Recitation 9: Externalities and Adverse Selection

Emma Zhu [1]

14.03 Fall 2025

Externalities

The behavior of others can affect your utility and your optimal decision, which we refer to as "social externalities." Consider a simple model in which each person can choose between two actions A and B. Each action yields a utility that depends on the percentage of people who take each action.

An **equilibrium** occurs when no player wants to change their action, keeping all other players' actions fixed. We can classify equilibria into two categories: stable and unstable.

We often want to achieve the social optimum. Sometimes, the social optimum is a unique stable equilibrium (the system will reach this stable equilibrium on its own) - other times, the social optimum may be one of many stable equilibria (may or may not reach the social optimum naturally, depending on initial conditions), an unstable equilibrium, or not an equilibrium at all.

Let's illustrate these concepts with an example problem (which is similar to an example we saw in class... and might be very similar to an example seen in psets!).

Problem: Tiktok

[2]

We are back to the pre-tiktok glory days of 2016, but unfortunately, it just launched. Initially, no one uses it, so the utility for using TikTok is -0.5, as it's just embarassing. However, as more people use Tiktok, the utility of using Tiktok increases. Let $U_N(p) = 0$ be the utility for not using Tiktok, and $U_L(p) = -0.5 + p$ be the utility for using Tiktok.

Q: Does this problem demonstrate a **positive** or **negative** externality?
A: Positive, because people benefit from greater Tiktok adoption.

Q What are the equilibrium points? Are they stable or unstable?
A: We can draw a graph depicting the two utility curves as well as the social utility to answer these questions.

[1] Adapted from notes by Sean Wang, 2023
[2] Problem about tiktok, though tiktok is probably also a problem

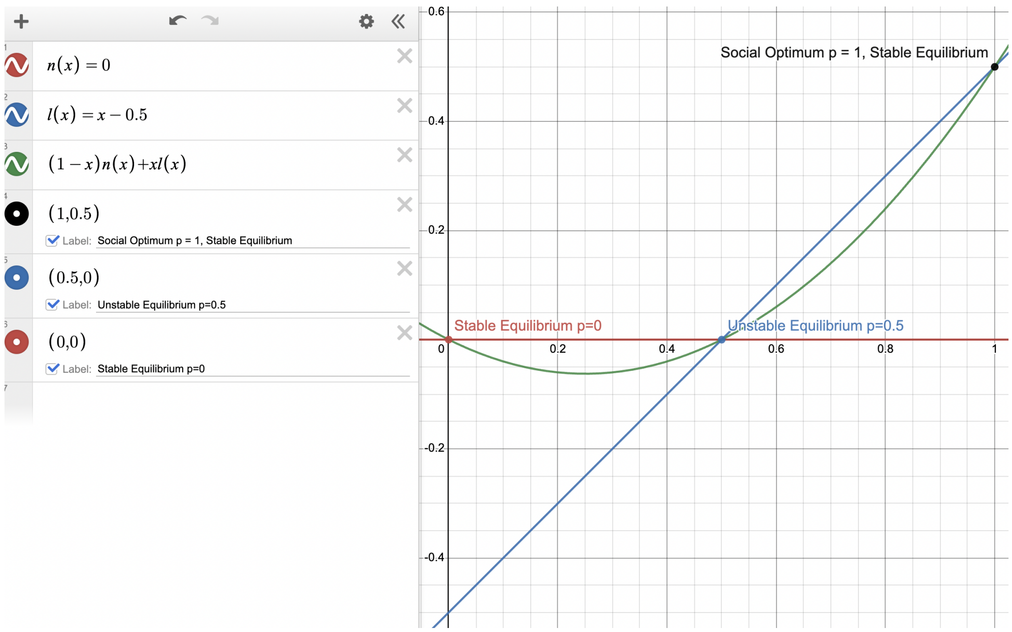

The green line is the average social utility. Observe how it is a weighted average of $n(x)$ and $l(x)$. At $x = 0$, no one is using Tiktok, so the social utility is zero. At $x = 1$, everyone is using TikTok, so the social utility is 0.5.

Q: What are the equilibrium points? Are they stable or unstable?

A: Stable equilibriua incur when small shocks/perturbations will just bring people back to the original point. Unstable ones mean that a small pertubation will bring the solution to a point.

- $p = 0$: Stable - if we perturb this equilibrium $p = 0 + \epsilon$, then the people using Tiktok will switch back to not using it, so we will be back to $p = 0$.
- $p = 0.5$: Unstable - if we perturb this equilibrium to $p = 0.5 + \epsilon$ such that $\epsilon < 0$, then not using Tiktok is better than using Tiktok, so people will fall back to the $p = 0$ stable equilibrium. But if $\epsilon > 0$, then, using Tiktok is better than not using Tiktok, so everyone will start using Tiktok and reach the $p = 1$ stable equilibrium.
- $p = 1$: Stable - for the same reason $p = 0$ is stable. This is also the social optimum (the argmax value of the green curve).

Q What about the example in class, where everyone is kinda meh about Tiktok?

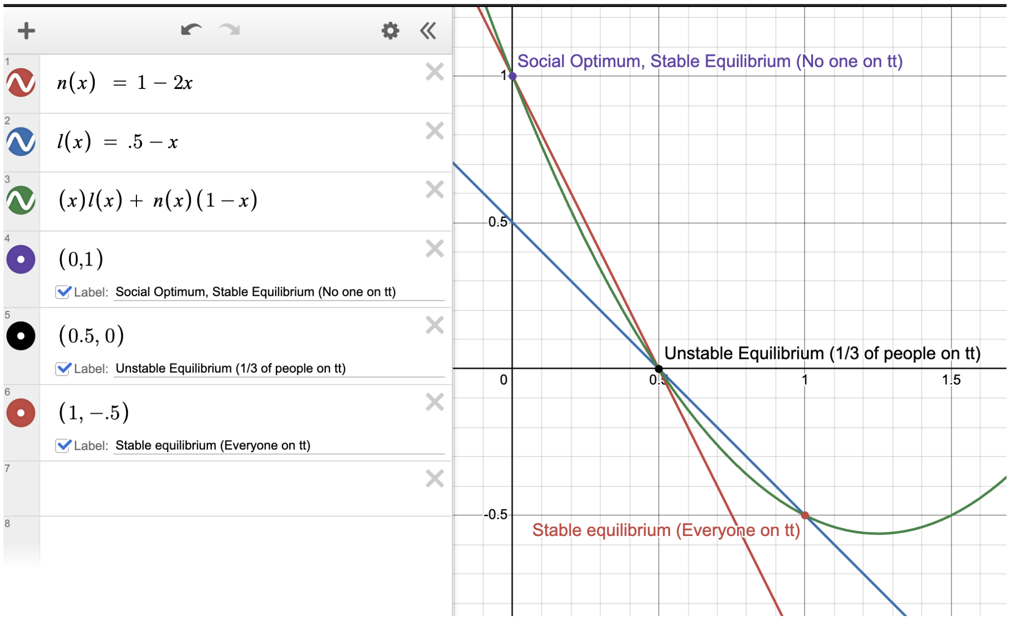

The green line is the average social utility. Observe how it is a weighted average of $n(x)$ and $l(x)$. At $x = 0$, no one is using Tiktok, so the social utility is 1. At $x = 1$, everyone is using TikTok, so the social utility is 0.

Adverse Selection/Moral Hazard

Many concepts in game theory and micro in general rely on information. Thus far in this class, we've been assuming full information in bargaining and such. In reality, most people only know their own valuations and we can only *infer* what the other actors' might be. One of these situations is in Adverse Selection

We will review the relevant concepts as they come up in the problem.

0.1 Problem: Car Insurance

Suppose that a car insurance company assesses people's risks of destroying their cars in an accident. Assume that there is a population of people with latent (unobserved) risk factor r_i, uniformly distributed $r_i \sim U[0, 1]$ - the risk denotes the probability of each person destroying their car (worth $10,000$). Each person suffers expected $U(r) = 0.1(r_i)(L_i)$ (on top of the lost car) from if they don't have insurance, but if they have insurance, their utility loss is $-p$ where p is the price of insurance. It costs the insurance company r_i in expectation for each person with risk factor r_i.

Recall from R2 the formulas for expected values. For discrete random variables like the value of buying a lemon, the formula is:

$$E[U(r_i)] = \sum^{i} P(r_i)U(r_i)$$

For continuous random variables like r_i in this problem, the formula is

$$E[U(r_i)] = \int_{-\infty}^{\infty} r_i(f(r_i))dr$$

where $f(r_i)$ is the probability density distribution of r.

Why does adverse selection occur

Q: What are people willing to pay?

A: People will be willing to pay at least their utility of not paying. Reminder that since they have to worry about things, this is on top of their utility loss from losing their car (call this L). Each person i has utility in expectation

$$E[U(L, r_i)] = r_i(-L) + (1 - r_i)(0) + 0.1(r_i L) = 1.1 r_i L = 11,000 r_i$$

So everyone is willing to pay 1.1x their expected loss to cover for this worry. Companies see this and realize there's a money making opportunity.

Q: What would happen in a world of perfect information?

A: Every individual would get their own price plan, between their cost and what they're willing to pay $(r_i, 1.1r_i)$.

The insurance company doesn't know everyone's types though, and thus has to set a price for everyone. They can do this with guess and check or any process, but since we are omnicient, we can solve this problem for them!

Say they are greedy and want to set the price $p = \$9900$.

Q: Who is the marginal buyer

A: People will continue buying until there's no incentive to buy, so until their willingness to pay is exactly the price p. Here, its $1.1 r_i L = 9900$ meaning $r_i = 9000$. Since we assumed a uniform distribution earlier, we know that 1) everyone that has $r_i \geq .9$ will buy insurance and that 2) this is $\frac{1}{10}$ of people and 3) due to the linear willingness to pay, the average cost of people in this range is the cost of the average risk of people in this range which is $(\frac{1+0.9}{2})$ = average risk person, so cost per person is \$9500.

Note that the cost per person is a lot lower than the last person willing to pay, so there's some surplus missing somewhere

Q: What is the equilibrium price?

A: When the average cost per person is exactly the cost the marginal buyer is willing to pay:

$$(\frac{1+r_i}{2}) \times (10,000) = 11,000r_i \Rightarrow r_i = \frac{5}{6}, p = \$9166.67$$

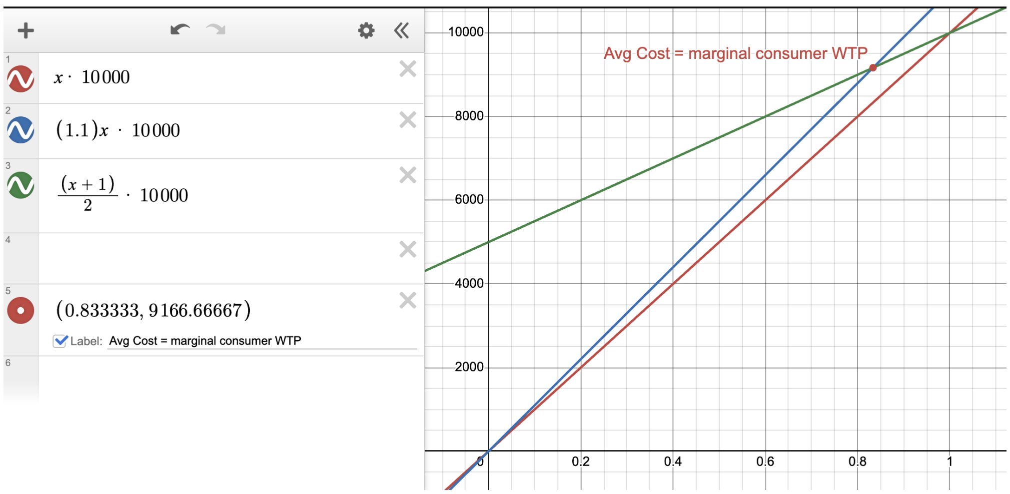

The green line is the insurance company's average cost per customer, the purple line is the cost incurred per customer, and the black line is the marginal customer's willingness to pay. Note that this is opposite of the example drawn in class, as these axis are increasing in x.

Consumers are worried because *they don't know if or if not* something will happen. They know they in expectation, they will pay $10,000r_i$ (and they know their own type so its chill), but any event, including a car crash, is binary! It happen or it doesn't happen. An event can't actually 75% happen. However, due to the *law of large numbers*, the insurance company isn't risk adverse because, if they insure enough people, by the Law of Large Numbers any function of r_i, which is a random variable, will converge to its expectation. $r_i \to E[r_i]$. Then, they will pay *in Expectation* $\$7,500$ each for people with risk r_i

Theorem 1 (Weak Law of Large Numbers)**.** *Let $X_1, X_2, \ldots$ be i.i.d. random variables with $\mathbb{E}[X_1] = \mu$ and $\mathbb{V}ar(X_1) < \infty$. Then the sample mean*

$$\overline{X}_n = \frac{1}{n}\sum_{i=1}^{n} X_i$$

satisfies

$$\overline{X}_n \xrightarrow{p} E[X]$$

Since separate people are independent, we can treat everyone as a random variable with their r_i, so if we have enough people with probability of wrecking their car $= r_i$, the insurance company will pay them each $r_i * L$

Reali ty

However, in practice, people are usually **risk adverse**, meaning they are more worried about the uncertainty than the actual risk. If someone is super super likely to get into a car crash, they know that they're going to get into a car crash so they don't have to worry about if they're getting into a car crash (confusing example, but think: are you more uncertain about the outcome of a coin flip or the outcome of a lottery ticket? You are probably more worried about a coin flip because you know that you're likeley not going to win the lottery. Or, when lottery numbers are called out, people get more and more stressed as numbers get read if they match their ticket, because they think it's more risky, even if the probability of winning would be the same as just getting a ticket). For example: utility could be $U(L, r_i) = r_i L + L\sqrt{r_i}$, meaning people increase a little bit more than linearly with risk.

Theorem 2 (Jensen's Inequality for Concave Functions)**.** *Let X be an integrable random variable. If $\varphi : \mathbb{R} \to \mathbb{R}$ is concave, then*

$$\varphi(\mathbb{E}[X]) \;\geq\; \mathbb{E}[\varphi(X)]\,.$$

By Jensen's Inequality, the expected willingness to pay (which is φ in the above), will always be smaller than the willingness to pay of the average customer (which is what the insurance company is find) so the insurance company can never truly make money. Yay?