Moreover, if $[w(\cdot), I(\cdot)]$ satisfies (13.B.16) and (13.B.17), then there exists a $t(\cdot)$ such that $[w(\cdot), t(\cdot), I(\cdot)]$ satisfies (13.B.14)–(13.B.16). Condition (13.B.17), however, is exactly the budget constraint faced by a central authority who runs the firms herself. Hence, we can restrict attention to schemes in which the authority runs the firms herself and uses a direct revelation mechanism $[w(\cdot), I(\cdot)]$ satisfying (13.B.16) and (13.B.17).

Now consider any two types θ' and θ'' for which $I(\theta') = I(\theta'')$. Setting $\theta = \theta'$ and $\hat{\theta} = \theta''$ in condition (13.B.16), we see that we must have $w(\theta') \geq w(\theta'')$. Likewise, letting $\theta = \theta''$ and $\hat{\theta} = \theta'$, we must have $w(\theta'') \geq w(\theta')$. Together, this implies that $w(\theta') = w(\theta'')$. Since $I(\theta) \in \{0, 1\}$, we see that any feasible mechanism $[w(\cdot), I(\cdot)]$ can be viewed as a scheme that gives each worker a choice between two outcomes, $(w_e, I = 1)$ and $(w_u, I = 0)$ and satisfies the budget balance condition (13.B.17). This is exactly the class of mechanisms studied above.

13.C Signaling

Given the problems observed in Section 13.B, one might expect mechanisms to develop in the marketplace to help firms distinguish among workers. This seems plausible because both the firms and the high-ability workers have incentives to try to accomplish this objective. The mechanism that we examine in this section is that of *signaling*, which was first investigated by Spence (1973, 1974). The basic idea is that high-ability workers may have actions they can take to distinguish themselves from their low-ability counterparts.

The simplest example of such a signal occurs when workers can submit to some costless test that reliably reveals their type. It is relatively straightforward to show that in any subgame perfect Nash equilibrium all workers with ability greater than $\underline{\theta}$ will submit to the test and the market will achieve the full information outcome (see Exercise 13.C.1). Any worker who chooses not to take the test will be correctly treated as being no better than the worst type of worker.

However, in many instances, no procedure exists that directly reveals a worker's type. Nevertheless, as the analysis in this section reveals, the potential for signaling may still exist.

Consider the following adaptation of the model discussed in Section 13.B. For simplicity, we restrict attention to the case of two types of workers with productivities θ_H and θ_L, where $\theta_H > \theta_L > 0$ and $\lambda = \text{Prob}\,(\theta = \theta_H) \in (0, 1)$. The important extension of our previous model is that before entering the job market a worker can get some education, and the amount of education that a worker receives is observable. To make matters particularly stark, we assume that education does *nothing* for a worker's productivity (see Exercise 13.C.2 for the case of productive signaling). The cost of obtaining education level e for a type θ worker (the cost may be of either monetary or psychic origin) is given by the twice continuously differentiable function $c(e, \theta)$, with $c(0, \theta) = 0$, $c_e(e, \theta) > 0$, $c_{ee}(e, \theta) > 0$, $c_\theta(e, \theta) < 0$ for all $e > 0$, and $c_{e\theta}(e, \theta) < 0$ (subscripts denote partial derivatives). Thus, both the cost and the marginal cost of education are assumed to be lower for high-ability workers; for example, the work required to obtain a degree might be easier for a high-ability individual. Letting $u(w, e \mid \theta)$ denote the utility of a type θ worker who chooses education level e and receives wage w, we take $u(w, e \mid \theta)$ to equal her wage less any educational costs incurred: $u(w, e \mid \theta) = w - c(e, \theta)$. As in Section 13.B, a worker of type θ can earn $r(\theta)$ by working at home.

In the analysis that follows, we shall see that this otherwise useless education may serve as a signal of unobservable worker productivity. In particular, equilibria emerge in which high-productivity workers choose to get more education than low-productivity workers and firms correctly take differences in education levels as a signal of ability. The welfare effects of signaling activities are generally ambiguous. By revealing information about worker types, signaling can lead to a more efficient allocation of workers' labor, and in some instances to a Pareto improvement. At the same time, because signaling activity is costly, workers' welfare may be reduced if they are compelled to engage in a high level of signaling activity to distinguish themselves.

To keep things simple, throughout most of this section we concentrate on the special case in which $r(\theta_H) = r(\theta_L) = 0$. Note that under this assumption the unique equilibrium that arises in the absence of the ability to signal (analyzed in Section 13.B) has all workers employed by firms at a wage of $w^* = E[\theta]$ and is Pareto efficient. Hence, our study of this case emphasizes the potential inefficiencies created by signaling. After studying this case in detail, we briefly illustrate (in small type) how, with alternative assumptions about the function $r(\cdot)$, signaling may instead generate a Pareto improvement.

A portion of the game tree for this model is shown in Figure 13.C.1. Initially, a random move of nature determines whether a worker is of high or low ability. Then, conditional on her type, the worker chooses how much education to obtain. After obtaining her chosen education level, the worker enters the job market. Conditional on the observed education level of the worker, two firms simultaneously make wage offers to her. Finally, the worker decides whether to work for a firm and, if so, which one.

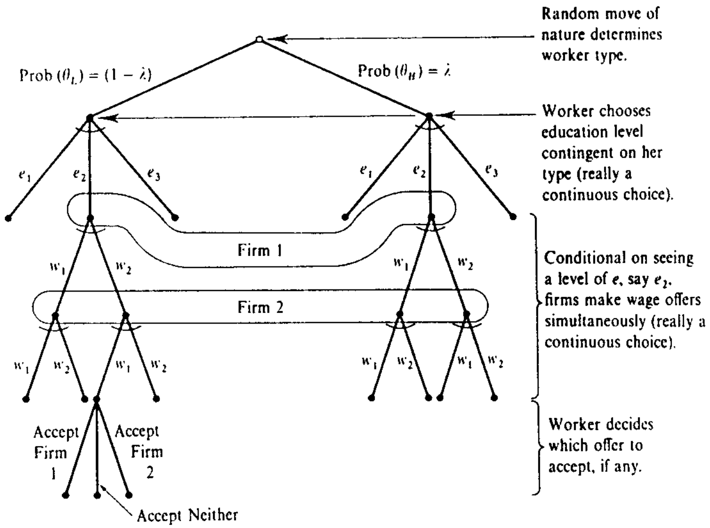

Figure 13.C.1 The extensive form of the education signaling game.

Note that, in contrast with the model of Section 13.B, here we explicitly model only a single worker of unknown type; the model with many workers can be thought of as simply having many of these single-worker games going on simultaneously, with the fraction of high-ability workers in the market being λ. In discussing the equilibria of this game, we often speak of the "high-ability workers" and "low-ability workers," having the many-workers case in mind.

The equilibrium concept we employ is that of a weak perfect Bayesian equilibrium (see Definition 9.C.3), but with an added condition. Put formally, we require that, in the game tree depicted in Figure 13.C.1, the firms' beliefs have the property that, for each possible choice of e, there exists a number $\mu(e) \in [0, 1]$ such that: (i) firm 1's belief that the worker is of type θ_H after seeing her choose e is $\mu(e)$ and (ii) after the worker has chosen e, firm 2's belief that the worker is of type θ_H and that firm 1 has chosen wage offer w is precisely $\mu(e)\sigma_1^*(w \mid e)$, where $\sigma_1^*(w \mid e)$ is firm 1's equilibrium probability of choosing wage offer w after observing education level e. This extra condition adds an element of commonality to the firms' beliefs about the type of worker who has chosen e, and requires that the firms' beliefs about each others' wage offers following e are consistent with the equilibrium strategies both on and off the equilibrium path.

We refer to a weak perfect Bayesian equilibrium satisfying this extra condition on beliefs as a *perfect Bayesian equilibrium* (PBE). Fortunately, this PBE notion can more easily, and equivalently, be stated as follows: A set of strategies and a belief function $\mu(e) \in [0, 1]$ giving the firms' common probability assessment that the worker is of high ability after observing education level e is a PBE if

(i) The worker's strategy is optimal given the firm's strategies.
(ii) The belief function $\mu(e)$ is derived from the worker's strategy using Bayes' rule where possible.
(iii) The firms' wage offers following each choice e constitute a Nash equilibrium of the simultaneous-move wage offer game in which the probability that the worker is of high ability is $\mu(e)$.[20]

In the context of the model studied here, this notion of a PBE is equivalent to the sequential equilibrium concept discussed in Section 9.C. We also restrict our attention throughout to pure strategy equilibria.

We begin our analysis at the end of the game. Suppose that after seeing some education level e, the firms attach a probability of $\mu(e)$ that the worker is type θ_H. If so, the expected productivity of the worker is $\mu(e)\theta_H + (1 - \mu(e))\theta_L$. In a simultaneous-move wage offer game, the firms' (pure strategy) Nash equilibrium wage offers equal the worker's expected productivity (this game is very much like the Bertrand pricing game discussed in Section 12.C). Thus, in any (pure strategy) PBE, we must have both firms offering a wage exactly equal to the worker's expected productivity, $\mu(e)\theta_H + (1 - \mu(e))\theta_L$.

20. Thus, the extra condition we add imposes equilibrium-like play in parts of the tree off the equilibrium path. See Section 9.C for a discussion of the need to augment the weak perfect Bayesian equilibrium concept to achieve this end.

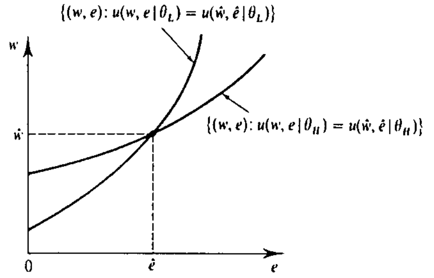

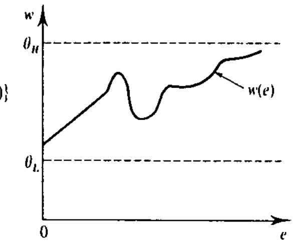

Figure 13.C.2 (left) Indifference curves for high- and low-ability workers: the single-crossing property.

Figure 13.C.3 (right) A wage schedule.

Knowing this fact, we turn to the issue of the worker's equilibrium strategy, her choice of an education level contingent on her type. As a first step in this analysis, it is useful to examine the worker's preferences over (wage rate, education level) pairs. Figure 13.C.2 depicts an indifference curve for each of the two types of workers (with wages measured on the vertical axis and education levels measured on the horizontal axis). Note that these indifference curves cross only once and that, where they do, the indifference curve of the high-ability worker has a smaller slope. This property of preferences, known as the *single-crossing property*, plays an important role in the analysis of signaling models and in models of asymmetric information more generally. It arises here because the worker's marginal rate of substitution between wages and education at any given (w, e) pair is $(dw/de)_{\bar{u}} = c_e(e, \theta)$, which is decreasing in θ because $c_{e\theta}(e, \theta) < 0$.

We can also graph a function giving the equilibrium wage offer that results for each education level, which we denote by $w(e)$. Note that since in any PBE $w(e) = \mu(e)\theta_H + (1 - \mu(e))\theta_L$ for the equilibrium belief function $\mu(e)$, the equilibrium wage offer resulting from any choice of e must lie in the interval $[\theta_L, \theta_H]$. A possible wage offer function $w(e)$ is shown in Figure 13.C.3.

We are now ready to determine the equilibrium education choices for the two types of workers. It is useful to consider separately two different types of equilibria that might arise: *separating equilibria*, in which the two types of workers choose different education levels, and *pooling equilibria*, in which the two types choose the same education level.

Separating Equilibria

To analyze separating equilibria, let $e^*(\theta)$ be the worker's equilibrium education choice as a function of her type, and let $w^*(e)$ be the firms' equilibrium wage offer as a function of the worker's education level. We first establish two useful lemmas.

Lemma 13.C.1: In any separating perfect Bayesian equilibrium, $w^*(e^*(\theta_H)) = \theta_H$ and $w^*(e^*(\theta_L)) = \theta_L$; that is, each worker type receives a wage equal to her productivity level.

Proof: In any PBE, beliefs on the equilibrium path must be correctly derived from the equilibrium strategies using Bayes' rule. Here this implies that upon seeing education level $e^*(\theta_L)$, firms must assign probability one to the worker being type θ_L. Likewise, upon seeing education level $e^*(\theta_H)$, firms must assign probability one

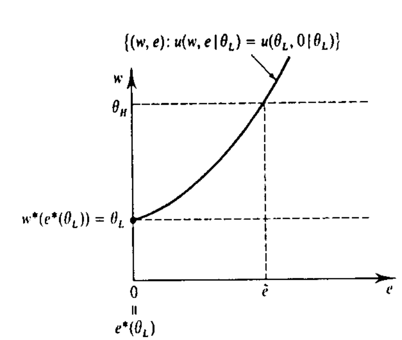

Figure 13.C.4 (left)
Low-ability worker's outcome in a separating equilibrium.

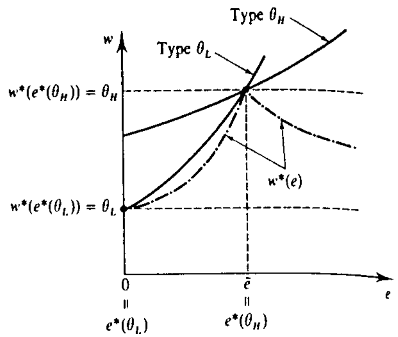

Figure 13.C.5 (right)
A separating equilibrium: Type is inferred from education level.

to the worker being type θ_H. The resulting wages are then exactly θ_L and θ_H, respectively. ■

Lemma 13.C.2: In any separating perfect Bayesian equilibrium, $e^*(\theta_L) = 0$; that is, a low-ability worker chooses to get no education.

Proof: Suppose not, that is, that when the worker is type θ_L, she chooses some strictly positive education level $\hat{e} > 0$. According to Lemma 13.C.1, by doing so, the worker receives a wage equal to θ_L. However, she would receive a wage of at least θ_L if she instead chose $e = 0$. Since choosing $e = 0$ would have save her the cost of education, she would be strictly better off by doing so, which is a contradiction to the assumption that $\hat{e} > 0$ is her equilibrium education level. ■

Lemma 13.C.2 implies that, in any separating equilibrium, type θ_L's indifference curve through her equilibrium level of education and wage must look as depicted in Figure 13.C.4.

Using Figure 13.C.4, we can construct a separating equilibrium as follows: Let $e^*(\theta_H) = \tilde{e}$, let $e^*(\theta_L) = 0$, and let the schedule $w^*(e)$ be as drawn in Figure 13.C.5. The firms' equilibrium beliefs following education choice e are $\mu^*(e) = (w^*(e) - \theta_L)/(\theta_H - \theta_L)$. Note that they satisfy $\mu^*(e) \in [0, 1]$ for all $e \geq 0$, since $w^*(e) \in [\theta_L, \theta_H]$.

To verify that this is indeed a PBE, note that we are completely free to let firms have any beliefs when e is neither 0 nor $\tilde{e}$. On the other hand, we must have $\mu(0) = 0$ and $\mu(\tilde{e}) = 1$. The wage offers drawn, which have $w^*(0) = \theta_L$ and $w^*(\tilde{e}) = \theta_H$, reflect exactly these beliefs.

What about the worker's strategy? It is not hard to see that, given the wage function $w^*(e)$, the worker is maximizing her utility by choosing $e = 0$ when she is type θ_L and by choosing $e = \tilde{e}$ when she is type θ_H. This can be seen in Figure 13.C.5 by noting that, for each type that she may be, the worker's indifference curve is at its highest-possible level along the schedule $w^*(e)$. Thus, strategies $[e^*(\theta), w^*(e)]$ and the associated beliefs $\mu(e)$ of the firms do in fact constitute a PBE.

Note that this is not the only PBE involving these education choices by the two types of workers. Because we have so much freedom to choose the firms' beliefs off the equilibrium path, many wage schedules can arise that support these education

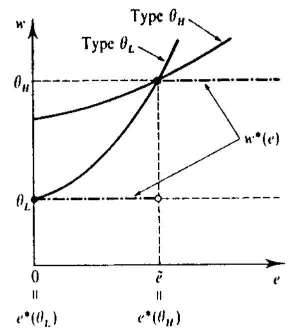

Figure 13.C.6 (left)
A separating equilibrium with the same education choices as in Figure 13.C.5 but different off-equilibrium-path beliefs.

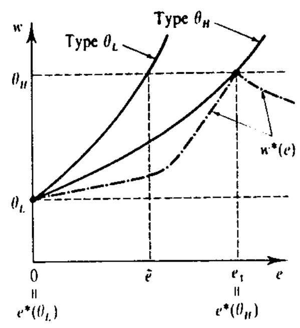

Figure 13.C.7 (right)
A separating equilibrium with an education choice $e^*(\theta_H) > \tilde{e}$ by high-ability workers.

choices. Figure 13.C.6 depicts another one; in this PBE, firms believe that the worker is certain to be of high quality if $e \geq \tilde{e}$ and is certain to be of low quality if $e < \tilde{e}$. The resulting wage schedule has $w^*(e) = \theta_H$ if $e \geq \tilde{e}$ and $w^*(e) = \theta_L$ if $e < \tilde{e}$.

In these separating equilibria, high-ability workers are willing to get otherwise useless education simply because it allows them to distinguish themselves from low-ability workers and receive higher wages. The fundamental reason that education can serve as a signal here is that the marginal cost of education depends on a worker's type. Because the marginal cost of education is higher for a low-ability worker [since $c_{e\theta}(e, \theta) < 0$], a type θ_H worker may find it worthwhile to get some positive level of education $e' > 0$ to raise her wage by some amount $\Delta w > 0$, whereas a type θ_L worker may be unwilling to get this same level of education in return for the same wage increase. As a result, firms can reasonably come to regard education level as a signal of worker quality.

The education level for the high-ability type observed above is not the only one that can arise in a separating equilibrium in this model. Indeed, many education levels for the high-ability type are possible. In particular, any education level between $\tilde{e}$ and e_1 in Figure 13.C.7 can be the equilibrium education level of the high-ability workers. A wage schedule that supports education level $e^*(\theta_H) = e_1$ is depicted in the figure. Note that the education level of the high-ability worker cannot be below $\tilde{e}$ in a separating equilibrium because, if it were, the low-ability worker would deviate and pretend to be of high ability by choosing the high-ability education level. On the other hand, the education level of the high-ability worker cannot be above e_1 because, if it were, the high-ability worker would prefer to get no education, even if this resulted in her being thought to be of low ability.

Note that these various separating equilibria can be Pareto ranked. In all of them, firms earn zero profits, and a low-ability worker's utility is θ_L. However, a high-ability worker does strictly better in equilibria in which she gets a lower level of education. Thus, separating equilibria in which the high-ability worker gets education level $\tilde{e}$ (e.g., the equilibria depicted in Figures 13.C.5 and 13.C.6) Pareto dominate all the others. The Pareto-dominated equilibria are sustained because of the high-ability worker's fear that if she chooses a lower level of education than that prescribed in the equilibrium firms will believe that she is not a high-ability worker. These beliefs can be maintained because in equilibrium they are never disconfirmed.

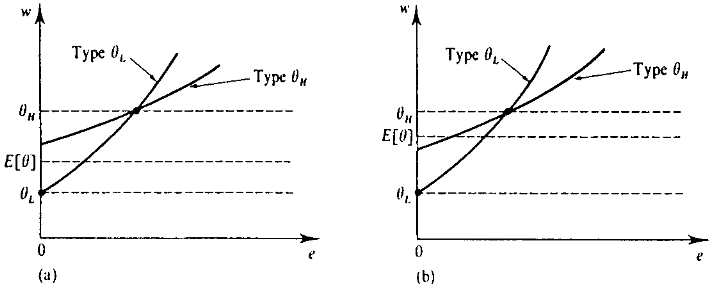

Figure 13.C.8
Separating equilibria may be Pareto dominated by the no-signaling outcome. (a) A separating equilibrium that is not Pareto dominated by the no-signaling outcome. (b) A separating equilibrium that is Pareto dominated by the no-signaling outcome.

It is of interest to compare welfare in these equilibria with that arising when worker types are unobservable but no opportunity for signaling is available. When education is not available as a signal (so workers also incur no education costs), we are back in the situation studied in Section 13.B. In both cases, firms earn expected profits of zero. However, low-ability workers are strictly worse off when signaling is possible. In both cases they incur no education costs, but when signaling is possible they receive a wage of θ_L rather than $E(\theta)$.

What about high-ability workers? The somewhat surprising answer is that high-ability workers may be either better or worse off when signaling is possible. In Figure 13.C.8(a), the high-ability workers are better off because of the increase in their wages arising through signaling. However, in Figure 13.C.8(b), even though high-ability workers seek to take advantage of the signaling mechanism to distinguish themselves, they are *worse* off than when signaling is impossible! Although this may seem paradoxical (if high-ability workers choose to signal, how can they be worse off?), its cause lies in the fact that in a separating signaling equilibrium firms' expectations are such that the wage–education outcome from the no-signaling situation, $(w, e) = (E[\theta], 0)$, is no longer available to the high-ability workers; if they get no education in the separating signaling equilibrium, they are thought to be of low ability and offered a wage of θ_L. Thus, they can be worse off when signaling is possible, even though they are choosing to signal.

Note that because the set of separating equilibria is completely unaffected by the fraction λ of high-ability workers, as this fraction grows it becomes more likely that the high-ability workers are made worse off by the possibility of signaling [compare Figures 13.C.8(a) and 13.C.8(b)]. In fact, as this fraction gets close to 1, nearly every worker is getting costly education just to avoid being thought to be one of the handful of bad workers!

Pooling Equilibria

Consider now pooling equilibria, in which the two types of workers choose the same level of education, $e^*(\theta_L) = e^*(\theta_H) = e^*$. Since the firms' beliefs must be correctly derived from the equilibrium strategies and Bayes' rule when possible, their beliefs when they see education level e^* must assign probability λ to the worker being type θ_H. Thus, in any pooling equilibrium, we must have $w^*(e^*) = \lambda\theta_H + (1 - \lambda)\theta_L = E[\theta]$.

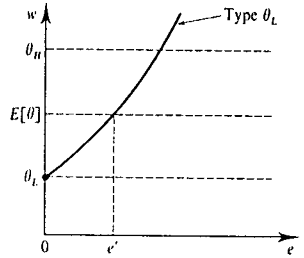

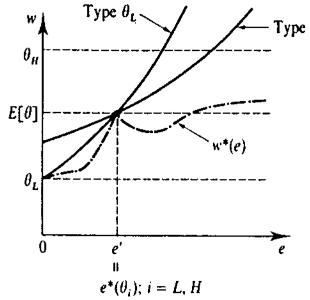

Figure 13.C.9 (left)
The highest-possible education level in a pooling equilibrium.

Figure 13.C.10 (right)
A pooling equilibrium.

The only remaining issue therefore concerns what levels of education can arise in a pooling equilibrium. It turns out that any education level between 0 and the level e' depicted in Figure 13.C.9 can be sustained.

Figure 13.C.10 shows an equilibrium supporting education level e'. Given the wage schedule depicted, each type of worker maximizes her payoff by choosing education level e'. This wage schedule is consistent with Bayesian updating on the equilibrium path because it gives a wage offer of $E[\theta]$ when education level e' is observed.

Education levels between 0 and e' can be supported in a similar manner. Education levels greater than e' cannot be sustained because a low-ability worker would rather set $e = 0$ than $e > e'$ even if this results in a wage payment of θ_L. Note that a pooling equilibrium in which both types of worker get no education Pareto dominates any pooling equilibrium with a positive education level. Once again, the Pareto-dominated pooling equilibria are sustained by the worker's fear that a deviation will lead firms to have an unfavorable impression of her ability. Note also that a pooling equilibrium in which both types of worker obtain no education results in exactly the same outcome as that which arises in the absence of an ability to signal. Thus, pooling equilibria are (weakly) Pareto dominated by the no-signaling outcome.

Multiple Equilibria and Equilibrium Refinement

The multiplicity of equilibria observed here is somewhat disconcerting. As we have seen, we can have separating equilibria in which firms learn the worker's type, but we can also have pooling equilibria where they do not; and within each type of equilibrium, many different equilibrium levels of education can arise. In large part, this multiplicity stems from the great freedom that we have to choose beliefs off the equilibrium path. Recently, a great deal of research has investigated the implications of putting "reasonable" restrictions on such beliefs along the lines we discussed in Section 9.D.

To see a simple example of this kind of reasoning, consider the separating equilibrium depicted in Figure 13.C.7. To sustain e_1 as the equilibrium education level of high-ability workers, firms must believe that any worker with an education level below e_1 has a positive probability of being of type θ_L. But consider any education level $\hat{e} \in (\tilde{e}, e_1)$. A type θ_L worker could never be made better off choosing such an education level than she is getting education level $e = 0$ *regardless of what* *firms believe about her as a result.* Hence, any belief by firms upon seeing education level $\hat{e} > \tilde{e}$ other than $\mu(\hat{e}) = 1$ seems unreasonable. But if this is so, then we must have $w(\hat{e}) = \theta_H$, and so the high-ability worker would deviate to $\hat{e}$. In fact, by this logic, the only education level that can be chosen by type θ_H workers in a separating equilibrium involving reasonable beliefs is $\tilde{e}$.

In Appendix A we discuss in greater detail the use of these types of reasonable-beliefs refinements. One refinement proposed by Cho and Kreps (1987), known as the *intuitive criterion*, extends the idea discussed in the previous paragraph to rule out not only the dominated separating equilibria but also all pooling equilibria. Thus, if we accept the Cho and Kreps (1987) argument, we predict a *unique* outcome to this two-type signaling game: the best separating equilibrium outcome, which is shown in Figures 13.C.5 and 13.C.6.

Second-Best Market Intervention

In contrast with the market outcome predicted by the game-theoretic model studied in Section 13.B (the highest-wage competitive equilibrium), in the presence of signaling a central authority who cannot observe worker types may be able to achieve a Pareto improvement relative to the market outcome. To see this in the simplest manner, suppose that the Cho and Kreps (1987) argument predicting the best separating equilibrium outcome is correct. We have already seen that the best separating equilibrium can be Pareto dominated by the outcome that arises when signaling is impossible. When it is, a Pareto improvement can be achieved simply by banning the signaling activity.

In fact, it may be possible to achieve a Pareto improvement even when the no-signaling outcome does not Pareto dominate the best separating equilibrium. To see how, consider Figure 13.C.11. In the figure, the best separating equilibrium has low-ability workers at point $(\theta_L, 0)$ and high-ability workers at point $(\theta_H, \tilde{e})$. Note that the high-ability workers would be worse off if signaling were banned, since the point $(E[\theta], 0)$ gives them less than their equilibrium level of utility. Nevertheless, note that if we gave the low- and high-ability workers outcomes of $(\hat{w}_L, 0)$ and $(\hat{w}_H, \hat{e}_H)$, respectively, both types would be better off. The central authority can achieve this outcome by mandating that workers with education levels below $\hat{e}_H$ receive a wage of $\hat{w}_L$ and that workers with education levels of at least $\hat{e}_H$ receive a

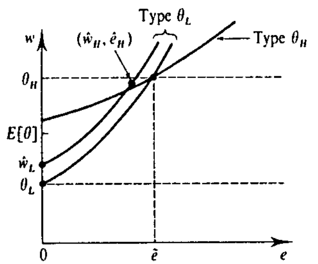

Figure 13.C.11 Achieving a Pareto improvement through cross-subsidization.

wage of $\hat{w}_H$. If so, low-ability workers would choose $e = 0$ and high-ability workers would choose $e = \hat{e}_H$. This alternative outcome involves firms incurring losses on low-ability workers and making profits on high-ability workers. However, as long as the firms break even on *average*, they are no worse off than before and a Pareto improvement has been achieved. The key to this Pareto improvement is that the central authority introduces *cross-subsidization*, where high-ability workers are paid less than their productivity level while low-ability workers are paid more than theirs, an outcome that cannot occur in a separating signaling equilibrium. (Note that the outcome when signaling is banned is an extreme case of cross-subsidization.)

Exercise 13.C.3: In the signaling model discussed in Section 13.C with $r(\theta_H) = r(\theta_L) = 0$, construct an example in which a central authority who does not observe worker types can achieve a Pareto improvement over the best separating equilibrium through a policy that involves cross-subsidization, but cannot achieve a Pareto improvement by simply banning the signaling activity. [*Hint*: Consider first a case with linear indifference curves.]

The case with $r(\theta_H) = r(\theta_L) = 0$ studied above, in which the market outcome in the absence of signaling is Pareto optimal, illustrates how the use of costly signaling can reduce welfare. Yet, when the market outcome in the absence of signaling is not efficient, signaling's ability to reveal information about worker types may instead create a Pareto improvement by leading to a more efficient allocation of labor. To see this point, suppose that we have $r = r(\theta_L) = r(\theta_H)$, with $\theta_L < r < \theta_H$ and $E[\theta] < r$. In this case, the equilibrium outcome without signaling has no workers employed. In contrast, any Pareto efficient outcome must have the high-ability workers employed by firms.

We now study the equilibrium outcome when signaling is possible. Consider, first, the wage and employment outcome that results after educational choice e by the worker. Following the worker's choice of educational level e, equilibrium behavior involves a wage of $w^*(e) = \mu(e)\theta_H + (1 - \mu(e))\theta_L$. If $w^*(e) \geq r$, then both types of workers would accept employment; if $w^*(e) < r$, then neither type would do so.

We now determine the equilibrium education choices of the two types of workers. Note first that any pooling equilibrium must have both types choosing $e = 0$ and neither type accepting employment. To see this, suppose that both types are choosing education level $\hat{e}$. Then $\mu(\hat{e}) = \lambda$ and $w^*(e) = E[\theta] < r$, and so neither type accepts employment. Hence, if $\hat{e} > 0$, both types would be better off choosing $e = 0$ instead. Thus, only an education level of zero is possible in a pooling equilibrium. In this zero education pooling equilibrium, the outcome is identical to the equilibrium outcome arising in the absence of the opportunity to signal.

The set of separating equilibria, on the other hand, is illustrated in Figure 13.C.12. In any separating equilibrium, a low-ability worker sets $e = 0$, is offered a wage of θ_L, and chooses to work at home, thereby achieving a utility of r. High-ability workers, on the other hand, select an education level in the interval $[\hat{e}, e_2]$ depicted in the figure, are offered a wage of θ_H, and accept employment. Note that no separating equilibrium can have $e^*(\theta_H) < \hat{e}$, since then low-ability workers would deviate and set $e = e^*(\theta_H)$; also, no separating equilibrium can have $e^*(\theta_H) > e_2$, since high-ability workers would then be better off setting $e = 0$ and working at home.

Note that in all these equilibria, both pooling and separating, the high-ability workers are weakly better off compared with the equilibrium arising without signaling opportunities and are strictly better off in separating equilibria with $e^*(\theta_H) < e_2$. Moreover, both the low-ability workers and the firms are equally well off. Thus, in the case with $\theta_L < r < \theta_H$ and $E[\theta] < r$, any pooling or separating signaling equilibrium weakly *Pareto dominates* the outcome arising

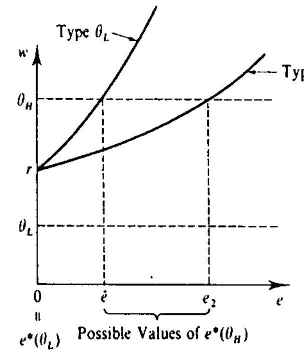

Figure 13.C.12 Separating equilibria when $r(\theta_L) = r(\theta_H) = r \in (\theta_L, \theta_H)$.

in the absence of signaling, and this Pareto dominance is *strict* for (essentially) all separating equilibria.

13.D Screening

In Section 13.C, we considered how signaling may develop in the marketplace as a response to the problem of asymmetric information about a good to be traded. There, individuals on the *more informed* side of the market (workers) chose their level of education in an attempt to signal information about their abilities to uninformed parties (the firms). In this section, we consider an alternative market response to the problem of unobservable worker productivity in which the *uninformed* parties take steps to try to distinguish, or *screen*, the various types of individuals on the other side of the market.[21] This possibility was first studied by Rothschild and Stiglitz (1976) and Wilson (1977) in the context of insurance markets (see Exercise 13.D.2).

As in Section 13.C, we focus on the case in which there are two types of workers, θ_L and θ_H, with $\theta_H > \theta_L > 0$ and where the fraction of workers who are of type θ_H is $\lambda \in (0, 1)$. In addition, workers earn nothing if they do not accept employment in a firm [in the notation used in Section 13.B, $r(\theta_L) = r(\theta_H) = 0$]. However, we now suppose that jobs may differ in the "task level" required of the worker. For example, jobs could differ in the number of hours per week that the worker is required to work. Or the task level might represent the speed at which a production line is run in a factory.

To make matters particularly simple, and to make the model parallel that in Section 13.C, we suppose that higher task levels add *nothing* to the output of the worker; rather, their *only* effect is to lower the utility of the worker.[22] The output of a type θ worker is therefore θ regardless of the worker's task level.

21. The setting analyzed here is one of *competitive screening* of workers, since we assume that there are several competing firms. See Section 14.C for a discussion of the *monopolistic screening* case, where a single firm screens workers.

22. As was true in the case of educational signaling, the assumption that higher task levels do not raise productivity is made purely for expositional purposes. Exercise 13.D.1 considers the case in which the firms' profits are increasing in the task level.

We assume that the utility of a type θ worker who receives wage w and faces task level $t \geq 0$ is

$$u(w, t \mid \theta) = w - c(t, \theta),$$

where $c(t, \theta)$ has all the properties assumed of the function $c(e, \theta)$ in Section 13.C. In particular, $c(0, \theta) = 0$, $c_t(t, \theta) > 0$, $c_{tt}(t, \theta) > 0$, $c_\theta(t, \theta) < 0$ for all $t > 0$, and $c_{t\theta}(t, \theta) < 0$. As will be clear shortly, the task level t serves to distinguish among types here in a manner that parallels the role of education in the signaling model discussed in Section 13.C.

Here we study the pure strategy subgame perfect Nash equilibria (SPNEs) of the following two-stage game:[23]

Stage 1: Two firms simultaneously announce sets of offered contracts. A contract is a pair (w, t). Each firm may announce any finite number of contracts.

Stage 2: Given the offers made by the firms, workers of each type choose whether to accept a contract and, if so, which one. For simplicity, we assume that if a worker is indifferent between two contracts, she always chooses the one with the lower task level and that she accepts employment if she is indifferent about doing so. If a worker's most preferred contract is offered by both firms, she accepts each firm's offer with probability $\frac{1}{2}$.

Thus, a firm can offer a variety of contracts; for example, it might have several production lines, each running at a different speed. Different types of workers may then end up choosing different contracts.[24]

It is helpful to start by considering what the outcome of this game would be if worker types were *observable*. To address this case, we allow firms to condition their offer on a worker's type (so that a firm can offer a contract (w_L, t_L) solely to type θ_L workers and another contract (w_H, t_H) solely to type θ_H workers).

Proposition 13.D.1: In any SPNE of the screening game with observable worker types, a type θ_i worker accepts contract $(w_i^*, t_i^*) = (\theta_i, 0)$, and firms earn zero profits.

Proof: We first argue that any contract (w_i^*, t_i^*) that workers of type θ_i accept in equilibrium must produce exactly zero profits; that is, it must involve a wage $w_i^* = \theta_i$. To see this, note that if $w_i^* > \theta_i$, then some firm is making a loss offering this contract and it would do better by not offering any contract to type θ_i workers. Suppose, on the other hand, that $w_i^* < \theta_i$, and let $\Pi > 0$ be the aggregate profits earned by the two firms on type θ_i workers. One of the two firms must be earning no more than $\Pi/2$ from these workers. If it deviates by offering a contract $(w_i^* + \varepsilon, t_i^*)$ for any

23. For this game, the set of subgame perfect Nash equilibria is identical to the sets of strategy profiles in weak perfect Bayesian equilibria or sequential equilibria.

24. The models in the original Rothschild and Stiglitz (1976) and Wilson (1977) analyses differ from our model in two respects. First, firms in those papers were restricted to offering only a single contract. This could make sense in the production line interpretation, for example, if each firm had only a single production line. Second, those authors allowed for "free entry," so that an additional firm could always enter if a profitable contracting opportunity existed. In fact, making these two changes has little effect on our conclusions. The only difference is in the precise conditions under which an equilibrium exists. (For more on this, see Exercise 13.D.4.)