Lecture 18 - Information, Adverse Selection, and Insurance Markets

14.03 Spring 2003

1 Lecture 18 - Information, Adverse Selection, and Insurance Markets

1.1 Introduction

- Risk is costly to bear (in utility terms). If we can defray risk through market mechanisms, we can potentially make many people better off without making anyone worse off.
- We gave three explanations for why and how insurance markets operate:
  1. Risk pooling – Law of large numbers.
  2. Risk spreading – Social insurance for non-diversifiable risks.
  3. Risk transfer (Lloyds of London) – Trading risk between more and less risk averse entities.
- (Note: risk spreading does not generate Pareto improvements, but it may still be economically efficient.)
- Hence, there is an exceedingly strong economic case for many types of insurance. Efficient insurance markets can unequivocally improve social welfare.
- Question: If the economic case for full insurance is so strong, do we see full insurance for:
  - Health
  - Loss of property: home, car, cash
  - Low wages
  - Bad decisions:
    - Marrying wrong guy/gal
    - Going to the wrong college

    * Eating poorly

- Obviously, the answer to these questions is no. We see instead:
  - Markets where not everyone is insured (health insurance, life insurance)
  - Incomplete insurance in every market where insurance exists at all:
    * Deductibles
    * Caps on coverage
    * Tightly circumscribed rules (e.g., must install smoke detectors in house, must not smoke to qualify for life insurance).
    * Coverage denied
    * Insurance markets that don't exist at all, even for major life risks
      · Low earnings
      · Bad decisions
- What is going wrong, given our belief in the economic case for insurance?
- Roughly 4 explanations:
  1. Credit constraints: People cannot afford insurance and hence must bear risk. Health insurance could be an example (i.e., if you already know you have an expensive disease, it may be too late to buy insurance).
  2. Non-diversifiable risk cannot be insured, e.g., polar ice cap melts, planet explodes. No way to buy insurance because we all face identical risk simultaneously.
  3. Adverse selection–Individuals' private information about their own 'riskiness' causes insurers not to want to sell policies to people who want to buy them.
  4. Moral hazard ('hidden action')–Once insured, people take risky/costly actions that they otherwise would not. This makes policies prohibitively costly.
- The model that we'll examine in this lecture concerns adverse selection.
- More generally, the issues of moral hazard and adverse selection open up a deep topic in economics: the economics of information.
- Information is not a standard market good:
  - Non-rivalrous (no marginal cost to each person knowing it)
  - Extremely durable (not consumed)
  - Quantities are hard to define

- Most goods can be described before purchased. In the case of information, the description is the good. That is, I cannot allow you to 'sample' my information without actually giving you my information.
- Strategic interactions often create incentives not to reveal information.

- A critical attribute of information in markets is that it is often asymmetric. That is, some parties to a transaction may be better informed than others (often for strategic reasons).
- A general point is that when buyers and sellers have asymmetric information about market transactions, the trades actually completed may be biased to favor the actor with better information.
- Equally critically, many potentially Pareto-improving trades will not be completed due to informational asymmetries (including trades that would voluntarily occur if all parties had full information).
- Consequently, economic models of information are often about the information environment–who knows what when. Specifying these features carefully in the model is critical to understanding what follows.
- The model we'll analyze today is by Rothschild-Stiglitz (1976). This paper was specifically cited by the Nobel committee in its 2001 award of the Nobel Prize in Economics to Joseph Stiglitz (along with George Akerlof and Michael Spence, whose work will feature in the next few lectures).

1.2 The environment

- Consider an insurance market where each potential insured faces two states of the world.
  1. No accident, in which case wealth is w.
  2. Accident, in which case wealth is $w - d$ (where $d < 0$ stands for damage).
- Hence, the wealth endowment is $W = (w, w - d)$.
- If a person is insured, their endowment is changed as follows $W = (w - \alpha_1, w - d + \alpha_2)$
- Hence, the vector $\alpha = (\alpha_1, \alpha_2)$ complete describes the insurance contract. Notice you can think of the insurance premium α_1 as paid in both the accident and no-accident states. Hence, α_2 is the net payout of the policy in event of accident.

- Denote the probability of an accident as p. An individual will buy insurance if the expected utility of being insured exceeds the expected utility of being uninsured, i.e.,

$$(1-p) \cdot u(w-\alpha_1) + p \cdot u(w-d+\alpha_2) > (1-p) \cdot u(w) + p \cdot u(w-d).$$

- An insurance company will sell a policy if expected profits are non-negative:

$$\alpha_1 - p(\alpha_1 + \alpha_2) \geq 0.$$

- Competition will insure that this equation holds with equality, and hence in equilibrium:

$$\alpha_1 - p(\alpha_1 + \alpha_2) = 0.$$

- We now need to define an equilibrium construct for this model. R-S use the following:

    1. No insurance contracts make negative profits.
    2. No contract outside of the set offered exists that, if offered, would make a non-negative profit.

- In plain language, this says that companies will not sell contracts to any individuals who are unprofitable in expectation, and further, there must not be a potential contract that could be offered that would be more profitable than the contracts offered in equilibrium.

1.3 Base case: Homogeneous risk pool

- To fix ideas, it is always useful to start with the simplest case.
- Assume for now that all potential insured have the same probability of loss, p. And note that we have already assumed that all losses are equal to d.
- (Fixing d is without loss of generality. We only need one free parameter here, either p or d. We'll be using p below.)
- See Figure 18#1, which recaps the state preference diagram from long, long ago (i.e., before the 2nd midterm).

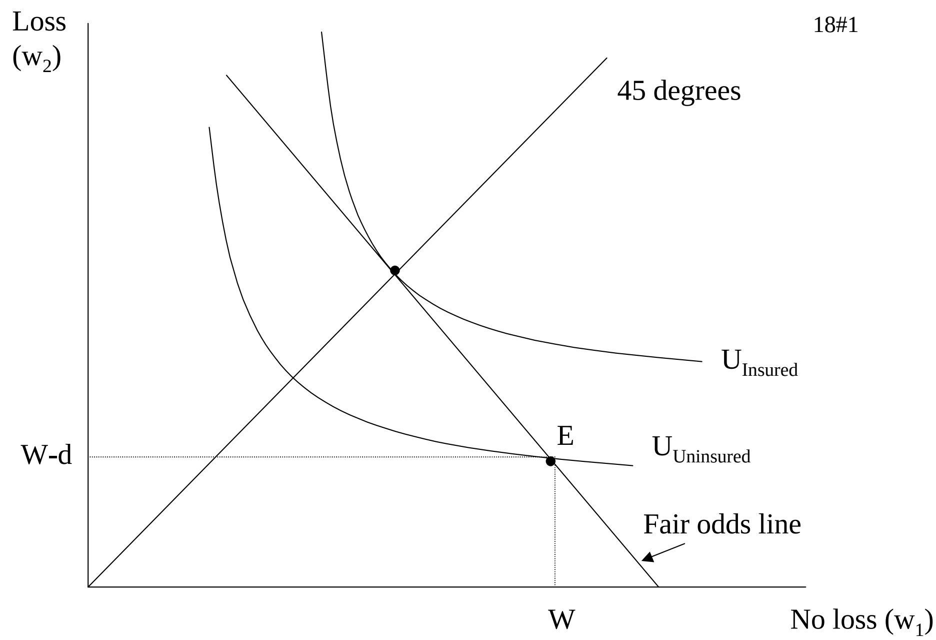

- Note that from the initial endowment $E = (w, w-d)$, the fair odds line extends with slope $\frac{-(1-p)}{p}$, reflecting the odds ratio between the accident and no-accident states.

- As we showed some weeks ago, a risk averse agent (and everyone is assumed to be risk averse here) will optimally purchase full insurance. Following from the Von Neumann Morgenstern expected utility property, the highest indifference curve tangent to the fair odds line has slope $\frac{(1-p)}{p}$ at its point of tangency with the the fair odds line, which is where it intersects the 45^0 line. At this point, wealth is equalized across states. So, the tangency condition is

$$\frac{(1-p)u'(w-\alpha_1)}{pu'(w-d+\alpha_2)} = \frac{(1-p)}{p} \Rightarrow w - \alpha_1 = w - d + \alpha_2.$$

- (You can demonstrate this to yourself by solving for α_1, α_2 in the Lagrangian for wealth allocation across states where the constraint is: $(1-p)w + p(w-d) - (1-p)(w-\alpha_1) + p(w-d+\alpha_2) = 0$. You should discover that $\alpha_1 = pd$, and $\alpha_2 = (1-p)d$, so that $w - \alpha_1 = w - d + \alpha_2 = w - pd = w - d + (1-p)d = w - pd$.)

- In this initial case, insurance companies will be willing to sell this policy $\alpha = (pd, (1-p)d)$ since they break even.

- This will be an equilibrium since no alternative profit-making policies could potentially be offered.

1.4 Adding heterogenous risk and private information

- Now, extend the model to the case where:

  1. Heterogeneity: The loss probability p varies across individuals. Specifically, assume two types of insurance buyers:

  $$\begin{aligned} H &: \text{Probability of loss } p_h, \\ L &: \text{Probability of loss } p_l, \\ \text{with } p_h &> p_l. \end{aligned}$$

     These buyers are otherwise identical in w and the amount of loss d in event of an accident (and their utility functions $u()$). Only their odds of a loss differ.

  2. Private information: Assume that individuals' i know their risk type p_i but this information is not known to insurance companies. (Note: private information without heterogeneity is not meaningful since if everyone is identical, there is no private information.)

- How realistic is the latter assumption? The gist is clearly correct: you know more about your 'riskiness' than your insurance companies. It is this informational advantage that is at the heart of the model. The model presents a particularly stark case, but the same results would hold with any degree of informational asymmetry.

- Given that there are two risk groups, H and L, there are two possible classes of equilibria in the model:

  1. A 'pooling equilibrium' – Both risk types buy the same policy.
  2. A 'separating equilibrium' – Each risk type (H, L) buys a different policy.

- We'll take these possibilities in order.

1.5 Candidate pooling equilibrium

- In a pooling equilibrium, both risk types buy the same policy.

- The equilibrium construct requires that this policy lie on the *aggregate* fair odds line (so that it earns neither negative nor positive profits).

- Define λ as the proportion of the population that is high risk.

- Hence, the expected share of the population experiencing a loss is

  $$\lambda p_h + (1 - \lambda) p_l.$$

  And the expected share not experiencing a loss is

  $$1 - (\lambda p_h + (1 - \lambda) p_l).$$

- Define
$$\overline{p} = \lambda p_h + (1-\lambda)p_l.$$
- The slope of the aggregate fair odds line is
$$-\frac{1-\overline{p}}{\overline{p}}.$$
- To simplify notation, define
$$\begin{aligned} w_1 &= w + \alpha, \\ w_2 &= w - d + \alpha_2. \end{aligned}$$
- See Figure 18#2

18#2

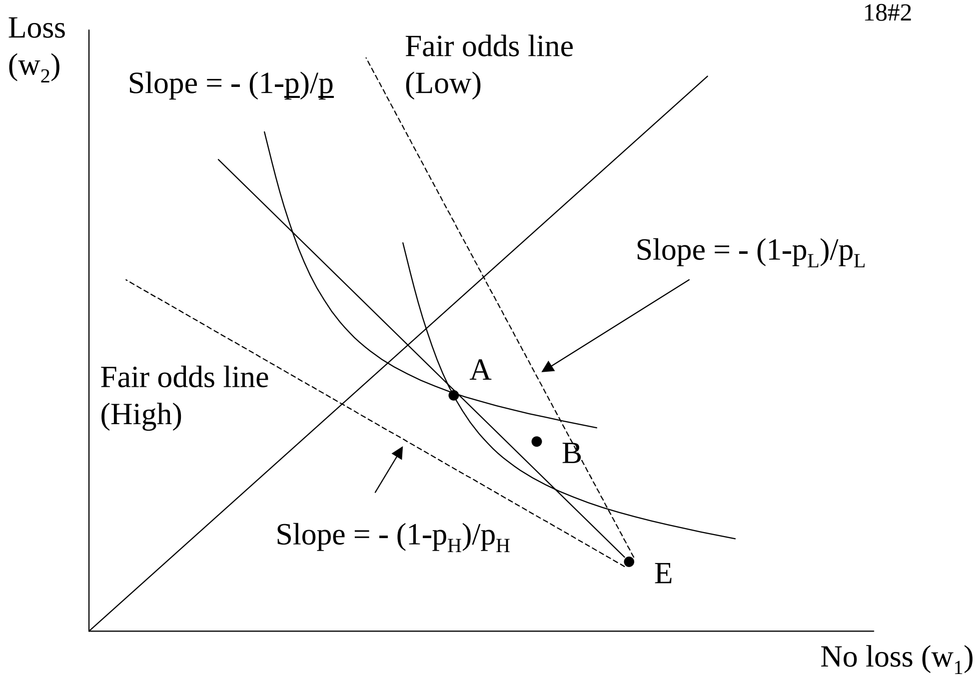

- Notice first that the 'pooling' policy A must lie on the aggregate fair odds line. If it lay above, it would be unprofitable and so would not exist in equilibrium. If it lay below, it would it would make positive profits and so would not exist in equilibrium.
- Notice 2nd that I've drawn the figure in a particular way with
$$MRS^H_{w_1,w_2} < MRS^L_{w_1,w_2}.$$
This is quite important. How do we know it's true?

- From the VNM property, we know the following

$$\begin{aligned} MRS^H_{w_1,w_2} &= -\frac{dw_2}{dw_1} = \frac{u'(w-\alpha_1)(1-p_h)}{u'(w-d+\alpha_2)p_h}, \\ MRS^L_{w_1,w_2} &= -\frac{dw_2}{dw_1} = \frac{u'(w-\alpha_1)(1-p_L)}{u'(w-d+\alpha_2)p_L}. \end{aligned}$$

  Since we've stipulated that High and Low types are otherwise identical, we know that $u_h(w) = u_l(w)$. This implies that

$$\frac{MRS_H}{MRS_L} = \frac{1-p_h}{p_h} \cdot \frac{p_l}{1-p_l} = \frac{1-p_h}{1-p_l} \cdot \frac{p_l}{p_h} < 1.$$

  This shows that the slope of the indifference curve for type H is less steep than for type L.

- This should also follow intuitively. Since the probability of loss is lower for type L, type L must get strictly more income than H in the loss state to compensate for income taken from the no loss state. This implies that the L types have steeper indifference curves for transfers of income between loss and no-loss states.

- Notice 3rd that the pooling equilibrium involves a cross-subsidy from L to H types. We know there is a cross-subsidy because H, L pay the same premium but H makes more claims. Herein lies the problem...

- Now consider what happens when another insurance company offers a policy like point B in the figure:

- How do H types react? They do not. As you can see, B lies strictly below $\overline{u}_H$, so clearly H types are happier with the current policy.

- However, L types strictly prefer this policy, as is clear from the fact that B is above $\overline{u}_L$. Why would this be? Point B is actuarially a better deal – it lies above the fair odds line for the pooling policy. On the other hand, it doesn't provide as much insurance – it lies closer to E than does point A. This is attractive to L types because they would rather have a little more money and a little less insurance since they are cross-subsidizing the H types. (For the opposite reasons, H types prefer the old policy.)

- So what occurs when policy B is offered is that all the L types change to B, and the H types stick with A. And notice that B is profitable if it only attracts L types because it lies below the fairs odds line for L types.

- But A cannot be offered without the L types participating–it depends on the cross-subsidy.

- Hence, the pooling equilibrium *cannot* exist. It will always be undermined by a 'separating' policy that skims off the L types from the pool.

- Free entry leads to ‘cream skimming’ of low-risk from the pool.
- Causes pooling policy to lose money b/c only high risk remain.
- Pooling policy disappears.

- An intuitive explanation is that cross-subsidies are not likely to exist in equilibrium. If a company loses money on one group but makes it back on another, there is a strong incentive to separate the profitable from the unprofitable group and charge them different prices (or just drop the unprofitable group), thereby undermining the cross-subsidy.
- We must consider a ‘separating equilibrium.’

1.6 Candidate separating equilibrium

- See Figure 18#3

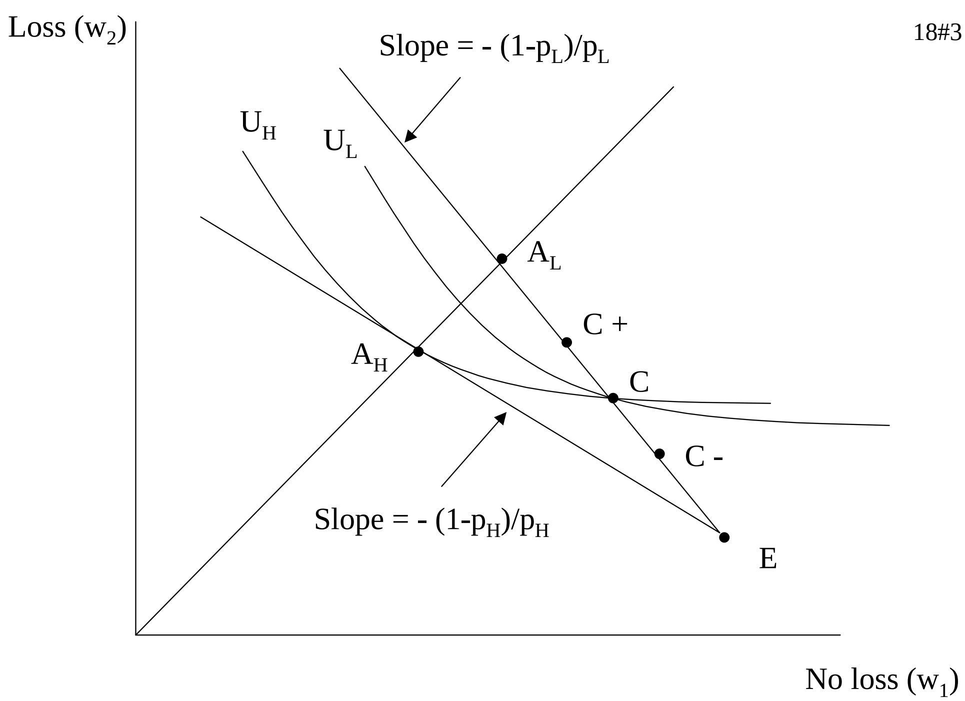

- Notice that there are two fair odds line corresponding to the two different risk groups.
- Points A_L and A_H are the full-insurance points for the two risk groups. Group L has higher wealth because its odds of experiencing a loss are lower.

- Note the point labeled C on the fair odds line for the L group. This is where the indifference curve from the full-insurance point for the H group crosses the fair odds line for the L group.
- Notice that the indifference curve $\overline{u}_L$ that intersects point C is steeper than the corresponding curve for the H group. In other words $MRS^H_{w_1,w_2} < MRS^L_{w_1,w_2}$, as we showed earlier.
- What is special about point C?
  - Observe that it is the best policy you could offer to the L types that would not also attract H types.
  - If a firm offered the policy C^+ on the figure, L types would strictly prefer it – but so would H types, which would put us back in the pooling equilibrium.
  - If a firm offered the policy C^-, H types would not select it, but L types would strictly prefer C, the original policy. So, any policy like $C-$ is strictly dominated by C.
- So, C is the point that defines the 'separating constraint' for types H, L. Any policy that is more attractive to H types would result in pooling.
- Notice by the way that any point like C^+ also involves cross-subsidy (if H types take it). We can see this because $\overline{u}_H$ is the indifference curve for the fully-insured type H. If there is a point that H types prefer to full-insurance (that is actuarially fair of course), it can only mean that the policy is subsidized.
- So we have a candidate equilibrium:
  - Policies A_H and C are offered.
  - Type H chooses A_H and type L chooses C.
  - Both policies break even since each lies on the fair odds line for the insured group.
- Before we ask whether this candidate pair of policies is in fact an equilibrium (according to the criteria above), let's look at its properties:
  1. High risk are fully insured.
  2. Low risk only partly insured! (Esp. ironic since they should be 'easier' to insure.) But if a company offered a policy that fully insured the L risk, it would also attract the high risk.
- Hence, preferences of H risk buyers act as a constraint on the market. Firms must maximize the well-being of L risk buyers subject to the constraint that they don't attract H risk buyers.

- Notice also that H risk are no better off for the harm they do to L risk. The externality is entirely 'dissipative,' meaning that one group loses but no group gains (opposite of Pareto improvement). This is potentially a large social cost.

1.7 Are the 'separating' policies an equilibrium?

- We now need to confirm whether the proposed set of policies does in fact constitute a separating equilibrium.
- See Figure 18#4

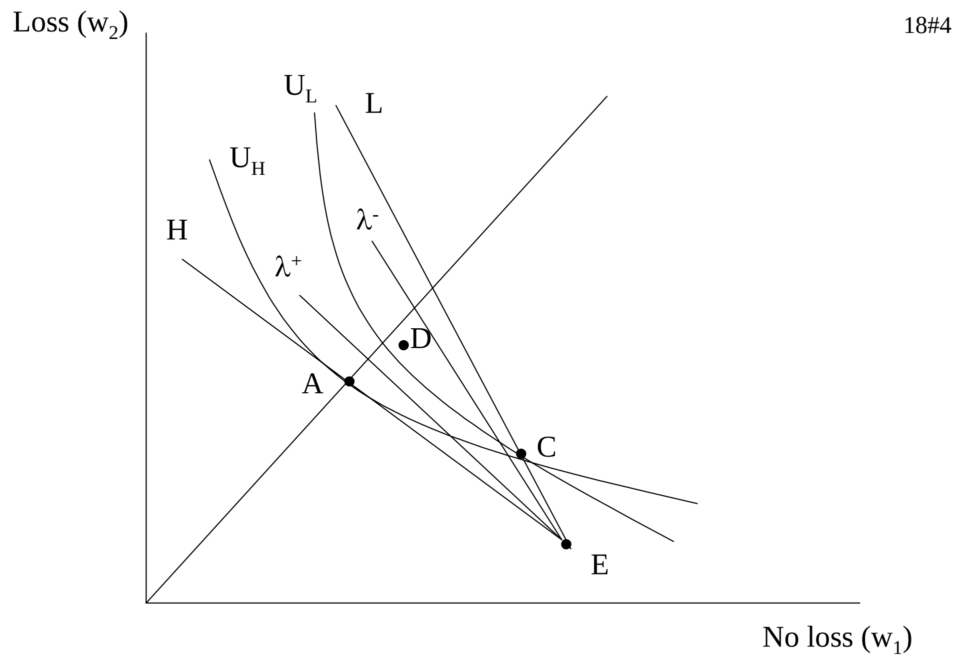

- Consider policy D in the figure. If offered, who would buy this policy? Clearly, both types H, L because the policy lies above each of their indifference curves when purchasing policies A_L and C.
- What's the potential problem with D? D is a pooling policy, which we already know cannot exist in equilibrium.
- What will determine whether D is offered? Clearly, whether it is more profitable than A_L, C both of which offer zero profits.
- What will determine the profitability of D is λ because λ determines the slope of the fair odds line for pooling policies. The larger is λ, the closer the pooling odds line lies to the H risk line. And the smaller is λ, the closer the pooling odds line lies to the L risk line.

- Notice in the figure the pooling odds lines corresponding to λ^+ and λ^- :
  - If the population is mostly high risk (λ^+), the pooling policy D that would break the separating equilibrium is unprofitable (lies above the fair odds line) and so will not be offered.
  - But if the population is mostly low risk (λ^-), the pooling policy D that would break the separating equilibrium is *profitable* (lies below the fair odds line) and so *will be offered.*
- *In the latter case, the model has no equilibrium.*
- Why is it the case that having λ low causes the separating equilibrium to fail?
  - At the separating equilibrium, the L risk types are not fully insured, and they are unhappy about this.
  - So, a pooling policy like D that requires just a little cross-subsidy to H types but offers more insurance is preferred to policy C by type $L's$.
  - Hence, if there are sufficiently few H types in the market, a firm could profitably offer this policy and it will dominate the two separating policies.
  - But as we know, once the pooling policy is offered, it will be broken by a separating policy that 'skims off' the L types.
  - In this case, the model has no stable equilibrium.
- To sum up:
  1. Welfare (efficiency) losses from adverse selection can be high.
  2. The costs appear to be born entirely by the low risk claimants (b/c of the need to get the H risk to select out of the pool).
  3. Pooling equilibria are unstable/non-existent
  4. Separating equilibrium may also not exist.

1.8 Implications of Rothschild-Stiglitz

- When information is private, the usual efficiency results for market outcomes can be radically undermined. Clearly, this model violates the first welfare theorem because the free market equilibrium is not Pareto efficient. And the only flaw in this market is that one set of parties is better informed than another. (Clarify for yourself that there would be no inefficiency if the insurance company could tell who is type L, H).
- How relevant is this model to 'real life' insurance markets?
- Consider life insurance policies (only small moral hazard problems here):

- Extensive health and background checks (deny coverage ex post if you lie).
- Will not sell you as much as you want to buy (quantity restrictions).
- Waiting and grace periods before coverage goes into place.
- Further oddities: what happens if you die while your application is being screened?

- Health insurance (more moral hazard issues in this case):
  - Health club benefits, maternity benefits – Why are these offered?
  - Why do individual policies cost so much more than group policies?

- There is something fundamentally correct about the insights of this model, though results are probably too stark (usually the case).