ADD-ON 21A

ADVERSE SELECTION AND SCREENING IN INSURANCE MARKETS

In this Add-On, we discuss the effects of adverse selection and the nature of competitive screening in insurance markets. This material parallels the discussion of labor market screening that appears in the main text, but draws heavily on Chapter 11 (Risk and Uncertainty). Here we assume a working knowledge of that chapter, particularly Section 11.3, “Insurance.”

A SIMPLE MODEL OF DISABILITY INSURANCE

All workers run the risk of incurring a disability that limits their ability to work and thereby reduces their income. Naturally, risk-averse workers would like to protect against that risk by purchasing insurance. Unfortunately, workers may know more about their risks of disability than insurers. For example, a worker may know whether she is clumsy or accident prone, characteristics that are not easily observed by those who evaluate risks for insurance companies. As a result, disability insurers may be exposed to adverse selection. Recognizing this problem, insurers may attempt to screen potential customers by offering policies designed to appeal to individuals facing different risks. Disability insurance therefore provides a suitable context for studying adverse selection and screening in insurance markets.

Let’s assume that the disability insurance market is competitive, and that the cost of operating an insurance company, aside from the payment of claims, is negligible. To keep matters relatively simple, we’ll assume that there are two types of individuals, those facing a high probability of disability, Π_H, and those facing a low probability, Π_L, where $\Pi_H > \Pi_L$. In either case, an individual earns an income of W_{Able} when healthy and $W_{Disabled}$ when disabled, where $W_{Disabled} < W_{Able}$. Each individual cares about the amounts of goods she will be able to consume if she is healthy and if she is disabled. We illustrate such consumption bundles in Figure 21A.1. Without insurance, she consumes whatever she earns (W_{Able} when healthy and $W_{Disabled}$ when disabled), so her consumption bundle corresponds to point W in the figure. Because the possibility of disability exposes her to risk, the point W lies below the guaranteed consumption line. (We introduced the concept of a guaranteed consumption line on page 374.)

Figure 21A.1 includes an indifference curve for a high-risk individual and one for a low-risk individual. These indifference curves reflect risk-averse preferences. The indifference curve of a low-risk individual is always steeper than that of a high-risk individual where they intersect, as shown in the figure. (Can you explain why? If not, review pages 375–376, including Figure 11.4.)

Figure 21A.1

Consumption bundles and preferences An individual earns W_{Able} when healthy and $W_{Disabled}$ when disabled. Without insurance, she consumes whatever she earns, so her consumption bundle corresponds to point W. Because the possibility of disability exposes her to risk, point W lies below the guaranteed consumption line. The indifference curve of a low-risk individual is always steeper than that of a high-risk individual where they intersect, as shown.

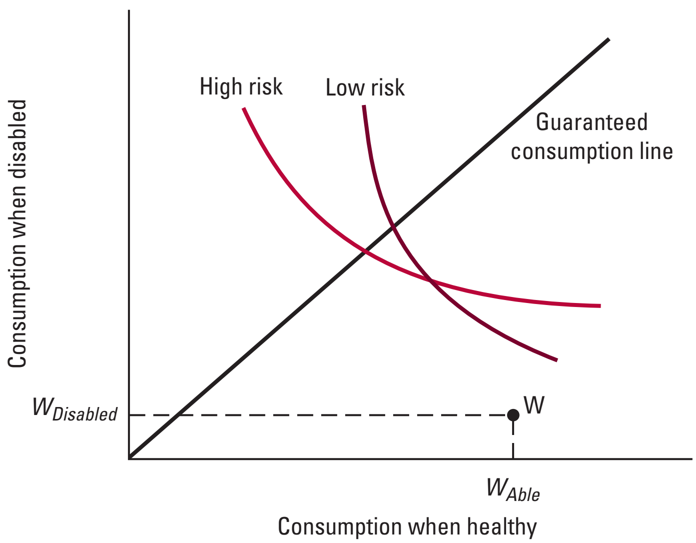

Let's suppose for the moment that insurers can accurately assess each individual's disability risk. If insurers charge a high-risk individual a premium of $\$M_H$ for a policy promising a benefit of $\$B$, the profit earned from the average policyholder will be $\$[M_H - \Pi_H B]$. Because competition drives profits to zero, the competitive insurance premium for high-risk individuals will be $M_H = \Pi_H B$. For similar reasons, the competitive insurance premium for low-risk individuals will be $M_L = \Pi_L B$. In both cases, the premium is actuarially fair. (We introduced the concept of actuarial fairness on page 386.)[1] With actuarially fair insurance, each consumer's budget line consists of the solid portion of the constant expected consumption line that runs between point W and the guaranteed consumption line. (Can you explain why? If not, reread the section titled "The Demand for Fair Insurance" on pages 386–387). In Figure 21A.2, the budget line is B_H for high-risk individuals and B_L for low risk individuals. Because everyone is risk averse and insurance is actuarially fair, everyone purchases full insurance. Low-risk individuals end up at point A and high-risk individuals end up at point B. That outcome is efficient.

The problem becomes interesting when each individual knows his own disability risk, but insurers do not (unless the worker reveals it through his actions). Asymmetric information prevents a competitive market from achieving the efficient outcome shown in Figure 21A.2 (the one that would prevail with symmetric information between policyholders and insurers). As we've drawn the figure, both types of individuals would choose the policy associated with point A over the one associated with point B, and the company offering the policy associated with point A would lose money.[2] That can't be a competitive equilibrium. What then does competition deliver, if not the efficient outcome? The next two subsections address that question. As in the case of labor market screening, our discussion will focus on separating equilibria and pooling equilibria.

[1] On page 386, we used the symbol Π to stand for the probability of avoiding a loss, rather than the probability of incurring a loss. That is why the formula for an actuarially insurance premium given on page 386 was $M = (1 - \Pi)B$, rather than $M = \Pi B$, as above.
[2] Companies price the policy corresponding to point A on the assumption that all policyholders will have low risks of disability. If some purchasers of the policy have high risk, the company will lose money.

Throughout this Add-On, whenever we refer to the *price* of insurance, we mean the amount paid per dollar of promised benefit. In other words, if the insurance company charges a premium of $\$P \times B$ for a policy that promises benefits of $\$B$, the price of that insurance is $\$P$ per dollar of coverage. With actuarial fair premiums, the price of insurance is $\$\Pi_H$ for high-risk individuals and $\$\Pi_L$ for low-risk individuals.

EQUILIBRIUM WITH SEPARATION

If an insurer offers people a choice between two insurance prices, one high and one low, and places no restrictions on the amount of insurance purchased, everyone will obviously choose the lower price. However, by restricting the amount of insurance that an individual can purchase at a given price, an insurer can induce different types of individuals to sort themselves into different types of policies. Figure 21A.3 illustrates this point. Here we assume that the insurer offers two types of policies. One shifts an individual's consumption bundle from point W to point C, the other from point W to point D. Because point D is closer to the guaranteed consumption line than is point C, the policy associated with point D offers more complete coverage; however, it entails a higher price per dollar of coverage. We can tell that the price of insurance is higher for the policy associated with point D than for the one associated with point C because the straight line drawn from point W to point D is flatter than the one drawn from point W to point C. Therefore, the individual trades consumption when healthy for consumption when disabled at a less favorable rate when moving from point W to point D than when moving from point W to point C. As shown, people self-select into different policies: low-risk individuals choose point C, while high-risk individuals choose point D. The high-risk individuals settle for high-priced insurance because they have greater exposure to a loss, and the available high-priced policy offers more complete coverage.

In a separating equilibrium, individuals with different risks sort themselves into different types of policies, much as in Figure 21A.3. Insurers offer a high-priced policy and

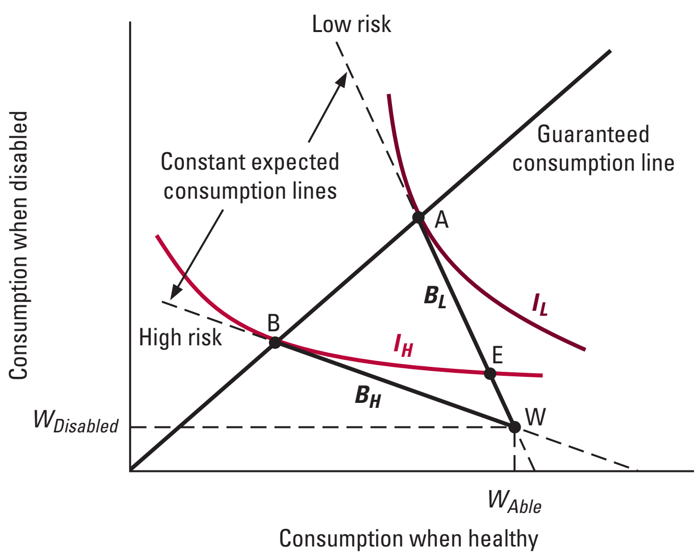

Figure 21A.2

Insurance Choices When Insurers Know Customers' Risks If competitive insurers know customers' risks, the budget line will be B_H for a high-risk individual and B_L for a low risk individual. Because everyone is risk averse and insurance is actuarially fair, everyone purchases full insurance. Low-risk individuals end up at point A and high-risk individuals end up at point B. That outcome is efficient.

Figure 21A.3

Voluntary Sorting by Insurance Customers Suppose two types of insurance policies are available, one that moves individuals from bundle W to bundle C, and another that moves individuals from bundle W to bundle D. Then low-risk individuals will purchase the first type of policy and high-risk individuals will purchase the second.

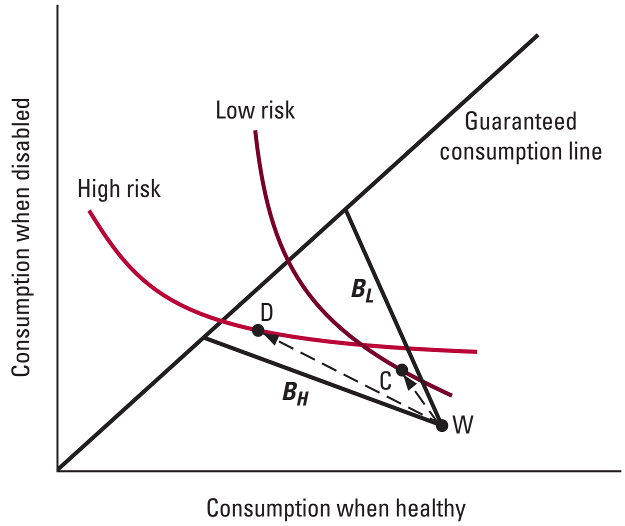

a low-priced policy, where the high-priced policy provides more complete coverage. They expect high-risk individuals to self-select into the high-priced policies and low-ability workers to self-select into the low-priced policies. Neither type of policy can generate positive profits; otherwise, new insurance companies would have an incentive to enter the market and offer policies of that type.[3] Nor can either type of policy generate negative profits; without an offsetting source of positive profits, any insurer offering the unprofitable type of policy would necessarily lose money and shut down. Therefore, both types of policies must break even; given the individuals who choose them, they must be actuarially fair. It follows that high-risk individuals will end up on the line labeled B_H in Figure 21A.2, and low-risk individuals will end up on the line labeled B_L. (Thus, while points C and D in Figure 21A.3 may induce high-risk and low-risk workers to make different choices, competitive firms will not offer the policies associated with those points.)

A SEPARATING EQUILIBRIUM

Look again at Figure 21A.2. Point E lies at the intersection of the line B_L and the indifference curve labeled "high risk" that runs through point B. Under certain conditions (which we identify below), competition between insurers leads to a separating equilibrium in which the policy chosen by high-risk individuals shifts their consumption bundle from point W to point B, and the policy chosen by low-risk individuals shifts their consumption bundle from point W to point E. High-risk individuals are willing to choose point B even when point E is available. Because low-risk individuals have steeper indifference curves than high-risk individuals, they will choose point E over point B. Given those choices, both types of policies generate zero profits for insurers. To determine whether

[3] This statement assumes that the entrant will attract some customers when offering a policy identical to one provided by an existing insurer. That assumption simplifies our analysis. If instead we assumed that the entrant would attract customers only if it offered a policy superior to those of existing insurers, then competition would only ensure that insurers break even across all policies, not necessarily on each policy. It turns out that, with this alternative assumption, the government cannot adopt a policy that makes everyone better off than in the competitive equilibrium, even though the equilibrium is inefficient compared to the outcome with perfect information.

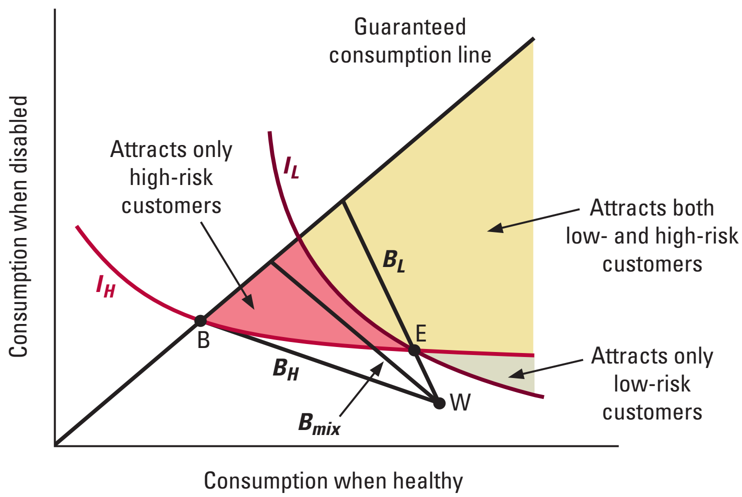

Figure 21A.4

A Separating Equilibrium In a separating equilibrium, low-risk individuals end up at point E and high-risk individuals end up at point B. Insurers break even on both types of customers and cannot earn positive profits by offering another type of policy. Policies corresponding to points in the unshaded area would attract no customers. Policies corresponding to points in the green-shaded area would attract only low-risk customers. Policies corresponding to points in the red-shaded area would attract only high-risk customers. Policies corresponding to points in the yellow-shaded area would attract both types of customers. Given the customers they attract, all of these policies would lose money.

this combination of policies survives against open competition, we need to assess whether a new insurer can enter this market and earn positive profits by offering some other type of policy. Sometimes it can, and sometimes it can't.

Figure 21A.4 reproduces the lines B_L and B_H, the points B and E, and the indifference curve of a high-risk individual that runs through those points (now labeled I_H) from Figure 21A.3. We have added the indifference curve for a low-risk individual that runs through point E (labeled I_L). Let's evaluate the profit opportunities available to a new insurer. Obviously, the policies associated with points B and E will attract insurance customers, but generate zero profits. What other alternatives are available?

First consider points in the unshaded portion of the figure. Because those points are below both I_H and I_L, the corresponding policies will not attract any insurance customers. Therefore, they aren't profitable.

Next consider points in the green-shaded area. Because those points are above I_L and below I_H, the corresponding policies will attract only low-risk customers. However, because those points also lie above B_L, the policies would pay out, on average, more than the associated premium, which means that the insurer would lose money.

Next consider points in the red-shaded area. Because those points are above I_H and below I_L, the corresponding policies will attract only high-risk customers. Because those points also lie above B_H, the policies would pay out, on average, more than the associated premium, which means that the insurer would lose money.

Finally, consider points in the yellow-shaded area. Because those points are above both I_L and I_H, the corresponding policies will attract all customers. Will they be profitable? Notice that we've added a new line to Figure 21A.4, labeled B_{mix}. The slope of B_{mix} reflects the average probability of disability across the entire population. A policy that attracts all customers isn't profitable unless it lies below B_{mix}. As the mix of customers shifts from high risk to low risk, B_{mix} rotates upward from B_H to B_L. Therefore, if high-risk individuals are sufficiently numerous, B_{mix} passes below the yellow-shaded area, as shown in the figure. In that case, an insurer who offered a policy corresponding to any point in the yellow-shaded area would lose money. However, if low-risk individuals were sufficiently numerous, B_{mix} would pass through the yellow-shaded area. In that case, points below B_{mix} and above both indifference curves would correspond to policies that would attract both types of customers while generating a profit.

What have we learned? If (and only if) high-risk individuals are sufficiently numerous, there is a separating equilibrium in which insurers offer the combination of policies shown in Figure 21A.4 (points B and E). High-risk individuals end up with the same consumption bundle regardless of whether employers know each worker's ability. The burden of asymmetric information falls on low-risk individuals, who can only purchase partial insurance when insurers are uninformed (point E lies below the guaranteed consumption line), at the same actuarially fair premium (points E and A both lie on the line B_L). In the separating equilibrium, insurance companies screen customers by presenting them with the following test: "if you want me to believe that you have a low risk of disability and sell you insurance at the price $\$\Pi_L$ rather than at the higher price $\$\Pi_H$, then prove that you have low risk by settling for partial insurance."

ARE THERE OTHER SEPARATING EQUILIBRIA?

As it turns out, there are no separating equilibria other than the one described in the preceding section. Figure 21A.5(a) shows why high-risk individuals must end up at point B. Let's suppose that the job chosen by those workers corresponds to some other point on

Figure 21A.5

Possibilities for Separation that Do Not Survive Competition Figure (a) shows that, in a competitive separating equilibrium, high-risk customers must end up at point B. If they ended up at another point on B_H such as F, an insurer could offer a policy corresponding to a point in the red-shaded area, attract high-risk customers (and possibly low-risk customers), and earn a profit. Figure (b) shows that low-risk customers must end up at point E. If they ended up at another point on B_L involving less insurance, such as point H, an insurer could offer a policy corresponding to a point in the green-shaded area, attract only low-risk customers, and earn a profit.

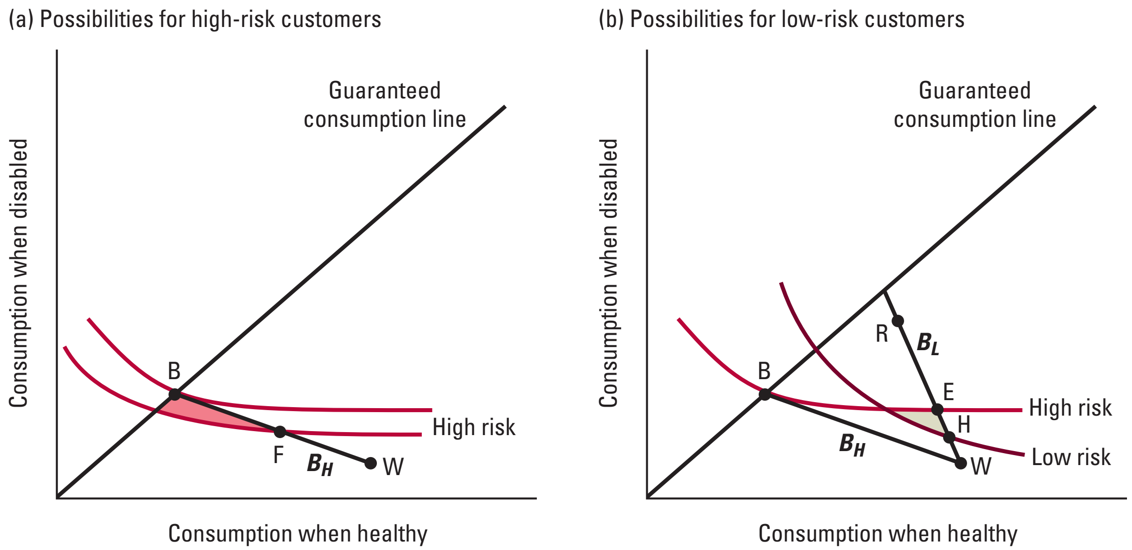

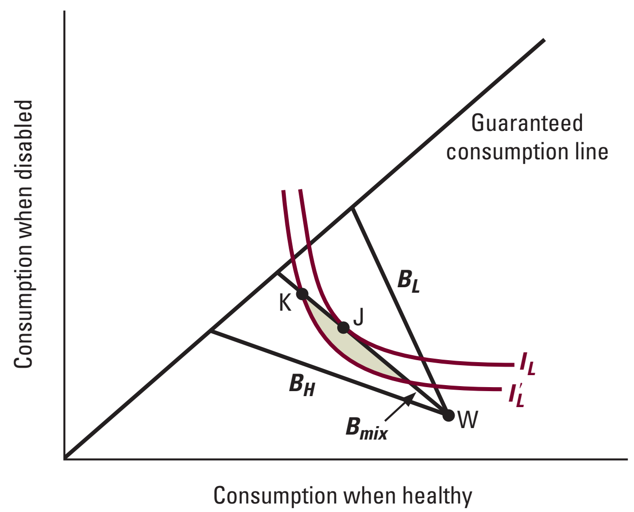

Figure 21A.6

Possibilities for Pooling that Do Not Survive Competition In a pooling equilibrium, all customers must end up at point J. If they instead ended up at another point such as K, an insurer could offer a policy corresponding to a point in the green-shaded area, attract low-risk customers (and perhaps high-risk customers as well), thereby earning a profit.

the line B_H, such as point F. The indifference curve of a high-risk customer that runs any such point must pass below point B, as shown. If a new insurer entered this market and offered a policy corresponding to any point in the red-shaded area of the figure (below B_H and above the indifference curve that runs through point F), it would attract customers—certainly those with high risks, and potentially those with low risks as well—while charging a price greater than Π_H. Because that strategy permits the new entrant to earn a profit, the market isn't in a competitive equilibrium.

What about low-risk customers? To convince ourselves that they must end up at point E in Figure 21A.5(b), let's rule out the alternatives. Recall that they must end up at a point on B_L. They cannot end up at a point to the left of point E, such as R, because then high-risk customers would choose the policy intended for low-risk customers. Neither can they end up at any point to the right of point E, like point H. Why not? The indifference curve of a low-risk customer through any such point must pass below point E, as shown.[4] If a new insurer entered this market and offered a policy corresponding to any point in the green-shaded area of the figure (below B_L and between the two indifference curves), it would be able to attract low-risk customers (because the point is above a low-risk customer's indifference curve through point H) but no high-risk customers (because the point is below a high-risk customer's indifference curve through point B), while charging a price less than Π_L. Because that strategy permits the new entrant to earn a profit, the market isn't in a competitive equilibrium.

ARE THERE POOLING EQUILIBRIA?

In a competitive pooling equilibrium, everyone chooses the same insurance policy. Let's narrow down the possibilities. Because insurance companies must break even, individuals end up at a point on the line B_{mix}, which we've reproduced in Figure 21A.6. At point

[4] We assume here that any reduction in risk with no change in expected consumption makes an individual better off. This property is not necessarily implied by risk aversion, which tells us only that the individual prefers a point on the guaranteed consumption line to all other points on the same constant expected consumption line.

J in the figure, B_{mix} is tangent to the indifference curve labeled I_L, which belongs to a low-risk individual. From the perspective of a low-risk individual, point J is therefore the best alternative on B_{mix}. In a competitive pooling equilibrium, insurers will not offer policies corresponding to any point on B_{mix} other than point J. Figure 21A.6 shows why. Suppose insurers offered policies corresponding to some other point on B_{mix}, such as point K. The indifference curve of a low-risk individual that runs through point K, labeled I'_L, necessarily passes below point J, as shown. If a new insurer entered this market and created a policy corresponding to any point in the green-shaded area of the figure (below B_{mix} and above the low-risk individual's indifference curve that runs through point K), it would certainly attract low-risk customers. Regardless of whether the policy would also attract high-risk customers, the insurer would definitely earn a positive profit, because the point in question lies below both B_L and B_{mix}. Therefore, the market isn't in a competitive equilibrium.

Is there a competitive pooling equilibrium in which insurers offer policies that correspond to the only remaining possibility, point J? Unfortunately, the answer to that question is ambiguous; it depends on whether insurers can observe and react to each others'

Figure 21A.7

A Possible Pooling Equilibrium If insurers can't observe each other's offers, then there is no pooling equilibrium. Assuming that all customers end up at point J, an entrant could offer a policy corresponding to a point in the green- or yellow-shaded areas of figure (a), attract only low-risk customers, and earn a profit. Even if insurers can observe each others' offers, there is no pooling equilibrium for the case shown in figure (a). By offering two policies, one corresponding to a point in the yellow-shaded area and another corresponding to point B, an entrant would earn a profit even if existing firms turned away customers. That entry strategy does not work for the case shown in figure (b). In that case, there is a pooling equilibrium in which all individuals end up at point J, provided that insurers can observe each others' offers.

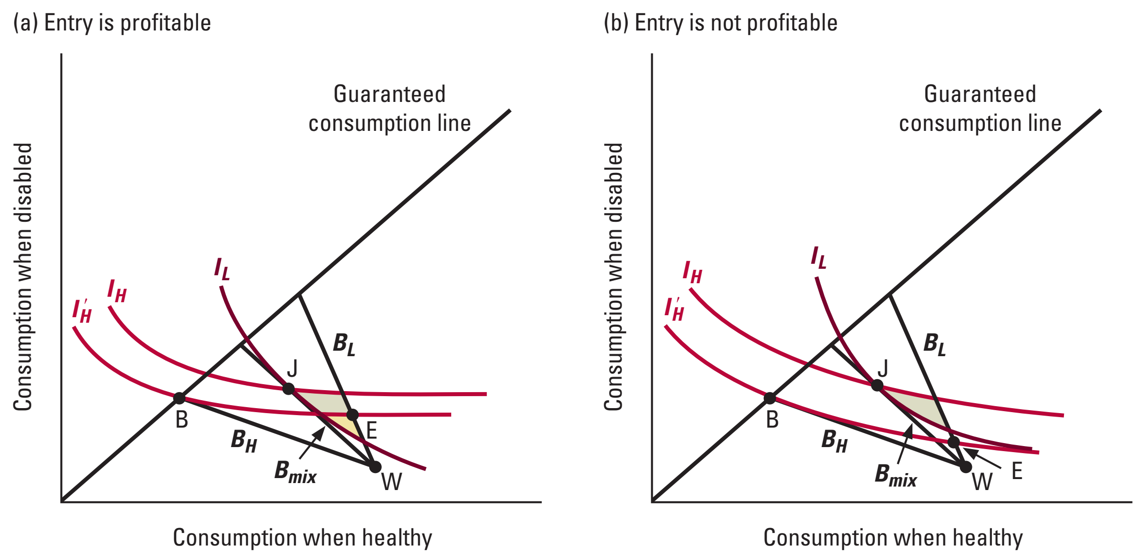

offers by turning away business. If they can't observe each others' offers, then there is definitely no competitive pooling equilibrium, just as we concluded in Section 21.3. Figure 21A.7(a) shows why. If an insurer offered a policy corresponding to any point in the green- or yellow-shaded areas of the figure (below B_L and the high-risk individual's indifference curve I_H, but above the low-risk individual's indifference curve I_L), it would attract only low-risk customers, while charging more than $\$\Pi_L$ per dollar of coverage, thereby earning a profit. Therefore, a competitive equilibrium cannot lead to point J.

If, however, insurers can observe and react to each others' offers by turning away customers, matters are rather different. (We mentioned this possibility in footnote 23 on page 21-29 of Section 21.3, but did not elaborate on it.) Upon observing that a competitor has offered a policy that will attract all of the low-risk customers, an insurer who offers the policy corresponding to point J will infer that its own applicants will have high risks. Rather than sell insurance to those applicants and lose money, the company will turn down all applications. Finding that the policies corresponding to point J are not actually available, the high-risk individuals will apply for the insurance offered by the new entrant. Because the entrant's offering corresponds to a point that lies above B_{mix}, the entrant will lose money after all.

For the case depicted in Figure 21A.7(a), the low-risk customers prefer point E to point J. As a result, the entrant can protect itself against the possibility that existing insurers might turn down applicants by offering two policies, one corresponding to a point in the yellow-shaded area (above I_L, below I'_H, and to the left of B_L), and the other corresponding to point B. Then, if high-risk customers find point J unavailable, they will choose the policy associated with point B, rather than the one that the entrant intends for low-risk customers. With that strategy, the entrant earns a profit on low-risk customers and breaks even on any high-risk customers it insures. Therefore, in Figure 21A.7(a), there is definitely no competitive pooling equilibrium.

In contrast, for the case depicted in Figure 21A.7(b), the low-risk customers prefer point J to point E. As a result, there is no yellow-shaded area (that is, there are no points above I_L, below I'_H, and to the left of B_L), and the entry strategy described in the last paragraph is not feasible. Because all policies in the green-shaded area are more attractive to high-risk customers than point B, offering the policy that corresponds to point B does not discourage high-risk customers from choosing the policy that the entrant intends for low-risk customers in the event that existing insurers turn away applicants. Because a new entrant cannot earn a profit, there is a competitive pooling equilibrium in which everyone purchases the policy associated with point J.

Notably, if insurers can observe and react to each others' offers, the pooling equilibrium exists exactly when the separating equilibrium does not, and vice versa. Why? We've seen that a separating equilibrium exists if the line B_{mix} does not pass through the yellow-shaded area in Figure 21A.4. That condition is equivalent to the statement that a low-risk individual prefers point E to point J, which in turn implies that a pooling equilibrium does not exist [as illustrated in Figure 21A.7(a)]. We've also seen that a separating equilibrium does not exist if the line B_{mix} does pass through the yellow-shaded area in Figure 21A.4. That condition is equivalent to the statement that a low-risk individual prefers point J to point E, which in turn implies that a pooling equilibrium exists [as illustrated in Figure 21A.7(b)]. Therefore, there is always one and only one competitive equilibrium outcome.

AN ADDITIONAL REMARK ON THE ROLE OF GOVERNMENT

At the end of Section 21.3, we explained that government intervention may be justified when competitive insurance companies encounter adverse selection and attempt to screen applicants. Here we illustrate a point to which we alluded in the main text: the government may be able to improve on the market outcome and make everyone better off by designing a social insurance system that relies on self-selection, pays for itself, and induces low-risk individuals to cross-subsidize high-risk individuals.

Consider Figure 21A.8. When competition leads to a separating equilibrium, high-risk individuals end up at point B and low-risk individuals end up at point E. Suppose the government bans private disability insurance and offers two policies, one corresponding to point M, and the other corresponding to point N (which lies at the intersection of I'_H, the indifference curve of a high-risk individual that runs through point M, and I_L, the indifference curve of a low-risk individual that runs through point E). With these offerings, high-risk individuals will be willing to choose the policy associated with point M, and low-risk individuals will prefer the policy associated with point N. High-risk individuals will be better off than with point B, and low-risk individuals will be just as well off as with point E. (We could make them better off as well by sliding point N slightly to the right along the indifference curve I'_H, without tempting high-risk individuals to choose that policy.) The government will earn profits on the policies sold to low-risk individuals (because point N lies below the line B_L) and lose money on the policies sold to high-risk individuals (because point M is above the line B_H). However, if low-risk individuals are sufficiently numerous relative to high-risk individuals, the program will at least break even.

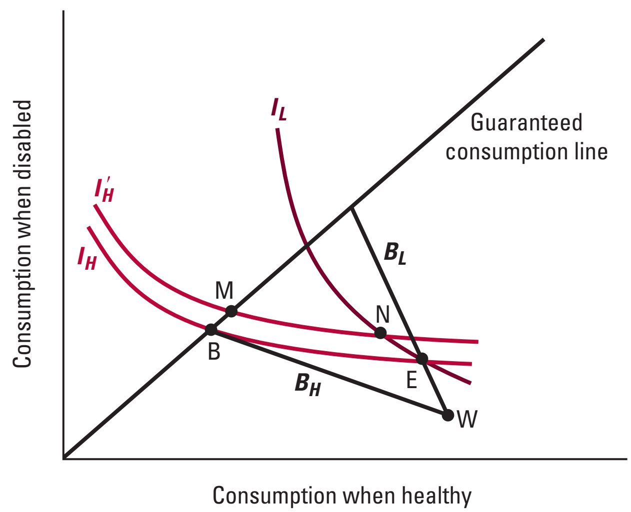

Figure 21A.8

The Benefits of Cross-Subsidization The government can potentially improve on a separating equilibrium by offering policies corresponding to the points M and N, and requiring every individual to choose one of them. High-risk individuals are willing to choose point M over point N and are better off with point M than with point B. Low-risk individuals prefer point N over point M and are just as well off with point N as with point E. The government earns a profit on policies sold to low-risk individuals and loses money on policies sold to high-risk individuals but may break even (or better) overall.