Notes on Adverse Selection

Jonathan Dworak

February 10, 1999

1 Introduction

It is evident that many real-world economic transactions are not conducted under conditions of symmetric information. To recall a few classic examples, sellers of used cars are better informed about the quality of the car than potential buyers, purchasers of insurance are likely to have a better assessment of the risks they face than insurance companies do, regulated firms know the costs they face better than government bureaucrats, and so on. In each case, at the time contracts are signed, one party to the trade is better informed about some issue of relevance to the transaction. We refer to such a case as one of *adverse selection* (as opposed to cases of *moral hazard*, in which information assymetries enter after the contract is established). In these notes I will illustrate the basic theoretical issues and types of models that deal with adverse selection.

2 The basic issue: the lemons problem

In models of adverse selection it is common to speak of an economic agent's type. The type of an agent will be a mathematical parameter which summarizes that agent's economic characteristics. For instance, an agent's type may index the intensity of his preference for a particular good. Assymetric information about an agent's type is what leads to the adverse selection problem. This can be illustrated in an example based on the classic statement of these issues, given by Akerlof.

We consider the used car market. There are buyers and sellers of used cars, and two quality levels of cars. We will identify the type of the seller with the quality of car he has to sell. Therefore denote a seller's type by θ, with $\theta \in \{\theta_H, \theta_L\}$, and of course $\theta_H > \theta_L > 0$. A seller gets θ dollars worth of utility if he keeps his car (that is, his reservation price is θ). There is only one type of buyer; a buyer buying from a seller of type θ receives 2θ dollars worth of utility, and buyers are risk-neutral expected utility maximizers. The proportion of high-type sellers in the population is p (and hence there are $1-p$ low-type sellers).

What will happen in this market? Let us consider first the case when everyone knows a seller's type. Assume that a buyer offers to pay what a car is worth to him. Then when a buyer is matched with a low-quality seller, the price offered is $2\theta_L$, the seller accepts (since $2\theta_L > \theta_L$) and trade occurs. Similarly high-quality cars are sold at price $2\theta_H$. This outcome, which occurs when there is no information problem, is referred to as the *first-best* outcome. Alternatively, the first-best outcome may be found by maximizing the total utility in the economy (and thus ignoring transfers of money). In this case, the first-best allocation has all the cars being transferred to the buyers, since any particular car is worth more to a buyer then a seller.

Now, let us introduce an information assymetry. Buyers no longer get to observe the quality of the car they are purchasing. A randomly drawn car is worth $2\left[p\theta_H + (1-p)\theta_L\right]$ to a buyer, so this is the most a buyer will offer. Sellers, on the other hand, know the value of their car. A high-type seller, for instance, will only accept this price if $2\left[p\theta_H + (1-p)\theta_L\right] \geq \theta_H$. If so, again all cars will be sold, just at different prices. First-best efficiency will still obtain (since the new prices just reflect a different transfer of money). However, if $2\left[p\theta_H + (1-p)\theta_L\right] < \theta_H$, high-quality sellers will withdraw their cars from the market since the car is worth more to them then they can get by selling it. If p is very close to zero, this could well happen. In that case, only the low-quality cars (the "lemons") will be offered for sale; "adverse selection" occurs in that this particular price causes only the worst goods to appear on the market. Of course, the buyers recognize this, and so offer only $2\theta_L$, and only the market for lemons clears. This *second-best* outcome is inefficient since the high-quality cars are also worth more to buyers than sellers. Because of the information assymetry, inefficiency arises.

To be concrete, suppose that $\theta_H = 3$ and $\theta_L = 1$. Then the expected value of a random car to a buyer is $2\left[p\theta_H + (1-p)\theta_L\right] = 4p+2$. If $4p+2 \geq 3$ ($p \geq \frac{1}{4}$), both types of car will be sold at this price. However, if $p < \frac{1}{4}$, the high-quality market will unravel, and only low-quality cars will be sold at a price of 2.

3 Screening

In models of adverse selection it is common to distinguish between screening and signalling models. In *screening* models, it is the uninformed parties that choose actions intended to distinguish among the different types of the informed agents. In *signalling* models, it is the informed agents who choose actions in an attempt to reveal to the uninformed parties their characteristics. Screening models are somewhat simpler so I discuss those now. Many of the issues encountered in the theory of contracts and organizations are dealt with in principal-agent models, in which an uninformed principal makes a take-it-or-leave-it offer to an informed agent, and thus screening models are also to some degree more prevelant.

One of the classic screening models deals with non-linear pricing. (Recall that a linear price p is charged whenever the total cost of q units of a good is pq; any deviation from this scheme results in non-linear pricing). The pricing of phone service, for example, is rife with non-linearities. Long-distance plans typically combine differing fixed fees with various per-minute rates, and consumers are asked to choose the fixed fee and marginal price that best suits their needs. One interpretation of this behavior is that the phone company is attempting to separate consumers by their valuation of phone service.

To be precise, suppose that their are two types of consumers, again indexed by θ, with $\theta \in \{\theta_H, \theta_L\}$ Consumers like phone service but they dislike paying money; and these preferences depend on type. In fact, assume that

$$U(q, t; \theta) = \theta q - t$$

where q is the quantity of phone service received and t is the total transfer paid to the phone company. A consumer may choose not to purchase any service in which case he receives zero utility. There are p high types and $1-p$ low types. The phone company wishes to maximize total expected profits; receiving t dollars for q units of service results in profits of $t - q^2$.

Let us first examine the case when the phone company observes the consumers type. The phone company must solve the following problem for each consumer type θ

$$\max t - q^2 \text{ subject to } \theta q - t \geq 0$$

The constraint that $\theta q - t \geq 0$ is referred to as the *individual rationality* (IR) or participation constraint. The phone company cannot offer terms that are so bad for the consumers that consumers choose not to purchase (the phone company can't force anyone to purchase anything). However, it is obvious that the phone company will choose q and t to make the IR bind with equality, since if he left any utility to the consumer he could suck it back in the form of an increased transfer t. Therefore we know that $\theta q - t = 0 \Rightarrow t = \theta q$. The problem becomes

$$\max_q \theta q - q^2$$

which has first-order condition $\theta - 2q = 0 \Rightarrow q^* = \frac{\theta}{2}$. Note that in the absence of assymetric information the monopolist has chosen to maximize the total social utility. The first-best outcome therefore has $q_H = \frac{\theta_H}{2}$, $q_L = \frac{\theta_L}{2}$. Also, the monopolist manages to suck back all the utility in society through the money he charges for service.

3.1 Incentive compatibility and the revelation principle

Now we introduce an information problem. The monopolist is no longer allowed to observe the type of potential consumers. It does not follow, however, that the monopolist will not offer different quantities of service and different transfers to the two consumer types. By allowing consumers to choose among different contracts, the phone company may be able to distinguish among the two groups. Of course, since consumer types are unobserved by the phone company, the phone company must take into account the incentives of the consumers to reveal their types through the choice of contract. As translated into theory, this leads us to the consideration of incentive compatibility constraints and the revelation principle.

One could imagine a number of schemes, potentially quite complicated, that the monopoilist could use to sell the good. The *revelation principle* states the (perhaps surpising) result that the outcome of any such scheme can be replicated by asking that consumers simply report their types to the monopolist, with outcomes depending on those reports. The only additional (but important) requirement is that the consumers find it in their interests to report their types truthfully[1]. This allows us to express the effect of asymmetric information on the principal's problem in a very convenient manner. In our case, the monopolist chooses (q_H, t_H, q_L, t_L), where q_H is the quantity of service an agent gets if he announces that he is a high type, and so on. The *incentive compatibility constraints* reflect the requirement that agents wish to truthfully reveal their types:

$$\begin{array}{ll} \mathrm{IC}_H & \theta_H q_H - t_H \geq \theta_H q_L - t_L \\ \mathrm{IC}_L & \theta_L q_L - t_L \geq \theta_L q_H - t_H \end{array}$$

In each case, the left-hand side of the IC is the utility an agent receives when he truthfully announces his type. The right-hand side is the utility an agent receives if he lies. For instance, a high type who lies receives (q_L, t_L); however, the utility he gets is $\theta_H q_L - t_L$ since his true type is θ_H.

What would happen if we ignored these constraints? Suppose the phone company continued to offer $\frac{\theta_H}{2}$ for price $\frac{\theta_H^2}{2}$, and $\frac{\theta_L}{2}$ for price $\frac{\theta_L^2}{2}$. Now, the high types do better for themselves by purchasing the low-type's bundle! If they stick to their "own" bundle, the receive zero utility, while if they purchase the low bundle, they receive $\frac{1}{2}\theta_L(\theta_H - \theta_L) > 0$. Again, the adverse selection problem manifests itself in that the contract that is offered selects only certain types in the market. In this case, the high-quantity market collapses and too many consumers purchase a low quantity.

It is now easy to set up the phone company's problem taking the constraints into account. The company seeks to maximize total expected profits, subject to the IC and IR constraints:

$$\max p\left(t_H - q_H^2\right) + (1-p)\left(t_L - q_L^2\right) \quad \text{subject to:}$$

$$\begin{array}{ll} \mathrm{IC}_H & \theta_H q_H - t_H \geq \theta_H q_L - t_L \\ \mathrm{IC}_L & \theta_L q_L - t_L \geq \theta_L q_H - t_H \end{array}$$

$$\begin{array}{ll} \mathrm{IR}_H & \theta_H q_H - t_H \geq 0 \\ \mathrm{IR}_L & \theta_L q_L - t_L \geq 0 \end{array}$$

[1] There are really two parts to the revelation principle. One is that it suffices to restrict the agents' actions to reporting their type; we don't need to consider alternative schemes. The second is that there is no loss in insisting that the agents truthfully reveal their types.

Now, one may simply set up a Lagrangian and solve the above problem with brute force. It is often instructive, however, to analyze the constraints, with the hope of eliminating some of them. First, note that IC_H and IR_L together imply IR_H:

$$\theta_H q_H - t_H \geq \theta_H q_L - t_L \geq \theta_L q_L - t_L \geq 0$$

where the middle inequality follows from the fact that $\theta_H > \theta_L$. Thus we can ignore IR_H; if we find a solution that satisfies the other three constraints, it will satisfy IR_H as well. As to the IC constraints, subtract IC_L from IC_H

$$(\theta_H - \theta_L)\, q_H \geq (\theta_H - \theta_L)\, q_L \Rightarrow (\theta_H - \theta_L)\,(q_H - q_L) \geq 0 \Rightarrow q_H \geq q_L$$

where again the last inequality follows from the fact that $\theta_H > \theta_L$. To satisfy the IC constraints, the monopolist will find it necessary to offer a higher quanity of service to the higher-valuation consumers. Furthermore, note that IC_H is an equality; if it were not, we would have the following (using IR_L):

$$\theta_H q_H - t_H > \theta_H q_L - t_L \geq \theta_L q_L - t_L \geq 0$$

which means that the monopolist could increase t_H without affecting incentive compatibility ore the IR constraints. Conveniently if IC_H is an equality, and using the fact that $q_H \geq q_L$ we have

$$\begin{aligned} t_H - t_L = \theta_H\,(q_H - q_L) &\geq \theta_L\,(q_H - q_L) \Rightarrow \\ \theta_L q_L - t_L &\geq \theta_L q_H - t_H \end{aligned}$$

so that IC_L can be ignored. Also, not surprisingly, it must be that IR_L is an equality as well (if not, both types are receiving positive utility, and the monopolist would suck that back through increased transfers). The point of all this discussion is that we may rewrite the phone company's problem as

$$\begin{aligned} \max p\left(t_H - q_H^2\right) + (1-p)\left(t_L - q_L^2\right) \quad & \text{subject to :} \\ & \theta_H q_H - t_H = \theta_H q_L - t_L \\ & \theta_L q_L - t_L = 0 \\ & q_H \geq q_L \end{aligned}$$

It is traditional to ignore the $q_H \geq q_L$ constraint, since it is typically automatically satisfied (in the first-best world the monopolist likes to choose $q_H \geq q_L$). Solving for t_H and t_L from the constraints, we have the simple problem

$$\max_{q_H, q_L} p(\theta_H q_H - \theta_H q_L + \theta_L q_L - q_H^2) + (1-p)\left(\theta_L q_L - q_L^2\right)$$

with solution

$$\begin{aligned} q_H^* &= \frac{\theta_H}{2} \\ q_L^* &= \frac{\theta_L}{2} - \frac{p}{1-p}\frac{\theta_H - \theta_L}{2} \end{aligned}$$

There are a number of things to take note of. First, the high types continue to receive the first-best quantity level. Essentially, the high types are the most efficient producers of social surplus, and the more social surplus, the more the monopolist can suck back. The information assymetry, however, leads the monoplist to distort the quantity offered the low types. As we saw before, at the first-best contract, it is the high types who face an incentive to switch to the low-quantity contract. By lowering the quantity offered the low types, the phone company makes the low contract relatively more unattractive to the high types. Alternatively, note that since $t_L = \theta_L q_L$, and $\theta_H q_H - t_H = \theta_H q_L - t_L$, it follows that the utility the high types get from their bundle is equal to $q_L(\theta_H - \theta_L)$. By lowering q_L, the monopolist lowers the rent left over for the high types. Finally, just to be explicit, we can easily calculate that

$$
\begin{aligned}
t_L^* &= \theta_L q_L^* = \theta_L \left(\frac{\theta_L}{2} - \frac{p}{1-p} \frac{\theta_H - \theta_L}{2} \right) \\
t_H^* &= t_L^* + \theta_H \left(q_H^* - q_L^* \right) = \theta_L \left(\frac{\theta_L}{2} - \frac{p}{1-p} \frac{\theta_H - \theta_L}{2} \right) + \theta_H \left(\frac{\theta_H - \theta_L}{2(1-p)} \right)
\end{aligned}
$$

These expressions are not that illuminating; note that from an efficiency standpoint the q's are much more important since the t's are just transfers.

Let's put some concrete numbers on this stuff. Take $p = \frac{1}{2}$, $\theta_H = 4$, $\theta_L = 3$. We have

$$
q_H^* = 2, q_L^* = 1, t_H^* = 7, t_L^* = 3
$$

Now, we solved this problem by using the revelation principle. Obviously, in the real world, we very rarely see schemes in which consumers are literally expected to announce their valuation for a good. A more realistic possibility is that the phone company would allow consumers their choice of calling plans. One plan offers 2 units of service for a price of 7, while the second offers 1 unit of service for the price of 3. The above considerations show that the high type consumers will choose the higher quantity of service, and so on[2].

3.2 Bargaining and buyer-seller models

Additional complications arise when there is private information on both sides of a market. Even more than in the classic lemons problem, bilateral assymetric information may make it impossible for efficient trades to take place; this possibility provides a theoretical justification for inefficient outcomes of bargaining processes. I will change notation somewhat. There are two agents, a buyer and a seller, and a good to be traded. The value to the buyer in dollars of the good is v, the value to the seller of keeping the good is c. There are two possible types of buyer, so that $v \in \{v_L, v_H\}$, and two types of sellers with $c \in \{c_L, c_H\}$. Assume that the proportion of high type buyer is π_b and of high type sellers is π_s, and that valuations are drawn independently. Since we ignore transfers,

[2]This pricing scheme is non linear in that the price of 2 units is not twice the price of 1.

the efficient outcome is for the buyer to get the good whenever $v > c$, and for the seller to keep the good whenever $c < v$. When can this be sustained? Once again, we can examine the possibilities by using the revelation principle: ask that the buyer and the seller each report their valuations, and devise a mechanism under which they find it advantageous to truthfully report their types. I shall restrict attention to schemes that are budget-balanced; the only payments allowed are transfers from the buyer to the seller (as opposed to transfers to both from some outside third party). I also assume that either party can refuse to trade. A buyer who refuses to trade gets zero. A seller who refuses to trade gets c.

Let's consider the two simplest cases: $v_H > v_L > c_H > c_L$ and $c_H > c_L > v_H > v_L$. In the first case, no matter what the valuations of each party, it is efficient for the buyer to own the good. In the second case, no matter what the valuations of each party, it is efficient for the seller to keep the good. In each case, it is easy to devise an incentive compatible scheme which implements this outcome. In the second case, the scheme requires no transfer of money or of the good no matter what the agents report. In the first, we can choose a price p such that $v_L > p > c_H$. Then, the scheme specifies that independent of the reports of the two agents, the buyer gets the good for price p. Since reports don't affect the outcome, no agent has an incentive to misrepresent his type. And, we have chosen p so that all types of both agents prefer to participate in the scheme than not. Note the similarity of this scheme to the standard neoclassical market mechanism, in which no agent can affect the price at which a good is sold.

The interesting case is when $v_H > c_H > v_L > c_L$. That is, it is no longer common knowledge that trade should occur (or not occur). When a high-cost seller is matched with a low-valuation buyer, the seller should keep the good. We shall examine this situation using the revalation principle. If we want efficient trade to take place, the buyer must receive the good whenever the type profiles are (v_H, c_H), (v_H, c_L), or (v_L, c_L). The seller should keep the good whenever the outcome is (v_L, c_H). The only thing left to determine are the transfers. Since their are four possible profiles of reports, there are four levels of payment from the buyer to the seller to consider. Rather than deal with the full-blown set of IR and IC constraints (of which their are 4 each), let's just consider the incentives facing the agents. High-valuation buyers have an incentive to pretend to have a low-valuation, and thus get a lower price. However, by such misrepresentation they face the possibility that trade will not occur, since they may be matched with a high-cost seller. Similarly, a low-cost buyer may try to extract a higher price at the risk of trade not occuring. We know that the price paid by a high-valuation buyer to the high-cost seller must be at least c_H; otherwise the high-cost seller will refuse to participate. Also, the price paid by a low-valuation buyer to a low-cost seller can be no higher than v_L, otherwise the low-valuation buyer will refuse to participate. In fact, by making the prices in these cases exactly c_H and v_L we mute the incentives of agents to misrepresent their valuations (since we have minimized the price you get when claiming to be a high-cost seller, and maximized the price you pay when claiming to be a low-valuation buyer). Thus the allocation rule as a function of reports must be:

$$\begin{aligned} p(v_H, c_H) &= c_H \\ p(v_L, c_H) &= \text{(no trade)} \\ p(v_L, c_L) &= v_L \\ p(v_H, c_L) &= p \end{aligned}$$

where p is a number we have yet to determine. We have two incentive constraints to consider; that of the high-valuation buyer and the low-cost seller.

$$\begin{aligned} v_H - \pi_s c_H - (1-\pi_s)p &\geq (1-\pi_s)(v_H - v_L) \\ \pi_b p + (1-\pi_b)v_L - c_L &\geq \pi_b (c_H - c_L) \end{aligned}$$

Is it possible to find a p that satisifes both of these constraints? Not surprisingly, it depends. Let's choose some numbers:

$$v_H = 3, v_L = 1, c_H = 2, c_L = 0$$

so the constraints become

$$\begin{aligned} 3 - 2\pi_s - (1-\pi_s)p &\geq 2(1-\pi_s) \\ \pi_b p + (1-\pi_b) &\geq 2\pi_b \end{aligned}$$

From the second constraint, it must be that $p \geq 3 - \frac{1}{\pi_b}$. From the first constraint it must be that $p \leq \frac{1}{1-\pi_s}$. To be able to find such a price, the distribution of buyer's and seller's types must be such that

$$3 - \frac{1}{\pi_b} \leq \frac{1}{1-\pi_s} \Rightarrow \pi_s \geq \frac{2\pi_b - 1}{3\pi_b - 1}$$

So that, for example, if $\pi_b = \frac{2}{3}$ and $\pi_s = \frac{1}{4}$, then no choice of price can satisfy the constraints.

What is going on? Again, the incentive for both agents is to misrepresent their valuations in the hopes of getting a more favorable price. The countervailing incentive is that misrepresentation increases the chances that trade will not occur. From the standpoint of the low-cost seller, a higher π_b (higher probability of a high-valuation buyer) means that the costs of misrepresenting are very low; similarly, a low π_s means a buyer's misrepresentation does not cause a very high probabiliy of a no trade outcome. Thus, we see that the above condition is most easily satisfied when π_s is high and π_b is low; otherwise, we need to use price as an incentive to truthfully report. Parameters may be such that no price can induce efficient trade. We may interpret this case as one of inefficient bargaining; even though there are gains from trade, the incentives of agents to capitalize on their private information may lead them to reject mutually beneficial exchanges.

4 Signalling

In signalling models we allow the privately informed agents to take actions which may allow them to reveal their type. The key difficulty in signalling models is that the benefit of sending a signal depends on how it will be interpreted by the receiver (the uninformed party). This is in contrast to screening models in which the uninformed party knows how the privately informed agents will react to the incentives he constructs. In any case, signalling is likely to be a very important real-world phenomenon. I shall illustrate the basic properties of signalling models in the classic Spence education model.

Suppose there are two types of workers, θ_H and θ_L, in proportion π and $(1-\pi)$ respectively. A worker's type is his productivity in a particular job; in fact, assume that the value of the worker's output is exactly θ. Worker's are informed of their productivity; firms are not. A worker may, if he chooses, choose to obtain some amount of education e. The cost of education depends on a worker's type: $C(e) = \frac{e}{\theta}$. Note that education adds nothing whatsoever to a worker's productivity, nor does the worker get any enjoyment out of consuming education. From the social standpoint, education is a pure waste. Nevertheless, we shall find that some workers may choose to get an education if it allows them to signal that they are high-productivity types. The key will be that high types face a lower marginal cost of education.

Assume that "competition" among the firms forces them to pay the expected productivity of a worker, conditional on his education level. Workers like higher wages and, of course, lower education. Social efficiency would have no worker getting any education. If both types of workers chose not to be educated, firms would not be able to infer anything about their productivities. We would have a single wage paid to a worker with no education with $w^P = \pi\theta_H + (1-\pi)\,\theta_L$. Is this an equilibrium? It depends on what the firms believe about workers with higher levels of education. For instance, they may simply assume that any worker who chose $e > 0$ had simply made a mistake, and that mistakes were equally likely to come from high and low types. Then again the firms will pay only w^P to any worker, and no worker will find it profitable to become educated. This type of equilibrium is called a *pooling equilibrium*, since both types of workers take the same action. Note, however, that the high types are getting a raw deal in the sense that they would be paid more (θ_H) if they could somehow reveal their true productivity. On the other hands, the low types benefit in that they are mixed in with the high types and get higher wages as a result.

It is possible, therefore, to sustain an equilibrium in which the high types distinguish themselves from the low types by obtaining education. I will look for an equilibrium with the lowest amount of education, since that is the most efficient. Low types will be assumed to get zero education, and high types will get some number e^H which is yet to be determined. Once again we will have to deal with incentive constraints. Firms will pay any worker who gets zero education θ_L, and any worker who chooses $e = e^H$ gets paid θ_H. A high worker may wish to choose zero education (avoiding those costs) even though he gets paid less; more likely, e^H must be high enough so that a low worker does not wish to imitate the high worker in the hopes of getting a higher wage. Thus we have:

$$\theta_H - \frac{e^H}{\theta_H} \geq \theta_L \Rightarrow e^H \leq \theta_H(\theta_H - \theta_L)$$
$$\theta_L \geq \theta_H - \frac{e^H}{\theta_L} \Rightarrow e^H \geq \theta_L(\theta_H - \theta_L)$$

Note that since $\theta_H > \theta_L$ it is always possible to satisfy both constraints. We see, in fact, that the lowest level of education that can sustain this *separating equilibrium* is $e^H = \theta_L(\theta_H - \theta_L)$. Of course, workers may choose levels of education other than e^H or zero. Whether they wish to do so or not depends on what firms believe about them when they observe these out-of-equilibrium choices. A simple specification of such beliefs is that firms assume any worker not getting education level e^H is a low type. Other beliefs could break the above equilibrium.

The point of all this is the value of education as a signal even when it has no intrinsic value. Getting more education can prove that you're a smart person, for example, even if what you actually learn in school doesn't help you in your job at all. Since high types face a lower marginal cost of education, they are more willing to trade education for higher wages than the low types; this fact allows workers to separate themselves by education level. Although it may be extreme to think that education has no value whatsoever, it does seem that the premium paid for a college education, for example, is in excess of the actual value of that education in workplace (especially, perhaps, in the case of liberal arts education).

A brief word is in order on the pathologies of the signalling models. Since the value of the signal depends on the firms' beliefs, we have seen already that multiple equilibria are possible. Here, in fact, we have illustrated the two most efficient such equilibria. On the other hand, perverse beliefs by the firms could cause really inefficient equilibria to arise, equilibria in which, for example, workers pool on a positive level of education. Similarly, there are separating equilibria with education levels higher than e^H. It may be an instructive exercise to think about which firm beliefs would generate such outcomes, and whether these seem plausible or not.

5 Further reading

Milgrom and Roberts (1992) have a nice overview of the issues associated with adverse selection in chapter 5, pages 149ff. The buyer-seller model of section 3.2, in particular, is derived from their discussion. The orginal lemons model was developed in Akerlof (1970), if you feel like digging it up. Spence (1973) is a reasonably non-technical book that summarizes his work on signalling. The standard screening model is dealt with in many places; Salanie (1997) is a recent, lucid presentation. And Mas-Colell, Whinston and Green (1995) give all of the gory details in chapter 13 and 14, but I caution you that their presentation is very technical and far beyond the level of this course.

References

[1] Akerlof, "The Market for Lemons", *Quarterly Journal of Economics*, 1970.

[2] Mas-Colell, Whinston, and Greene, *Microeconomic Theory*, 1995.

[3] Milgrom and Roberts, *Economics, Organization, and Management*, 1992.

[4] Salanie, *The Economics of Contracts*, 1997

[5] Spence, *Market Signalling*, 1973