Lecture 19 - Adverse Selection: Applications and Extensions

14.03 Spring 2003

1 Adverse Selection: Applications and Extensions

- This lecture extends the discussion of adverse selection beyond the insurance case presented in Rothschild-Stiglitz. The core results in the R-S model depend upon risk aversion (as well as private information); risk aversion is what gives rise to the existence of a separating equilibrium if one exists at all. This raises the question: is adverse selection only important in markets for risk?

- The answer it turns out is a resounding no. The paper by Akerlof (1973) discussed below develops the case of adverse selection in markets for heterogenous goods (in this case, used cars). This example has nothing to do with risk; it's simply asymmetric information. We'll discuss two principles that arise from the Akerlof analysis:

  1. The "Lemons Principle"
  2. The"Full Disclosure Principle"

- It turns out that these principles are roughly inverses.

1.1 Example: Altman, Cutler, Zeckhauser (1998)

- Before this, let's apply the R-S adverse selection idea to an actual health insurance market.

- Group Insurance Commission of Massachusetts: Data visible to individuals *but not known to their insurance plans at the time of enrollment.*

- Three policies offered:

  1. Traditional indemnity: Most generous. You pay for whatever services you want, the insurance company pays you back.

2. PPO (Preferred Provider Organization): Less generous, but few restrictions on utilization.
3. HMO (Health Maintenance Organization): Most restrictive. They limit who you can see and what care is administered.

- Question 1: Why is indemnity plan so much more expensive than HMO plan?
  Table I

	Premium	Unadjusted Benefit Cost	Age/Sex Adjusted Benefit Cost
Indemnity	$2,670	$2,176	$1,908
HMO	$1,631	$1,115	$1,320
Difference	$984	$943	$588

- It appears that about $\frac{2}{3}rds$ of the difference in the premium of the indemnity versus HMO plan is *not* explained by the age and sex composition of the plan clientele. Why?
  - One explanation is moral hazard: people spend more when they get on this plan.
  - Another is adverse selection: people with greater taste for spending on health care choose the Indemnity plan.
  - Most likely explanation is a combination of the two.
- Table II:
  - People leaving indemnity plan for the HMO have average cost of $\$1,444$.
  - People staying in indemnity have average cost of $\$2,252$.
  - This is a 37% difference!
- Similarly:
  - People leaving HMO for indemnity have cost of $\$1,615$.
  - People staying in HMO have cost of $\$1,125$.
  - This is a 47% difference.
- However, notice that plan 'immigrants' spent about average for the new plan they enter once they are in it.

	Indemnity	HMO
Stayers in 2nd year	$1,960	$1,344
Immigrants in 2nd year	$2,095	$1,181

- Table III
  - Accounting for adverse selection:

    * If no switching were allowed: Indemnity cost would be $16 lower per person.
    * If no switching were allowed: HMO costs would be $9 higher per person.
  - New enrollees:
    * By contrast, new enrollees *reduce* average costs by about $20 per person per plan per year ($20 in indemnity plan, $19 in HMO). Why?
  - Adverse retention:
    * If costs are convex in age (that is, rise nonlinearly as people age)
    * And if stayers are older at indemnity plans than HMOs
    * Then plan premium will rise faster at indemnity than HMO plan.
    * Over time, this will cause more young people to depart HMOs, bringing premiums up further.
    * The effect of "adverse retention" is $23 at indemnity plan versus $9 at HMO.

- So, program costs rise about $40 faster at indemnity than HMO, which may not be stable. Adverse selection spiral of death is conceivable...

1.2 Adverse Selection: The Market for Lemons (Akerlof, 1970)

- Now, we extend the case of adverse selection outside of insurance markets. It turns out, the problem is far more general than the R-S model might suggest.
- The fundamental problem:
  1. Goods of different quality exist in the marketplace.
  2. Owners/sellers of goods know more about their goods' quality than do buyers.
  3. Critical insight of Akerlof: *Potential buyers know that sellers know more about the quality of goods than they do.*
- This information asymmetry dramatically changes the market.
- It can easily be the case that there is *no trade* whatsoever for a given good even though:
  1. At any given price, there are traders willing to sell their products.
  2. At this price, there are buyers willing to buy the product.

- Akerlof (1970) was the first economist to analyze this paradox rigorously in a paper that eventually led to a Nobel prize (simultaneously with the Rothschild-Stiglitz and Spence papers that we're covering).

- Akerlof's paper was nominally about the market for used cars. It's always been folk wisdom that it's a bad idea to buy used cars–that 'you are buying someone else's problem.' But why should this be true? If used cars are just like new cars only a few years older, why should someone else's used car be any more problematic than your new car after it ages a few years?

- Let's take a simple example.

  – There are 2 types of *new* cars available at dealerships: good cars and lemons (which break down often).
  – The fraction of lemons at a dealership is λ.
  – Dealers do not distinguish among good cars versus lemons – they sell what's on the lot at the sticker price.
  – Buyers cannot tell good cars from lemons apart, but they know that some fraction λ will be lemons.
  – After buyers have owned the car for any period of time, they can tell whether or not they have bought a lemon.
  – Assume that good cars are worth $\$2,000$ and lemons are worth $\$1,000$.
  – Finally, for simplicity (and without loss of generality), assume that cars do not deteriorate and that buyers are risk neutral.

- First, what is the equilibrium price for new cars? Clearly this will be

$$P^N = (1 - \lambda) \cdot 2,000 + \lambda \cdot 1,000.$$

- Since dealers sell all cars at the same price, buyers pay the expected value of a new car.

- Now, let's consider the used car market. Assume that used car sellers are willing to part with their cars at 20% below their new value (perhaps they need cash to buy looted Iraqi art treasures). So,

$$P_G^U = \$1,600 \text{ and } P_L^U = \$800.$$

- Since cars don't deteriorate, used car buyers will be willing to pay $\$2,000$ and $\$1,000$ respectively for used good cars and lemons. Hence, there is a surplus of \$400 or \$200 gain from trade from each sale. Selling either a good car or a lemon yields a Pareto improvement.

- Question: What will be the equilibrium price of used cars?

- The natural answer is

$$P^U = (1 - \lambda) \cdot 1,600 + \lambda \cdot 800,$$

  but this is not necessarily correct.

- Recall that buyers cannot distinguish good cars from lemons whereas owners of used cars know which is which. Assuming sellers are profit maximizing, this means that at any $P^U \geq 800$, owners of lemons will gladly sell them. But at $P^U < 1,600$, owners of good cars will *keep* their cars.

- So, the equilibrium price of used cars will depend on λ.

- Take the case where $\lambda = 0.4$. In this case, the expected value (to a buyer) of a randomly chosen used car–assuming both good cars and lemons are sold–would be

$$P^U_{\lambda=.4} = (1 - .4) \cdot 2000 + .4 \cdot 1000 = 1,600.$$

  Here, used cars sell at exactly the average price at which potential sellers value them. Owners of good cars are indifferent and owners of lemons get a \$800 surplus.

- But now take the case where $\lambda = 0.5$. If both good cars and lemons were sold, would the expected value of a randomly chosen used car be

$$P^U_{\lambda=.5} = (1 - .5) \cdot 2000 + .5 \cdot 1000 = 1,500?$$

- The answer is clearly no. Since owners demand \$1,600 for their good used cars, they will clearly not sell them. This means the only used cars on the market at $P^U_{\lambda=.5} = 1,500$ are lemons.

- But if this is so, then the equilibrium price of used cars cannot be \$1,500. Instead, it must be $P^U_{\lambda=.5} = \$1,000$.

- In other words, if the share of lemons in the overall car population is high enough ($\lambda > .4$ in this example), the bad products drive the good ones out of the market. Although buyers would gladly pay \$2,000 for a good used car, their inability to distinguish good cars from lemons means that they are not willing to pay more than \$1,500 for *any* used car. In this case, no good cars are sold, and so the equilibrium price must fall further to exclusively reflect the value of used lemons.

- The key insight here is that market quality is *endogenous,* it depends on price. When sellers have private information about products' intrinsic worth, they will only bring good products to market when prices are high.

- Buyers understand this, and so must adjust the price they are willing to pay to reflect the quality of the goods they expect to buy at that price.

- *In equilibrium, goods available at a given price must be worth that price.* If they are not, then there will be no equilibrium price and it's possible that no trade will occur (which is the case in the lemons model in the Akerlof paper.)

1.3 A richer example

- Now that you've seen a stylized example, let's go through the same logic with a richer example. In particular, we want to use a continuous distribution of product quality (rather than just two quality types).
- Consider the market for 'fine' art. Imagine that paintings are worth between \$0 and \$100,000 dollars to sellers (denote this as V_s), and that they are uniformly distributed between these two values, so the average painting is worth \$50,000 to a seller.
- Assume that the only way to know the value of a painting is to buy it and have it appraised. So, buyers cannot tell masterpieces from junk, but sellers can.
- Assume that buyers value paintings at 50% above the seller's price (denote this as V_b). Hence, if a painting has $V_s = \$1,000$ then $V_b = \$1,500$.
- What is the equilibrium price of paintings in this market? An equilibrium price must satisfy the condition that the goods that sellers are willing to sell at this price are worth that price to buyers.
- Take the sellers' side first. A seller will sell a painting if $V_s \leq P$. What is the expected seller's value of paintings for sale as a function of P? Given that paintings are distributed uniformly, this implies that

$$E(V_s|V_s \leq P) = \frac{0+P}{2}.$$

  So, if $P = 100,000$ then *all* paintings are available for sale and their expected value to sellers is \$50,000. Or if $P = 50,000$, the expected seller value of paintings for sale is \$25,000.

- Now take the buyer's side. Since the $V_b = 1.5 \cdot V_s$, buyers' willingness to pay for paintings as a function of their price is

$$E(V_b|P) = 1.5 \cdot E(V_s|V_s \leq P) = 1.5\left(\frac{0+P}{2}\right) = \frac{3}{4}P.$$

- So, the equilibrium price satisfies

$$\frac{1}{2}P^* = \frac{3}{4}P^*,$$

  which has only one solution: $P^* = 0$. The equilibrium price in this market is \$0, and so there is no trade.

- Given that buyers value paintings strictly above the seller's price, this result seems rather ironic. What's going wrong?

- The answer is similar to the Rothschild-Stiglitz model. The sellers of low-quality goods generate a negative externality for sellers of high quality goods. For every \$1.00 the price rises, seller value only increases by \$0.50 because additional low-quality sellers crowd into the market $\left(\frac{\partial E(V_s|V_s \leq P)}{\partial P} = 0.5\right)$.

- And for every dollar that the price rises, buyer value only increases by \$0.75 $\left(\frac{\partial E(V_b|P)}{\partial P} = 0.75\right)$. And so the twain do not meet.

- This is in effect the "Lemons Principle" – The goods available at a given price are worth less than or equal to that price (to sellers).

- Thus in this extreme example, there is no trade.

1.4 Reversing the Lemons equilibrium

- Is there any way around this result? Intuition should suggest that the answer is yes. Sellers of good products clearly have a strong incentive to their products' quality – otherwise, they cannot sell the products at their true value (and may not be able to sell them all).

- So why does this type of disclosure not occur in the example above? Because I have stipulated that the value of a piece of art can only be assessed ex-post by appraisal. So sellers of good paintings have no credible means to convey that their paintings are of good quality (b/c low quality seller will also claim that they have high quality paintings).

- So what is needed is a means to disclose information credibly. Intuition should now suggest that *if* there is an inexpensive (or free) means to credibly disclose the quality of paintings, sellers of above average paintings will probably want to do this. In fact, we will find the result is much stronger than this.

- Imagine now that a seller of a painting can spend \$1,000 to get a certified appraisal. This appraisal will credibly convey the true seller's value of the painting (and so the buyer's willingness to pay will be 1.5 times this value). Who will pay to get this appraisal?

- Your first guess might be that since buyers are willing to pay \$75,000 for a painting of average quality, any seller with a painting that would sell for at least \$76,000 if appraised (that is, the average sale price of non-appraised paintings plus the appraisal cost) would choose to get an appraisal. This is correct but incomplete.

- The gain to getting an appraisal α is the following:

$$\alpha = 1.5 \cdot V_s - 1000 - P^*.$$

  That is, the seller's gain is 1.5 times their value of the work of art minus the appraisal fee minus the price they would have gotten for a non-appraised painting.

- So, clearly if $P^* = 75,000$, then a seller with painting of $V_s \geq \frac{75000+1000}{1.5} =$ \$50,667 should pay for an appraisal.

- Are we done? No. Because we now need to determine P^*.

- Assume initially that $P^* = 75,000$. Could this be an equilibrium?

- Note that for any P^*, all paintings with

$$V_s \geq \frac{P^* + 1000}{1.5}$$

  will be appraised. So, the price P^* only applies to *non-appraised* paintings. A seller who has a painting valued $V_s \geq V_s^* = \frac{P^*+1000}{1.5}$, will appraise it and sell it for $1.5 \cdot V_s$. Only for $V_s < V_s^*$ will the seller will not bother to appraise the painting.

- Now clearly, buyers must take this fact into account when they decide how much to pay for a non-appraised painting.

- So, the equilibrium price of a non-appraised painting must be

$$P^* = 1.5 \cdot E(V_s | V_s < V_s^*) = 1.5 \cdot \frac{0 + V_s^*}{2}.$$

  And we have already established that

$$V_s^* = \frac{P^* + 1000}{1.5}.$$

- So combining these equations, we have

$$\begin{aligned} P^* &= 1.5 \cdot \frac{\frac{P^*+1000}{1.5}}{2} = \frac{P^* + 1000}{2} \\ &\Rightarrow P^* = 1,000 \\ &\Rightarrow V_s^* = \frac{1000 + 1000}{1.5} = 1,333.3. \end{aligned}$$

- In other words:

  1. Sellers with $V_s < 1,333$ do not have their paintings appraised and these paintings sell for \$1,000.

2. Sellers with paintings $V_s \geq 1,333$ have their paintings appraised, and these paintings sell for $V_s \cdot 1.5$.

- This result is quite counter-intuitive, so let's check it. If this is an equilibrium, then a seller with a painting equal to V_s^* must be indifferent between selling it at P^* or instead appraising it for \$1,000 and selling it at $V_s^* \cdot 1.5 = \$2,000$.

- Confirm. If the seller does not have the painting appraised, consumer's will value it at:

$$V_b = 1.5 \cdot E(V_s | V_s \leq 1,333.3) = 1.5 \cdot 666.7 = 1,000,$$

  so a non-appraised painting will sell for \$1,000.

- If the seller *does* have the painting appraised, his profits net of the appraisal cost will be

$$1.5 \cdot 1333.3 - 1000 = 2,000 - 1,000 = 1,000.$$

- So, the seller with $V_s = 1,333.3$ does equally well by appraising or not. And if $V_s < 1,333.3$, the seller does better by not appraising since P^* for a non-appraised painting is \$1,000. And for a seller with $V_s > 1,333.3$, the price of an appraised painting is always above \$1,000 and so these paintings will be appraised. So, we seemingly have an equilibrium price of $P^* = \$1,000$.

- However, as was pointed out by a student at the end of class, there is a problem with this equilibrium too. The problem is that at $P^* = \$,1000$, sellers with paintings of $V_s = 1,333.3$ will not want we willing to sell them; their own private valuation is higher.

- As in the adverse selection case above, the sellers in the market at $P^* = 1,000$ have $V_s \leq 1,000$. And so the average buyer value of paintings sold for \$1,000 is $1.5 \cdot 500 = \$750$. So, as above, there will be *no trade* in unappraised paintings, meaning that $P^* = 0$.

- Given this, which painting owners will get their paintings appraised. It turns out, we need two conditions:

$$\begin{aligned} V_s^* &\geq \frac{P^* + 1000}{1.5} \text{ and} \\ 1.5 \cdot V_s^* - 1,000 &\geq V_s^*. \end{aligned}$$

  The second condition stipulates that the market price for the painting minus the appraisal cost is greater than the seller's value–otherwise the seller doesn't sell.

- Solving (at last), we obtain that $V_s^* = 2,000$. This is an equilibrium because for a seller with a painting of $V_s = 2,000$ can sell it, if appraised, for \$3,000. Since the appraisal cost is \$1,000, the seller's net revenue is \$2,000, which is exactly equal to V_s. For $V_s > \$2,000$, the seller will also be willing to get an appraisal, and for $V_s < \$2,000$, the seller would rather keep the painting.

- What is going on? What is operative here is the 'Full-Disclosure Principle.' Roughly stated: If there is a credible means for an individual to disclose that he is above the average of a group, she will do so. This disclosure will implicitly reveal that other non-disclosers were below the average, which will give them the incentive to disclose, and so on... In equilibrium, everyone will explicitly or implicitly disclose their private information. (But if there is a cost to disclosure as in the appraisal example above, there will typically be a subset of sellers who do not find it worthwhile to disclose.)

- The Full Disclosure Principle is essentially the inverse of the Lemons Principle. In the Lemons case, the bad products drive down the price of the good ones. In the Full Disclosure case, the good products drive down the price of the bad ones. What distinguishes these cases is simply whether or not there is a credible disclosure mechanism (and what the direct discolure costs are).

1.5 A simpler example of the full-disclosure principle

- The example above is more complicated than need to make the point. Consider the following simple case.

- 100 bullfrogs are arrayed around a pond on a moonless night. Females choose mates according to their croaks. The frog with the deepest croak will attract the best mate, and so on... Each of the male frogs has a different croaking depth, and they all know where they stand in the ranking. If a frog doesn't croak, females will take their best guess at his croaking depth (i.e., the expected value). The question: Which frog(s) will croak, thereby revealing their type?

- Consider the decision of the frog with the 49th deepest croak (i.e., the frog just above average). Clearly, he should croak since otherwise females will assume that he is only as good as the average frog. Similarly, frogs 1 - 48 should croak since they are all above average too.

- What about the below average frogs? Now that the first 51 frogs have croaked, females should assume that the silent frogs are all below the average of the croaking frogs. Recognizing this, all of the frogs that are above the average of the previously silent group should now croak. So frogs 51 - 74 should also croak.

- Since frogs 75 - 100 are clearly below the average of the frogs that have croaked, frogs numbered 76 - 83 should now croak to show that they are above the average of the previously silent group...
- In the end, all of the frogs will croak, except perhaps for the highest pitched frog, who needn't bother since his type is implicitly revealed.

1.6 Conclusions

- These models suggest that unobservable quality heterogeneity creates significant problems for market efficiency.
- If this hypothesis is correct, there should be market mechanisms *already in place* to partly solve the problem. If these mechanisms did not exist, we would have reason to doubt whether the Lemons analysis identifies a genuine economic problem.
- What are some of these mechanisms?
  - Private mechanisms: Information provision, warranties, brand names, specialists.
  - Licensing.
  - Mandated information provision.
  - Legal liability.
  - Regulation.
- Each of these topics is worthy of hours of discussion. If we are fortunate, we'll have 15 minutes to talk about them in class.