7. ADVERSE SELECTION, SIGNALING, AND SCREENING

In many economic situations, there exists asymmetric information between the different agents. Examples are abundant: A seller has better information about the quality of a product than a buyer; A job applicant has better information about her ability than the potential employer; etc. In the presence of asymmetric information, market equilibria often fail to be Pareto optimal. This is so especially because of a phenomenon called adverse selection: the choices of an informed agent depend on her private information in a manner that adversely affects uninformed market participants. Under adverse selection, the market may completely break down.

To certain extent, markets and institutions develop in response to the problems of asymmetric information and adverse selection. Two possible ways of such responses are called market signaling and market screening.

The analysis of markets under asymmetric information represents one of the most active research areas in microeconomics in the last 20-30 years, and it continues to present important new challenges for economists.

7.1 Asymmetric information and Adverse Selection

Consider a labor market model with many identical firms wanting to hire workers. The firms all produce a homogeneous product with a constant-returns-to-scale technology and act as price takers (we use this as the numeraire, whose price is 1). A worker, if hired by a firm, can produce θ units of output, where θ differs across workers and is distributed on $[\underline{\theta}, \overline{\theta}]$, where $0 \leq \underline{\theta} < \overline{\theta} < \infty$, with c.d.f. $F(\theta)$. The total number (measure) of workers is N.

A worker can choose to work either at a firm or at home, and a worker of type θ earns $r(\theta)$ if she works at home. If w is the wage a worker earns at a firm, then she will choose to work at the firm if and only if $w \geq r(\theta)$.

To start, consider the benchmark where θ is public observable (there is no asymmetric information). In a competitive equilibrium, we must have $w^*(\theta) = \theta$ for all θ. Thus, in equilibrium, the set of workers who are employed by firms is $\{\theta : r(\theta) \leq \theta\}$. It is easy to see that this competitive equilibrium outcome is Pareto efficient. Since each person is allocated to a position (in a firm or at home) that utilizes her labor most efficiently, aggregate surplus must be maximized at this equilibrium.

Now turn back to the case of asymmetric information. Since θ is not observable, there must be a single wage rate for all workers in the market. First, consider labor supply. For any market wage rate w, a worker is willing to work for a firm if and only if $r(\theta) \leq w$. Thus, the set of worker types who are willing to accept employment at wage rate w is

$$\Theta(w) = \{\theta : r(\theta) \leq w\}.$$

Next, consider labor demand as a function of w. If a firm believes that the average productivity of the workers who accept employment is μ, its demand for labor is given by

$$z(w) = \begin{cases} 0 & if \quad \mu < w \\ [0, \infty) & if \quad \mu = w \\ \infty & if \quad \mu > w \end{cases}$$

Definition 1 *In the competitive labor market model with unobservable worker productivity levels, a competitive equilibrium is a wage rate* w^* *and a set* Θ^* *of worker types who accept employment such that*

$$\Theta^* = \{\theta : r(\theta) \leq w^*\}$$

and

$$w^* = E[\theta \mid \theta \in \Theta^*].$$

This definition requires that in equilibrium each worker and firm behave optimally, the beliefs of firms be correct in equilibrium, and each firm earn zero expected profit. Note that if $\Theta^* = \emptyset$, $E[\theta \mid \theta \in \Theta^*]$ is not well defined, in which case we shall take $w^* = E[\theta]$ in equilibrium.

We now argue that the competitive equilibrium under unobservable θ is typically not Pareto efficient. First, consider perhaps the simplest case: $r(\theta) = r$ for all θ and $F(r) \in (0, 1)$. So all workers have the same productivity at home (equal to r), while their productivity at firms could be higher or lower than r. Pareto efficiency would require that all workers with $\theta \geq r$ be employed by firms while workers with $\theta < r$ work at home. At a competitive equilibrium, however, at w^*, it must be that either all workers are employed if $r \leq w^* = E[\theta]$, or all workers work at home if $r > w^* = E[\theta]$. Thus, in equilibrium, either too many or too few people are employed at firms. The intuition for this inefficiency is straightforward: because firms cannot distinguish among workers of differing productivity, the market is unable to allocate workers efficiently between firms and home production.

When there exists adverse selection, the problem of inefficiency may become even severe. To see this, suppose that $r(\theta) \leq \theta$ for all θ and that $r(\theta)$ is strictly increases in θ. The first assumption implies that Pareto efficiency would require every worker be employed by firms, and the second assumption implies that workers with higher productivities in firms are also more productive at home (have higher opportunity costs of working for firms). We can see how adverse selection may occur here: given any market wage, only the less capable workers tend to accept jobs with firms, since they also have lower productivities at home.

Assume that there is a p.d.f. $f(\theta) > 0$ on $[\underline{\theta}, \overline{\theta}]$. Then

$$
\begin{aligned}
E[\theta \quad | \quad & r(\theta) \leq w] \\
&= \frac{\int_{r(\theta)\leq w} \theta f(\theta) d\theta}{\int_{r(\theta)\leq w} f(\theta) d\theta}.
\end{aligned}
$$

It is increasing in w for wages between $r(\underline{\theta})$ and $r(\overline{\theta})$, has a minimum value $\underline{\theta}$ at $w = r(\underline{\theta})$, and has a maximum value $\overline{\theta}$ at $w = r(\overline{\theta})$. In equilibrium, we have

$$w^* = E[\theta \mid r(\theta) \leq w^*].$$

In a diagram where the horizontal axis is w, w^* occurs at the point where curve $E[\theta \mid r(\theta) \leq w]$ intersects the 45^o line.

The equilibrium need not be efficient: those with $r(\theta) > w^*$ will not work for firms although they are more productive in firms than at home. The problem here is that to attract these high-θ workers, the wages need to be increased, but that attracts all the low-θ workers as well, which may then make the average productivity lower than the wage.

In fact, the adverse selection may drive the wage down to $\underline{\theta}$, which implies that no one (measure zero, to be precise) gets employed: The labor market completely breaks down. To see how this can happen, suppose that $r(\theta) = \alpha\theta$, where $\frac{1}{2} < \alpha < 1$, and θ is uniformly distributed on [0,2]. Thus $r(\underline{\theta}) = \underline{\theta}$, and $r(\theta) < \theta$ for all $\theta > 0$. In this case

$$\begin{aligned} E[\theta \quad | \quad r(\theta) \leq w] &= \frac{\int_0^{\frac{w}{\alpha}} \theta \frac{1}{2} d\theta}{\int_0^{\frac{w}{\alpha}} \frac{1}{2} d\theta} = \frac{(\frac{w}{2\alpha})^2}{\frac{w}{2\alpha}} = \frac{w}{2\alpha} \\ &\leq w, \text{ with equality only if } w = 0. \end{aligned}$$

Therefore the only competitive equilibrium in this case is $w^* = 0$ and all workers work at home.

Several other issues that you need to be aware of: (1) There can be multiple competitive equilibria in the labor market model with unobservable worker abilities. (2) To see whether government intervention can improve social welfare under asymmetric information, the government should in general be imposed the same informational constraint as other uninformed market participants. This relates to the concept of constrained Pareto optimal. (3) If firms are not price takers (they make wage offers), then a game-theoretic approach may be needed. One model, for instance, could be that two firms first make competing wage offers and workers then decide whether to work for a firm and if so for which firm.

7.2 Signaling

In the presence of asymmetric information, market mechanisms may be developed that would at least partially reveal information. One such mechanism is signaling, some informed market participants taking certain actions that would reveal private information.

In the job market model, if there exists a costless test that would determine a person's productivity θ, and everyone can decide whether or not to take this test before applying to a firm for employment. Then in any SPNE all workers with $\theta > \underline{\theta}$ will take the test and the market will achieve the full information outcome. (You are asked to show this in one of your problems, but the intuition is very simple: if you have the highest θ, you surely want to take the test to show it to the firm. On the other hand, if you are not tested, then you will be treated according to the average of the workers who are not tested. But then if you have the highest θ among those who are not tested, you must want to take the test. Therefore in equilibrium there cannot be a set of workers with different θ who are not tested. In other words, anyone with $\theta > \underline{\theta}$ will take the test.)

More often, however, signaling is costly. Nevertheless, signaling can occur in equilibrium. Consider the job market signaling model with education. Suppose that there are two types of workers, those with $\theta = \theta_H$ and those with $\theta = \theta_L$, with $\text{prob}(\theta = \theta_H) = \lambda$. A worker can acquire education before going to look for a job. Suppose that the cost of obtaining education level e for a type-θ worker is $c(e, \theta)$, where $c(0, \theta) = 0$, $c_e(e, \theta) > 0$, $c_{ee}(e, \theta) \geq 0$, $c_\theta\ (e, \theta) < 0$, and $c_{e\theta}(c, \theta) < 0$. Thus both the cost and the marginal cost of education is lower for the high-quality workers. The level of education a worker receives is observable. Let the utility of a type θ worker be

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

if she works for a firm and receives wage rate w. She can receive $r(\theta)$ if working at home, and For convenience, we shall assume that $r(\theta) = 0$ for both θ types.

To focus on the study of signaling, we shall assume that education has no effect on a worker's productivity. Thus with perfect information each worker will choose 0 education. Under asymmetric information, as we shall see shortly, a θ_H-type worker may want to choose to have more education in order to signal her high ability in equilibrium. One would expect then that the welfare effects of signaling could be ambiguous: to the extent that signaling reveals information, it improves resource allocation; but at the same time signaling may involve wasteful expenditures that reduce welfare.

The game is as follows: First, nature determines the type of a worker, which is learned by the worker herself. Second, a worker chooses the level of education to receive. Third, after observing the education level of a worker, two firms simultaneously make offers to her. Finally, the worker decides whether to work for a firm and, if so, which one.

The equilibrium concept we use here is perfect Bayesian equilibrium, Let the worker's strategy be e, the common belief function of both firms that the worker's type is θ_H be $\mu(e)$, and the strategy of firm j be $w_j(e)$. A set of strategies and a belief function is a PBE if

(i) The worker's strategy is optimal given the firms' 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 e constitute a Nash equilibrium of the simultaneous-move wage offer game in which the probability of the worker having θ_H

is $\mu(e)$.

This concept adds to the concept of WPBE the requirement that each firm's beliefs about the other firm's wage offers following e are consistent with the equilibrium strategies both on and off the equilibrium path.

We start our analysis from the end of the game. In the simultaneous-move game of offering wages, given belief $\mu(e)$, the expected productivity of the worker is $\mu(e)\theta_H + [1-\mu(e)]\theta_L$. The only Nash equilibrium in this subgame must be that both firms offer wages equal to $\mu(e)\theta_H + [1-\mu(e)]\theta_L$ (the same logic as in the Bertrand model).

Now, let's turn to the worker's equilibrium strategy, her choice of e contingent on her type. It is useful to examine the worker's preference over (w, e). For a given utility level u_0, we can draw an indifference curve for the worker on the e-w space (e on the horizontal axis), which is upward slopping and is

$$w - c(e, \theta) = u_0.$$

Thus, the worker can be kept indifferent with a higher e if she can obtain a higher w. Since $c(e, \theta_H) < c(e, \theta_L)$ and $c_e(e, \theta_H) < c_e(e, \theta_L)$ for all e, the indifference curve for the θ_H-type must be flatter (it is willing to trade an increase in e for a relatively smaller increase in w), and thus there will be a unique crossing point for any two indifference curves of the two types. This property of preferences is knows as the single-crossing property, which plays an important role in signaling models and more generally models of asymmetric information.

We can also graph the equilibrium wage offer function, $w(e)$, which must lie between $w = \theta_L$ and $w = \theta_H$.

There are potentially two types of equilibria: separating equilibrium and pooling equilibrium. In a separating equilibrium, the two types choose different levels of e, while in a pooling equilibrium both types choose the same e.

We first characterize separating equilibria. Let $e^*(\theta)$ be the worker's equilibrium education choice, and $w^*(e)$ be the firms' equilibrium wage offer.

Lemma 2 *In any separating PBE, $w^*(e^*(\theta_H)) = \theta_H$ and $w^*(e^*(\theta_L)) = \theta_L$; that is, each worker type receives a wage equal to her productivity level.*

Proof. In any PBE, the beliefs of the firms must be consistent with the strategy of the worker. Therefore, if the equilibrium strategy of the worker is $e^*(\theta_H)$ and $e^*(\theta_L)$, firms must assign probability 1 to the worker being type θ_H if the worker chooses $e^*(\theta_H)$, and assign probability 1 to the worker being type θ_L if the worker chooses $e^*(\theta_L)$. The resulting equilibrium wage offers must then be θ_H and θ_L respectively. ■

Lemma 3 *In any separating PBE, $e^*(\theta_L) = 0$; that is, a low-ability worker chooses to get zero education.*

Proof. Suppose instead that $e^*(\theta_L) > 0$. From the previous Lemma, a θ_L-type worker's wage is θ_L in any separating PBE. Thus the equilibrium payoff to the θ_L-type worker would be $\theta_L - c(\theta_L, e^*(\theta_L)) < \theta_L$. But if she chooses $e = 0$, her wage and her payoff is at least θ_L, contradicting the assumption that $e^*(\theta_L) > 0$ is her equilibrium choice. ■

In any separating equilibrium, type θ_L's indifference curve on the e-w space is upward-sloping and starts at $e = 0$ and $\theta = \theta_L$. This then allows us to construct a separating equilibrium as follows: $e^*(\theta_H) = \widetilde{e}$, $e^*(\theta_L) = 0$, where $\widetilde{e}$ is at the point where $\theta_L = \theta_H - c(\theta_H, \widetilde{e})$; $w^*(e)$ is any function whose curve is below both types' indifference curves at all e except at $e = 0$ and $e = \widetilde{e}$, where $w^*(0) = \theta_L$ and $w^*(\theta_H) = \theta_H$; the firms' beliefs are $\mu^*(e) = \frac{w^*(e)-\theta_L}{\theta_H-\theta_L}$.

Intuitively, in the equilibrium so constructed, the high-ability type chooses the level of eduction that is just high enough that the low-ability type would not want to imitate and would thus simply choose $e = 0$. The beliefs of the firms are $\mu^*(0) = 0$ and $\mu^*(\widetilde{e})$, which are correct given the strategy of the worker. Off the equilibrium path, however, there can be many different beliefs that would support the equilibrium outcome, and so can be many different wage offers. It is thus clear that there are many equilibria at which $e^*(\theta_H) = \widetilde{e}$, $e^*(\theta_L) = 0$. Two examples of these are given in the book.

In these separating equilibria of the model here, high-ability workers are willing to get otherwise useless education because it allows them to signal their high abilities and receive higher wages. This is possible because a high-ability worker has lower marginal cost of acquiring education.

There can also be separating equilibria that involve $e^*(\theta_L) = 0$ but $e^*(\theta_H) = e_1 > \widetilde{e}$. How high can e_1 be? It can be as high as $\theta_H - c(\theta_H, e_1) = \theta_L$. One possible equilibrium strategy of the firms is

$$w^*(e) = \begin{cases} \theta_L & \text{if} \quad e < e_1 \\ \theta_H & \text{if} \quad e \geq e_1 \end{cases}$$

with

$$\mu(e) = \begin{array}{ll} 0 \;\; \text{if} & e < e_1 \\ 1 \;\; \text{if} & e \geq e_1 \end{array}$$

To see that the strategies and beliefs constitute a PBE, notice that each type's choice of e is optimal given $w^*(e)$, $w^*(e)$ is the Nash equilibrium offers by the firms given their beliefs, and $\mu(e)$ is derived from the worker's strategy using Bayes' rule where possible. Note that these different separating equilibria can be Pareto ranked. Since firms all earn zero expected profit and the θ_L type's payoff is always θ_L, the separating equilibrium that gives the θ_H type the highest payoff, the one that has $e^*(\theta_H) = \widetilde{e}$, must Pareto dominate all other separating equilibria.

In any separating equilibrium, the θ_L type is worse off than if no signaling is possible. If signaling is not possible, then the equilibrium wage must be $E(\theta)$, which is higher than θ_L. Interestingly, being able to signal need not always make the θ_H type better off. This can be seen in the following two diagrams:

In the first case, in the separating equilibrium the θ_H type's indifference curve is above $E(\theta)$, which implies that it is better off when signaling is possible. In the second case, in the separating equilibrium the θ_H type's indifference curve intersects the w axis below $E(\theta)$, which implies that the θ_H type is worse off when signaling is possible. Thus the fact that signaling is possible can actually make everybody worse off. The intuition is that when signaling is possible, one who does not choose a high level of e could be perceived to have low ability and to be paid θ_L instead of $E(\theta)$. This makes the θ_H type choose $\widetilde{e}$ even though she would prefer to receiving $E(\theta)$ with $e = 0$. This is another example that having more choices could make a person worse off.

We now consider pooling equilibrium, in which two types of workers choose the same level of education, $e(\theta_H) = e(\theta_L) = e^*$. The equilibrium beliefs of the firms, which should be derived from the worker's strategy and Bayes' rule where possible, must be $\mu(e^*) = \lambda$. Thus the wage offers at any pooling equilibrium must be $w^*(e^*) = \lambda\theta_H + (1-\lambda)\theta_L = E(\theta)$.

What are the possible equilibrium values of e^*? The answer is that any $e \in [0, e']$ can be e^*, where e' is found by equation

$$E(\theta) - c(e', \theta_L) = \theta_L,$$

as is shown in the following diagram, where one possible $w^*(e)$ curve is drawn. Given $w^*(e)$, both types would be optimal by choosing e', the beliefs of the firms are derived from the Bayes' rule on the equilibrium path, and given these beliefs the wage offers constitute a Nash equilibrium in the wage-offering stage game.

Here, choosing a positive amount of e is a pure waste since no information is revealed, and yet it could be an equilibrium outcome because one fears that a deviation could be interpreted as coming from a low-ability type.

The fact that there can be multiple equilibria is disturbing, and it is considered a general problem in signaling models. This problem arises largely because the beliefs of the uninformed are not restricted off the equilibrium path. A large part of recent research has gone to studying what would be "reasonable" beliefs off the equilibrium path. These are called theories of equilibrium refinements.

Another area of research about signaling games is to consider multiple audience for the signals sent by some informed agents. When a firm tries to signal to the capital market, for instance, the signal may also be observed by the product market competitors.

Still another area of research about signaling games is to consider private information in multidimensions and also signaling in multidimensions.

Another mechanism through which private information may be revealed in the marketplace is screening. In this case, the uninformed market participants can design offers in such a way that an informed agent would self-select the offer that is best for her, which then reveals her type in equilibrium. In the labor market, for example, if firms first announce the wage offers for different education levels, then we could have an equilibrium where more able people choose higher level of education. When firms do not know consumers' valuations for a product, they may offer the product in different forms to have consumers self-select. In airline pricing, for instance, first-and economy-class tickets are offered.

One way to think of the difference between signaling and screening models is the order of moves: In a signaling model, the informed moves first, while in a screening model the uninformed moves first. One advantage of working with a screening model is the set of possible equilibria is typically smaller in a screening model than in a signaling model.