Week 10: Inefficient Trade and Adverse Selection

Dr Daniel Sgroi

Reading: Snyder and Nicholson, pp. 650–9.

With thanks to Peter J. Hammond.

Inefficient Trade and Adverse Selection

A central result of microeconomic theory
is that perfectly competitive markets generate
Pareto efficient allocations of goods and services,
relative to the preferences that drive
consumers' demand (and supply) behaviour in those markets.

For this result to have normative or ethical significance, however, market participants should have preferences that are well informed — i.e., ideally there should be perfect information.

It is therefore important to understand
the extent to which these arguments stand or fall
in the face of incomplete information,
particularly when some people are more informed
about the true value of some goods than others are.

Example: Would You Buy An Old Orange Grove?

To address this issue, we follow important ideas due to George Akerlof (1970) concerning *adverse selection*.

Akerlof wrote about the market for "lemons", which is colloquial American English for a useless product, especially a used car.

We discuss another kind of citrus fruit.

For many years player 1 has owned an orange grove, but the trees (and the farmer) are in poor shape, nearing the end of their productive lives.

It is hard to judge the quality of the land based on the crop that these old trees are producing.

But player 1 remembers how fruitful (or not) the orange grove used to be when the trees were in their prime.

Three Possible Types of Land

Expert agronomists report that the quality of land can be either poor, mediocre or good, each with probability $\frac{1}{3}$.

Thus, player 1's type $t_1 \in T_1 = \{t_1^L, t_1^M, t_1^H\}$ is the land quality that only he knows for sure.

Assume that, if player 1 does not sell the land, he can rejuvenate the orange grove, in which case the respective values for player 1 would be $v_1 \in \{10, 20, 30\}$.

A potential buyer (player 2) plans to remove the trees and grow soybeans.

Suppose the respective values to player 2 of the three different types of land would be $v_2 \in \{14, 24, 34\}$. That is, $v_2(t_1) = t_1 + 4$ for each type $t_1 \in T_1 = \{t_1^L, t_1^M, t_1^H\}$.

But player 2 only knows what the agronomists tell him.

One Round of Bargaining

Consider the game where player 2, the buyer,
makes just one take-it-or-leave-it offer to player 1, the seller,
which he can either accept (A) or reject (R).

The game ends with either a sale at the offer price, or no deal.

A strategy for player 2, the buyer, is a single price offer p.

A strategy for player 1, the seller,
is a mapping $t_1 \mapsto s_1(t_1) \in \{A, R\}$
from the seller's type $t_1 \in T_1$, which only he knows,
to a response which either accepts (A) or rejects (R) the offer.

If the quality of land could become commonly known,
player 2's value is 4 more than player 1's in every case,
so a Pareto improving trade could take place
at a price between t_1 and $t_1 + 4$.

A Modified Ultimatum Game

The game with just one take-it-or-leave-it offer by the buyer reduces to the one round bargaining or ultimatum game in which the unique subgame perfect equilibrium has player 2, the buyer, offer the lowest price t_1 acceptable to player 1, the seller, who then does accept that offer.

The Bayesian Nash Equilibrium Outcome

Player 1 will accept an offer p if $p > v_1(t_1)$, and only if $p \geq v_1(t_1)$.

If player 2 offers $p \geq 30$, all three possible types of player 1 accept;
player 2 winds up with land
worth only $\frac{1}{3}(14 + 24 + 34) = 24$ on average.

If player 2 offers $p \in [20, 30)$,
then only the two "lowest" types of player 1 accept;
player 2 winds up with land
worth only $\frac{1}{2}(14 + 24) = 19$ on average.

If player 2 offers $p \in [10, 20)$,
then only the one "lowest" type of player 1 accepts;
player 2 winds up with land worth only 14.

Unravelling

We have shown:

Proposition

An outcome of the game can be sustained
as a Bayesian Nash Equilibrium
if and only if player 2 offers $p^* \in [10, 14]$*;*
then the land changes hands
if and only if it is of the lowest possible quality.

A fully specified equilibrium strategy for player 1
involves three different reservation prices:
$p^L = v_1(t_1^L) = 10$; $p^M = v_1(t_1^M) = 20$; $p^H = v_1(t_1^H) = 30$.
Namely, the seller's value given his known land quality.
In other words, when of type t_1 player 1 accepts iff $p \geq v_1(t_1)$;
otherwise he rejects.

Only Lemons are Traded

Because of "adverse selection",
only the lowest quality of land gets traded.

Were the buyer to offer his average value,
only below average seller types would accept.

Then, were the buyer to offer his average value
of these below average types ("truncating the distribution"),
only even lower average seller types would accept. And so on.

In the example, this unravelling causes traded quality
to drop to the bottom, preventing the market
from implementing efficient trade of any but the worst kind of land.

Adverse Selection

This is another case of common values,
since the seller's type directly affects the buyer's payoff.

This is what causes the adverse selection
of the set of quality types that are traded in equilibrium
when there is this kind of "asymmetric" information.

Remark

Both players would benefit if the seller could somehow uncover old records which reveal publicly the land's true quality, before it was planted with orange trees.
However, mere announcements like "my soil is of high quality" are useless: if any buyer were to believe them, the seller would always make this claim, regardless of the land's true type.

Instead, some credible revelation of the land's type is needed in order to change the outcome.

An Insurance market with symmetric information

Let us now consider asymmetric information
in the market for insurance.

Assume that preferences take the expected utility form.

We begin with the symmetric information case
— when both buyers and sellers know
the quality of the good being traded
— in this case the quality is the riskiness of the risk being insured.

Suppose, as before, that there are two states of the world,
labelled state 1 and state 2 respectively.

Let us assume that their respective probabilities are p and $1 - p$.

Let c_1 and c_2 denote consumption in the two states 1 and 2.

A Fair Insurance Market

With fair (or perfect) insurance, the budget line facing the person who is thinking about taking out insurance takes the form $pc_1 + (1-p)c_2 =$ constant.

This has a slope equal to $-p/(1-p)$ where p and $1-p$ are the probabilities of the two states.

At any point along this budget line the insurance company expects to break even — on average the premium it collects for sure equals the average amount that it pays out when it has to.

To draw a graph, we consider a numerical example where $p = 0.4$ and the individual has an *ex ante* uncertain income of 10 in State 1 and 70 in State 2, as at the point marked W.

The budget line facing this individual under fair insurance is $0.4 \cdot c_1 + 0.6 \cdot c_2 = 0.4 \cdot 10 + 0.6 \cdot 70 = 4 + 42 = 46$, or $2c_1 + 3c_2 = 230$.

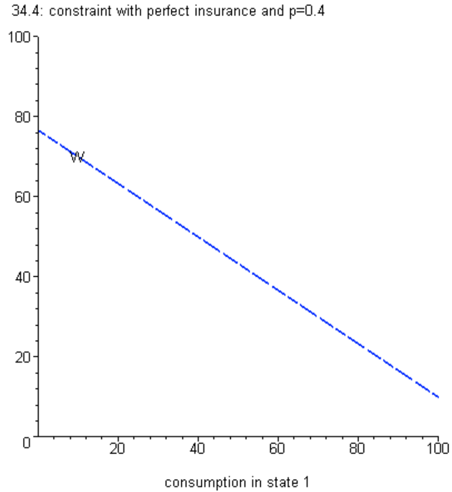
34.4: constraint with perfect insurance and p=0.4
100
80
60
40
20
0
W
20
40
60
80
100
consumption in state 1

Expected Utility Indifference Curves

Suppose now that our individual expected utility maximiser is risk-averse, so has preferences for pairs (c_1, c_2) represented by $pu(c_1) + (1 - p)u(c_2)$ where $u'(c) > 0$ and $u''(c) < 0$ for all $c \geq 0$.

The individual's indifference curves look as in the following graph.

Remember the important result that along the certainty line the slope of each expected utility indifference curve is

$$-\frac{pu'(c)}{(1-p)u'(c)} = -\frac{p}{(1-p)}.$$

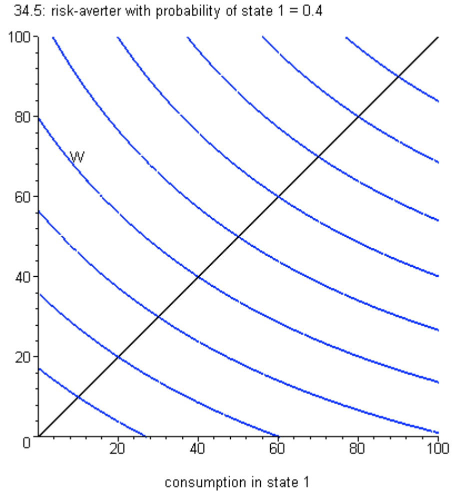
34.5: risk-averter with probability of state 1 = 0.4
100
80
60
40
20
0
20
40
60
80
100
W
consumption in state 1

Insurance Demand

If we put these two graphs together,
we get the following optimising solution shown next.

The slope of the budget line $pc_1 + (1-p)c_2 = \text{constant}$ (under fair insurance) is equal to $-p/(1-p)$.

The slope of each indifference curve along the certainty line is also equal to $-p/(1-p)$.

So we get the familiar result
that the risk-averter chooses to become fully insured
by choosing a point where the budget line intersects the 45° line.

The insurance company, of course, breaks even.

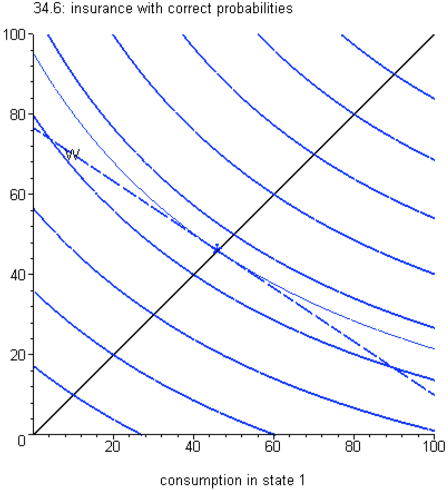
34.6: insurance with correct probabilities
100
80
60
40
20
0
20
40
60
80
100
W
consumption in state 1

High and Low Risk Types

Suppose policy holders may be one of two risk types:

type H with high probability p^H of adverse state 2;

type L with low probability $p^L < p^H$ of adverse state 2.

Otherwise all consumers are identical, and risk averse.

Let E denote the initial endowment point for each type,
where wealth W_2 in state 2 is lower than wealth W_1 in state 1.

To avoid duplicating notation,
let $t \in T = \{H, L\}$ denote the space of (two) possible risk types,
with associated probabilities p^t of adverse state 2.

We begin by considering what happens
when the two types can be distinguished.

Competitive Equilibrium and Efficiency

Complete markets imply separate insurance markets for each type, with different insurance premiums.

Suppose there is a perfectly competitive insurance industry, where free entry drives each firm's expected profit to zero.

Equilibrium will be Pareto efficient, with full insurance for each type of consumer.

Individuals of type $t \in \{H, L\}$ face a budget constraint of the form $(1 - p^t)W_1 + p^t W_2 = y^t$, with slope $-(1 - p^t)/p^t$. Because $p^H > p^L$, individuals of type L face a steeper budget line.

As in the usual competitive market model, the two budget lines would pass through the endowment point E.

So $y^t = (1 - p^t)W_1^E + p^t W_2^E$ for both types $t \in \{H, L\}$.

First Best Competition

Because the two types can be distinguished, they can be faced with the blue line $B^H E$ and the red line $B^L E$
as the two type H and L's separate budget constraints.
Their equations are $(1 - p^t)W_1 + p^t W_2 = (1 - p^t)W_1^E + p^t W_2^E$,
with respective slopes $-(1 - p^H)/p^H$ and $-(1 - p^L)/p^L$.
Because $p^H > p^L$, the line $B^L E$ is steeper than $B^H E$.
Equilibrium is at A^H and A^L on the full insurance 45° line.

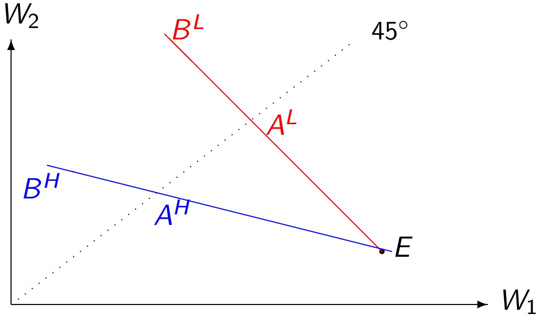

First Best Monopoly

Consider the first-best case when a monopoly insurer can costlessly observe the individual's true risk type, and charge discriminatory prices accordingly.

To maximize its expected profit, the monopolist minimizes each insuree's expected consumption subject to the participation constraint that neither type can be worse off than at the endowment point E.

Individuals of each type $t \in \{H, L\}$ face a budget constraint of the form $(1 - p^t)W_1 + p^t W_2 = y^t$, having slope $-(1 - p^t)/p^t$.

Each type of insuree is offered an inflexible full insurance contract, for a premium that extracts as much profit as possible so only the monopolist gains from trade.

So the insurer puts type H consumers at A and type L consumers at B, where the two relevant indifference curves through E intersect the 45° line.

First Best Monopoly Optimum Illustrated

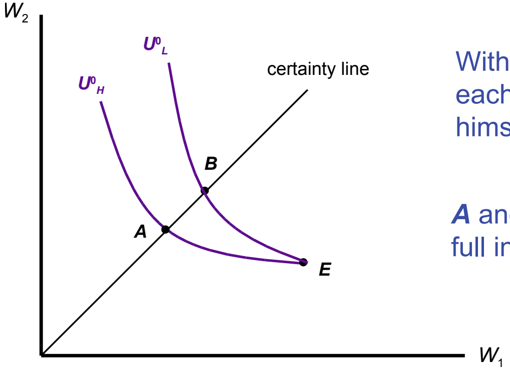

Without insurance each type finds himself at ***E***

A and ***B*** represent full insurance

Adverse Selection in Insurance

Adverse selection in insurance occurs when riskier types cannot be distinguished from less risky types.

Contracts select adversely because risker types are more likely to accept any policy offering full insurance (with the same level of consumption in each state) but are more expensive to serve if they do.

Suppose a monopoly insurer cannot observe the consumer's type.

If the monopoly insurer were to offer the two different full insurance contracts A and B, then the high-risk type would choose B instead of A.

So the first-best contracts are not be incentive compatible

The monopoly insurer would not be maximizing its profit, and could even make a loss with both low and high-risk types at B.

Second Best Monopoly Optimum

The insurer can earn more profit by not offering at all
the cheaper full insurance contract B for type L individuals.

Instead, it offers a partial insurance contact like D,
which moves type L individuals along their indifference curve
passing through the endowment point E.

To deter type H individuals from choosing D,
it offers an alternative contract C that minimizes expected value
subject to the incentive constraint requiring
that they must not strictly prefer type L's contract D
to the contract C the insurance company wants them to choose.

Specifically, the monopolist offers full insurance at C,
where the indifference curve U_H^1 through D intersects the 45° line.

Second Best Monopoly Optimum Illustrated

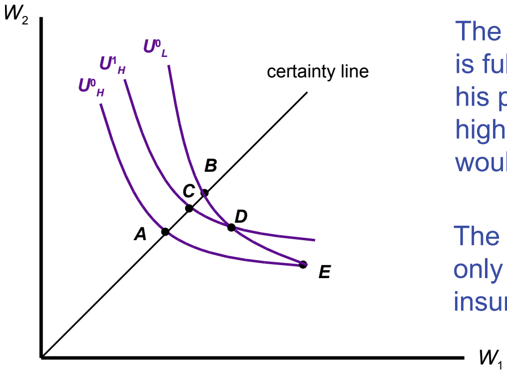

The high-risk type is fully insured, but his premium is higher (than it would be at ***B***)

The low-risk type is only partially insured

Second Best Monopoly Optimum

In this second best the monopoly insurer offers all consumers a choice between the two contracts C and D.

Contract C offers full insurance,
at a premium which is lower than at B but higher than at A.

Contract D offers just enough partial insurance
so that the high-risk types do not strictly prefer it.

Compared to the pair of contracts A and B,
the high-risk type extracts some "rent" from being hidden.

The monopolist has to give up some profit
in order to distinguish between types H and L.

Let $\alpha \in (0, 1)$ denote the proportion of type H individuals.
As α increases from 0 to 1, the second-best contracts C and D will move away from B toward A and E respectively.

Linear Pricing

Suppose the two types face a fixed price ratio $p/(1-p)$, with $p^L < p < p^H$, at which they can buy whatever insurance they like, up to full insurance.
They will face the black line segment joining the endowment E to the full insurance contract F, on the 45° line.

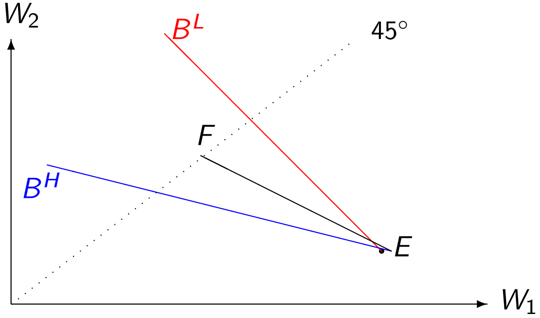

Demands at Intermediate Prices

Because $p < p^H$, for type H insurees the price ratio $p/(1-p)$ is below their actuarially fair price ratio $p^H/(1-p^H)$.

They would therefore demand excess insurance, if allowed.

We assume, however, they are limited to full insurance, not going beyond point F with coordinates $(\bar{W}, \bar{W})$, which is what they choose.

Because $p > p^L$, for type L insurees the price ratio $p/(1-p)$ is above their actuarially fair price ratio $p^L/(1-p^L)$.

They will therefore demand only partial insurance at some point (W_1^L, W_2^L) in the interior of the line segment EF.

That point is below the 45° line, so $W_1^L > \bar{W} > W_2^L$.

It is also on the budget line, so $(1-p)W_1^L + pW_2^L = (1-p)W_1^E + pW_2^E$.

Profits from the Two Types

Let $C^t = (W_1^t, W_2^t)$ ($t \in \{H, L\}$) be any pair of contracts which:

1. both lie on the same intermediate budget line segment EF with equation $(1-p)W_1 + pW_2 = (1-p)W_1^E + pW_2^E$, where $p^L < p < p^H$;
2. offer both types at least partial insurance, because $W_1^E > W_1^t \geq W_2^t > W_2^E$.

Then the expected profit made from contract $C^t = (W_1^t, W_2^t)$ is

$$\pi^t = (1-p^t)W_1^E + p^t W_2^E - (1-p^t)W_1^t - p^t W_2^t.$$

But the budget constraints imply that

$$0 = (1-p)W_1^E + pW_2^E - (1-p)W_1^t - pW_2^t.$$

Because $W_1^E - W_1^t + W_2^t - W_2^E > 0 + 0 = 0$,
subtracting this from the previous equation gives

$$\pi^t = (p - p^t)(W_1^E - W_1^t + W_2^t - W_2^E) \gtrless 0$$

according as $p \gtrless p^t$.

Cross Subsidization

We have shown that, if both types face contracts C^t on the same budget line with slope $-p/(1-p)$, then profits satisfy $\pi^t \gtrless 0$ according as $p \gtrless p^t$.

So if $p^L < p < p^H$, then $\pi^H < 0$ and $\pi^L > 0$.

Thus, every insurance company makes profits on type L consumers and makes losses on type H consumers.

Suppose free entry competes away any positive expected profit.

Then the zero expected profit condition $\alpha\pi^H + (1-\alpha)\pi^L = 0$ must be satisfied, given that α is the proportion of high risk types.

The profit $\pi^L > 0$ earned on each low risk insuree is being used as a cross subsidy in order to cover the loss $-\pi^H > 0$ earned on each high risk insuree.

A special case of this is at the pooled contract $C = C^H = C^L$.

Pooled Contracts

Given the proportions α and $1 - \alpha$ of the two types H and L, let $p^* := \alpha p^H + (1 - \alpha) p^L$ denote the pooled probability that any individual experiences the adverse state 2.

In the following diagram, a single pooling contract C^* is traded.

It earns zero profit because it lies on the pooled budget line $B^* E$ whose equation is $(1 - p^*) W_1 + p^* W_2 = (1 - p^*) W_1^E + p^* W_2^E$.

The blue indifference curve through C^* labelled U^H is for type H.
At the pooled price ratio $p^*/(1 - p^*)$,
which is below the actuarially fair price ratio $p^H/(1 - p^H)$,
type H consumers would like to buy more insurance if they could.

The red indifference curve through C^* labelled U^L is for type L.
At the pooled price ratio $p^*/(1 - p^*)$,
which is above the actuarially fair price ratio $p^L/(1 - p^L)$
type L consumers would not like to buy any more insurance.

Disrupted Pooling

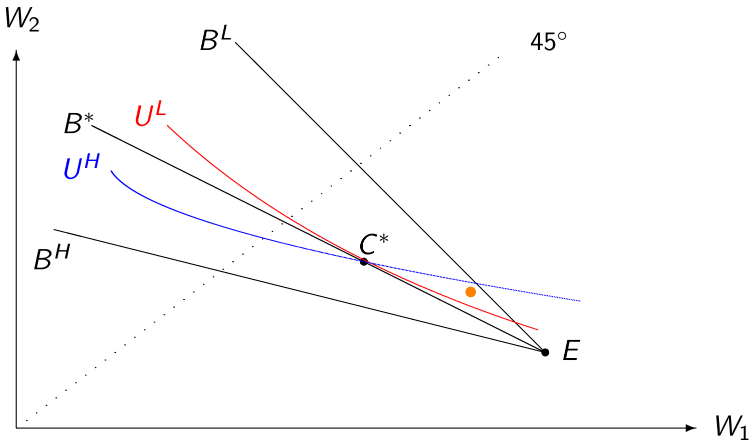

Cream Skimming

An entering insurance company could "cherry pick"
or skim the cream off the top of the market.

It does so by offering a new contract, like that at the red dot.

This is below the high risk consumers' indifference curve U^H,
but above the low risk consumers' indifference curve U^L.

It is designed, in fact, to attract only "good" type L consumers.

Because the red dot lies strictly below the type L line,
and is bought only by type L consumers, it earns a profit.

By contrast, the pooled contract C^* only breaks even.

Separating Equilibrium

The next diagram illustrates a separating equilibrium where all types of consumer are offered a choice between just two different contracts, represented by the blue and red dots.

Type H consumers are willing to buy full insurance on the 45° line at what is an actuarially fair price $p^H/(1-p^H)$ for them.

At this price, type L consumers find full insurance too expensive.

Instead, type L consumers are allowed to buy only partial insurance at an actuarially fair price $p^L/(1-p^L) < p^H/(1-p^H)$ for them.

The position of the red dot is just close enough to point E to deter type H consumers from rushing in to choose a contract which otherwise they would all strictly prefer to the blue dot.

Separating Equilibrium Illustrated

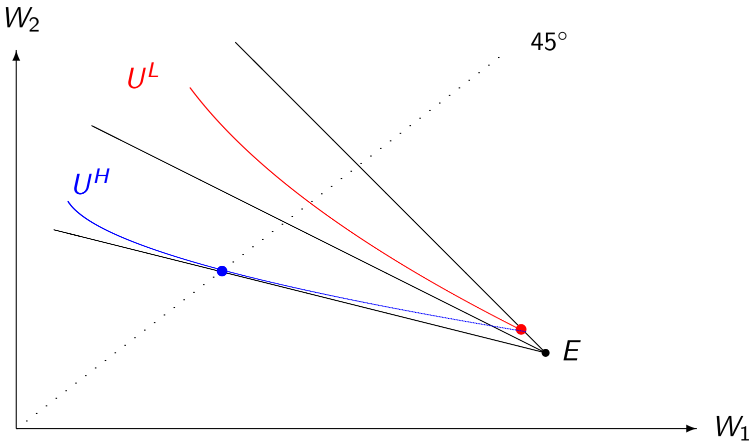

No Separating Equilibrium

The separating equilibrium may also be vulnerable to entry if there is a sufficiently small proportion α of type H consumers.

In the previous diagram, the red indifference curve lies above the middle line through E on which an enforced pooling contract breaks even. This allows a separating equilibrium to exist.

In the next diagram, however, the red indifference curve has points below this middle line.

The orange dot then represents a new pooling contract which both types prefer to the separating allocation represented by the blue and red dots.

It also earns a profit because it lies below the break-even line.

There is no equilibrium in this case. The only candidate for a separating equilibrium is vulnerable to a pooling contract, but this can never be an equilibrium either.

Illustrating the Case with no Separating Equilibrium

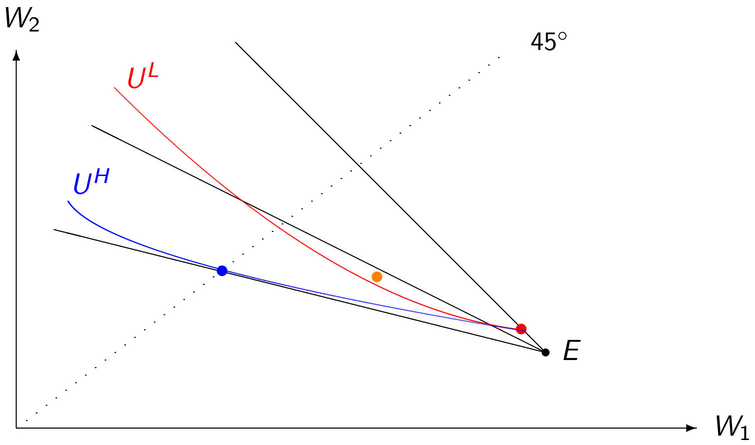

Final Thoughts

Imperfect information puts severe limits on market efficiency.

When there are hidden actions, as in the case of moral hazard, they amount to externalities which may have to be controlled directly if efficiency is to be restored.

Adverse selection, or hidden types, can be even more troublesome.

Not only are market allocations inefficient.
There may not even be any satisfactory equilibrium in the market.

Superficially, this creates a strong case for appropriate regulation, such as excluding the kind of cream skimming that destroys any possibility of a pooling equilibrium.