CHAPTER 8

The Theory of Training Investments

1. General Vs. Specific Training

In the Ben-Porath model, an individual continues to invest in his human capital after he starts employment. We normally think of such investments as "training", provided either by the firm itself on-the-job, or acquired by the worker (and the firm) through vocational training programs. This approach views training just as schooling, which is perhaps too blackbox for most purposes.

More specifically, two complications that arise in thinking about training are:

(1) Most of the skills that the worker acquires via training will not be as widely applicable as schooling. As an example, consider a worker who learns how to use a printing machine. This will only be useful in the printing industry, and perhaps in some other specialized firms; in this case, the worker will be able to use his skills only if he stays within the same industry. Next, consider the example of a worker who learns how to use a variety of machines, and the current employer is the only firm that uses this exact variety; in this case, if the worker changes employer, some of his skills will become redundant. Or more extremely, consider a worker who learns how to get along with his colleagues or with the customers of his employer. These skills are even more "specific", and will become practically useless if he changes employer.

(2) A large part of the costs of training consist of forgone production and other costs borne directly by the employer. So at the very least, training investments have to be thought as joint investments by the firm and the worker, and in many instances, they may correspond to the firm's decisions more than to that of the worker.

The first consideration motivates a particular distinction between two types of human capital in the context of training:

(1) Firm-specific training: this provides a worker with *firm-specific skills*, that is, skills that will increase his or her productivity *only* with the current employer.
(2) General training: this type of training will contribute to the worker's general human capital, increasing his productivity with a range of employers.

Naturally, in practice actual training programs could (and often do) provide a combination of firm-specific and general skills.

The second consideration above motivates models in which firms have an important say in whether or not the worker undertakes training investments. The extreme but not show case is the one where training costs are borne by the firm (for example, because the process of training reduces production), and in this case, the firm directly deciding whether and how much training the worker will obtain may be a good approximation to reality and a good starting point for our analysis.

2. The Becker Model of Training

Let us start with investments in general skills. Consider the following stylized model:

- At time $t = 0$, there is an initial production of y_0, and also the firm decides the level of training τ, incurring the cost $c(\tau)$. Let us assume that $c(0) = 0$, $c'(0) = 0$, $c'(\cdot) \geq 0$ and $c''(\cdot) > 0$. The second assumption here ensures that it is always socially beneficial to have some amount of positive training.
- At time $t = 1/2$, the firm makes a wage offer w to the worker, and other firms also compete for the worker's labor. The worker decides whether to quit and work for another firm. Let us assume that there are many identical firms who can use the general skills of the worker, and the worker does not incur any cost in the process of changing jobs. This assumption makes the

labor market essentially competitive. (Recall: there is no informational asymmetry here).

- At time $t = 1$, there is the second and final period of production, where output is equal to $y_1 + \alpha(\tau)$, with $\alpha(0) = 0$, $\alpha'(\cdot) > 0$ and $\alpha''(\cdot) < 0$. For simplicity, let us ignore discounting.

First, note that a social planner wishing to maximize net output would choose a positive level of training investment, $\tau^* > 0$, given by

$$c'(\tau^*) = \alpha'(\tau^*).$$

The fact that τ^* is strictly positive immediately follows from the fact that $c'(0) = 0$ and $\alpha'(0) > 0$.

Before Becker analyzed this problem, the general conclusion, for example conjectured by Pigou, was that there would be underinvestment in training. The reasoning went along the following lines. Suppose the firm invests some amount $\tau > 0$. For this to be profitable for the firm, at time $t = 1$, it needs to pay the worker at most a wage of

$$w_1 < y_1 + \alpha(\tau) - c(\tau)$$

to recoup its costs. But suppose that the firm was offering such a wage. Could this be an equilibrium? No, because there are other firms who have access to exactly the same technology, they would be willing to bid a wage of $w_1 + \varepsilon$ for this worker's labor services. Since there are no costs of changing employer, for ε small enough such that

$$w_1 + \varepsilon < y_1 + \alpha(\tau),$$

a firm offering $w_1 + \varepsilon$ would both attract the worker by offering this higher wage and also make positive profits. This reasoning implies that in any competitive labor market, we must have

$$w_1 = y_1 + \alpha(\tau).$$

But then, the firm cannot recoup any of its costs and would like to choose $\tau = 0$. Despite the fact that a social planner would choose a positive level of training investment, $\tau^* > 0$, the pre-Becker view was that this economy would fail to invest in training.

The mistake in this reasoning was that it did not take into account the worker's incentives to invest in his own training. In effect, the firm does not get any of the returns from training because the worker is receiving all of them. In other words, the worker is the *full residual claimant* of the increase in his own productivity, and in the competitive equilibrium of this economy without any credit market or contractual frictions, he would have the right incentives to invest in his training.

Let us analyze this equilibrium now. As is the case in all games of this sort, we are interested in the subgame perfect equilibria. So we have to solve the game by backward induction. First note that at $t = 1$, the worker will be paid $w_1 = y_1 + \alpha(\tau)$. Next recall that τ^* is the efficient level of training given by $c'(\tau^*) = \alpha'(\tau^*)$. Then in the unique subgame perfect equilibrium, in the first period the firm will offer the following package: training of τ^* and a wage of

$$w_0 = y_0 - c(\tau^*).$$

Then, in the second period the worker will receive the wage of

$$w_1 = y_1 + \alpha(\tau^*)$$

either from the current firm or from another firm.

To see why no other allocation could be an equilibrium, suppose that the firm offered (τ, w_0), such that $\tau \neq \tau^*$. For the firm to break even we need that $w_0 \leq y_0 - c(\tau)$, but by the definition of τ^*, we have

$$y_0 - c(\tau^*) + y_1 + \alpha(\tau^*) > y_0 - c(\tau) + y_1 + \alpha(\tau) \geq w_0 + y_1 + \alpha(\tau)$$

So the deviation of offering $(\tau^*, y_0 - c(\tau^*) - \varepsilon)$ for ε sufficiently small would attract the worker and make positive profits. Thus, the unique equilibrium is the one in which the firm offers training τ^*.

Therefore, in this economy the efficient level of training will be achieved with firms bearing none of the cost of training, and workers financing training by taking a wage cut in the first period of employment (i.e, a wage $w_0 < y_0$).

There are a range of examples for which this model appears to provide a good description. These include some of the historical apprenticeship programs where young individuals worked for very low wages and then "graduated" to become master craftsmen; pilots who work for the Navy or the Air Force for low wages, and then obtain much higher wages working for private sector airlines; securities brokers, often highly qualified individuals with MBA degrees, working at a pay level close to the minimum wage until they receive their professional certification; or even academics taking an assistant professor job at Harvard despite the higher salaries in other departments.

3. Market Failures Due to Contractual Problems

The above result was achieved because firms could commit to a wage-training contract. In other words, the firm could make a credible commitment to providing training in the amount of τ^*. Such commitments are in general difficult, since outsiders cannot observe the exact nature of the "training activities" taking place inside the firm. For example, the firm could hire workers at a low wage pretending to offer them training, and then employ them as cheap labor. This implies that contracts between firms and workers concerning training investments are naturally *incomplete*.

To capture these issues let us make the timing of events regarding the provision of training somewhat more explicit.

- At time $t = -1/2$, the firm makes a training-wage contract offer (τ', w_0). Workers accept offers from firms.
- At time $t = 0$, there is an initial production of y_0, the firm pays w_0, and also unilaterally decides the level of training τ, which could be different from the promised level of training τ'.
- At time $t = 1/2$, wage offers are made, and the worker decides whether to quit and work for another firm.

- At time $t = 1$, there is the second and final period of production, where output is equal to $y_1 + \alpha(\tau)$.

Now the subgame perfect equilibrium can be characterized as follows: at time $t = 1$, a worker of training τ will receive $w_1 = y_1 + \alpha(\tau)$. Realizing this, at time $t = 0$, the firm would offer training $\tau = 0$, irrespective of its contract promise. Anticipating this wage offer, the worker will only accept a contract offer of the form (τ', w_0), such that $w_0 \geq y_0$, and τ does not matter, since the worker knows that the firm is not committed to this promise. As a result, we are back to the outcome conjectured by Pigou, with no training investment by the firm.

A similar conclusion would also be reached if the firm could write a binding contract about training, but the worker were subject to credit constraints and $c(\tau^*) > y_0$, so the worker cannot take enough of a wage cut to finance his training. In the extreme case where $y_0 = 0$, we are again back to the Pigou outcome, where there is no training investment, despite the fact that it is socially optimal to invest in skills (which one of these problems, contractual incompleteness or credit market constraints, appears more important in the context of training?).

4. Training in Imperfect Labor Markets

4.1. Motivation. The general conclusion of both the Becker model with perfect (credit and labor) markets and the model with incomplete contracts (or severe credit constraints) is that there will be no firm-sponsored investment in general training. This conclusion follows from the common assumption of these two models, that the labor market is competitive, so the firm will never be able to recoup its training expenditures in general skills later during the employment relationship.

Is this a reasonable prediction? The answer appears to be no. There are many instances in which firms bear a significant fraction (sometimes all) of the costs of general training investments.

The first piece of evidence comes from the German apprenticeship system. Apprenticeship training in Germany is largely general. Firms training apprentices have to follow a prescribed curriculum, and apprentices take a rigorous outside exam in their trade at the end of the apprenticeship. The industry or crafts chambers certify whether firms fulfill the requirements to train apprentices adequately, while works councils in the firms monitor the training and resolve grievances. At least in certain technical and business occupations, the training curricula limit the firms' choices over the training content fairly severely. Estimates of the net cost of apprenticeship programs to employers in Germany indicate that firms bear a significant financial burden associated with these training investments. The net costs of apprenticeship training may be as high as DM 6,000 per worker (in the 1990s, equivalent of about $6,000 today).

Another interesting example comes from the recent growth sector of the US, the temporary help industry. The temporary help firms provide workers to various employers on short-term contracts, and receive a fraction of the workers' wages as commission. Although blue-collar and professional temporary workers are becoming increasingly common, the majority of temporary workers are in clerical and secretarial jobs. These occupations require some basic computer, typing and other clerical skills, which temporary help firms often provide before the worker is assigned to an employer. Workers are under no contractual obligation to the temporary help firm after this training program. Most large temporary help firms offer such training to all willing individuals. As training prepares the workers for a range of different assignments, it is almost completely general. Although workers taking part in the training programs do not get paid, all the monetary costs of training are borne by the temporary help firms, giving us a clear example of firm-sponsored general training. This was first noted by Krueger and is discussed in more detail by David Autor.

Other evidence is not as clear-cut, but suggests that firm-sponsored investments in general skills are widespread. A number of studies have investigated whether workers who take part in general training programs pay for the costs by taking lower wages. The majority of these studies do not find lower wages for workers in training programs, and even when wages are lower, the amounts typically appear too small to compensate firms for the costs. Although this pattern can be explained within the paradigm of Becker's theory by arguing that workers selected for training were more skilled in unobserved dimensions, it is broadly supportive of widespread firm-sponsored-training.

There are also many examples of firms that send their employees to college, MBA or literacy programs, and problem solving courses, and pay for the expenses while the wages of workers who take up these benefits are not reduced. In addition, many large companies, such as consulting firms, offer training programs to college graduates involving general skills. These employers typically pay substantial salaries and bear the full monetary costs of training, even during periods of full-time classroom training.

How do we make sense of these firm-sponsored investments in general training? We will now illustrate how in frictional labor markets, firms may also be willing to make investments in the general skills of their employees.

4.2. A Basic Framework. Consider the following two-period model. In period 1, the worker and/or the employer choose how much to invest in the worker's general human capital, τ. There is no production in the first period. In period 2, the worker either stays with the firm and produces output $y = f(\tau)$, where $f(\tau)$ is a strictly increasing and concave function. The worker is also paid a wage rate, $w(\tau)$ as a function of his skill level (training) τ, or he quits and obtains an outside wage. The cost of acquiring τ units of skill is again $c(\tau)$, which is again assumed to be continuous, differentiable, strictly increasing and convex, and to satisfy $c'(0) = 0$. There is no discounting, and all agents are risk-neutral.

Assume that all training is *technologically general* in the sense that $f(\tau)$ is the same in all firms.

If a worker leaves his original firm, then he will earn $v(\tau)$ in the outside labor market. Suppose

$$v(\tau) < f(\tau).$$

That is, despite that fact that τ is general human capital, when the worker separates from the firm, he will get a lower wage than his marginal product in the current firm. The fact that $v(\tau) < f(\tau)$ implies that there is a surplus that the firm and the worker can share when they are together. Also note that $v(\tau) < f(\tau)$ is only possible in labor markets with frictions—otherwise, the worker would be paid his full marginal product, and $v(\tau) = f(\tau)$.

Let us suppose that this surplus will be divided by asymmetric Nash bargaining with worker bargaining power given by $\beta \in (0,1)$. Recall from above that asymmetric Nash bargaining and risk neutral preferences imply that the wage rate as a function of training is

$$w(\tau) = v(\tau) + \beta \left[f(\tau) - v(\tau) \right]. \tag{8.1}$$

An important point to note is that the equilibrium wage rate $w(\tau)$ is independent of $c(\tau)$: the level of training is chosen first, and then the worker and the firm bargain over the wage rate. At this point the training costs are already sunk, so they do not feature in the bargaining calculations (bygones are bygones).

Assume that τ is determined by the investments of the firm and the worker, who independently choose their contributions, c_w and c_f, and τ is given by

$$c(\tau) = c_w + c_f.$$

Assume that \$1 investment by the worker costs \$$p$ where $p \geq 1$. When $p = 1$, the worker has access to perfect credit markets and when $p \to \infty$, the worker is severely constrained and cannot invest at all.

More explicitly, the timing of events are:

- The worker and the firm simultaneously decide their contributions to training expenses, c_w and c_f. The worker receives an amount of training τ such that $c(\tau) = c_w + c_f$.

- The firm and the worker bargain over the wage for the second period, $w(\tau)$, where the threat point of the worker is the outside wage, $v(\tau)$, and the threat point of the firm is not to produce.
- Production takes place.

Given this setup, the contributions to training expenses c_w and c_f will be determined noncooperatively. More specifically, the firm chooses c_f to maximize profits:

$$\pi(\tau) = f(\tau) - w(\tau) - c_f = (1-\beta)\left[f(\tau) - v(\tau)\right] - c_f.$$

subject to $c(\tau) = c_w + c_f$. The worker chooses c_w to maximize utility:

$$u(\tau) = w(\tau) - pc_w = \beta f(\tau) + (1-\beta)v(\tau) - pc_f$$

subject to the same constraint.

The first-order conditions are:

$$(1-\beta)\left[f'(\tau) - v'(\tau)\right] - c'(\tau) = 0 \quad \text{if } c_f > 0 \tag{8.2}$$

$$v'(\tau) + \beta\left[f'(\tau) - v'(\tau)\right] - pc'(\tau) = 0 \quad \text{if } c_w > 0 \tag{8.3}$$

Inspection of these equations implies that generically, one of them will hold as a strict inequality, therefore, one of the parties will bear the full cost of training.

The result of no firm-sponsored investment in general training by the firm obtains when $f(\tau) = v(\tau)$, which is the case of perfectly competitive labor markets. (8.2) then implies that $c_f = 0$, so when workers receive their full marginal product in the outside labor market, the firm will never pay for training. Moreover, as $p \to \infty$, so that the worker is severely credit constrained, there will be no investment in training. In all cases, the firm is not constrained, so one dollar of spending on training costs one dollar for the firm.

In contrast, suppose there are labor market imperfections, so that the outside wage is less than the productivity of the worker, that is $v(\tau) < f(\tau)$. Is this gap between marginal product and market wage enough to ensure firm-sponsored investments in training? The answer is no. To see this, first consider the case with no *wage compression*, that is the case in which a marginal increase in skills is valued appropriately in the outside market. Mathematically this corresponds to $v'(\tau) = f'(\tau)$ for all τ. Substituting for this in the first-order condition of the firm, (8.2), we immediately find that if $c_f > 0$, then $c'(\tau) = 0$. So in other words, there will be no firm contribution to training expenditures.

Next consider the case in which there is wage compression, i.e., $v'(\tau) < f'(\tau)$. Now it is clear that the firm may be willing to invest in the general training of the worker. The simplest way to see this is again to consider the case of severe credit constraints on the worker, that is, $p \to \infty$, so that the worker cannot invest in training. Then, $v'(0) < f'(0)$ is sufficient to induce the firm to invest in training.

This shows the importance of *wage compression* for firm-sponsored training. The intuition is simple: wage compression in the outside market translates into wage compression inside the firm, i.e., it implies $w'(\tau) < f'(\tau)$. As a result, the firm makes greater profits from a more skilled (trained) worker, and has an incentive to increase the skills of the worker.

To clarify this point further, the figure draws the productivity, $f(\tau)$, and wage, $w(\tau)$, of the worker. The gap between these two curves is the sector-period profit of the firm. When $f'(\tau) = w'(\tau)$, this profit is independent of the skill level of the worker, and the firm has no interest in increasing the worker's skill. A competitive labor market, $f(\tau) = v(\tau)$, implies this case. In contrast, if $f'(\tau) > w'(\tau)$, which follows is a direct implication of $f'(\tau) > v'(\tau)$ given Nash bargaining, the firm makes more profits from more skilled workers, and is willing to invest in the general skills of its employees.

Let τ_w be the level of training that satisfies (8.3) as equality, and τ_f be the solution to (8.2). Then, it is clear that if $\tau_w > \tau_f$, the worker will bear all the cost of training. And if $\tau_f > \tau_w$, then the firm will bear all the cost of training (despite the fact that the worker may have access to perfect capital markets, i.e. $p = 1$).

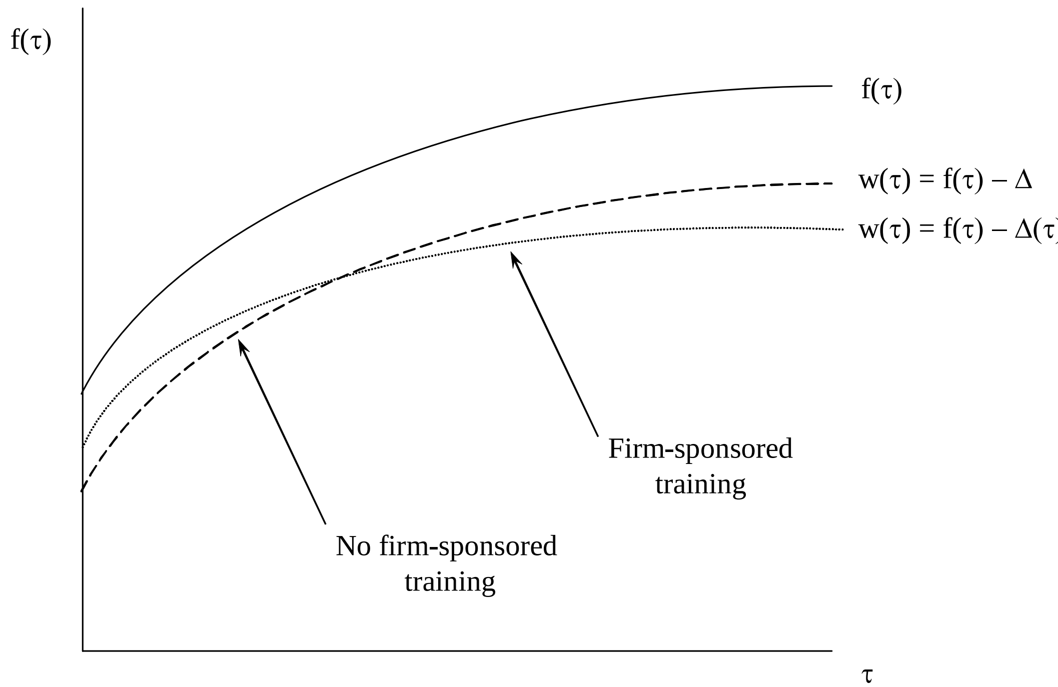

Figure 8.1

To derive the implications of changes in the skill premium on training, let $v(\tau) = af(\tau) - b$. A decrease in a is equivalent to a decrease in the price of skill in the outside market, and would also tilt the wage function inside the firm, $w(\tau)$, decreasing the relative wages of more skilled workers because of bargaining between the firm and in the worker, with the outside wage $v(\tau)$ as the threat point of the worker. Starting from $a = 1$ and $p < \infty$, a point at which the worker makes all investments, a decrease in a leads to less investment in training from (8.3). This is simply an application of the Becker reasoning; without any wage compression, the worker is the one receiving all the benefits and bearing all the costs, and a decline in the returns to training will reduce his investments.

As a declines further, we will eventually reach the point where $\tau_w = \tau_f$. Now the firm starts paying for training, and a further decrease in a increases investment in general training (from (8.2)). Therefore, there is a U-shaped relation between the skill premium and training—starting from a compressed wage structure, a further decrease in the skill premium may increase training. Holding $f(\tau)$ constant a tilting up of the wage schedule, $w(\tau)$, reduces the profits from more skilled workers, and the firm has less interest in investing in skills.

Changes in labor market institutions, such as minimum wages and unionization, will therefore affect the amount of training in this economy. To see the impact of a minimum wage, consider the next figure, and start with a situation where $v(\tau) = f(\tau) - \Delta$ and $p \to \infty$ so that the worker cannot invest in training, and there will be no training. Now impose a minimum wage as drawn in the figure. This distorts the wage structure and encourages the firm to invest in skills up to τ^*, as long as $c(\tau^*)$ is not too high. This is because the firm makes higher profits from workers with skills τ^* than workers with skills $\tau = 0$.

This is an interesting comparative static result, since the standard Becker model with competitive labor markets implies that *minimum wages should always reduce training*. The reason for this is straightforward. Workers take wage cuts to finance their general skills training, and minimum wages will prevent these wage cuts, thus reducing training. We will discuss this issue further below.

5. General Equilibrium with Imperfect Labor Markets

The above analysis showed how in imperfect labor market firms will find it profitable to invest in the general skills of their employees as long as the equilibrium wage structure is compressed. The equilibrium wage structure will be compressed, in turn, when the outside wage structure, $v(\tau)$, is compressed—that is, when $v'(\tau) < f'(\tau)$. The analysis was partial equilibrium in that this outside wage structure was taken as given. There are many reasons why in frictional labor markets we may expect this outside wage structure to be compressed. These include adverse selection, bargaining, and efficiency wages, as well as complementarity between general and specific skills. Here we will discuss how adverse selection leads to wage compression.

5.1. The Basic Model of Adverse Selection and Training. This is a simplified version of the model in Acemoglu and Pischke (1998). Suppose that fraction p of workers are high ability, and have productivity $\alpha(\tau)$ in the second period if they receive training τ in the first period. The remaining $1-p$ are low ability and produce nothing (in terms of the above model, we are setting $y=0$).

No one knows the worker's ability in the first period, but in the second period, the current employer learns this ability. Firms never observe the ability of the workers they have not employed, so outsiders will have to form beliefs about the worker's ability.

The exact timing of events is as follows:

- Firms make wage offers to workers. At this point, worker ability is unknown.
- Firms make training decisions, τ.
- Worker ability is revealed to the current employer and to the worker.
- Employers make second period wage offers to workers.
- Workers decide whether to quit.
- Outside firms compete for workers in the "secondhand" labor market. At this point, these firms observe neither worker ability nor whether the worker has quit or was laid off.
- Production takes place.

Since outside firms do not know worker ability when they make their bids, this is a (dynamic) game of incomplete information. So we will look for a Perfect Bayesian Equilibrium of this game, which is defined in the standard manner. We will characterize equilibria using backward induction conditional on beliefs at a given information set.

First, note that all workers will leave their current employer if outside wages are higher. In addition, a fraction λ of workers, irrespective of ability, realize that they form a bad match with the current employer, and leave whatever the wage is. The important assumption here is that firms in the outside market observe neither worker ability nor whether a worker has quit or has been laid off. However, worker training is publicly observed (what would happen to the model is training was not observed by outside employers?).

These assumptions ensure that in the second period each worker obtains his expected productivity *conditional* on his training. That is, his wage will be independent of his own productivity, but will depend on the average productivity of the workers who are in the secondhand labor market.

By Bayes's rule, the expected productivity of a worker of training τ, is

$$v(\tau) = \frac{\lambda p \alpha(\tau)}{\lambda p + (1-p)} \tag{8.4}$$

To see why this expression applies, note that all low ability workers will leave their initial employer, who will at most pay a wage of 0 (since this is the productivity of a low ability worker), and as we will see, outside wages are positive, low ability workers will quit (therefore, the offer of a wage of 0 is equivalent to a layoff; can there exist in equilibrium in which workers receive zero wage and stay at their job?). Those workers make up a fraction $1-p$ of the total workforce. In addition, of the high ability workers who make up a fraction p of the total workforce, a fraction λ of them will also leave. Therefore, the total size of the secondhand labor market is $\lambda p + (1-p)$, which is the denominator of (8.4). Of those, the low ability ones produce nothing, whereas the λp high ability workers produce $\alpha(\tau)$, which explains this expression.

Anticipating this outside wage, the initial employer has to pay each high ability worker $v(\tau)$ to keep him. This observation, combined with (8.4), immediately implies that there is wage compression in this world, in the sense that

$$v'(\tau) = \frac{\lambda p \alpha'(\tau)}{\lambda p + (1-p)} < \alpha'(\tau),$$

so the adverse selection problem introduces wage compression, and via this channel, will lead to firm-sponsored training.

To analyze this issue more carefully, consider the previous stage of the game. Now firm profits as a function of the training choice can be written as

$$\pi(\tau) = (1-\lambda)\,p\,[\alpha(\tau) - v(\tau)] - c(\tau)\,.$$

The first-order condition for the firm is

$$\begin{aligned} \pi'(\tau) &= (1-\lambda)\,p\,[\alpha'(\tau) - v'(\tau)] - c'(\tau) = 0 \\ &= \frac{(1-\lambda)\,p\,(1-p)\,\alpha'(\tau)}{\lambda p + (1-p)} - c'(\tau) = 0 \end{aligned} \tag{8.5}$$

There are a number of noteworthy features:

(1) $c'(0) = 0$ is sufficient to ensure that there is firm-sponsored training (that is, the solution to (8.5) is interior).

(2) There is underinvestment in training relative to the first-best which would have involved $p\alpha'(\tau) = c'(\tau)$ (notice that the first-best already takes into account that only a fraction p of the workers will benefit from training). This is because of two reasons: first, a fraction λ of the high ability workers quit, and the firm does not get any profits from them. Second, even for the workers who stay, the firm is forced to pay them a higher wage, because they have an outside option that improves with their training, i.e., $v'(\tau) > 0$. This reduces profits from training, since the firm has to pay higher wages to keep the trained workers.

(3) The firm has *monopsony power* over the workers, enabling it to recover the costs of training. In particular, high ability workers who produce $\alpha(\tau)$ are paid $v(\tau) < \alpha(\tau)$.

(4) Monopsony power is not enough by itself. Wage compression is also essential for this result. To see this, suppose that we impose there is no wage compression, i.e., $v'(\tau) = \alpha'(\tau)$, then inspection of the first line of (8.5) immediately implies that there will be zero training, $\tau = 0$.

(5) But wage compression is also not automatic; it is a consequence of some of the assumptions in the model. Let us modify the model so that high ability workers produce $\eta + \alpha(\tau)$ in the second period, while low ability workers

produce $\alpha(\tau)$. This modification implies that training and ability are no longer complements. Both types of workers get exactly the same marginal increase in productivity (this contrasts with the previous specification where only high ability workers benefited from training, hence training and ability were highly complementary). Then, it is straightforward to check that we will have

$$v(\tau) = \frac{\lambda p \eta}{\lambda p + (1-p)} + \alpha(\tau),$$

and hence $v'(\tau) = \alpha'(\tau)$. Thus no wage compression, and firm-sponsored training. Intuitively, the complementarity between ability and training induces wage compression, because the training of high ability workers who are contemplating to leave their firm is judged by the market as the training of a relatively low ability worker (since low ability workers are overrepresented in the secondhand labor market). Therefore, the marginal increase in a (high ability) worker's productivity due to training is valued less in the outside market, which views this worker, on average, as low ability. Hence the firm does not have to pay as much for the marginal increase in the productivity of a high ability worker, and makes greater profits from more trained high-ability workers.

(6) What happens if

$$\pi(\tau) = (1-\lambda)\, p\, [\alpha(\tau) - v(\tau)] - c(\tau) > 0,$$

that is, if firms are making positive profits (at the equilibrium level of training)? If there is free entry at time $t = 0$, this implies that firms will compete for workers, since hiring a worker now guarantees positive profits in later periods. As a result, firms will have to pay a positive wage at time $t = 0$, precisely equal to

$$W = \pi(\tau)$$

as a result of this competition. This is because once a worker accepts a job with a firm, the firm acquires monopsony power over this worker's labor services at time $t = 1$ to make positive profits. Competition then implies that these profits have to be transferred to the worker at time $t = 0$. The interesting result is that not only do firms pay for training, but they may also pay workers extra in order to attract them.

5.2. Evidence. How can this model be tested? One way is to look for evidence of this type of adverse selection among highly trained workers. The fact that employers know more about their current employees may be a particularly good assumption for young workers, so a good area of application would be for apprentices in Germany.

According to the model, workers who quit or are laid off should get lower wages than those who stay in their jobs, which is a prediction that follows simply from adverse selection (and Gibbons and Katz tested in the U.S. labor market for all workers by comparing laid-off workers to those who lost their jobs as a result of plant closings). The more interesting implication here is that if the worker is separated from his firm for an exogenous reason that is clearly observable to the market, he should not be punished by the secondhand labor market. In fact, he's "freed" from the monopsony power of the firm, and he may get even higher wages than stayers (who are on average of higher ability, though subject to the monopsony power of their employer).

To see this, note that a worker who is exogenously separated from his firm will get to wage of $p\alpha(\tau)$ whereas stayers, who are still subject of the monopsony power of their employer, obtain the wage of $v(\tau)$ as given by (8.4), which could be less than $p\alpha(\tau)$. In the German context, workers who leave their apprenticeship firm to serve in the military provide a potential group of such exogenous separators. Interestingly, the evidence suggests that although these military quitters are on average lower ability than those who stay in the apprenticeship firm, the military quitters receive higher wages.

5.3. Mobility, training and wages. The interaction between training and adverse selection in the labor market also provides a different perspective in thinking about mobility patterns. To see this, change the above model so that $\lambda = 0$, but workers now quit if

$$w(\tau) - v(\tau) < \theta$$

where θ is a worker-specific draw from a uniform distribution over $[0, 1]$. θ, which can be interpreted as the disutility of work in the current job, is the worker's private information. This implies that the fraction of high ability workers who quit their initial employer will be

$$1 - w(\tau) + v(\tau),$$

so the outside wage is now

$$v(\tau) = \frac{p[1 - w(\tau) + v(\tau)]\,\alpha(\tau)}{p[1 - w(\tau) + v(\tau)] + (1 - p)} \tag{8.6}$$

Note that if $v(\tau)$ is high, many workers leave their employer because outside wages in the secondhand market are high. But also the right hand side of (8.6) is increasing in the fraction of quitters, $[1 - w(\tau) + v(\tau)]$, so $v(\tau)$ will increase further. This reflects the fact that with a higher quit rate, the secondhand market is not as adversely selected (it has a better composition).

This implies that there can be multiple equilibria in this economy. One equilibrium with a high quit rate, high wages for workers changing jobs, i.e. high $v(\tau)$, but low training. Another equilibrium with low mobility, low wages for job changers, and high training. This seems to give a stylistic description of the differences between the U.S. and German labor markets. In Germany, the turnover rate is much lower than in the U.S., and also there is much more training. Also, in Germany workers who change jobs are much more severely penalized (on average, in Germany such workers experience a substantial wage loss, while they experience a wage gain in the U.S.).

Which equilibrium is better? There is no unambiguous answer to this question. While the low-turnover equilibrium achieves higher training, it does worse in terms of matching workers to jobs, in that workers often get stuck in jobs that they do not like. In terms of the above model, we can see this by looking at the average disutility of work that workers receive (i.e., the average θ's).

5.4. Adverse selection and training in the temporary help industry. An alternative place to look for evidence is the temporary help industry in the U.S. Autor (2001) develops an extended version of this model, which also incorporates self-selection by workers, for the temporary help industry. Autor modifies the above model in four respects to apply it to the U.S. temporary help industry. These are:

(1) The model now lasts for three periods, and in the last period, all workers receive their full marginal products. This is meant to proxy the fact that at some point temporary-help workers may be hired into permanent jobs where their remuneration may better reflect their productivity.
(2) Workers have different beliefs about the probability that they are high ability. Some workers receive a signal which makes them believe that they are high ability with probability p, while others believe that they are high ability with probability $p' < p$. This assumption will allow self-selection among workers between training and no-training firms.
(3) Worker ability is only learned via training. Firms that do not offer training will not have superior information relative to the market. In addition, in contrast to the baseline version of the above model, it is also assumed that firms can offer different training levels and commit to them, so firms can use training levels as a method of attracting workers.
(4) The degree of competitiveness in the market is modeled by assuming that firms need to make a certain level of profits π, and a higher π corresponds to a less competitive market.

Autor looks for a "separating"/self-selection equilibrium in which p' workers select into no-training firms, whereas p workers go to training firms. In this context, self-selection equilibrium is one in which workers with different abilities (different beliefs) choose to accept jobs in different firms, because ability is rewarded differentially in different firms. This makes sense since training and ability are complements as before. Since firms that do not train their employees do not learn about employability, there is no adverse selection for workers who quit from no-training firms. Therefore, the second-period wage of workers who quit from no-training firms will be simply

$$v(0) = p'\alpha(0)$$

In contrast, the secondhand labor market wage of workers from training firms will be given by $v(\tau)$ from (8.4) above.

In the third period, all workers will receive their *expected* full marginal product. For workers who were employed by the non-training firms (and thus would did not receive training), this is $p'\alpha(0)$, whereas for workers with training, it is $p\alpha(\tau)$.

In the second-period, all workers receive their outside option in the secondhand market, so $v(0)$ for workers in no-training firms, and $v(\tau)$ for workers in training firms.

The condition for a self-selection equilibrium is

$$p(\alpha(\tau) - \alpha(0)) > v(0) - v(\tau) > p'(\alpha(\tau) - \alpha(0)),$$

that is, expected gain of third-period wages for high-belief workers should outweigh the loss (if any) in terms of second period wages (since there are no costs in the first-period by the assumption that there are no wages in the first-period). Otherwise, there could not be a separating equilibrium.

This immediately implies that if $v(0) - v(\tau) < 0$, that is, if workers with training receive higher wages in the second period, then there cannot be a self-selection equilibrium—all workers, irrespective of their beliefs, would like to take a job with training firms. Therefore, the adverse selection problem needs to be strong enough to ensure that $v(0) - v(\tau) > 0$. This is the first implication that Autor investigates empirically using data about the wages of temporary help workers in firms that offer free training compared to the wages of workers in firms that do not offer training. He finds that this is generally the case.

The second implication concerns the impact of greater competition on training. To see this more formally, simply return to the basic model, and look at the profits of a typical training firm. These are

$$\pi(\tau) = \frac{(1-\lambda)\,p\,(1-p)\,\alpha(\tau)}{\lambda p + (1-p)} - c(\tau).$$

Therefore, if in equilibrium we must have $\pi(\tau) = \pi$ for some exogenous level of profits π, and π increases exogenously, the training level offered by training firms must increase. To see this, note that in equilibrium we could never have $\pi'(\tau) > 0$, since then the firm can increase both its profits and attract more workers by simply increasing training. Therefore, the equilibrium must feature $\pi'(\tau) \leq 0$, and thus a decline in π, that is, increasing competitiveness, will lead to higher training.

Autor investigates this empirically using differences in temporary help firms concentration across MSAs, and finds that in areas where there is greater concentration, training is lower.

5.5. Labor market institutions and training. The theory developed here also implies that changes in labor market institutions, such as minimum wages and unionization, will therefore affect the amount of training in this economy. To see the impact of a binding minimum wage on training, let us return to the baseline framework and consider the next figure, and start with a situation where $v(\tau) = f(\tau) - \Delta$ and $p \to \infty$ so that the worker cannot invest in training, and there will be no training. Now impose a minimum wage as drawn in the figure. This distorts the wage structure and encourages the firm to invest in skills up to τ^*, as long as $c(\tau^*)$ is not too high. This is because the firm makes higher profits from workers with skills τ^* than workers with skills $\tau = 0$.

This is an interesting comparative static result, since the standard Becker model with competitive labor markets implies that *minimum wages should always reduce training*. The reason for this is straightforward. Workers take wage cuts to finance their general skills training, and minimum wages will prevent these wage cuts, thus reducing training.

Therefore, an empirical investigation of the relationship between minimum wage changes and worker training is a way of finding out whether the Becker channel

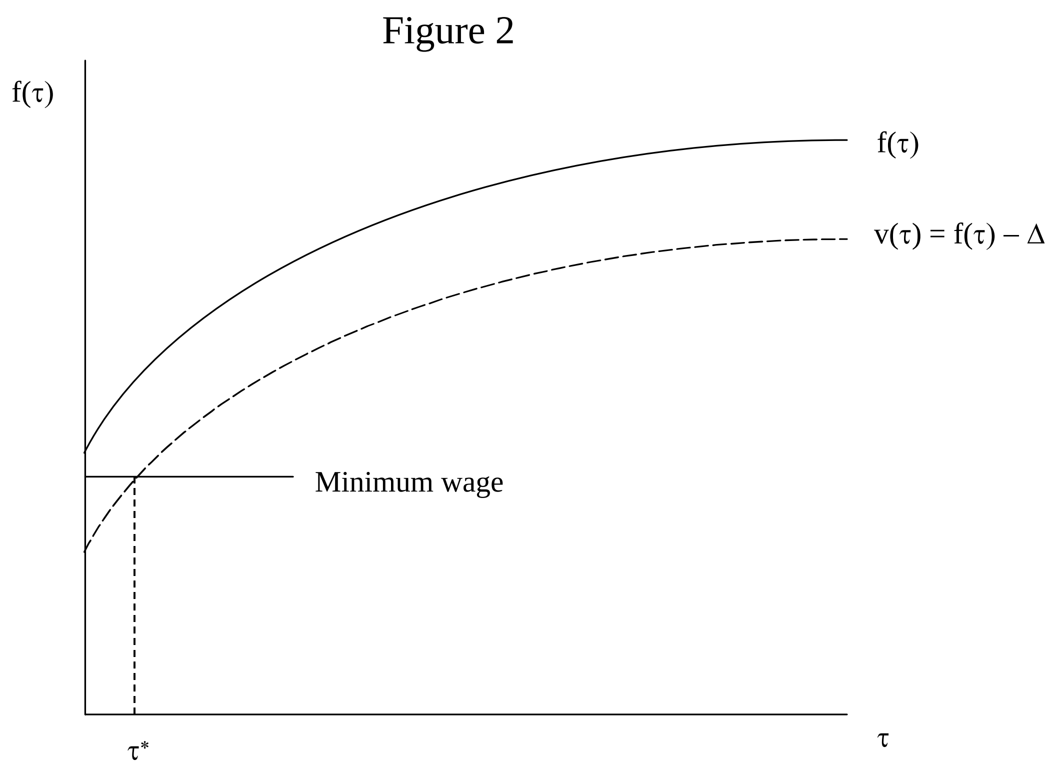

FIGURE 8.2

or the wage-compression channel is more important. Empirical evidence suggests that higher minimum wages are typically associated with more training for low-skill workers (though this relationship is not always statistically significant).

Table 9
The Effect of Minimum Wage Increases on Affected Workers

Independent Variable	*(1)*	*(2)*	*(3)*
Minimum wage increased and wage in prior year is below the current minimum wage	0.010 (0.014)	--	--
Minimum increased and wage in prior year is below the current minimum and above prior year minimum	--	0.016 (0.019)	--
Minimum wage increased and wage in prior year is below 150 % of the current minimum wage	--	--	0.005 (0.008)
Minimum wage increased and wage in prior year is below 130% of the current minimum wage	--	--	--
Change in high school graduation status	0.070 (0.053)	0.070 (0.054)	0.070 (0.054)
Change in new job status	0.032 (0.008)	0.032 (0.008)	0.032 (0.008)

Notes: Basic sample with all workers with high school education or less, who do not move between stat
the next. Dependent variable is the change in training incidence between two consecutive years. All re
include a constant and year dummies. Regressions are weighted by NSLY sampling weights. Standard
for the presence of state*time effecıs in the error term, and therefore robust to the MA structure of the e
17101

FIGURE 8.3

CHAPTER 9

Firm-Specific Skills and Learning

The analysis so far has focused on general skills, acquired in school or by investments in general training. Most labor economists also believe that there are also important firm-specific skills, acquired either thanks to firm-specific experience, or by investment in firm-specific skills, or via "matching". If such firm-specific skills are important we should observe worker productivity and wages to increase with *tenure*—that is, a worker who has stayed longer in a given job should earn more than a comparable worker (with the same schooling and experience) who has less tenure.

1. The Evidence On Firm-Specific Rents and Interpretation

1.1. Some Evidence. The empirical investigation of the importance of firm-specific skills and rents is a difficult and challenging area. There are two important conceptual issues that arise in thinking about the relationship between wages and tenure, as well as a host of econometric issues. The conceptual issues are as follows:

(1) We can imagine a world in which firm-specific skills are important, but there may be no relationship between tenure and wages. This is because, as we will see in more detail below, productivity increases due to firm-specific skills do not necessarily translate into wage increases. The usual reasoning for why high worker productivity translates into higher wages is that otherwise, competitors would bid for the worker and steal him. This argument does not apply when skills are firm-specific since such skills do not contribute to the worker's productivity in other firms. More generally, the

relationship between productivity and wages is more complex when firm-specific skills are a significant component of productivity. For example, we might have two different jobs, one with faster accumulation of firm-specific skills, but wages may grow faster in the other job because the outside option of the worker is improving faster.

(2) An empirical relationship between tenure and wages does not establish that there are imported from-specific effects. To start with, wages may increase with tenure because of *backloaded* compensation packages, which, as we saw above, are useful for dealing with moral hazard problems. Such a relationship might also result from the fact that there are some jobs with high "rents," and workers who get these jobs never quit, creating a positive relationship between tenure and wages. Alternatively, a positive relationship between tenure and wages may reflect the fact that high ability workers stay in their jobs longer (selection).

The existing evidence may therefore either overstate or understate the importance of tenure and firm-specific skills, and there are no straightforward ways of dealing with these problems. In addition, there are important econometric problems, for example, the fact that in most data sets most tenure spells are uncompleted (most workers are in the middle of their job tenure), complicating the analysis. A number of researchers have used the usual strategies, as well as some creative strategies, to deal with the selection and omitted variable biases, pointed out in the second problem. But is still requires us to ignore the first problem (i.e., be cautious in inferring the tenure-productivity relationship from the observed tenure-wage relationship).

In any case, the empirical relationship between tenure and wages is of interest in its own right, even if we cannot immediately deduce from this the relationship between tenure and firm-specific productivity.

With all of these complications, the evidence nevertheless suggests that there is a positive relationship between tenure and wages, consistent with the importance of firm-specific skills. Here we will discuss two different types of evidence.

The first type of evidence is from regression analyses of the relationship between wages and tenure exploiting within job wage growth. Here the idea is that by looking at how wages grow within a job (as long as the worker does not change jobs), and comparing this to the experience premium, we will get an estimate of the tenure premium. In other words, we can think of wages as given by the following model

$$\ln w_{it} = \beta_1 X_{it} + \beta_2 T_{it} + \varepsilon_{it} \tag{9.1}$$

where X_{it} this total labor market experience of individual i, and T_{it} is his tenure in the current job. Then, we have that his wage growth on this job is:

$$\Delta \ln w_{it} = \beta_1 + \beta_2 + \Delta\varepsilon_{it}$$

If we knew the experience premium, β_1, we could then immediately compute the tenure premium β_2. The problem is that we do not know the experience premium. Topel suggests that we can get an upper bound for the experience premium by looking at the relationship between entry-level wages and labor market experience (that is, wages in jobs with tenure equal to zero). This is an upper bound to the extent that workers do not randomly change jobs, but only accept new jobs if these offer a relatively high wage. Therefore, whenever $T_{it} = 0$, the disturbance term ε_{it} in (9.1) is likely to be positively selected. According to this reasoning, we can obtain a *lower bound* estimate of β_2, $\hat{\beta}_2$, using a two-step procedure—first estimate the rate of within-job wage growth, $\hat{\beta}_1$, and then subtract from this the estimate of the experience premium obtained from entry-level jobs (can you see reasons why this will lead to an upwardly biased estimate of the importance of tenure rather than a lower bound on tenure affects as Topel claims?).

Using this procedure Topel estimates relatively high rates of return to tenure. For example, his main estimates imply that ten years of tenure increase wages by about 25 percent, over and above the experience premium.

It is possible, however, that this procedure might generate tenure premium estimates that are upward biased. For example, this would be the case if the return to tenure or experience is higher among high-ability workers, and those are under-represented among the job-changers. Alternatively, returns to experience may be non-constant, and they may be higher in jobs to which workers are a better match. If this is the case, returns experience for new jobs will understate the average returns to experience for jobs in which workers choose to stay.

On the other hand, the advantage of this evidence is that it is unlikely to reflect simply the presence of some jobs that offer high-rents to workers, unless these jobs that provide high rents also have (for some reason) higher wage growth (one possibility might be that, union jobs pay higher wages, and have higher wage growth, and of course, workers do not leave union jobs, but this seems unlikely).

The second type of evidence comes from the wage changes of workers resulting from job displacement. A number of papers, most notably Jacobson, LaLonde and Sullivan, find that displaced workers experience substantial drop in earnings. This is shown in the next figure.

Part of this is due to non-employment following displacement, but even after three years a typical displaced worker is earning about $1500 less (1987 dollars). Econometrically, this evidence is simpler to interpret than the tenure-premium estimates. Economically, the interpretation is somewhat more difficult than the tenure estimates, since it may simply reflect the loss of high-rent (e.g. union) jobs.

In any case, these two pieces of evidence together are consistent with the view that there are important firm-specific skills/expertises that are accumulated on the job.

1.2. What Are Firm-Specific Skills? If we are going to interpret the above evidence as reflecting the importance of firm-specific skills, then we have to be more specific about what constitutes firm-specific skills. Here are four different views:

(1) Firm-specific skills can be thought to result mostly from firm-specific training investments made by workers and firms. Here it is important to distinguish between firms' and workers' investments, since they will have different incentives.

VOL. 83 NO. 4 JACOBSON ET AL.: EARNINGS LOSSES OF DISPLACED WORKERS 697

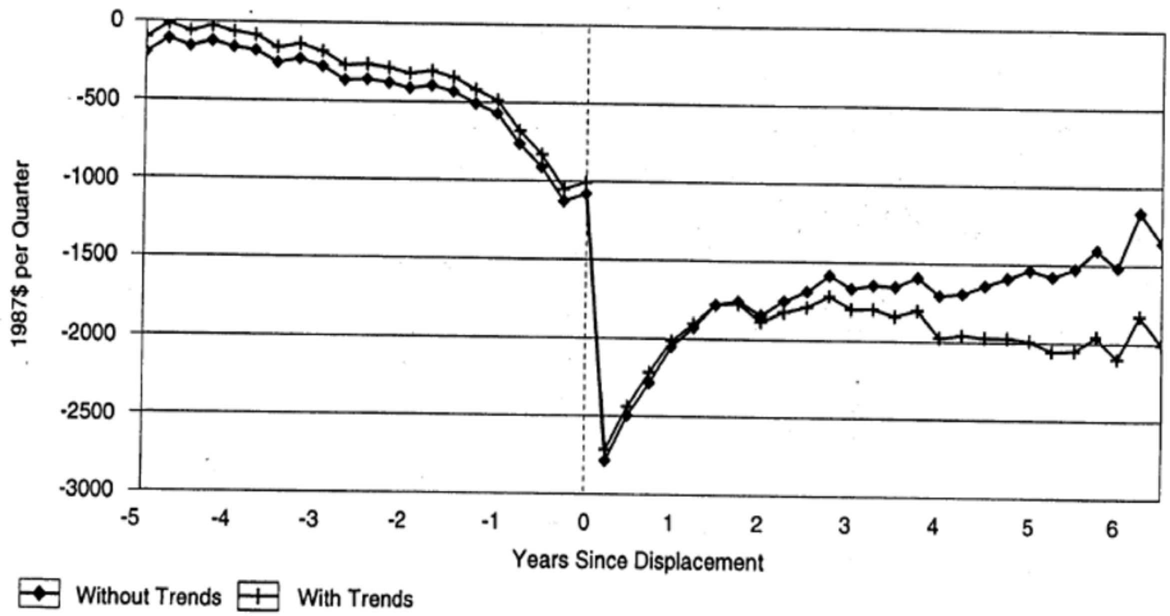

FIGURE 2. EARNINGS LOSSES FOR SEPARATORS IN MASS-LAYOFF SAMPLE

FIGURE 9.1

(2) Firm-specific skills simply reflect what the worker learns on-the-job without making any investments. In other words, they are simply unintentional byproducts of working on the job. The reason why it is useful to distinguish this particular view from the firm-specific investments view is that according to this view, we do not need to worry about the incentives to acquire firm-specific skills. However, most likely, even for simple skills that workers can acquire on-the-job, they need to exert some effort, so this view may have relatively little applicability.

(3) Firm-specific skills may reflect "matching" as in Jovanovic's approach. Here, there is no firm-specific skill, but some workers are better matches to some firms. Ex ante, neither the firm nor the worker knows this, and the information is revealed only slowly. Only workers who are revealed to be good

matches to a particular job will stay on that job, and as a result, they will be more productive in this job than a randomly chosen worker. We can think of this process of learning about the quality of the match as the "accumulation of firm-specific skills".

(4) There may be no technologically firm-specific skills. Instead, you may think of all skills as technologically general, in the sense that if the worker is more productive in a given firm, another firm that adopts exactly the same technologies and organizational structure, and hires the same set of co-workers will also be able to benefit from this high productivity. These technologically general skills are transformed into *de facto* firm-specific skills because of market imperfections. For example, if worker mobility is costly, or if it is difficult or unprofitable for firms to copy some other firms' technology choices, these skills will be de facto specific to the firm that has first made the technology/organizational choices. But if this is the case, we are back to the model of general training investments under imperfect markets we studied above. The reason why it is important to distinguish this view of de facto firm-specific skills from the first view above is that now changes in technology/market organization will affect which skills are specific and how much of a given bundle of technologically-determined skills are "specific".

2. Investment in Firm-Specific Skills

2.1. The basic problem. The problem with general training investments was that part of the costs had to be borne by the firm, but, at least in competitive labor markets, the worker was the residual claimant. The worker, in turn, was the residual claimant because the skills were general, and other firms could compete for this worker's labor services. In contrast, with specific skills, the current employer is the only (or at least the main) "consumer," so there is no competition from other firms to push up the worker's wages. As a result, firm-specific skills will make the firm the ex post monopsonist. This creates the converse problem. Now the worker also bears some (perhaps most) of the costs of investment, but may not have the right incentives to invest, since the firm will get most of the benefits.

To capture these problems, consider the following very simple model:

- At time $t = 0$, the worker decides how much to invest in firm-specific skills, denoted by s, at the cost $\gamma(s)$. $\gamma(s)$ is strictly increasing and convex, with $\gamma'(0) = 0$.
- At time $t = 1$, the firm makes a wage offer to the worker.
- The worker decides whether to accept this wage offer and work for this firm, or take another job.
- Production takes place and wages are paid.

Let the productivity of the worker be $y_1 + f(s)$ where y_1 is also what he would produce with another firm. Since s is specific skills, it does not affect the worker's productivity in other firms.

First, note that the first-best level of firm-specific skills is given by

$$\gamma'(s^*) = f'(s^*).$$

Here s^* is strictly positive since $\gamma'(0) = 0$.

Let us next solve this game by backward induction again, starting in the last period. The worker will accept any wage offer $w_1 \geq y_1$, since this is what he can get in an outside firm. Knowing this, the firm simply offers $w_1 = y_1$. In the previous period, realizing that his wage is independent of his specific skills, the worker makes no investment in specific skills, even though the first best level of firm-specific skills s^* is strictly positive.

What is the problem here? By investing in his firm-specific skills, the worker is increasing the firm's profits. Therefore, the firm would like to encourage the worker to invest. However, given the timing of the game, wages are determined by a take-it-leave-it offer by the firm after the investment. Therefore, it will always be in the interest of the firm to offer a low wage to the worker after the investment, in other words, the firm will *hold* the worker *up*. The worker anticipates this holdup problem and does not invest in his firm-specific skills.

Why is there not a contractual solution to this underinvestment problem? For example, the firm could write a contract ex ante promising a certain payment to the worker. Leaving aside the problems of enforcing such contracts (the firm could always try to fire the worker, or threaten to fire him), there is and more fundamental problem. If the employment contract does *not* make the wage of the worker conditional on his firm-specific skills, it will not encourage investment. So the only contracts that could help with the underinvestment problem are those that make the worker's wages contingent on his firm-specific skills. However, such skills are very difficult to observe or verify by outside parties. This motivates the assumption in this literature, as well as in the incomplete contracts literature, that such contingent contracts cannot be written (they cannot be enforced, and hence are useless). Therefore, contractual solutions to the underinvestment problem are difficult to devise

As a result, there is a severe underinvestment problem here, driven by exactly the converse of the underinvestment problem in general training. The worker will not undertake the required investments, because he's afraid of being held up by the firm.

2.2. Worker power and investment. How can we improve the worker's investment incentives?

At a very general level, the answer is simple. The worker's earnings have to be conditioned on his specific skills. There are a number of ways of achieving this. Perhaps the simplest is to give the worker some "power" in the employment relationship. This power may come simply because the worker can bargain with his employer effectively (either individually or via unions—though the latter would probably be not useful in this context, since union bargaining does not typically will link a worker's wage to his productivity). The worker may be able to bargain with the firm, in turn, for a variety of reasons. Here are some:

(1) Because of regulations, such as employment protection legislation, or precisely because of his specific skills, the firm needs the worker, hence we are

in the bilateral monopoly situation, and the rents will be shared (rather than the firm making a take-it-leave-it offer).

(2) The firm may purposefully give access to some important assets of the firm to the worker, so that the worker may feel secure that he will not be held up. This is basically the insight that follows from the incomplete contracting approach to property rights, which we discussed previously. Recall that in the Grossman-Hart-Moore approach to the internal organization of the firm, the allocation of property rights determine who can use assets and the use of the firm's assets is a way of manipulating ex post bargaining and via this channel ex ante investment incentives.

(3) The firm may change its organizational form in order to make a credible commitment not to hold up the worker.

(4) The firm may develop a reputation for not holding up workers who have invested in firm-specific human capital.

Here let us consider a simple example of investment incentives with bargaining power, and show why firms may preferred to give more bargaining power to their employees in order to ensure high levels of firm-specific investments. In the next section, we discuss alternative "organizational" solutions to this problem.

Modify the above game simply by assuming that in the final period, rather than the firm making a take-it-leave-it offer, the worker and the firm bargain over the firm-specific surplus, so the worker's wage is

$$w_1(s) = y_1 + \beta f(s)$$

Now at time $t = 0$, the worker maximizes

$$y_1 + \beta f(s) - \gamma(s),$$

which gives his investment as

$$\beta f'(\hat{s}) = \gamma'(\hat{s}) \tag{9.2}$$

Here $\hat{s}$ is strictly positive, so giving the worker bargaining power has improved investment incentives. However, $\hat{s}$ is strictly less than the first-best investment level s^*.

To investigate the relationship between firm-specific skills, firm profits and the allocation of power within firms, now consider an extended game, where at time $t = -1$, the firm chooses whether to give the worker access to a key asset. If it does, ex post the worker has bargaining power β, and if it does not, the worker has no bargaining power and wages are determined by a take-it-leave-it offer of the firm. Essentially, the firm is choosing between the game in this section and the previous one. Let us look at the profits of the firm from choosing the two actions. When it gives no access, the worker chooses zero investment, and since $w_1 = y_1$, the firm profits are $\pi_0 = 0$. In contrast, with the change in organizational form giving access to the worker, the worker undertakes investment $\hat{s}$, and profits are

$$\pi_\beta = (1 - \beta) f(\hat{s}).$$

Therefore, the firm would prefer to give the worker some bargaining power in order to encourage investment in specific skills.

Notice the contrast in the role of worker bargaining power between the standard framework and the one here. In the standard framework, worker bargaining power always reduces profits and causes inefficiency. Here, it may do the opposite. This suggests that in some situations reducing worker bargaining power may actually be counterproductive for efficiency.

Note another interesting implication of the framework here. If the firm could choose the bargaining power of the worker without any constraints, it would set $\bar{\beta}$ such that

$$\frac{\partial \pi_\beta}{\partial \beta} = 0 = -f\left(\hat{s}\left(\bar{\beta}\right)\right) + \left(1 - \bar{\beta}\right) f'\left(\hat{s}\left(\bar{\beta}\right)\right) \frac{d\hat{s}\left(\bar{\beta}\right)}{d\beta},$$

where $\hat{s}(\beta)$ and $d\hat{s}/d\beta$ are given by the first-order condition of the worker, (9.2).

One observation is immediate. The firm would certainly choose $\bar{\beta} < 1$, since with $\bar{\beta} = 1$, we could never have $\partial \pi_\beta / \partial \beta = 0$ (or more straightforwardly, profits would be zero). In contrast, a social planner who did not care about the distribution of income between profits and wages would necessarily choose $\beta = 1$. The reason why the firm would not choose the structure of organization that achieves the best investment outcomes is that it cares about its own profits, not total income or surplus.

If there were an ex ante market in which the worker and the firm could "transact', the worker could make side payments to the firm to encourage it to choose $\beta = 1$, then the efficient outcome would be achieved. This is basically the solution that follows from the analysis of the incomplete contracts literature discussed above, but this literature focuses on vertical integration, and attempts to answer the question of who among many entrepreneurs/managers should own the firm or its assets. In the context of worker-firm relationships, such a solution is not possible, given credit constraints facing workers. Perhaps more importantly, such an arrangement would effectively amount to the worker buying the firm, which is not possible for two important reasons:

- the entrepreneur/owner of the firm most likely has some essential knowledge for the production process and transferring all profits to workers or to a single worker is impractical and would destroy the value-generating capacity of the firm;
- in practice there are many workers, so it is impossible to improves their investment incentives by making each worker the residual claimant of the firm's profits.

2.3. Promotions. An alternative arrangement to encourage workers to invest in firm-specific skills is to design a promotion scheme. Consider the following setup. Suppose that there are two investment levels, $s = 0$, and $s = 1$ which costs c.

Suppose also that at time $t = 1$, there are two tasks in the firm, difficult and easy, D and E. Assume outputs in these two tasks as a function of the skill level are

$$y_D(0) < y_E(0) < y_E(1) < y_D(1)$$

Therefore, skills are more useful in the difficult task, and without skills the difficult task is not very productive.

Moreover, suppose that

$$y_D(1) - y_E(1) > c$$

meaning that the productivity gain of assigning a skilled worker to the difficult task is greater than the cost of the worker obtaining skills.

In this situation, the firm can induce firm-specific investments in skills if it can commit to a wage structure attached to promotions. In particular, suppose that the firm commits to a wage of w_D for the difficult task and w_E for the easy task. Notice that the wages do not depend on whether the worker has undertaken the investment, so we are assuming some degree of commitment on the side of the firm, but not modifying the crucial incompleteness of contracts assumption.

Now imagine the firm chooses the wage structures such that

$$y_D(1) - y_E(1) > w_D - w_E > c, \tag{9.3}$$

and then ex post decides whether the worker will be promoted.

Again by backward induction, we have to look at the decisions in the final period of the game. When it comes to the promotion decision, and the worker is unskilled, the firm will naturally choose to allocate him to the easy task (his productivity is higher in the easy task and his wage is lower). If the worker is skilled, and the firm allocates him to the easy task, his profits are $y_E(1) - w_E$. If it allocates him to the difficult task, his profits are $y_D(1) - w_D$. The wage structure in (9.3) ensures that profits from allocating him to the difficult task are higher. Therefore, with this wage structure the firm has made a credible commitment to pay the worker a higher wage if he becomes skilled, because it will find it profitable to promote the worker.

Next, going to the investment stage, the worker realizes that when he does not invest he will receive w_E , and when he invests, he will get the higher wage w_D. Since, again by (9.3), $w_D - w_E > c$, the worker will find it profitable to undertake the investment.

2.4. Investments and layoffs—The Hashimoto model. Consider the following model which is useful in a variety of circumstances. The worker can invest in $s = 1$ at time $t = 0$ again at the cost c. The investment increases the worker's productivity by an amount $m + \eta$ where η is a mean-zero random variable observed only by the firm at $t = 1$. The total productivity of the worker is $x + m + \eta$ (if he does not invest, his productivity is simply x). The firm unilaterally decides whether to fire the worker, so the worker will be fired if

$$\eta < \eta^* \equiv w - x - m,$$

where w is his wage. This wage is assumed to be fixed, and cannot be renegotiated as a function of η, since the worker does not observe η. (There can be other more complicated ways of revealing information about η, using stochastic contracts, whereby workers and firms make direct reports about the values of η and θ, and different values of these variables map into a wage level and a probability that the relationship will continue; using the Revelation Principle we can restrict attention to truthful reports subject incentive compatibility constraints and solve for the most efficient contracts of this form; nevertheless, to keep the discussion simple, we ignore these stochastic contracts here).

If the worker is fired or quits, he receives an outside wage v. If he stays, he receives the wage paid by the firm, w, and also disutility, θ, only observed by him. The worker unilaterally decides whether to quit or not, so he will quit if

$$\theta > \theta^* \equiv w - v$$

Denoting the distribution function of θ by Q and that of η by F, and assuming that the draws from these distributions are independent, the expected profit of the firm is

$$Q(\theta^*)\left[1 - F(\eta^*)\right]\left[x + m - w + E(\eta \mid \eta \geq \eta^*)\right]$$

The expected utility of the worker is

$$v + Q(\theta^*)\left[1 - F(\eta^*)\right]\left[w - v - E(\theta \mid \theta \leq \theta^*)\right]$$

In contrast, if the worker does not invest in skills, he will obtain

$$\begin{aligned} v \text{ if } w &> x \\ v + Q(\theta^*)\left[w - v - E(\theta \mid \theta \leq \theta^*)\right] \text{ if } w &\leq x \end{aligned}$$

So we can see that a high wage promise by the firm may have either a beneficial or an adverse effect on investment incentives. If $w = x + \varepsilon > v$, the worker realizes that he can only keep his job by investing. But on the other hand, a high wage makes it more likely that $\eta < \eta^*$, so it may increase the probability that given the realization of the productivity shock, profits will be negative, and the worker will be fired. This will reduce the worker's investment incentives. In addition, a lower wage would make it more likely that the worker will quit, and through this channel increase inefficiency and discourage investment.

According to Hashimoto, the wage structure has to be determined to balance these effects, and moreover, the ex post wage structure chosen to minimize inefficient separations may dictate a particular division of the costs of firm-specific investments.

An interesting twist on this comes from Carmichael, who suggests that commitment to a promotion ladder might improve incentives to invest without encouraging further layoffs by the firm. Suppose the firm commits to promote N_h workers at time $t = 1$ (how such a commitment is made is an interesting and difficult question). Promotion comes with an additional wage of B. So the expected wage of the worker, if he keeps his job, is now

$$w + \frac{N_h}{N}B,$$

where N is employment at time $t = 1$, and this expression assumes that a random selection of the workers will be promoted. A greater N_h or B, holding the layoff rate of the firm constant, increases the incentive of the worker to stay around, and encourages investment.

Next think about the layoff rate of the firm. The total wage bill of the firm at time $t = 1$ is then

$$W = Nw + N_h B.$$

The significance of this expression is that if the firm fires a worker, this will only save the firm w, since it is still committed to promote N_h workers. Therefore, this commitment to (an absolute number of) promotions, reduces the firm's incentive to fire, while simultaneously increasing the reward to staying in the firm for the worker.

This is an interesting idea, but we can push the reasoning further, perhaps suggesting that it is not as compelling as it first appears. If the firm can commit to promote N_h workers, why can it not commit to employing N' workers, and by manipulating this number effectively make a commitment not to fire workers? So if this type of commitment to employment level is allowed, promotions are not necessary, and if such a commitment is not allowed, it is not plausible that the firm can commit to promoting N_h workers.

3. A Simple Model of Labor Market Learning and Mobility

An important idea related to firm-specific skills is that these skills are (at least in part) a manifestation of the quality of the match between a worker and his job. Naturally, if workers could costlessly learn about the quality of the matches between themselves and all potential jobs, they would immediately choose the job for which they are most suited to. In practice, however, jobs are "experience goods," meaning that workers can only find out whether they are a good match to a job (and to a firm) by working in that firm and job. Moreover, this type of learning does not take place immediately.

What makes these ideas particularly useful for labor economics is that a simple model incorporating this type of match-specific learning provides a range of useful results and also opens up even a larger set of questions for analysis. Interestingly, however, after the early models on these topics, there has been relatively little research.

The first model to formalize these ideas is due to Jovanovic. Jovanovic considered a model in which match-specific productivity is the draw from a normal distribution, and the output of the worker conditional on his match-specific productivity is also normally distributed. Though, as we have seen, normal distributions are often very convenient, in this particular context the normal distribution has a disadvantage, which is that as the worker learns about his match-specific productivity, we need to keep track of both his belief about the level of the quality and also the precision of his beliefs. This makes the model somewhat difficult to work with.

Instead, let us consider a simpler version of the same model.

Each worker is infinitely lived in discrete time and maximizes the expected discounted value of income, with a discount factor $\beta < 1$. There is no ex ante heterogeneity among the workers. But worker-job matches are random.

In particular, the worker may be a good match for a job (or a firm) or a bad match. Let the (population) probability that the worker is a good match be $\mu_0 \in (0, 1)$. A worker in any given job can generate one of two levels of output, high, y_h, and low $y_l < y_h$. In particular, suppose that we have

$$\text{good match} \rightarrow \begin{array}{cc} y_h & \text{with probability } p \\ y_l & \text{with probability } 1-p \end{array}$$

and

$$\text{bad match} \rightarrow \begin{array}{cc} y_h & \text{with probability } q \\ y_l & \text{with probability } 1-q \end{array}$$

where, naturally,

$$p > q.$$

Let us assume that all learning is symmetric (as in the career concerns model). This is natural in the present context, since there is only learning about the match quality of the worker and the firm will also observe the productivity realizations of the worker since the beginning of their employment relationship. This implies that the firm and the worker will share the same *posterior probability* that the worker is a good match to the job. For worker i job j and time t, we can denote this posterior probability (belief) as μ_{ijt}. When there is no risk of confusion, we will denote this simply by μ.

Jovanovic assumes that workers always receive their full marginal product in each job. This is a problematic assumption, since match-specific quality is also firm specific, thus there is no reason for the worker to receive this entire firm-specific surplus. As in the models with firm-specific investments, the more natural assumption would be to have some type of wage bargaining. Let us assume the simplest bargaining structure in which a firm will pay the worker a fraction $\phi \in (0, 1]$ of his expected productivity at that point. In particular, the wage of a worker whose posterior of a good match is μ will be

$$w(\mu) = \phi [\mu y_h + (1 - \mu) y_l] .$$

Note that this is different from the Nash bargaining solution, which would have to take into account the outside option and also the future benefits to the worker from being in this job (which result from learning). But having such a simple expression facilitates the analysis and the exposition here. [Alternatively, we could have assume that bargaining takes place after the realization of output, in which case the wage would be equal to ϕy_h with probability μ and to ϕy_l with probability $1 - \mu$; since both the worker and affirm our risk neutral, there is no difference between these two cases].

To make progress, let us consider a worker with belief μ. If this worker produces output y_h, then Bayes's rule implies that his posterior (belief) next period should be

$$\mu_h'(\mu) \equiv \frac{\mu p}{\mu p + (1 - \mu) q} > \mu,$$

where the fact that this is greater than μ immediately follows from the assumption that $p > q$. Similarly, following an output realization of y_l, the belief of the worker will be

$$\mu_l'(\mu) \equiv \frac{\mu (1 - p)}{\mu (1 - p) + (1 - \mu)(1 - p)} < \mu.$$

Finally, let us also assume that every time a worker changes jobs, he has to incur a training or mobility cost equal to $\gamma \geq 0$.

Under these assumptions, we can write the net present discounted value of a worker with belief μ recursively using simple dynamic programming arguments. In particular, this is

$$\begin{aligned} V(\mu) &= w(\mu) + \beta[(\mu p + (1-\mu) q) V(\mu_h'(\mu)) \\ &\quad + (\mu(1-p) + (1-\mu)(1-q)) \max\{V(\mu_l'(\mu)); V(\mu_0) - \gamma\}. \end{aligned}$$

Intuitively, the worker receives the wage $w(\mu)$ as a function of the (symmetrically held) belief about the quality of his match at the moment.

The continuation value, which is discounted with the discount factor $\beta < 1$, has the following explanation: with probability μ, the match is indeed good and then the worker will produce an output equal to y_h with probability p. With probability $1-\mu$, the match is not good and the worker will produce high output with probability q. In either case, the posterior about match quality will be $\mu_h'(\mu)$, and using the recursive reasoning, his value will be $V(\mu_h'(\mu))$. Since he was happy to be in this job with belief μ, $\mu_h'(\mu) > \mu$ as stated above, and clearly (can you prove this?) $V(\mu)$ is increasing in μ, he will not want to quit after a good realization and thus his value is written as $V(\mu_h'(\mu))$.

With probability $(\mu(1-p) + (1-\mu)(1-q))$, on the other hand, he will produce low output, y_l, and in this case the posterior will be $\mu_l'(\mu)$. Since $\mu_l'(\mu) < \mu$, at this point the worker may prefer to quit and take another job. Since a new job is a new draw from the match-quality distribution, the probability that he will be good at this job is μ_0. Subtracting the cost of mobility, γ, the value of taking a new job is therefore $V(\mu_0) - \gamma$. The worker chooses the maximum of this and this continuation value in the same firm, $V(\mu_l'(\mu))$.

An immediate result from dynamic programming is that if the instantaneous reward function, here $w(\mu)$, is strictly increasing in the state variable, which here is the belief μ, then the value function $V(\mu)$ will also be strictly increasing. This implies that there will exist some cutoff level of belief μ^* such that workers will stay in their job as long as

$$\mu \geq \mu^*,$$

and they will quit if $\mu < \mu^*$.

Let $\bar{\mu} = \inf \{\mu{:}\mu'_l(\mu) < \mu^*\}$. Then a worker with beliefs $\mu > \bar{\mu}$ will not quit irrespective of the realization of output. Workers with $\mu < \mu^*$ should have quit already. Therefore, the only remaining range of beliefs is $\mu \in [\mu^*, \bar{\mu}]$. A worker with beliefs in this range will quit the job if he generates low output.

Now a couple of observations are immediate.

(1) Provided that $\mu_0 \in (0,1)$, μ will never converge to 0 or 1 in finite time. Therefore, a worker who generates high output will have higher wages in the following period, and a worker who generates low output will have lower wages in the following period. Thus, in this model worker *wages will move with past performance.*

(2) It can be easily proved that if $\gamma = 0$, then $\mu^* = \mu_0$. This implies that when γ is equal to 0 or is very small, a worker who starts a job and generates low output will quit immediately. Therefore, as long as γ is not very high, there will be *a high likelihood of separation in new jobs.*

(3) Next consider a worker who has been in a job for a long time. Such workers will on average have high values of μ, since they have never experienced (on this job) a belief less than μ^*. This implies that the average value of their beliefs must be high. Therefore, *workers with long tenure are unlikely to quit or separate from their job.* [Here average refers to the average among the set of workers who have been in a job for a given length of time; for example, the average value of μ for all workers who have been a job for T periods].

(4) With the same argument,workers who have been in a job for a long time will have high average μ and thus high wages. This implies that in equilibrium *there will be a tenure premium.*

(5) Moreover, because Bayesian updating immediately implies that the gaps between $\mu'_h(\mu)$ and μ and between $\mu'_l(\mu)$ and μ are lowest when μ is close to 1 (and symmetrically when it is close to 0, but workers are never in jobs where their beliefs are close to 0), workers with long tenure will not

experience large wage changes. In contrast, *workers at the beginning of their tenure will have higher wage variability.*

(6) What will happen to wages when workers quit? If $\gamma = 0$, wages will necessarily fall when workers quit (since before they quit $\mu > \mu_0$, whereas in the new job $\mu = \mu_0$). If, on the other hand, γ is non-infinitesimal, workers will experience a wage gain when they change jobs, since in this case $\mu^* < \mu_0$ because they are staying in their current job until this job is sufficiently unlikely to be a good match. This last prediction is also consistent with the data, where on average workers who change jobs experience in increase in wages. [But is this a reasonable explanation for wage increases when workers change jobs?].

What is missing from this model is differential learning opportunities in different jobs. If we assume that output and underlying job quality are normally distributed, we already obtain some of amount of differential learning, since the value of learning is higher in new jobs because the precision of the posterior is smaller. Another possibility will be to have heterogeneous jobs, where some jobs have greater returns to match quality, or perhaps some jobs enable faster learning (e.g., more informative signals). Even more interesting would be to allow some amount of learning about general skills. For example, an academic will not be learning and revealing only about his match-specific quality but also about his industry-specific quality (e.g., his research potential). When this is the case, some jobs may play the role of "stepping stones" because they reveal information about the skills and productivity of the worker in a range on other jobs.

Finally, if instead of the reduced-form wage equation, we incorporate competition among firms into this model, some of the predictions change again. For example, we can consider a world in which a finite number of firms with access to the same technology compete a la Bertrand for the worker. Clearly the worker will start working for the firm where the prior of a good match is greatest. Bertrand competition implies that this firm will pay the worker his value at the next best job. Once the worker receives bad news and decides to quit, then he will switch to the job that was previously his next best option. But this implies that his wage, which will now be determined by the third best option (which may in fact be his initial employer) is necessarily smaller, thus job changes will always be associated with wage *declines*. This discussion shows that wage determination assumptions in these models are not innocuous, and more realistic wage determination schemes may lead to results that are not entirely consistent with the data (wage declines rather than wage increases upon job changes). This once again highlights the need for introducing some amount of general skills and job heterogeneity, so that workers quit not only because they have received bad news in their current job but also because they have learned about their ability and can therefore go and work for "higher-quality" jobs.