Externalities and Public Goods

Lecture 9

Div Bhagia

Introduction

When Markets Fail

Under the conditions of the First Welfare Theorem, competitive markets deliver Pareto-efficient outcomes. But those conditions are demanding: complete markets, well-defined property rights, price-taking behavior, and no spillover effects between agents.

We studied one source of market failure: **asymmetric information**. In this lecture we study two more classic sources of market failure:

- **Externalities:** one agent's action affects another's welfare without passing through the price system
- **Public goods:** goods whose consumption is non-rival and non-excludable, so markets under-provide them

In both cases, private incentives diverge from social incentives and there is a potential role for policy.

Externalities

What Is an Externality?

- An **externality** arises when an economic actor does not face the correct price for taking a specific action. The correct price is the **marginal social cost** of that action.
- When markets work properly, they align private costs and benefits with social costs and benefits. When private benefits differ from social benefits (either higher or lower), externalities result.
- Put differently: externalities occur when one agent's action directly affects another's welfare, and this effect is **not reflected in market prices**.

Classic Examples

- **Traffic congestion:** When I take the highway, I increase congestion for other drivers. But my decision to drive is based on my private cost (time, fuel) and does not account for the additional delay I impose on others.
- **Disease transmission:** When deciding whether to vaccinate, individuals are most likely to weigh only private costs and benefits, ignoring the positive externality of reduced disease spread to others.
- **Pollution:** Because clean air is not priced, we pay essentially no cost to pollute. The marginal cost of our polluting actions do not incorporate the social cost of additional pollution, so we pollute more than is socially optimal.

A Key Question

Can Externalities Be Resolved by the Market?

Until the publication of Ronald Coase's 1960 paper, most economists would have answered **no**.

Coase made them reconsider that view. Coase gave the example of a **doctor and a baker** who share an office building. The baker's loud machinery disturbed the doctor's medical practice.

The standard reasoning (at the time): the baker should compensate the doctor for the harm, since the baker was "causing" the externality.

But Coase pointed out a re-framing: suppose the doctor sets up his office *after* the baker is already there, and then demands the baker shut down. Who is responsible for the externality now?

The Reciprocal Nature of Externalities

One can argue that the *doctor* is creating an externality by requiring the baker to bake in silence. The baker's noise is an "input" into his production of baked goods.

From a legal point of view, the answer may be clear. From an economic point of view, the answer is **indeterminate** based only on this information. We need to know the *costs* of the different solutions.

Suppose:

- The baker could buy **quieter machinery** for **$50**
- The doctor could **soundproof** his walls for **$100**

Economic efficiency demands the lowest-cost solution. The baker should buy quieter machinery.

But does this mean the baker should *pay* to abate? **Not necessarily.**

Two Scenarios

Scenario 1: The town council assigns the doctor the right to control noise. The baker spends $50 for quieter machinery.

Scenario 2: The town council assigns the baker the right to make noise. Will the doctor spend $100 to soundproof? If the parties can negotiate, they should arrange for the doctor to pay the baker $50 for quieter machinery instead.

In **both** cases, the efficient outcome occurs: quieter machinery is purchased. The assignment of rights determines only **who pays**, not what happens.

The Coase Theorem (Formal Statement)

If (1) **property rights are complete** (one party clearly owns the relevant right) and (2) **parties can negotiate costlessly**, then the parties will always negotiate an **efficient solution** to the externality.

The law determines who pays the cost, but the **outcome is the same**.

Note the parallel with the Welfare Theorems: **efficiency and distribution are separable problems**.

When Does Coase Fail?

The Coase theorem implies the market will solve externalities *unless*:

- **Property rights are incomplete** (e.g., no one owns the air)
- **Negotiation is costly** (e.g., the entire population cannot simultaneously negotiate about pollution)

The theorem is often misinterpreted to suggest markets will solve *all* externalities. This is not true. The market can *potentially* solve externalities **if** property rights are clearly assigned and negotiation is feasible.

And in many cases, Coasian bargaining is infeasible due to large numbers of affected parties, high transaction costs, or asymmetric information. In those cases, government intervention may be necessary to correct the externality.

Reducing Pollution: A Numerical Example

Consider two oil refineries that both produce fuel.

- Market **price of fuel is $3** per gallon
- Each refinery uses **$2 in raw inputs** to produce 1 gallon of fuel
- Each plant also produces **smog**, which causes **$0.01** of environmental damage per cubic foot
- Each plant can produce a maximum of **200 gallons**

Smog Production Differs Across Plants

The amount of smog *per gallon* differs at the two plants:

$$s_1 = y_1^2, \qquad s_2 = \frac{1}{2} y_2^2$$

where y_1, y_2 are gallons of fuel produced at each plant.

Plant 2 pollutes only $\frac{1}{2}$ as much as Plant 1 for a given level of production.

Competitive Outcome (No Regulation)

Without any legal framework for resolving the externality, each firm solves:

$$\max_{y_1} \pi_1 = y_1 \cdot (3 - 2) \quad s.t. \quad y_1 \leq 200$$

$$\max_{y_2} \pi_2 = y_2 \cdot (3 - 2) \quad s.t. \quad y_2 \leq 200$$

Result: $y_1^* = y_2^* = 200$. Each firm produces at capacity.

Pollution: $s_1 = 40{,}000, s_2 = 20{,}000$ cubic feet. The externality costs are $400 and $200 respectively.

The Socially Efficient Level

To find the efficient level, equate marginal social benefit to marginal social cost:

$$MB_s = MC_s$$

The social benefit of a gallon of fuel is $3 (from the demand curve). The social cost is $2 in inputs **plus** the pollution damage.

At the margin, no plant should produce more than **$1 of environmental damage** per gallon. So no plant should produce more than **100 cubic feet of smog** per gallon.

Welfare-Maximizing Output

If each plant faced private *plus* social costs:

$$\max_{y_1} \pi_1 = y_1(3-2) - 0.01 \cdot y_1^2 \quad s.t. \quad y_1 \leq 200$$

$$\max_{y_2} \pi_2 = y_2(3-2) - 0.01 \cdot \tfrac{1}{2} y_2^2 \quad s.t. \quad y_2 \leq 200$$

Efficient output: $y_1^{**} = 50, \quad y_2^{**} = 100$

Plant 1 produces less because it is the dirtier plant. Optimal pollution is **not zero**. It reflects the tradeoff between production benefits and environmental harm.

What regulation can achieve this outcome? We will discuss three approaches: command and control, Pigouvian taxes, and cap-and-trade.

Approach 1: Command and Control

The traditional approach: set numerical **quantity limits** on polluting activities.

If the regulator knows each plant's cost structure, they could mandate: "Plant 1 may produce 50 gallons, Plant 2 may produce 100 gallons." This achieves the efficient outcome.

Problems with Command and Control

This kind of regulation is clumsy:

- It's difficult to write laws that regulate each plant **individually**
- Once passed, such laws are difficult to **modify** as technology or costs change
- If the regulator makes a mistake (e.g., reverses the allocations), there is **no self-correcting mechanism**, so the incentives are inefficient

If the law cannot differentiate across plants and must impose a **uniform cap**, further inefficiencies result. (The optimal uniform cap is *not* simply 75 gallons; work this out as an exercise.)

Despite these weaknesses, command and control is the **most common** approach to regulating externalities.

Approach 2: Pigouvian Tax

An alternative: use the **price system** to internalize the externality.

If we set a tax of $t = \$0.01$ per cubic foot of smog, each plant's problem becomes:

$$\max_{y_1} \pi = y_1(3-2) - t \cdot y_1^2 \implies y_1^p = 50$$

$$\max_{y_2} \pi = y_2(3-2) - t \cdot \tfrac{1}{2} y_2^2 \implies y_2^p = 100$$

The tax achieves the efficient result with **less information**: we don't need to know firms' production functions, only the **marginal social damage** of pollution.

Advantages and Risks of Pigouvian Taxes

Advantage: Each firm optimally chooses its own level of pollution given the tax. No need to write a separate law for each plant.

Risk: If marginal damage varies with the *quantity* of pollution (e.g., pollution above a threshold causes mass extinction), setting the tax slightly wrong could be **catastrophic**.

When marginal damage is constant, the Pigouvian tax is straightforward. When it varies, setting the right tax schedule becomes much harder.

Approach 3: Cap and Trade

The Pigouvian tax doesn’t fully use the Coase theorem. The state sets the price, not the market. Can we do better by **assigning property rights** to pollution?

The optimal total pollution is $50^2 + \frac{1}{2}(100^2) = 7{,}500$ cubic feet. The government could issue **7,500 permits**, each allowing 1 cubic foot of smog. Permits can be traded.

Case 1: Permits Given to Plant 2

Plant 2 (the more efficient refinery) receives all 7,500 permits. It could:

1. **Use all permits itself:** produce 122.4 gallons (profit: $122.40)
2. **Produce 100 gallons, sell 2,500 permits to Plant 1:** Plant 1 produces 50 gallons and pays up to $50 for permits. Plant 2 earns $150 total.

Plant 2 **cannot do better** than option 2. The market leads it to the efficient allocation.

Case 2: Permits Given to Plant 1

Now Plant 1 (the dirtier refinery) receives all 7,500 permits. It could:

1. Use all permits: produce 86.6 gallons (profit: $86.60)
2. Sell all permits to Plant 2 (profit: $122.40)
3. **Keep 2,500 permits, sell 5,000 to Plant 2:** Plant 1 produces 50 gallons, Plant 2 produces 100 gallons, and Plant 1 earns up to $150

Again, the efficient allocation emerges. **Regardless of who receives the permits**, the key economic outcome (fuel produced, pollution produced, and the allocation across plants) is **identical**.

What Differs?

Only **which plant makes the profits** (a transfer among plant owners).

This is the power of the Coase theorem in action. By assigning property rights to pollution, the government allows the **market** to correct the externality.

The initial allocation of permits is a major **political** question, but it has no effect on **economic efficiency**.

Comparing the Three Approaches

Feature	Command & Control	Pigouvian Tax	Cap & Trade
Sets	Quantity for each firm	Price of pollution	Total quantity
Information needed	Each firm’s costs	Marginal damage	Total efficient pollution
Self-correcting?	No	Yes	Yes
Firms optimize?	No	Yes	Yes

Why Cap and Trade Has Special Virtues

1. Like command and control, it allows the regulator to set the **total amount** of pollution
2. Like the Pigouvian tax, it's simple to implement (no separate law per firm needed)
3. Unlike either alternative, it causes firms to **optimally reallocate pollution** through trade, even if the regulator doesn't know firms' cost structures

The SO_2 cap-and-trade program in the U.S. dramatically exceeded expectations: it reduced emissions at **far lower cost** than anticipated, because the market revealed that firms' true abatement costs were lower than what they had told regulators under command-and-control.

A Caveat

Cap and trade may not work well if the regulator cares about **which** plant does the polluting.

If all the low-cost polluters happen to be concentrated in one geographic area, cap and trade could lead to **localized pollution hotspots**, even if total pollution is efficient.

Key Takeaways

1. **Externalities** arise when private and social costs diverge, so agents face the "wrong" price
2. The **Coase theorem** shows that well-defined property rights + costless bargaining lead to efficient outcomes regardless of who holds the rights
3. Three regulatory instruments address externalities: **command and control**, **Pigouvian taxes**, and **cap and trade**, each with different information requirements and incentive properties
4. Cap and trade has special virtues: it sets total pollution, lets firms optimize, and generates efficient reallocation through trade
5. Rectifying an externality often means finding a way to restore market conditions so the Coase theorem can hold; when that isn't feasible, external regulation may be needed

Public Goods

Defining Public Goods

Goods can be classified along two dimensions:

1. **Rivalry:** whether one person's consumption reduces availability to others
2. **Excludability:** whether it is possible to prevent anyone from consuming the good

	Excludable	**Non-excludable**
Rival	Private good	Common good
Non-rival	Club good	*Public good*

Can You Classify These Goods?

1. A slice of pizza
2. National defense
3. Street lighting
4. Fish in the ocean
5. A Netflix subscription
6. A public park bench
7. A toll road
8. Clean air
9. A podcast episode
10. Open-source software
11. Fireworks displays
12. Public Wi-Fi at a cafe

Public Goods Model: Setup

- Economy with N households, indexed $i = 1, \ldots, N$
- Two goods: a private good X and a public good G
- Each household consumes private good X_i, with $X = \sum_{i=1}^{N} X_i$
- Public good is **non-rival**: every household consumes the same G
- Public good is funded by individual contributions g_i, so $G = \sum_{i=1}^{N} g_i$
- Utility: $U_i(X_i, G)$, increasing in both arguments
- Each household has income Y_i and faces the budget constraint: $X_i + g_i = Y_i$

Social Planner's Problem

The social planner chooses X_i for each household and G to maximize a weighted sum of utilities subject to the resource constraint:

$$\max_{X_1,\ldots,X_N,G} \sum_{i=1}^{N} \beta_i U_i(X_i, G) \quad s.t. \quad \sum_{i=1}^{N} X_i + G = \sum_{i=1}^{N} Y_i$$

First-order conditions:

$$\beta_i \frac{\partial U_i}{\partial X_i} = \lambda \quad \text{for each } i$$

$$\sum_{i=1}^{N} \beta_i \frac{\partial U_i}{\partial G} = \lambda$$

Social Planner's Optimal

Plugging $\beta_i = \lambda / \frac{\partial U_i}{\partial X_i}$ from the first condition into the second condition gives:

$$\sum_{i=1}^{N} \underbrace{\frac{\partial U_i / \partial G}{\partial U_i / \partial X_i}}_{\text{MRS}_i} = 1$$

This is the **Samuelson condition** for optimal provision of a public good.

Remember: MRS_i is the amount of private good X that household i is willing to give up for one more unit of public good G.

The Samuelson condition says that the total willingness to pay for an additional unit of G across all households must equal the relative cost of providing that unit (which is 1 in this case, since G uses up one unit of resources that could have been used for X).

Decentralized Solution

With no intervention, individuals choose their own contributions to the public good, taking others' contributions as given (**Nash equilibrium**). Each household solves:

$$\max_{X_i, g_i} U_i(X_i, g_1 + g_2 + \ldots + g_N) \quad s.t. \quad X_i + g_i = Y_i$$

Nash equilibrium outcome will satisfy:

$$MRS_i = \frac{\partial U_i / \partial G}{\partial U_i / \partial X_i} = 1$$

Since $MRS_i = 1$, total willingness to pay is $\sum_{i=1}^{N} MRS_i = N > 1$. This means the Nash equilibrium will under-provide the public good compared to the social optimum. This is the **free-rider problem**.

Numerical Example

- Two roommates $i = 1, 2$ sharing an apartment
- Each has \$100 to allocate between:
  - **Private consumption** X_i: food, clothes, personal stuff
  - **Contribution to shared goods** g_i: WiFi, cleaning supplies, Netflix
- Budget constraint: $X_i + g_i = 100$ for each i
- Total public good quality: $G = g_1 + g_2$
  - Both roommates enjoy G equally, regardless of who paid
- Utility: $U_i = X_i \cdot G$
- **The question:** How much does each roommate voluntarily contribute?

Nash Equilibrium

Each roommate takes the other's contribution as given and chooses their own contribution to maximize utility. For roommate 1:

$$\max_{X_1, g_1} X_1 \cdot (g_1 + g_2) \quad s.t. \quad X_1 + g_1 = 100$$

First-order conditions: (1) $g_1 + g_2 = \lambda$, (2) $X_1 = \lambda$, (3) $X_1 + g_1 = 100$.

Solving these gives us the **best response function** for roommate 1:

$$g_1^* = 50 - \frac{g_2}{2}$$

By symmetry, $g_2^* = 50 - \frac{g_1}{2}$. So **Nash equilibrium** contributions are given by:

$$g_1^* = g_2^* = 33.33$$

Social Planner's Problem

The planner chooses X_1, X_2, G to maximize total welfare:

$$\max_{X_1, X_2, G} \; X_1 G + X_2 G \quad s.t. \quad X_1 + X_2 + G = 200$$

Note, this is equivalent to maximizing XG subject to $X + G = 200$. The solution is $X = 100, G = 100$.

Intuition: The planner internalizes the fact that one unit of G benefits both roommates, so the efficient level of G is higher than in the Nash equilibrium.

Comparing Outcomes

Nash Equilibrium

- Each roommate contributes $g_i = 33.33$, so $G = 66.67$
- Private consumption: $X_i = 100 - g_i = 66.67$
- Utility: $U_i = 66.67 \times 66.67 \approx 4{,}445$
- $MRS_i = X_i/G = 1 $ (price ratio)

Social Optimum

- Planner sets $G = 100$, with $X_i = 50$
- Utility: $U_i = 50 \times 100 = 5{,}000$
- $MRS_i = X_i/G = 0.5$, so $_i MRS_i = 1 $ (**Samuelson condition**)

Free-Rider Problem

Free-riding occurs because each roommate benefits from the public good regardless of their own contribution, which leads to under-provision of the public good compared to the social optimum.

Are individuals really free-riders?

Early **lab experiments testing free-rider behavior** (Marwell & Ames, 1981)

- Groups of 5 subjects, each given 10 tokens.
- Can invest in either individual account (private good) or group account (public good).
  - Individual account: 1 token = $1 for the individual.
  - Group account: 1 token = $0.50 for each group member.
- What's the Nash equilibrium? What's the social optimum?

Free-Riding Evidence

- They found 40-60% of token were still invested in the public good.
- The experiment was run on various groups of high school and college students.
- Only one group free-rode alot: 1st-year Econ graduate students (donation rate: 20%).
- Andreoni (1988, 1993) implements experiments with repeated contributions and finds that contributions remain positive but decline over time.

Key question: How can we solve the underprovision problem?

Addressing Underprovision

Government provision: Set G directly, funded by mandatory taxes. Solves free-riding but requires knowing preferences to choose G efficiently.

Subsidizing private contributions: Reduce the private cost of contributing (e.g., tax deductions for charitable giving). Narrows the gap but may not eliminate underprovision.

Lindahl pricing: Charge each person a personalized price equal to their MRS_i. Achieves efficiency in theory, but people have incentives to misrepresent their willingness to pay.

All three face the same fundamental challenge: the government needs information about individual preferences that people have incentives to hide.