitive human capital vector and show that their measures of noncognitive skills (drawn from test scores in the NLSY79) are important for schooling choices, wages, employment, work experience, and occupational choice. A standard Mincer-style regression of wages regressed on cognitive and noncognitive test scores (along with the standard covariates) indicates that cognitive test scores explain approximately 12% of variation in wages, but only 0.4% of wage variation is explained by noncognitive test scores. However, using their factor structure, they find that the true marginal effects on wages of moving individuals around in the noncognitive skill distribution are the same as moving them in the cognitive skill distribution. The essence of their empirical methodology is the use of multiple test scores for cognitive and noncognitive skills linked with labor market outcomes and a model allowing for endogeneity of schooling choice and wages.

There is no consensus on what exactly these "noncognitive" skills measure; Heckman et al. (2006b) remain agnostic about what exactly these skills are and let the data back out the factor loadings on different test scores that correspond to noncognitive skills. Another approach is to look at the effect of one particular component of noncognitive skills.[13] For example, Borghans et al. (2006) take a specific set of skills, what they call "people skills," and argue that the ability to relate to and influence others is becoming a larger factor in wages and occupational choices over time. In their empirical work, they use a variety of data sources (including the DOT, British Skills Survey, and German Qualification and Career Survey discussed above) to document that people skills are systematically related to occupational choices and wages. For example, they show that early lifetime measures of increased sociability are associated with workers taking jobs that require them to interact more with people later in life. In addition, although jobs that require interaction with people tend to pay lower wages, the wage effects of sociability are largest for those who are in jobs with higher amounts of interaction.

Rather than looking at direct measures of skills, researchers often find evidence of task-specific human capital indirectly. For example, worker transitions can provide evidence that task-specific human capital is important for worker outcomes. Task-specific human capital has predictions about patterns of job transitions that the other models do not. For example, occupation-specific human capital models do not have much to say about what happens after a worker leaves an occupation: If a worker switches occupations, her occupation-specific human capital should be destroyed. However, with task-specific human capital, it makes sense that conditional on switching occupations, workers will tend to work in occupations that have similar overall task requirements so they can make the most efficient use of their current skills. Some of the strongest evidence that occupation-specific skills are actually proxies for task-specific skills comes from looking at the nature of transitions across occupations.

Poletaev & Robinson (2008) use four occupation-level measures of skills taken from the DOT and create a skill portfolio measure they use to measure the distance between jobs. They focus on workers whose firms close and document the wage losses associated with mobility across different distances in the task space. They find that workers who switch to occupations that are more "distant" in terms of tasks have larger wages losses. They interpret this pattern as evidence for a task-based component of human capital. Robinson (2010) extends the use of these distance measures to look at the evolution of mobility across tasks over time. He finds that since the 1970s, a decreasing fraction of occupational moves involve large changes in the skill portfolio. Decomposing by the type of switch, the distance of voluntary switches is increasing and the distance of involuntary searches is declining. Both are consistent with a pattern of increased access to upward mobility and reduced search costs. Using a similar method of characterizing occupations as a point in a multidimensional task space, Yamaguchi (2010) estimates a model in which workers can switch occupations over time to sort into different tasks based on their comparative advantage and shows that workers with high experience and education are better off in

[13]Interpreting the components of either cognitive or noncognitive skills in terms of other underlying factors is an active area of research, often drawing from the psychology literature (see Almlund et al. 2011, Borghans et al. 2011).

high-cognitive and interpersonal task jobs and that estimates of the returns to skills that do not correct for selection suffer from biases.

Consistent with these studies, Gathmann & Schoenberg (2010) use task data from Germany (the German Qualification and Career Survey) and also find that individuals are more likely to move to occupations with similar task requirements. They also construct a distance measure between occupations that is based on the similarities of reported tasks and use it to create a direct uni-dimensional measure of "task tenure." They show that the distance of moves across tasks decreases over time as predicted by a task-specific human capital model. In addition, wages after occupation changes are more highly correlated when the two occupations share more tasks. This is in stark contrast to the occupation-specific capital framework, which would predict no correlation in wages after a switch. They then quantify the importance of task-specific human capital for wage growth using wage changes of displaced workers within a control function approach. When they include task tenure along with occupational tenure and overall labor market experience and instrument for tenure as in Parent (2000), they find that the returns to task tenure exceed those of the other types of tenure in a wide variety of ordinary-least-squares and instrumental variables specifications. They also directly estimate the importance of task tenure for wage growth and find that task-specific human capital accounts for up to 52% of wage growth. As discussed above, Kambourov & Manovskii's (2009) and Parent's (2000) instrumental variables methods have been shown to fail in some settings by Pavan (2011). However, this is just one piece of evidence in their paper. The majority of evidence for the importance of task-specific human capital in Gathmann & Schoenberg (2010) is not based on this procedure.

There has been relatively less work attempting to develop formal tests of the task-specific human capital hypothesis. This may be partially because the complexity of the full theoretical framework shows that it is difficult to find clear empirical predictions because of the large number of typically unobservable variables, such as underlying skills and time spent investing. Autor & Handel (2009) use individual-specific task data developed by Princeton Data Improvement Initiative to test the implications of a task-specific self-selection model. Their data allow them to look at heterogeneity in tasks at the individual level within occupations, and they find that there is a significant amount of within-occupation task heterogeneity. They also develop an empirical model of self-selection based off the Roy (1951) model and test some of its reduced-form implications and find the data to be broadly consistent with the model.

3.4.3. Task-specific human capital accumulation. Given the existence and importance of multidimensional skills, a natural question is how workers develop the different types of human capital over their lifetime. As is clear from our model, the motivation of much of the standard human capital literature is to understand investment, but with multidimensional skills, all types of investment are no longer the same. Almost certainly, college augments cognitive skills more than noncognitive skills (for nonathletes at least), weight training increases noncognitive over cognitive skills, and surgical training that requires both precision and problem-solving increases both. To understand the effects of skill-augmenting technologies on labor market outcomes, one must estimate the skill-specific rates of human capital accumulation.

There have been two primary approaches to this problem, one using direct measurement of worker skills and the other using measurement of worker tasks to back out their underlying skill levels. The primary study for the first approach is Cunha et al. (2010), who use multiple test scores on cognitive and noncognitive skills and a dynamic factor structure to measure the effects of parental investment in the skill human capital of children. They estimate full production functions for the childhood development of their skills and link them to later labor market outcomes and behavioral problems such as crime. They find that it is easiest to rectify bad environments early in life rather than later when it comes to cognitive skills, but it is easier to fix a disadvantage in noncognitive skills later in childhood. Because they estimate a full model of investment and outcomes, they are able to run simulations to uncover the investment that a social planner would choose to maximize human capital accumulation or to minimize the crime rate.

The second approach is estimation of a dynamic hidden Markov model in which the worker's underlying skills are unobserved but the choice of tasks is observed. The observed choices over time, along with additional assumptions, are used to back out the underlying skills and their transitions over time. The first paper to estimate a model of this type is Yamaguchi (2012). He finds that the importance of human capital accumulation across skills differs by education level. Cognitive skill accumulation has large effects on wages for both college-educated workers and high school graduates. However, cognitive skills grow faster for the college educated than high school graduates, which can account for the divergence in wages between the two groups over the life cycle. By contrast, noncognitive skill accumulation has only a small influence on wages for college graduates and high school graduates but is responsible for a large amount of the wage growth of high school dropouts. Sanders (2012) adds worker uncertainty about their skills to the model and finds similar effects.

Pitt et al. (2010) use a similar framework but are able to access data on both individual-level skills and occupational tasks performed. They study the effects of investments in health versus schooling in developing countries. In particular, if men have a comparative advantage in manual skills versus cognitive skills due to biological factors, the choice of whether to invest in schooling or caloric consumption (which would increase manual skills through strength) leads to a problem of multidimensional human capital accumulation. Pitt et al. create a model of investment and self-selection and estimate it through instrumental variable procedures. They show that, as predicted by the Roy (1951) model, the measured marginal returns to investment differ greatly in the population, and the standard log-linear wage specification with a fixed coefficient on experience cannot accurately fit the data.

Little work has connected the early lifetime investments by parents and the later worker investments in their own human capital. Conceptually, the early lifetime focus of Cunha et al.'s (2010) methods, using early life measurements and test scores to measure individual-level skills, could be connected with Yamaguchi's (2012) approach, using observed task choices to back out unobserved individual skills. In addition, almost no empirical literature has looked at the relationship between firms and skill accumulation. Although the literature documenting the effects of skill-specific human capital has had success explaining a large number of short-run outcomes (such as what job to take after a plant closing) or long-term choices (such as career choices), the dynamics of skill accumulation combined with firm transitions are still not well understood and make up the most promising area for future research.

4. CONCLUSION

In this article we describe both the theory and empirical work behind life-cycle wage growth and heterogeneous skills. We begin with a model of general human capital based on the Ben-Porath (1967) model. We extend it to include search frictions and a multi-dimensional vector of human capital that allows for arbitrary substitution patterns across jobs. We then survey the empirical work beginning with canonical models of general and firm-specific human capital. We discuss the extensions of occupation-, industry-, and task-specific human capital. We show that once tenure and occupation-specific human capital are accounted for, there is little evidence of firm-specific human capital. Furthermore, the task-based approach fits patterns of wage growth and mobility that cannot be explained by human capital that is either purely occupation specific or purely industry specific.

The current literature's focus on task-specific human capital is the natural extension of older models that focused on firms, occupations, and industries, and the empirical work performed to date shows that it is potentially an important component of lifetime wage growth. However, there are still a number of interesting directions to be investigated. Typically task-specific models have abstracted completely from firms, but we show that the interaction between firms and tasks can lead to interesting behavior in the presence of frictions. Empirical work that focuses on the relationship among frictions, firms, and skill investment has promise for understanding the importance of skills on wage growth during firm transitions. Another active line of research involves breaking down the skill vector even further, beyond the common two-dimensional cognitive/manual or cognitive/noncognitive distinction. Modern data sets have a significant amount of more detailed information on worker tasks, but typically for simplicity and clarity, studies have reduced the skill vector to two dimensions. However, it seems plausible that there are many types of cognitive skills that are not fully substitutable across jobs (e.g., math skills and writing skills). Looking more carefully at these distinctions can better explain worker job choices and wage growth.

SUMMARY POINTS

1. Although pedagogically useful, the treatment of skills as either purely general or purely firm specific is not realistic.
2. We develop a model with search frictions and heterogeneous human capital. We show that workers can sometimes specialize in specific investments or, in other cases, prefer broader investments.
3. Task-specific human capital seems more intuitively appealing than either purely industry-specific or occupation-specific human capital. Task-specific human capital fits the data on wage growth and mobility patterns better than previous models.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We thank Jim Heckman, John Kennan, Derek Neal, Nicholas Roys, and several referees for comments. An earlier draft of this article constituted the third chapter of Sanders' dissertation at the University of Wisconsin-Madison. We take responsibility for any remaining errors.

LITERATURE CITED

Abraham K, Farber H. 1987. Job duration, seniority, and earnings. *Am. Econ. Rev.* 77:278–97

Acemoglu D, Autor D. 2011. Skills, tasks and technologies: implications for employment and earnings. In *Handbook of Labor Economics*, Vol. 4, ed. O Ashenfelter, D Card, pp. 1043–171. Amsterdam: North Holland

Acemoglu D, Pischke J. 1998. Why do firms train? Theory and evidence. *Q. J. Econ.* 113:79–118

Almlund M, Duckworth A, Heckman J, Kautz T. 2011. Personality psychology and economics. In *Handbook of the Economics of Education*, Vol. 4, ed. EA Hanushek, S Machin, L Woessmann, pp. 1–181. Amsterdam: North Holland

Altonji J, Shakotko R. 1987. Do wages rise with seniority? *Rev. Econ. Stud.* 543:437–59

Altonji J, Williams N. 2005. Do wages rise with job seniority? A reassessment. *Ind. Labor Relat. Rev.* 58:370–97

Antonovics K, Golan L. 2012. Experimentation and job choice. *J. Labor Econ.* In press

Autor D, Handel M. 2009. *Putting tasks to the test: human capital, job tasks and wages*. NBER Work. Pap. 15116

Autor D, Levy F, Murnane R. 2003. The skill content of recent technological change: an empirical exploration. *Q. J. Econ.* 118:1279–333

Becker G. 1964. *Human Capital*. Chicago: Univ. Chicago Press

Ben-Porath Y. 1967. The production of human capital and life cycle of earnings. *J. Polit. Econ.* 75:352–65

Borghans L, Golsteyn B, Heckman J, Humphries J. 2011. Identification problems in personality psychology. *Pers. Individ. Differ.* 51:315–20

Borghans L, ter Weel B, Weinberg B. 2006. *People people: social capital and the labor-market outcomes of underrepresented groups*. NBER Work. Pap. 11985

Browning M, Hansen L, Heckman J. 1999. Micro data and general equilibrium models. In *Handbook of Macro Economics*, Vol. 1, ed. JB Taylor, M Woodford, pp. 525–602. Amsterdam: North Holland

Burdett K, Coles M. 2010. *Tenure and experience effects on wages: a theory.* CESifo Work. Pap. 3202, Munich

Burdett K, Mortensen D. 1998. Wage differentials, employer size, and unemployment. *Int. Econ. Rev.* 39:257–73

Cahuc P, Postel-Vinay F, Robin J. 2006. Wage bargaining with on-the-job search: theory and evidence. *Econometrica* 74:323–64

Carrington W. 1993. Wage losses for displaced workers: Is it really the firm that matters? *J. Hum. Resour.* 28:435–62

Carrington W, Zaman A. 1994. Interindustry variation in the costs of job displacement. *J. Labor Econ.* 12:243–75

Cunha F, Heckman J, Schennach S. 2010. Estimating the technology of cognitive and noncognitive skill formation. *Econometrica* 78:883–931

Dustmann C, Meghir C. 2005. Wages, experience and seniority. *Rev. Econ. Stud.* 72:77–108

Felli L, Harris C. 1996. Learning, wage dynamics, and firm-specific human capital. *J. Polit. Econ.* 104:838–68

Friedman M, Kuznets S. 1954. *Income from Independent Professional Practice*. New York: NBER

Fu C. 2011. Training, search and wage dynamics. *Rev. Econ. Dyn.* 14:650–66

Gathmann C, Schoenberg U. 2010. How general is human capital? A task-based approach. *J. Labor Econ.* 28:1–50

Gibbons R, Waldman M. 2004. Task-specific human capital. *Am. Econ. Rev.* 94:203–7

Gibbons R, Waldman M. 2006. Enriching a theory of wage and promotion dynamics inside firms. *J. Labor Econ.* 24:59–107

Haley W. 1976. Estimation of the earnings profile from optimal human capital accumulation. *Econometrica* 44:1223–88

Heckman J. 1976. A life-cycle model of earnings, learning, and consumption. *J. Polit. Econ.* 84:S11–44

Heckman J, Lochner L, Cossa R. 2003. Learning-by-doing versus on-the-job training: using variation induced by the EITC to distinguish between models of skill formation. In *Designing Inclusion: Tools to Raise Low-End Pay and Employment in Private Enterprise*, ed. E Phelps, pp. 74–130. New York: Cambridge Univ. Press

Heckman J, Lochner L, Taber C. 1998a. Explaining rising wage inequality: explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents. *Rev. Econ. Dyn.* 1:1–58

Heckman J, Lochner L, Taber C. 1998b. General equilibrium treatment effects: a study of tuition policy. *Am. Econ. Rev.* 88:381–86

Heckman J, Lochner L, Taber C. 1998c. Tax policy and human-capital formation. *Am. Econ. Rev.* 88:293–97

Heckman J, Lochner L, Todd P. 2006a. Earnings functions, rates of return, and treatment effects: the Mincer equation and beyond. In *Handbook of the Economics of Education*, Vol. 1, ed. EA Hanushek, F Welch, pp. 307–58. Amsterdam: North Holland

Heckman J, Rubinstein Y. 2001. The importance of noncognitive skills: lessons from the GED testing program. *Am. Econ. Rev.* 91:145–49

Heckman J, Sedlacek G. 1985. Heterogeneity, aggregation, and market wage functions: an empirical model of self-selection in the labor market. *J. Polit. Econ.* 93:1077–125

Heckman J, Stixrud J, Urzua S. 2006b. The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *J. Labor Econ.* 24:411–82

Hu L, Taber C. 2011. Displacement, asymmetric information and heterogeneous human capital. *J. Labor Econ.* 29:113–52

Imai S, Keane M. 2004. Intertemporal labor supply and human capital accumulation. *Int. Econ. Rev.* 45:601–41

Ingram B, Neumann G. 2006. The returns to skill. *Labour Econ.* 13:35–59

Jovanovic B. 1979. Job matching and the theory of turnover. *J. Polit. Econ.* 87:972–90

Kambourov G, Manovskii I. 2009. Occupational specificity of human capital. *Int. Econ. Rev.* 50:63–115

Keane M, Johnson M. 2012. A dynamic equilibrium model of the U.S. wage structure, 1968–1996. *J. Labor Econ.* In press

Keane M, Wolpin K. 1997. The career decisions of young men. *J. Polit. Econ.* 105:473–522

Kim D. 1998. Reinterpreting industry premiums: match-specific productivity. *J. Labor Econ.* 16:479–504

Kletzer L. 1989. Returns to seniority after permanent job loss. *Am. Econ. Rev.* 79:536–43

Kuruscu B. 2006. Training and lifetime income. *Am. Econ. Rev.* 96:832–46

Lazear E. 2009. Firm-specific human capital: a skill-weights approach. *J. Polit. Econ.* 117:914–40

Lee D. 2005. An estimable dynamic general equilibrium model of work, schooling, and occupational choice. *Int. Econ. Rev.* 46:1–34

Lee D, Wolpin K. 2006. Intersectoral labor mobility and the growth of the service sector. *Econometrica* 74:1–46

Miller AR, Treiman DJ, Cain PS, Roos PA, eds. 1980. *Work, Jobs, and Occupations: A Critical Review of the* "Dictionary of Occupational Titles." Washington, DC: Natl. Acad.

Miller R. 1984. Job matching and occupational choice. *J. Polit. Econ.* 92:1086–120

Mincer J. 1958. Investment in human capital and personal income distribution. *J. Polit. Econ.* 66:281–302

Mincer J. 1974. *Schooling, Experience, and Earnings*. New York: Columbia Univ. Press

Mortensen D. 2005. *Wage Dispersion: Why Are Similar Workers Paid Differently?* Cambridge, MA: MIT Press

Murphy K, Welch F. 1990. Empirical age-earnings profies. *J. Labor Econ.* 8:202–29

Neal D. 1995. Industry-specific human capital: evidence from displaced workers. *J. Labor Econ.* 13:653–77

Neal D. 1998. The link between ability and specialization: an explanation for observed correlations between wages and mobility rates. *J. Hum. Resour.* 33:173–200

Neal D. 1999. The complexity of job mobility among men. *J. Labor Econ.* 17:237–61

Parent D. 2000. Industry-specific capital and the wage profile: evidence from the National Longitudinal Survey of Youth and the Panel Study of Income Dynamics. *J. Labor Econ.* 18:306–23

Pavan R. 2011. Career choice and wage growth. *J. Labor Econ.* 29:549–87

Pitt M, Rosenzweig M, Hassan N. 2010. *Human capital investment and the gender division of labor in a Brawn-based economy.* Work. Pap., Econ. Growth Cent., Yale Univ.

Poletaev M, Robinson C. 2008. Human capital specificity: evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984–2000. *J. Labor Econ.* 26:387–420

Robinson C. 2010. *Occupational mobility, occupation distance and basic skills: evidence from job-based skill measures*. Work. Pap., Univ. Western Ontario

Rosen S. 1972. Learning and experience in the labor market. *J. Hum. Resour.* 7:326–42

Rosen S. 1976. A theory of life earnings. *J. Polit. Econ.* 84:S45–67

Rosen S. 1983. Specialization and human capital. *J. Labor Econ.* 1:43–49

Roy AD. 1951. Some thoughts on the distribution of earnings. *Oxf. Econ. Pap.* 3(2):135–46

Rubinstein Y, Weiss Y. 2006. Post schooling wage growth: investment, search and learning. In *Handbook of the Economics of Education*, Vol. 1, ed. EA Hanushek, F Welch, pp. 1–67. Amsterdam: North Holland

Sanders C. 2012. *Skill uncertainty, skill accumulation, and occupational choice*. Work. Pap., Washington Univ. in St. Louis

Shaw K. 1984. A formulation of the earnings function using the concept of occupational investment. *J. Hum. Resour.* 19:319–40

Smith A. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. London: Methuen & Co.

Spitz-Oener A. 2006. Technical change, job tasks, and rising educational demands: looking outside the wage structure. *J. Labor Econ.* 24:235–70

Stevens M. 1994. A theoretical model of on-the-job training with imperfect competition. *Oxf. Econ. Pap.* 46:537–62

Sullivan P. 2010a. A dynamic analysis of educational attainment, occupational choices, and job search. *Int. Econ. Rev.* 51:289–317

Sullivan P. 2010b. Empirical evidence on occupation and industry specific human capital. *Labour Econ.* 17:567–80

Taber C. 2002. Tax reform and human capital accumulation: evidence from an empirical general equilibrium model of skill formation. *B.E. J. Econ. Anal. Policy* 2(1):3

Topel R. 1991. Specific capital, mobility, and wages: wages rise with job seniority. *J. Polit. Econ.* 99:145–76

Topel R, Ward M. 1992. Job mobility and the careers of young men. *Q. J. Econ.* 107:445–79

Weinberg B. 2001. Long-term wage fluctuations with industry-specific human capital. *J. Labor Econ.* 19:231–64

Yamaguchi S. 2010. Career progression and comparative advantage. *Labour Econ.* 17:679–89

Yamaguchi S. 2012. Tasks and heterogeneous human capital. *J. Labor Econ.* 30:1–53

Annual Review of
Economics

Volume 4, 2012

Contents

Indexes

Errata

An online log of corrections to *Annual Review of Economics* articles may be found at http://econ.annualreviews.org