Lecture Note 1: A Review of Perfectly Competitive Markets

A major goal of this course is to understand how **market structure**—the number of firms participating in a market and how they compete with each other—affects economic outcomes like prices, quantities, and social surplus.

We'll analyze three different market structures in this course:

- Perfect competition
- Monopoly
- Oligopoly (after Midterm 1)

Before turning to monopoly, it's useful to start with a quick review of perfect competition. (This should be familiar from ECN 100A.)

(Perfectly) competitive firms are price-takers

Commodities like soybeans are sold in highly competitive global markets.

- Soybeans are *homogeneous*: there isn't much product differentiation.
- There are many "small" soybean farms competing with each other.
- If one seller overcharges, it's easy for buyers to find another seller.

It's reasonable to assume that each soybean farm is a **price-taker**.

- Each farm can sell as much or as little as it wants at price p per unit.
- We can pretend that the firm is "forced" to charge p.
  - Why not charge more than p? Nobody would buy.
  - Why not charge less than p? You'd just lose money.
- For a price-taking firm, the demand curve is horizontal (or "flat").

Price-taking is the key, defining feature of a perfectly competitive market. If we say a market is competitive, we mean that firms treat p as a constant.

Maximizing profits under perfect competition

Suppose that a farm can produce q units of soybeans at a total cost $C(q)$.

- We'll assume $C(q)$ is increasing, continuous, and differentiable.
- In many cases, $C(q)$ will also be convex (decreasing returns to scale).

The farm's **profits** equal its **total revenues** minus its **total costs**.

$$\pi(q) = R(q) - C(q)$$

"Revenue" is all the money a firm gets from its customers. Since the soybean market is competitive, our farm receives p for each of the q units it sells:

$$R(q) = pq \implies \pi(q) = pq - C(q)$$

For a competitive firm, the **marginal revenue** simply equals the price:

$$MR = \frac{d}{dq}R(q) = \frac{d}{dq}(pq) = p$$

We write the firm's **profit-maximization problem** as follows:

$$\max_q \ \pi(q) = pq - C(q)$$

The notation here means: "Choose the value of q that maximizes profits."

To find the optimal quantity, we derive the **first-order condition** (FOC):

$$\frac{d}{dq}\left[pq - C(q)\right]\Bigg|_{q=q^*} = p - C'(q^*) = 0 \implies p = C'(q^*)$$

You've seen this before: "price equals marginal cost" (or "$p = MC$").

More precisely: in a competitive market, each firm chooses its level of output so that the price it faces exactly equals its marginal cost of production. Why?

If $p > MC$, then the firm can increase its profits by increasing its output. If $p < MC$, it should lower its output. Either way, it ends up at $p = MC$.

Solving for q^* with "real numbers"

The optimal quantity is implicitly defined by the FOC:

$$p = C'(q^*)$$

If we know the function $C(q)$, we can solve this equation explicitly for q^*.

Example: Suppose that $p = 20$ and $C(q) = q^2$. Find q^* and $\pi(q^*)$.

$$\max_q \ \pi(q) = 20q - q^2 \implies 20 - 2q^* = 0 \implies q^* = 10$$

At the optimum, $\pi(q^*) = 20 \cdot 10 - 10^2 = 100$.

You might remember from calculus that a solution to the first-order condition is only a "candidate optimum". So there are some additional details we need to think about. We'll discuss these details in next week's video lecture on optimization. For now, we'll trust that $q^* = 10$ really is the optimum (it is).

Exercises for self-study: Consider a farm operating in a competitive market.

(A) Let $C(q) = q^2 + 5q$. Find q^* if $p = 11$. Find q^* if $p = 3$. (Careful!)

Profits are $\pi(q) = pq - q^2 - 5q$, so the FOC is $p - 2q - 5 = 0$.

If $p = 11$, then $q^* = 3$. If $p = 3$, the FOC tells us that $q^* = -1$. But that's no good! The farm can't produce negative output, so the best it can do is set $q^* = 0$. We'll discuss corner solutions in the next lecture.

(B) Let $C(q) = 12\sqrt{q}$. For what price p will the firm choose $q^* = 9$?

Profits are $\pi(q) = pq - 12\sqrt{q}$, and the FOC is $p - \frac{6}{\sqrt{q^*}} = 0$. Setting $q^* = 9$, this simplifies to $p = 2$.

From farm to market: building supply curves

If we solve a soybean farm's profit-maximization problem for different values of p, we obtain the **firm supply curve**. This tells us how much the firm would be willing to supply at each price it might potentially face.

Example: A competitive firm with costs $C(q) = q^2$ faces price p. Find $q^*(p)$.

$$\max_q \; pq - q^2 \implies p - 2q^* = 0 \implies q^*(p) = \frac{1}{2}p$$

Practice your economic intuition: Why is q^* an increasing function of p?

When p is small, the firm only produces units that are cheap to produce. When p rises, it becomes profitable to produce more expensive units, too.

We construct the **market supply curve** by adding up the quantities supplied by all firms. If there are N firms in the market, numbered from 1 to N, then market supply is

$$Q_S(p) = q_S^1(p) + q_S^2(p) + \ldots q_S^N(p)$$

where $q_S^i(p)$ is the amount supplied by firm i. (Throughout the course, we'll use "Q" for market quantities and "q" for each firm's quantity.)

Graphically, this means we add up the firms' supply curves *horizontally*. For example, with three identical firms:

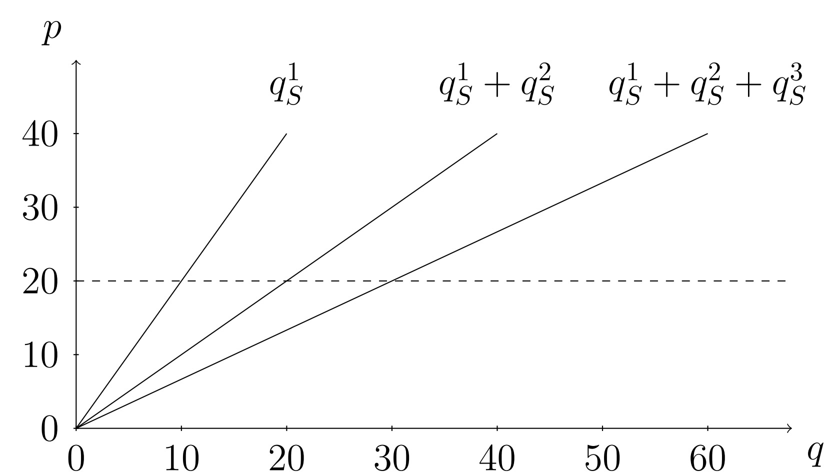

Supply, demand, and equilibrium

Of course, we also have buyers—in this case, probably wholesale merchants rather than individual consumers—interested in purchasing soybeans.

Each buyer has an individual demand curve indicating how many units of soybeans they'd be willing to buy at each price. As with supply, we construct the **market demand curve** by adding up the individual demands.

Under perfect competition, the **equilibrium price** and **equilibrium quantity** are determined by the intersection of the (market) supply curve and the (market) demand curve.

Example: Suppose that there are 100 soybean farms each with $q_S(p) = \frac{1}{2}p$, so that the market supply curve is $Q_S(p) = 100 \cdot \frac{1}{2}p = 50p$. Suppose that the market demand curve is $Q_D(p) = 3000 - 50p$.

Assuming that the market is perfectly competitive—which seems reasonable since 100 is a lot of firms—find p^* and Q^*.

Solution: Setting $Q_S = Q_D$, we obtain

$$50p = 3000 - 50p \implies p^* = 30 \implies Q^* = 1500$$

We can also solve this problem graphically. Since economists plot price against quantity, we should start by rearranging terms (solving each equation for p) to obtain the **inverse demand curve** and **inverse supply curve**.

- inverse demand curve: $p_D(Q) = 60 - \frac{1}{50}Q$
- inverse supply curve: $p_S(Q) = \frac{1}{50}Q$

Economists often use the terms "demand curve" and "inverse demand curve" interchangeably. It's just a question of whether we express Q as a function of p or p as a function of Q. Either way, it's the same curve.

Social welfare under perfect competition

Here's what our soybean market looks like:

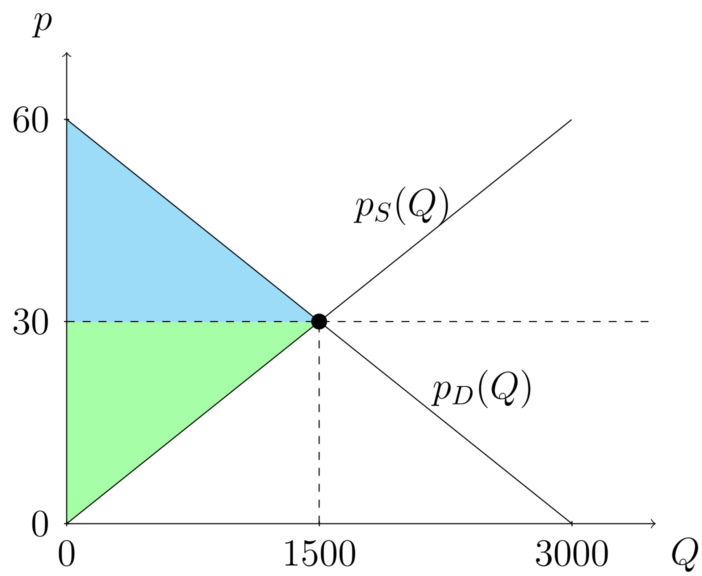

Consumer surplus (CS), in blue, is the amount, in dollars, that the buyers benefit from the existence of this market.

- If you're thinking "triangle", be careful! The consumer surplus doesn't have to be a triangle (though it often is in "ECN100" problems).
- The CS is the area above the price and below the demand curve.
- In this case, CS $= \frac{1}{2}(60 - 30)(1500) = 22{,}500$.

Producer surplus (PS), shown in green, is the amount that the sellers benefit from the existence of this market. In other words: profit.

- The PS is the area below the price and above the supply curve.
- Here, PS $= \frac{1}{2}(30 - 0)(1500) = 22{,}500$ as well.

Thanks to perfect competition, this market is working well: the outcome is **Pareto efficient**—we also use the terms "Pareto optimal", "socially efficient", and "socially optimal"—and there is **no deadweight loss**.

Taxing a competitive market

The United States has historically exported huge quantities of soybeans to China each year, but these exports have fallen amid the ongoing trade war.

Imagine that the market we're analyzing reflects US soybean sales to China. Suppose that China imposes a tariff t (i.e., a tax) on US soybean producers for each unit of soybeans sold to China. What are the new p^* and Q^*?

Taking the market price p as given, each producer now solves

$$\max_q \; pq - q^2 - tq \implies p - 2q^* - t = 0 \implies q^*(p) = \frac{1}{2}p - \frac{1}{2}t$$

With 100 identical producers in the market, market supply is

$$Q_S(p) = 100 \cdot \left(\frac{1}{2}p - \frac{1}{2}t\right) = 50p - 50t$$

At any given price, market demand is unchanged:

$$Q_D(p) = 3000 - 50p$$

So the new equilibrium is given by

$$50p - 50t = 3000 - 50p \implies p^* = 30 + \frac{1}{2}t \implies Q^* = 1500 - 25t$$

Since we've left t unspecified (rather than plugging in some particular value), it's easy for us to perform **comparative statics**:

- The equilibrium price is increasing in the size of the tax.
- The equilibrium quantity is decreasing in the size of the tax.

Even though the **nominal incidence** of this tax (who "writes the check") is on producers, the **economic incidence** (who "really" pays for it) is shared by consumers and firms: both producer and consumer surplus have fallen relative to the no-tax equilibrium.

A graphical look at the soybean tax

Under the soybean tax, the market supply curve is

$$Q_S(p) = 50p - 50t \implies p_S(Q) = \frac{1}{50}Q_S + t$$

In words: the supply curve has shifted up by the full amount of the tax.

For example, suppose that the tax rate is $t = 20$. Plugging this into our expressions for equilibrium price and quantity, we find that

$$p^* = 30 + \frac{1}{2} \cdot 20 = 40 \quad \text{and} \quad Q^* = 1500 - 25 \cdot 20 = 1000$$

Let's confirm this graphically:

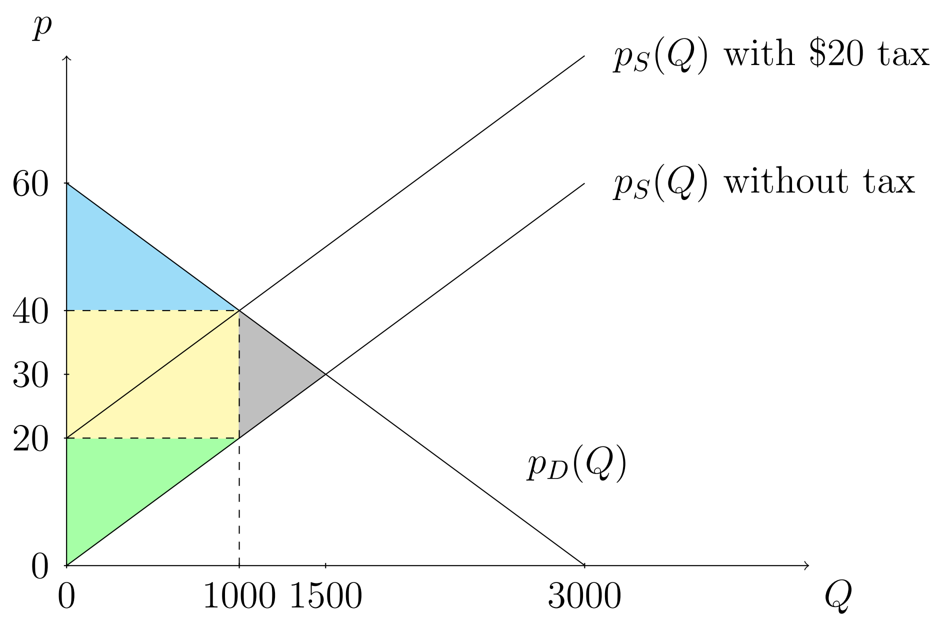

Exercise: Compute the consumer surplus, producer surplus, tax revenue, and deadweight loss without a tax, then under a \$20 tax on producers.

- Consumer surplus (blue) falls from 22,500 to 10,000.
- Producer surplus (green) falls from 22,500 to 10,000.
- Tax revenue (yellow) rises from 0 to 20,000.
- Deadweight loss (gray) rises from 0 to 5,000.