Lecture Note 13 - Private Information, Adverse Selection and Market Failure

David Autor, MIT Economics and NBER

14.03/14.003 Microeconomic Theory and Public Policy, Fall 2025
(Compiled on November 25, 2025)

Private Information, Adverse Selection and Market Failure

- Where there is private information, there is an incentive for agents to engage in strategic behavior. For example, if you are selling a product, and your buyer knows the distribution of product quality but not the quality of the individual product that you posses, how much should the buyer be willing to pay? The intuitive answer might be the *expected value* of the product, or perhaps the certainty equivalent of this lottery.

- But this answer ignores an important consideration: the choice of what product you sell may depend on what price the buyer offers. And the price that the buyer offers may depend on what product she thinks you'll sell at that price. The equilibrium outcome in which buyer and seller expectations are aligned—that is, the buyer gets what she wants at the price she offers—may be far from efficient.

A bit of background

- Economists had historically conjectured that markets for information were well-behaved, just like markets for other goods and services. One could optimally decide how much information to buy, and hence equate the marginal returns to information purchases with the marginal returns to all other goods.

- In the 1970s, economists were brought to reevaluate this belief by a series of papers by Akerlof, Rothschild-Stiglitz, and Spence. Many of these economists all went on to share the 2001 Nobel for their work on the economics of information.

- Information is not a standard market good:

  - It is non-rivalrous (no marginal cost to each person knowing it)
  - It is extremely durable (not consumed)
  - It not a typical experience good where you can 'try before you buy.' A seller cannot readily allow you to 'sample' information without actually giving you information.
  - Unlike other goods (or their attributes), information is extremely difficult to measure, observe, and verify

- This combination of odd properties often gives rise to settings where information is—at least potentially—*asymmetric*. That is, some agents in a market are better informed than others about the attributes of a product or transaction.

- The most natural (and surely ubiquitous) way in which this occurs is that buyers may have general information about the 'average' characteristics of a product that they wish to purchase whereas sellers will have *specific* information about the individual product that they are selling.

- When buyers and sellers have asymmetric information about market transactions, the trades that are transacted are likely to be a subset of the feasible, welfare-improving trades. Many trades that would voluntarily occur *if all parties had full information* will not take place.

- Economic models of information focus on the information environment—that is, who knows what when. Specifying these features carefully in the model is critical to understanding what follows.

- These papers were written before game theory was widely understood and broadly adopted in economics. Their limitations presage the need for careful formalization of strategic environments. That's game theory delivers.

- This lecture note discusses a key concept in the literature on asymmetric information: adverse selection. Two general manifestations of adverse selection appear below:

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

- It turns out that these principles are roughly inverses.

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

1.1 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 can substantially affect the market equilibrium. It is possible that there will be *no trade* whatsoever for a given good, although:

  1. At any given price p_0, there are traders willing to sell their products.
  2. At price p_0, there are buyers willing to pay strictly above p_0 for the good that traders would like to sell.

- George Akerlof (1970) was the first economist to analyze this paradox rigorously. His 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?

1.2 A simple example: The market for used cars

- Setting
  - 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 (perhaps by law) between good cars versus lemons; they sell what's on the lot at the sticker price.
  - Buyers cannot tell apart good cars and lemons. But they know that some fraction $\lambda \in [0, 1]$ of cars are lemons.
  - After buyers have owned a car for any period of time, they also can tell whether or not they have bought a lemon.
  - Assume that good cars are worth $B_N^G = \$2,000$ to buyers and lemons are worth $B_N^L = \$1,000$ to buyers. The superscripts G and L will be used for good cars and lemons respectively, and a subscript N is used for new cars.
  - For simplicity (and without loss of generality), assume that cars do not deteriorate and that buyers are risk neutral.

- What is the equilibrium price for *new* cars? This will be

$$P_N = (1 - \lambda) \cdot 2,000 + \lambda \cdot 1,000.$$

  [Here, we assume that the supply at the dealership is fixed (inelastic)]

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

- Now, consider the used car market. Assume that used car sellers know the quality of their car, and are willing to part with their cars at 20 percent below their new value. For simplicity, assume that cars don't deteriorate. So,

$$S_U^G = \$1,600 \text{ and } S_U^L = \$800.$$

- Buyers of used cars are willing to pay 1.25 times the value that used car sellers place on their cars. Specifically, used car buyers will be willing to pay $B_U^G = \$2,000$ and $B_U^L = \$1,000$ respectively for used good cars and lemons. Hence, there is a surplus of \$400 or \$200 (that is, a gain from trade) from each sale. Selling either a good car or a lemon is potentially Pareto improving.

- **Q**: What will be the equilibrium price of used cars? The intuitive answer is

$$P_U = E\left[S_U\right] = (1-\lambda)\cdot 1,600 + \lambda \cdot 800.$$

  But this is not necessarily correct.

- Recall that buyers cannot distinguish good cars from lemons, but owners of used cars know whether they are selling a good car or a lemon. Assuming (logically) that sellers will only part with their cars if offered a price that is greater than or equal to their reservation price ($S_U^L = 800$, $S_L^G = 1,600$). So, for $P_U \geq 800$ owners of lemons will gladly sell their cars. However, for $P_U < 1,600$, owners of good cars will keep their cars.

- An additional assumption we're making is that sellers of lemons will not disclose that their cars are lemons. Instead, they'll sell them at any price exceeding \$800. (If sellers of lemons voluntarily disclosed that their cars are lemons there would be no adverse selection problem.)

- Given sellers' reservation prices, the quality of cars available depends on the price. In particular, the share of Lemons is as follows:

$$\Pr(\text{Lemon}|P_u) = \begin{cases} 1 & \text{if } P_U < 1,600 \\ \lambda & \text{if } P_U \geq 1,600 \end{cases}$$

- That is, quality depends on price; more formally, quality is endogenous. We can say that expected reservation selling price reflects quality of cars available for sale:

$$E[S_U|P_U] = \begin{cases} 800 & \text{if } P_U < 1,600 \\ 800\cdot\lambda + (1-\lambda)\cdot 1,600 & \text{if } P_U \geq 1,600 \end{cases}$$

- Denote buyers' willingness to pay for a used car as B_u. It is a function of quality of cars, or the expected reservation selling price as we just defined it. Therefore:

$$B_U = 1.25 \times (E[S_U|P_U])$$

- If trade is going to take place, buyers' willingness to pay for a used car must satisfy the following inequality: $B_U\left(E[S_U|P_U]\right) \geq P_U$. That is, at price P_U, the quality of cars available for sale, as reflected in $S_u|P_u$ must be worth at least that price to buyers.

- The value to buyers of cars for sale as a function of price is:

$$B_U\left(E[S_U|P_U]\right) = \begin{cases} 1,000 & \text{if } P_U < 1,600 \\ 1,000\cdot\lambda + (1-\lambda)\cdot 2,000 & \text{if } P_U \geq 1,600 \end{cases}$$

- The willingness of buyer's to pay for used cars depends upon the market price (a result we have not previously seen in consumer theory).

- Take the case where $\lambda = 0.4$. Consider the price $P_U = 1,600$. At this price, the expected value (to a buyer) of a randomly chosen used car—assuming both good cars and lemons are sold—would be

$$B_U (P_U = 1,600, \lambda = 0.4) = (1 - 0.4) \cdot 2000 + 0.4 \cdot 1000 = 1,600.$$

  Here, used *good* 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. This equation therefore satisfies the condition that $B_U (E[S_U|P_U]) \geq P_U$.

- But now take the case where $\lambda = 0.5$. At price $P_U = 1,600$, the expected value of a randomly chosen used car is:

$$B_U (P_U = 1,600, \lambda = 0.5) = (1 - .5) \cdot 2000 + .5 \cdot 1000 = 1,500.$$

  Hence, $B_U (E[S_U|P_U]) < P_U$. This cannot be an equilibrium. Since owners of good used cars demand \$1,600, they will not sell their cars at \$1,500. Yet, $P_U = 1,500$ is the maximum price that buyers will be willing to pay for a used car, given that half of all cars are lemons. Consequently, good used cars will *not* be sold in equilibrium, despite the fact that they are worth more to buyers than to sellers. Thus, *only* lemons sell.

- More generally, if $\lambda > 0.4,$ then good used cars are not sold and $P_u \in [800, 1000]$. In this price range, $B_u (E[S_u|P_u]) \geq P_u$.

- Bottom line: If the share of lemons in the overall car population is high enough, the bad cars will drive out the good ones. Although buyers would be willing to pay \$2,000 for a good used car, their inability to distinguish good cars from lemons means that they will not be willing to pay more than \$1,500 for *any* used car.

- With λ high enough, no good cars are sold, and the equilibrium price falls to exclusively reflect the value of lemons. There is no market for good used cars.

1.3 Summing up the Akerlof adverse selection model

- The key insight of Akerlof's paper 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).

- The underlying economics of adverse selection are very nicely exposited in the 2011 paper in the Canvas readings module, "Selection in Insurance Markets: Theory and Empirics in Pictures," by Einav and Finkelstein. The examples in their paper are geared towards the health insurance market, but they apply equally well to any market setting where adverse selection is present.

- Here's a useful brief excerpt from their paper: "The link between the demand and cost curve is arguably the most important distinction of insurance markets (or selection markets more generally) from traditional product markets. The shape of the cost curve is driven by the demand-side customer selection. In most other contexts, the demand curve and cost curve are independent objects; demand is determined by preferences and costs by the production technology. The distinguishing feature of selection markets is that the demand and cost curves are tightly linked, because the individual's risk type not only affects demand but also directly determines cost."

  In this paragraph, the "demand side" of the market refers to consumers who differ in their expected health costs—i.e., some are sicker than others—and the "cost curve" refers to insurers' cost of providing health insurance to these consumers. Adverse selection arises because, in most realistic cases, consumers are better informed about their expected health costs than are insurers. The insight of the paragraph is that insurers' costs of providing policies depend on which consumers buy the policies, which is itself determined by what the policies cost. Specifically, the insurer's cost depends on the health of the consumers who choose to purchase the policy at a given price.

  Figures 1, 2a and 2b of Einav-Finkelstein provide a nice graphical depiction of these insights. Figure 1 corresponds to a case where adverse selection causes an efficiency loss in an insurance market, but the market does not shut down entirely. Figure 2a corresponds to a case where adverse selection causes no efficiency loss (though it does result in transfers of consumer surplus from low risk consumers to high risk consumers). Figure 2b depicts a case where adverse selection leads the insurance market to shut down; no one buys insurance even though all consumers are risk averse and hence willing to pay *more than* the actuarially fair cost of an insurance policy.

2 Equilibrium in models with asymmetric information

A core feature that makes models with asymmetric information different from models we've studied previously is that the equilibrium is not primarily determined by a set of marginal conditions. e.g. marginal profit is zero, or the marginal rate of substitution equals the price ratio. Instead, equilibrium reflects choices of strategy that are, in combination, stable—meaning that no party has an incentive to deviate from them. We think of parties on the different sides of the market (e.g., buyers v. sellers) as choosing strategies (feasible actions) that maximize their payoffs given the chosen strategies of the players on the other side of the market. But of course, the players on the other side of the market are likewise choosing strategies to maximize their payoffs given the actions (or anticipated actions) of the other players. An equilibrium in this setting is a set of complementary strategies such that neither side wants to unilaterally change its strategy given the strategy of the other side. This notion is what is called a Nash Equilibrium after John Forbes Nash, who developed the idea and proved its existence in a 28 page 1950 Princeton doctoral dissertation, which eventually won him the Nobel prize in Economics in 1994—and more importantly, led to a hit Hollywood biographical movie called *A Beautiful Mind* in 2002 in which the starring role of Nash is played by Russell Crowe.

Here's an informal definition of the Nash Equilibrium (paraphrased from Wikipedia): *A Nash Equilibrium is a solution concept for a game involving two or more players, in which each player is assumed to know the equilibrium strategies of the other players, and no player has anything to gain by changing only his or her own strategy unilaterally. If each player has chosen a strategy and no player can benefit by changing his or her strategy while the other players keep theirs unchanged, then the current set of strategy choices and the corresponding payoffs constitute a Nash equilibrium.*[1]

We saw this idea in the used car example above:

- We first worked out the strategies of used car sellers, taking as given the price offered by buyers. We concluded that if the offer price was less than $1,600$, only lemons were sold, whereas if the offer price was $\geq 1,600,$ both lemons and good cars were sold.
- We observed that buyers have two primary strategies available: offering 800 and offering $1600.$
- We then asked which of these buyer strategies constituted a Nash equilibrium given the strategies of sellers.
- Clearly, offering 800 is always a Nash equilibrium. If a buyer offers $800,$ sellers will offer only lemons. Since these cars are worth 800, the buyer and seller strategies constitute a Nash equilibrium. Neither party wishes to deviate from their strategy (e.g., offer more, sell a good car) given the strategy of the other player.
- We next asked whether offering $1,600$ could also be a Nash equilibrium. The answer, as we saw, depends upon λ, the population share of lemons. At offer price $1,600$, both lemons and good cars are sold (that's the sellers' strategy). For this to be a Nash equilibrium, buyers must be happy to pay a price of $1,600$ given the sellers' strategies when facing this price. We calculated that the value to buyers of cars available at offer price $1,600$ is:

$$E\left[S_U|P_U = 1,600\right] = (1-\lambda)\cdot 2000 + \lambda \cdot 1000.$$

[1] And here's a plot summary from IMDB.com, "After a brilliant but asocial mathematician accepts secret work in cryptography, his life takes a turn for the nightmarish."

For $P_u = 1,600$ to be a Nash equilibrium, it must be the case that $B_U = 1.25 \times E[S_U|P_U = 1,600] \geq 1,600$, which requires that $\lambda \leq 0.4$

We'll be using the notion of Nash equilibrium to solve asymmetric information models.

We can apply the Nash equilibrium logic to the used car example above by plotting $S_U(P_U)$ and $B_U(S_U(P_U))$ for different values of λ, as in Figures 1 and 2 below, which plot $S_U(P_U)$ and $B_U(S_U(P_U))$ on the $y-$axis as a function of P_U on the $x-$axis. Nash equilibria are visible as ranges (highlighted in dark red) where $B_U(S_U(P_U))$ lies above the 45^0 line, implying that $B_U(S_U(P_U)) \geq P_U$ at price P_U.[2]

In Figure 1, $\lambda = 0.3$, and there are two ranges of Nash equilibria: $P_U \in [800, 1000]$ and $P_U \in [1,600, 1,700]$. (Why is P_U capped at $1,700$ not $1,800$? Because $S_U(P_U \geq 1,600) = 0.3(800) + 0.7(1,600) = 1,360$, so $B_U(S_U(P_U(1,600))) = 1,700$.)

Figure 1: Potential Nash Equilibria of Used Car Market with $\lambda = 0.3$

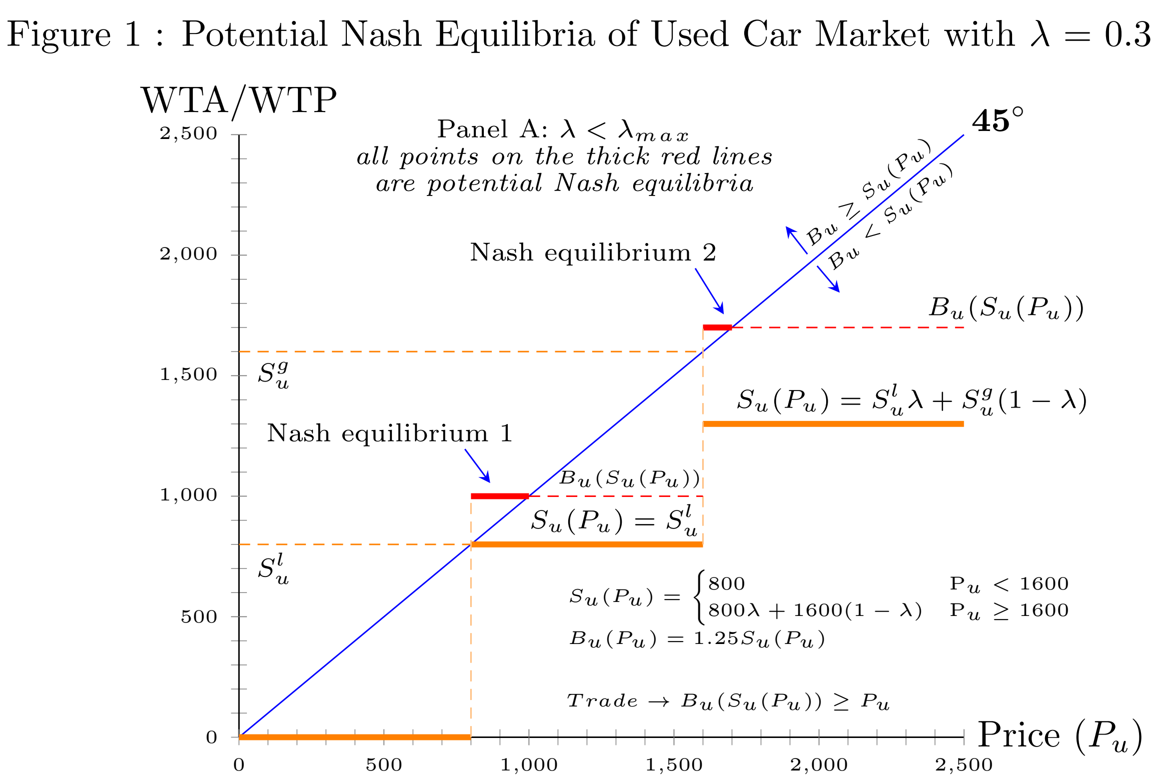

In Figure 2, $\lambda = 0.5$, and there is only one range of Nash equilibria: $P_U \in [800, 1000]$. For $P_U \geq$

[2]Many thanks to Sergey Naumov for plotting these illuminating figures.

1,600 we have $S_U(P_U >= 1,600) = 0.5(800) + 0.5(1,600) = 1,200$ and $B_U(S_U(P_U >= 1,600)) = 1,500$, which is less than P_U. Thus, there is no Nash equilibrium above $P_U = 1,000$.

Figure 2: Potential Nash Equilibria of Used Car Market with $\lambda = 0.5$

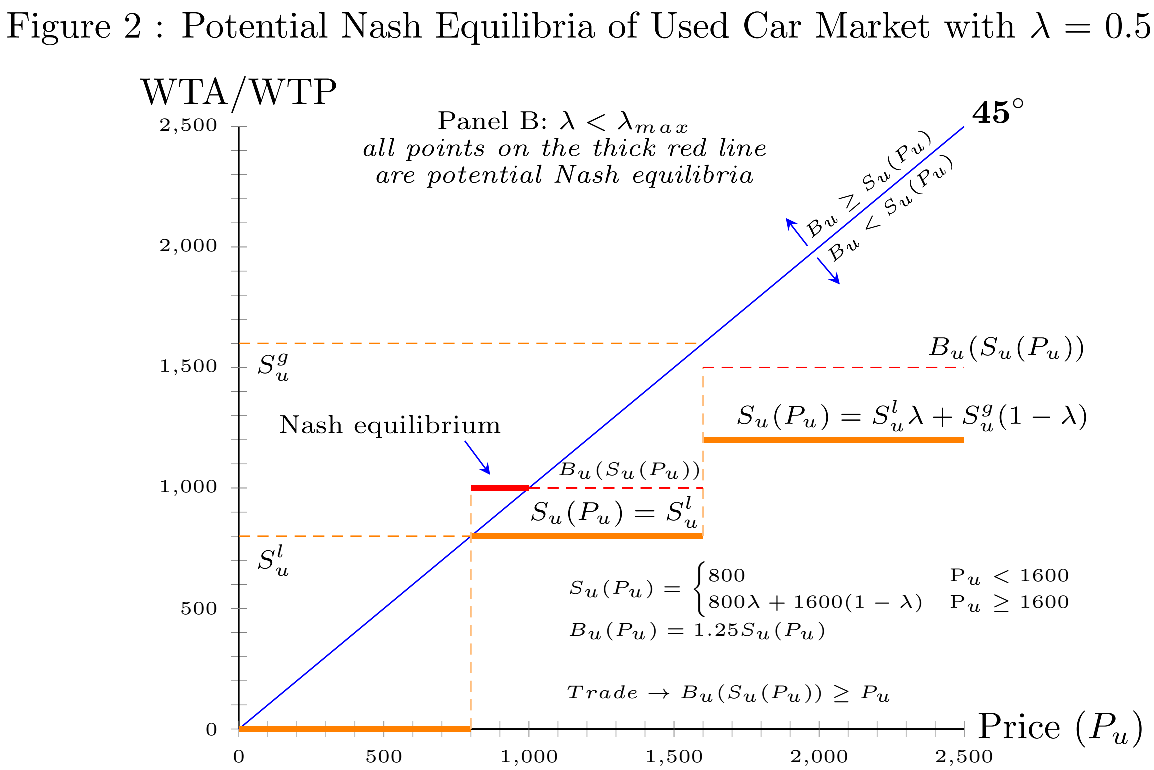

3 Adverse selection: A richer example

- Now that we have seen a stylized example, let's go through the same logic with a slightly richer example. We will consider a continuous distribution of product quality rather than just two types.

- Consider the market for 'fine' art. Imagine that sellers value paintings at between \$0 and \$100,000, denoted as V_s, and these values are uniformly distributed, so the average painting is worth \$50,000 to a seller.

- Assume that buyers value paintings at 50% above the seller's price. Denote this valuation as V_b. If a painting has $V_s = \$1,000$ then $V_b = \$1,500$.

- The only way to know the value of a painting is to buy it and have it appraised. Buyers cannot tell masterpieces from junk. Sellers can.

- 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: $V_b|E\left[V_s\left(P\right)\right] \geq P$.

- Take the sellers' side first. A seller will sell a painting if $P \geq V_s$.

- There is a range of sellers, each of whom will put their painting on the market if $P \geq V_s$.

- What is the *expected* seller's value of paintings for sale as a function of P? Given that paintings are distributed uniformly, it is:
$$E\left[V_s|P\right] = \frac{0+P}{2}.$$
So, if $P = 100,000$ then *all* paintings are available for sale and their expected value to sellers is $\$50,000$. 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
$$E\left[V_b|E\left[V_s|P\right]\right] = 1.5 \cdot E\left[V_s|P\right] = 1.5\left(\frac{0+P}{2}\right) = \frac{3}{4}P.$$
Clearly $E\left[V_b|E\left[V_s|P\right]\right] < P$. No trade occurs

- Since that buyers' valuation of paintings lies strictly above sellers' valuations, this outcome is economically inefficient—that is, the gains from trade are unrealized. What's wrong?

- 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 $\left(\frac{\partial E[V_s|P_s]}{\partial P} = 0.5\right)$.

- Consequently, for every dollar that the price rises, buyers' valuations only increases by $\$0.75,$ $\left(\frac{\partial E[V_b|E[V_s|P]]}{\partial P} = 0.75\right)$. There is no equilibrium point where the market price 'calls forth' a set of products that buyers are willing to buy at that price.

- 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).

- In this example, there is no trade.

4 Reversing the Lemons equilibrium: The Full Disclosure Principle

- Is there a way around this result? Intuition would suggest that the answer is yes. Sellers of *good* products have an incentive to demonstrate the quality of their products so that they can sell them at their true value. (Sellers of bad products have an incentive to not disclose quality, and this is what 'spoils' the market.)

- In the example above, sellers of good products do not disclose their products' quality because we have stipulated that the value of a piece of art can only be assessed *ex-post* by appraisal. Sellers of good paintings therefore have no credible means to convey their products' quality.

- Needed: A means to disclose information credibly. If there is an inexpensive (or free) means to credibly disclose the quality of paintings, you might expect that sellers of *above average* quality paintings will probably want to do this. In actuality, the result is much stronger than this: all sellers will choose to disclose.

4.1 Simplest case: Costless verification

- Imagine now that a seller of a painting can get a free 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 choose to get their paintings appraised?

- Your first instinct 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 \$75,000 if appraised would choose to get an appraisal.

- This intuition is on the right track but incomplete. It neglects the fact that the decision by some sellers to have their paintings appraised affects buyers' willingness to pay for non-appraised paintings.

- If only sellers with $V_s \geq 75,000$ had their paintings appraised, what would be the market price of non-appraised paintings?

$$1.5 \cdot E\left[V_s | V_s < 75,000\right] = 56,250.$$

- But if the market price is only \$56,250, then sellers with paintings at or above this price will also get them appraised. What is the new market price of non-appraised paintings?

$$1.5 \cdot E\left[V_s | V_s < 56,250\right] = 42,888.$$

- And so on...

- You can keep working through this example until you eventually conclude that *all* sellers will wish to have their paintings appraised. Why? Because each successive seller who has his painting appraised devalues the paintings of those who do not. This in turn causes additional sellers to wish to have their paintings appraised. In the limit, the only seller who doesn't have an incentive to obtain an appraisal is the seller with $V_s = 0$. This seller is indifferent.

- This example demonstrates 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 are below the average, which will give them the incentive to disclose, and so on... If disclosure is costless, in equilibrium all parties will explicitly or implicitly disclose their private information. If there is a cost to disclosure, 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 costs of disclosure are).

4.2 Costly verification

- Imagine now that a seller of a painting must pay $\$5,000$ for a appraisal. Which paintings will be appraised? If there are non-appraised paintings, will they be sold and at what price?

- We now need to consider three factors simultaneously:

  1. The net price of a painting that the seller would obtain if the painting were appraised (net of the appraisal fee) $(A = 1)$
  2. The net price if not appraised $(A = 0)$
  3. The value of the painting to the seller (remember that sellers won't sell for a net price less than V_s).

- The following conditions must be satisfied in equilibrium:

  1. Buyer's willingness to pay for an appraised painting is greater than or equal to seller's value of painting:
  $$V_b\,(A = 1) \geq V_s + 5000$$
  We will refer to this as Individual Rationality Constraint (IR).
  2. Seller cannot do better by appraising a non-appraised painting or v.v. We will refer to this as the Self-Selection Constraint (SS). Consider a cutoff value V_s^*. In equilibrium Paintings with $V_s \geq V_s^*$ are appraised and paintings with $V_s < V_s^*$ are not:
  $$V_b\,(V_s \geq V_s^*, A = 1) - 5,000 \geq V_b\,(V_s \geq V_s^*, A = 0)$$

and

$$V_b\left(V_s < V_s^*, A = 1\right) - 5,000 \leq V_b\left(V_s < V_s^*, A = 0\right).$$

- Let's go through these:

1. IR condition:

$$\begin{aligned} V_b\left(A = 1\right) &\geq V_s + 5000 \\ 1.5V_s &\geq V_s + 5000 \\ V_s^{IR*} &\geq 10,000 \end{aligned}$$

   This condition says that sellers who value their painting at less than 10,000 will choose not to get them appraised. This is because the market price less the appraisal cost is less than their private value of the painting.

2. SS condition

   This condition simply says that an individual seller must not be able to do better by switching their painting from appraised to non-appraised status or vice versa given the market equilibrium. Remember that the market value of a non-appraised painting is equal to 1.5 times their expected value to sellers. So, the self-selection constraint is:

$$\begin{aligned} V_b\left(V_s \geq V_s^{SS*}, A = 1\right) - 5,000 &\geq 1.5 \times \frac{V_s^{SS*}}{2} \\ 1.5 \times V_s^{SS*} - 5,000 &\geq 0.75 \times V_s^{SS*} \\ 0.75 \times V_s^{SS*} &\geq 5,000 \\ V_s^{SS*} &\geq 6,666 \end{aligned}$$

   and of course, the second inequality is also exactly satisfied at $V_s^{SS*} = 6,666$.

- Combining these results, we have

$$V_s^* = \max\left[V_s^{IR*}, V_s^{SS*}\right] = 10,000$$

Thus, the operative constraint is not that the *market price* for an non-appraised painting is higher than the *market price* for an appraised painting (which is SS) but that the *seller's own valuation* of an appraised painting net of appraisal cost must be greater than the *market price* of that painting when appraised (which is IR). Stated differently, when the IR constraint is satisfied, the SS constraint is also satisfied. But satisfaction of the SS constraint is necessary but not sufficient for satisfaction of the IR constraint. Hence, $V_s^* = V_s^{IR*} = 10,000$ is the threshold at which paintings are appraised.

- Meta-conclusion from this example:

  If verification of information is costly, the market equilibrium will not be Pareto efficient. In particular, there are two distortions evident in this market equilibrium. First, most sellers are spending $\$5,000$ to appraise and sell their paintings, even though this investment does nothing to improve the painting (so, this is a deadweight loss). Second, paintings with $V_s < \$10,000$ are not sold, even though these paintings are worth $1.5 \cdot V_s$ to buyers.

5 Insurance Markets

"The link between the demand and cost curve is arguably the most important distinction of insurance markets (or selection markets more generally) from traditional product markets. The shape of the cost curve is driven by the demand-side customer selection. In most other contexts, the demand curve and cost curve are independent objects; demand is determined by preferences and costs by the production technology. The distinguishing feature of selection markets is that the demand and cost curves are tightly linked, because the individualâs risk type not only affects demand but also directly determines cost." (Einav and Finkelstein, 2010, p117-118).

6 Example: Public Insurance Provision

Consider a unit mass of consumers i indexed from $i \in [0,1]$, each of whom has a VNM expected utility function of the form $U(w_i) = \ln(w_i)$. Each consumer i has initial wealth $w_{0i} = 150$ and faces a 50% probability of suffering loss $L_i = i \times 100$. Thus, for consumer $i = 0.60$:

$$E[w_i] = 150 - 0.5 \times 60 = \$120$$

$$U(E[w_i]) = \ln(120) = 4.79$$

$$E[U(w_i)] = 0.5 \times \ln(150) + 0.5 \ln(90) = 4.76$$

$$CE(E[U(w_i)]) = \exp(4.66) = \$116$$

Generalizing to any consumer i, we can write that i has the following wealth, expected utility, certainty equivalent income, etc.:

$$E[w_i] = 150 - 0.5 \times 100 \times i$$

$$U(E[w_i]) = \ln(150 - 0.5 \times 100 \times i)$$

$$E[U(w_i)] = 0.5 \times \ln(150) + 0.5 \ln(150 - 100 \times i)$$

$$CE\left(E\left[U\left(w_i\right)\right]\right)=\exp\left(E\left[U\left(w_i\right)\right]\right).$$

We can also calculate consumer $i's$ willingness to pay for insurance *in excess of* its actuarially fair value:

$$\begin{aligned} WTP_i &= E[w_i]-\exp\left(E\left[U\left(w_i\right)\right]\right)\\ &= \ln\left(150-0.5\times 100\times i\right)-\exp\left(0.5\times\ln\left(150\right)+0.5\ln\left(150-100\times i\right)\right) \end{aligned}$$

As per the example in section (6), assume that each consumer knows his or her type i but that insurers cannot distinguish individual types.

6.1 The naive policy

Consider a policy that pays each consumer L_i in the event of loss. That is, if i loses $L_i = 100\times i$, the insurer pays L_i to i to compensate her for the loss. A naive insurer decides to offer this policy at the price of $\$25$ since the average expected loss across the full population is $\$25$ per consumer (since $L_i \sim U\left[0,100\right]$ and each consumer faces a 50% probability of loss). Which consumers will purchase insurance in this case, and what would be the expected profits or losses of the policy?

To solve this problem, you need to identify the consumer i' who is indifferent between buying this policy and having no insurance. Consumers who have greater expected losses than i' will buy the policy (since the premium is the same for all consumers) whereas consumers who lower expected losses than i' will not purchase the policy. Formally, we want to find i' such that:

$$E\left[U\left(w_{i'}\right)\right]=0.5\times\ln\left(150\right)+0.5\ln\left(150-100\times i'\right)=\ln\left(150-25\right).$$

Notice that the lefthand side of this equation is the expected utility of i' if uninsured whereas the righthand side is wealth of i' if insured (in that case, i' pays the $\$25$ premium and so faces no risk of losing L_i).

$$\begin{aligned} 0.5\times\ln\left(150\right)+0.5\ln\left(150-100\times i'\right) &= \ln\left(125\right)\\ \ln\left(150-100\times i'\right) &= \left(2\times\ln\left(125\right)-\ln\left(150\right)\right)\\ 150-100\times i' &= \exp\left[2\times\ln\left(125\right)-\ln\left(150\right)\right]\\ i' &= \left[\exp\left[2\times\ln\left(125\right)-\ln\left(150\right)\right]-150\right]/\left(-100\right)\\ i' &= 0.46 \end{aligned}$$

Thus consumers on the interval $i'\in\left[0.46,1\right]$ will buy the policy. Notice that consumers with $i\in\left[0.46,50\right)$ would expect to lose money on the policy—since their expected losses are *less than* $\$25$ whereas the policy cost is $\$25$. These consumers choose to buy the insurance despite it being actuarially unfair (for them) because their risk aversion makes it worthwhile to pay a risk premium to obtain insurance.

This policy will lose money on average. Expected costs per insured consumer on this policy are $E\left[L_i | i \geq 0.46\right] = 100 \times 0.5 \times \left(\frac{1+0.46}{2}\right) = \$36.50$. This "naive" policy cannot be an equilibrium policy.

6.2 The free market policy

The insurer has learned its lesson. It will charge a premium that breaks even *given* that the population that enrolls for the policy will be adversely selected—the higher the premium, the less healthy the set of consumers that chooses to ensure. How can we find that break-even free market policy? Let i'' be the consumer who is just indifferent between the free market policy and no insurance. All consumers with greater expected losses than i'' will also buy the policy. Thus, the expected costs per insured of a policy sold to consumers with $i \geq i''$ is

$$E\left[L_i | i \geq i''\right] = 100 \times 0.5 \times \left(\frac{1+i''}{2}\right) = 25 \times \left(1 + i''\right).$$

Following the logic above, we can solve for i''

$$0.5 \times \ln(150) + 0.5 \ln\left(150 - 100 \times i''\right) = \ln\left(150 - 25 \times \left(1 + i''\right)\right).$$

After a bunch of algebra (or some spreadsheeting), the solution to this problem is $i'' = 0.75$. Thus, only one quarter of the population purchases insurance. This low rate of insurance is due to adverse selection. Those who have highest demand for insurance are those with the greatest expected losses. Since consumers with the greatest expected losses always enroll in the policy, the premium will be relatively high as compared to the average across the full population. This deters lower cost consumers from buying insurance, and hence only a subset insure, even though *all* are risk averse and would benefit from actuarially fair insurance.

Why does the market not completely unravel—leading to no one buying insurance? Since everyone is risk averse, consumers with highest risk are willing to pay a substantial premium in excess of their expected cost to obtain insurance. Thus, some consumers will sign up for the policy even though it is actuarially unfair for them because they prefer a 'bad deal' on insurance to no insurance at all.

6.3 An efficient mandatory policy

What would be an efficient insurance solution in this case? Refer to Figure 3. The key insight is that marginal social cost of providing insurance is below average cost for all but the riskiest consumer. And all consumers place positive value on insurance (except consumer $i = 0$, who has zero risk). Therefore an efficient market solution involves all consumers obtaining insurance. Since the marginal cost of insuring each consumer is less than or equal to her willingness to pay for this insurance, all consumers should be insured. As with the naive policy, the efficient policy has a premium of $\$25$, but this policy is *mandatory*.

Notice that as in section (6), not every consumer is better off under the mandatory policy. As we saw with the naive policy, consumers with $i' < 0.46$ would prefer not to buy the $\$25$ policy. So, in

Figure 3: Demand and Supply for Insurance

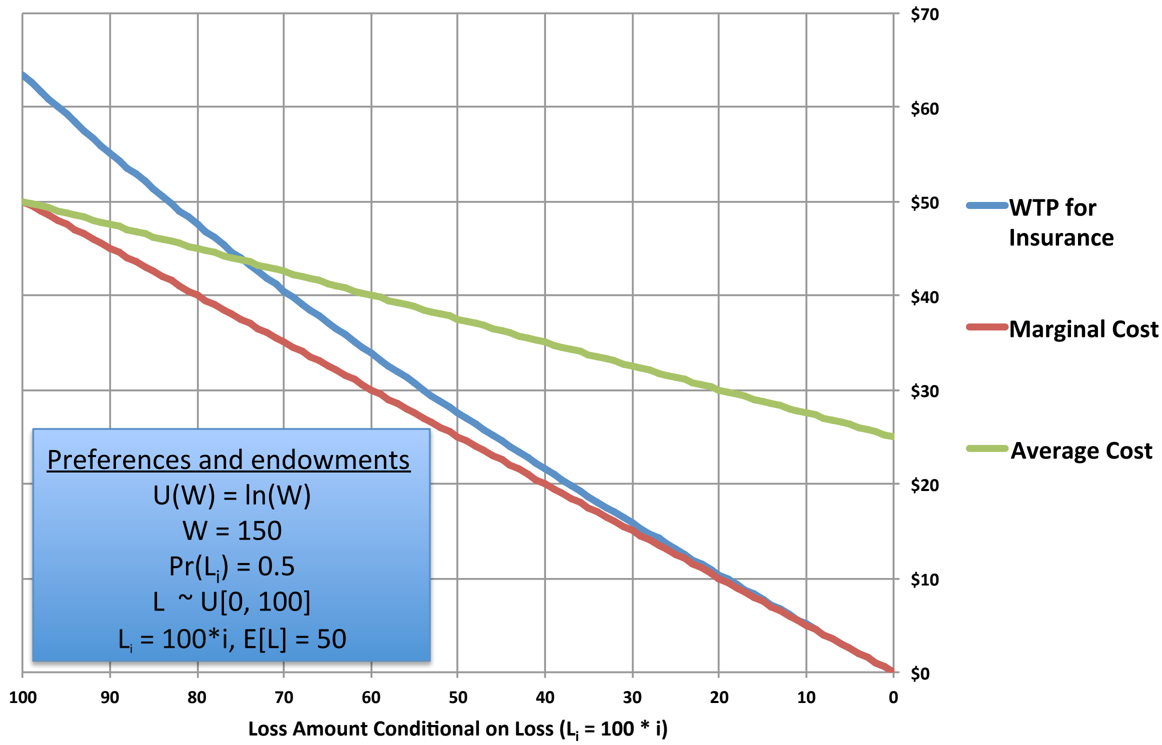

what sense is it efficient to require them to do so? This is tricky. You should think of the mandatory policy as having two parts: an insurance value and a transfer value. The transfer is from low cost to high cost consumers. Consumers with $i < 0.50$ effectively subsidize consumers with $i > 0.50$. While the insurance component makes consumers better off, the transfer component makes consumers with $i < 0.50$ worse off (and for consumers with $i < 0.46$, the net effect of the insurance and transfer is to lower utility relative to a case with no insurance). But remember that this transfer is just that: a *transfer* to other consumers. So we view it as a wash. (Notice that in the mandatory insurance case, each consumer has the same wealth and hence the same marginal utility of wealth in equilibrium. Transfers at the margin therefore do not effect social welfare; one person's loss is exactly offset by another's gain.)

As is implied by this logic, some simple calculations (perhaps made using a spreadsheet) will demonstrate that average consumer welfare is higher under the *mandatory* insurance policy than either in the *no-insurance* or the *free market insurance* case.

6.4 Introducing free screening

Imagine that a free *voluntary* test is introduced that will reveal the risk type of each consumer i who takes the test. Once i is tested, insurers will offer i an actuarially fair insurance policy at cost $0.5L_i = 50 \times i$.

Which consumers will volunteer to take this free screening test? Following the logic above, the

Figure 4: Expected Utility Under Three Insurance Schemes

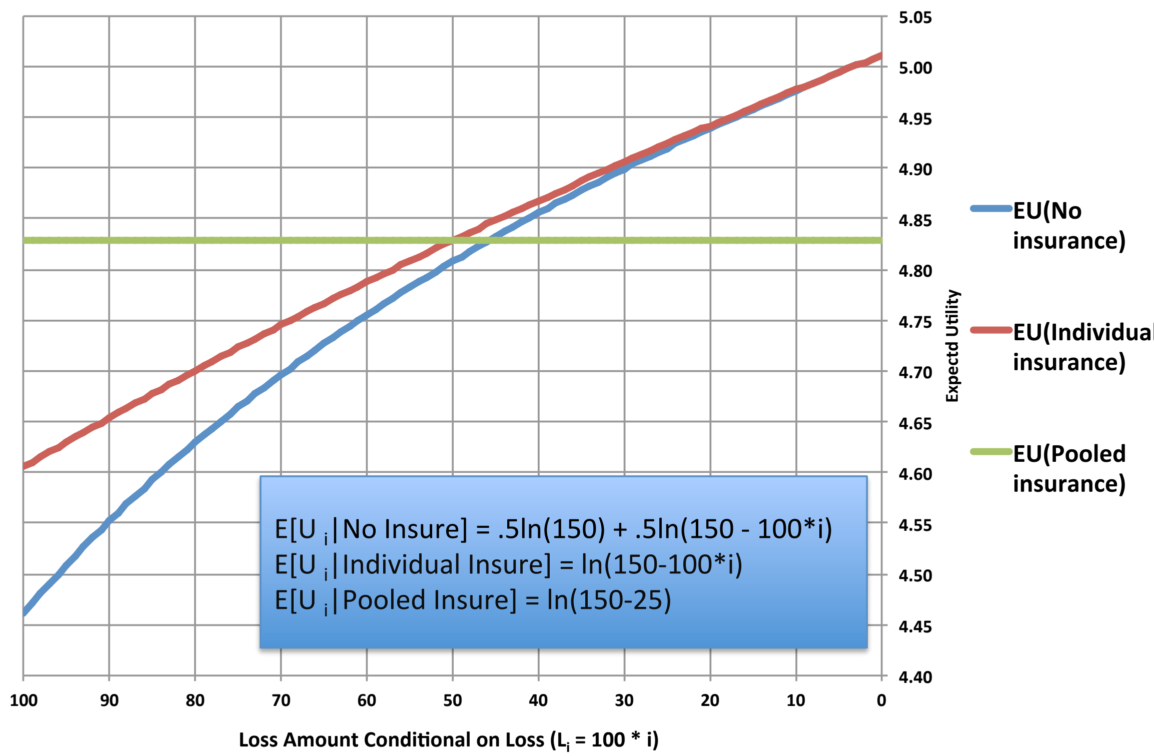

answer will be *everyone*. This is the Full Disclosure principle at work. The healthiest half of the population will first volunteer for the test. But then the healthiest half of the *remaining* untested population will volunteer for the test. And then the healthiest half of that remainder will volunteer for test. And so on.Pretty soon, everyone is tested.

Now, insurers can offer each and every consumer a break-even policy with premium $50 \times i$. There is no more adverse selection in this Full Disclosure case, and each consumer is fully insured. This equilibrium will *not* require an insurance mandate. The mandate was needed to resolve the adverse selection problem, which is no longer present.

Here's an interesting question: Is average consumer welfare higher under the mandatory policy in part (6.3) or under the *new* free market policy where each consumer pays an individualized premium of $50 \times i$? Perhaps surprisingly, the answer is that average social welfare is *higher* under the mandatory policy. Why? The answer is that the mandatory policy does two things: it provides insurance *and* it transfers income from rich to poor (that is, from those with low expected losses to those with high expected losses; we know this because every policyholder i pays the same premium, and the policy breaks even). The mandatory policy provides both risk pooling *and* risk spreading. The break-even policy with testing provides risk pooling but not risk spreading. You can again quickly demonstrate this result to yourself (perhaps again using a spreadsheet). Average social welfare is lowest in this example with no insurance, highest with the *mandatory* \$25 policy, and somewhere in between with the free market policy with testing. (See Figure 4)

7 Summary

- Unobservable quality heterogeneity creates important problems for market efficiency—market failures or incomplete markets quite likely.
- The problem is not the uncertainty *per se*. As we demonstrated during the lectures on risk, uncertainty and the market for insurance, there are market mechanisms for trading efficiently in risk. In these models, the uncertainty is *exogenous*—it is not under the control of economic agents.
- The fundamental problem in the adverse selection models above is that asymmetric information leads to a market equilibrium where sellers use their informational advantage strategically. Buyers respond strategically to sellers' choices. And the equilibrium of these strategic games are not likely to be first-best efficient. Quality is endogenous to price.

7.1 Market responses to asymmetric information

- If Lemons (adverse selection) hypothesis is correct, there should be some market mechanisms *already in place* to ameliorate the problem. Conversely, if no economic agent was observed attempting to solve the problem, we would have reason to doubt that the Lemons problem is relevant.
- What are some of these mechanisms?
  - Private mechanisms: Information provision, warranties, brand names, specialists and testers.
  - Licensing.
  - Mandated information provision.
  - Legal liability.
  - Regulation.
  - Example: Health insurance 'open enrollment' periods. Life insurance applications.
  - Lemon laws.
- Are there any markets that simply don't exist because of adverse selection (or moral hazard)?
  - Lifetime income insurance
  - Why is health insurance so expensive for the unemployed?
  - Why doesn't my life insurance policy cover suicide during the first five years after purchase?
  - Why can't I buy insurance against getting a low GPA at MIT?