Econ 802

Lecture Notes on Chapter 13

Greg Dow November 25, 2020

Up to This point, we have been studying price-taking firms. But if every firm (and consumer) is a price-taker, Then where do prices come from?

The Theory of perfectly competitive markets provides one answer to This question (There are other answers involving monopoly, oligopoly, etc).

In This chapter we consider perfect competition in a partial equilibrium context. Chapters 17 and 18 consider general equilibrium.

There are 3 new elements relative to what we have done earlier:

1. we have more Than one firm producing The same output.
2. we have a market demand curve for This output.
3. we have a market supply curve (which may involve either The short run or The long run)

The model of perfect competition is most appropriate when There are many buyers and sellers, The firms are producing a homogeneous output, and The buyers and sellers are well informed about The characteristics of The good or service.

We do not need to assume The firms are identical; They could have different technologies or levels of fixed inputs (in The SR).

The usual idea behind price taking behavior is that a firm believes it can sell as much output as it wants at the current market price (its output level has no noticeable effect on this price).

The firm believes that if it charged a higher price than this, it would have zero demand, because all consumers would go to other firms. It also believes that if it charged a lower price, it would get the entire market demand, but there is no gain from doing this, because it can sell as much as it wants already at the current market price.

One implication is that in a competitive market, there is a single price. If two firms charged different prices, the one with the higher price would have zero demand, so it may as well match the lower price (it can always choose to sell zero output at this price if it wants to).

Consider an individual firm with cost function $c(y)$ and facing the market price p. Such a firm chooses output to solve

$$\max_{y \geq 0} \{py - c(y)\}$$

The FOC says that if the solution has $y^* > 0$ we must have $p = c'(y^*)$.

If the solution has $y^* = 0$ we must have $p \leq c'(0)$.

This follows from Kuhn-Tucker conditions but the intuition should be clear: if it is optimal to produce zero output, then the derivative of profit with respect to output cannot be positive at $y = 0$, otherwise the firm would prefer some $y > 0$.

The SOC is $c''(y^*) \geq 0$ (necessary)
or $c''(y^*) > 0$ (sufficient)

When there is an interior solution $y^* > 0$, the FOC says that price equals marginal cost, and the SOC says MC is rising.

Note: sometimes people get sloppy in their use of language and say "the firm sets price equal to marginal cost." Technically this is not correct. The firm is a price taker so it does not set price equal to anything. What it really does is to look at a given market price and then adjust output until $c'(y)$ is equal to the price p.

There is one complication with the calculus approach used above. Even if there is some $y^* > 0$ with $p = c'(y^*)$ and the SOC holds, it could be that this local solution is not a global solution. Consider a short run situation where $c(y) = c_v(y) + F$. Maybe $p = c_v'(y^*)$ and $c_v''(y^*) > 0$, but

$$py^* - c_v(y^*) - F < 0 \quad \text{so profit is negative.}$$

Then the question is whether it is better to accept this negative profit, or shut down (set $y = 0$) and just pay the fixed cost, which gives $-F$. It is better to choose $y^* > 0$ if
(or at least as good)

$$py^* - c_v(y^*) - F \geq -F \quad \text{or} \quad py^* \geq c_v(y^*)$$
$$\text{or} \quad p \geq \frac{c_v(y^*)}{y^*}$$

The last inequality says price is at least equal to average variable cost.

Here is a graphical situation where this problem arises.

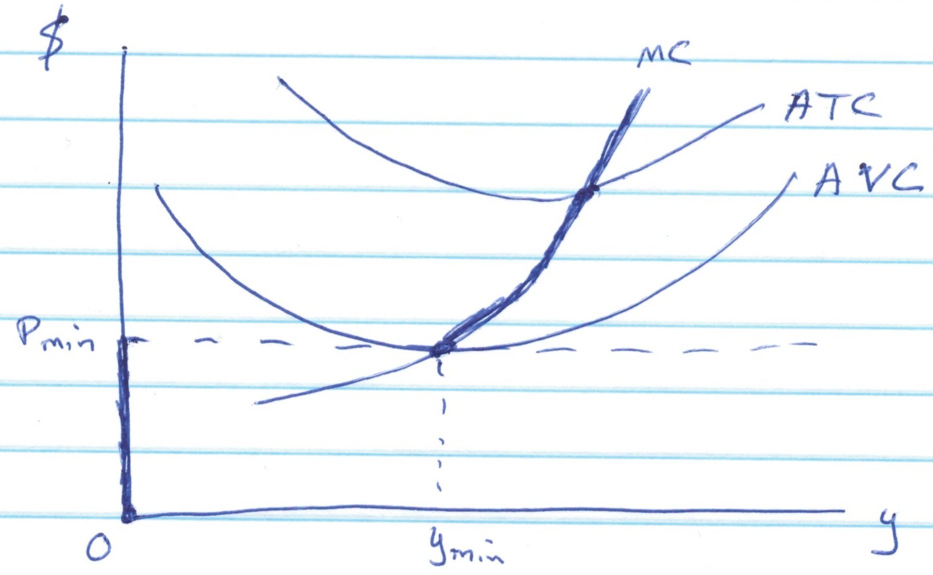

Note: the fixed cost F does not appear in MC or AVC, so it has no effect on p_{min} or y_{min}.

If $p < p_{min}$, then $p < AVC$ for all $y > 0$ and the firm prefers to shut down (set $y = 0$)
If $p > p_{min}$, then the firm chooses $y^* > 0$ such that $p = c'(y^*)$ and it satisfies the condition $p > AVC(y^*)$.
If $p = p_{min}$, the firm is indifferent between producing $y = 0$ or $y = y_{min}$ (but will not produce any intermediate y)

In this case (which generally arises in the short run whenever AVC is U-shaped), the firm's output supply function is

$y = 0$ for $p < p_{min}$

$y = y(p)$ for $p \geq p_{min}$ where $y(p)$ solves the FOC $p = c'(y)$.

Note that we often assume the firm produces positive output at p_{min}, despite being indifferent.

The firm's output supply curve here is discontinuous. It is the heavy part of the vertical axis for $p \leq p_{min}$ and the heavy part of the marginal cost curve for $p \geq p_{min}$.

More generally, we define the firm's output supply function $y(p)$ to be the optimal level of y determined by a given p. However, this assumes that there is a unique solution to the firm's profit maximization problem at each level of p. As we have just seen, this is not always true (there could be multiple solutions at certain values of p) so we have to be careful.

If the number of firms is fixed and they all have well defined output supply functions, the market supply function is

$$S(p) = \sum_{i=1}^{n} y_i(p) \quad \text{where price is the independent variable.}$$

Note: you need to make sure you are summing over the output levels produced at a given price. Do not sum over prices, this makes no sense!

An equilibrium price is defined to be a price at which the aggregate quantity supplied equals the aggregate quantity demanded, or $S(p) = D(p)$. This leads to the most popular graph in economics:

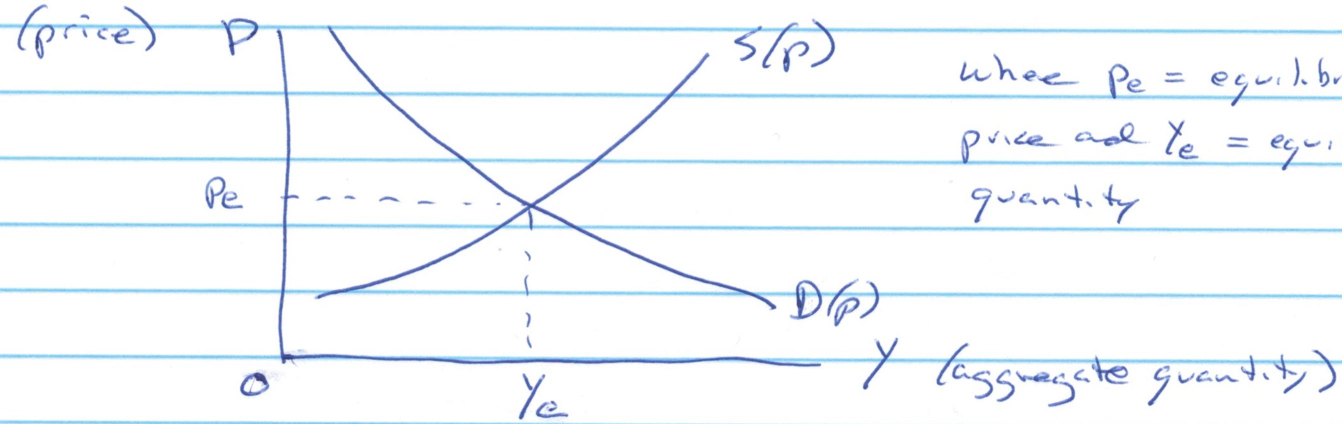

where P_e = equilibrium price and Y_e = equilibrium quantity

Algebraically, in a situation like this we would set quantity supplied = quantity demanded and solve for p_e:

$$S(p_e) = D(p_e)$$

Once we know p_e, it is usually easy to find Y_e by plugging p_e into either the supply or demand function.

Now let's consider whether a market equilibrium must exist. Suppose we have

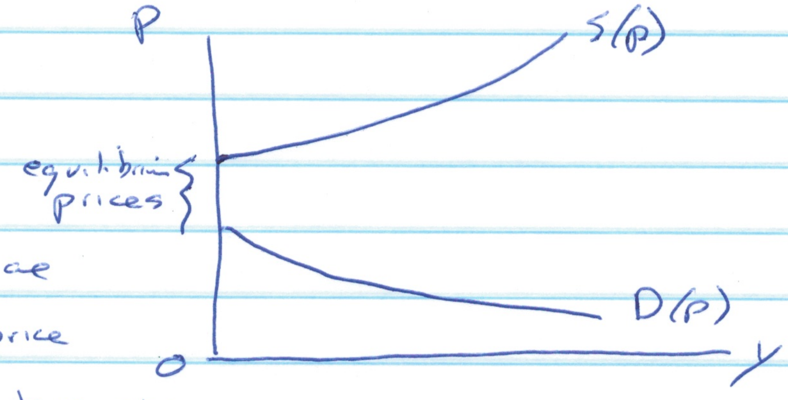

There actually is an equilibrium in such cases; in fact there are usually many. Any price between the vertical intercepts of $S(p)$ and $D(p)$ is an equilibrium, because at such prices, quantity supplied is equal to quantity demanded (both are zero).

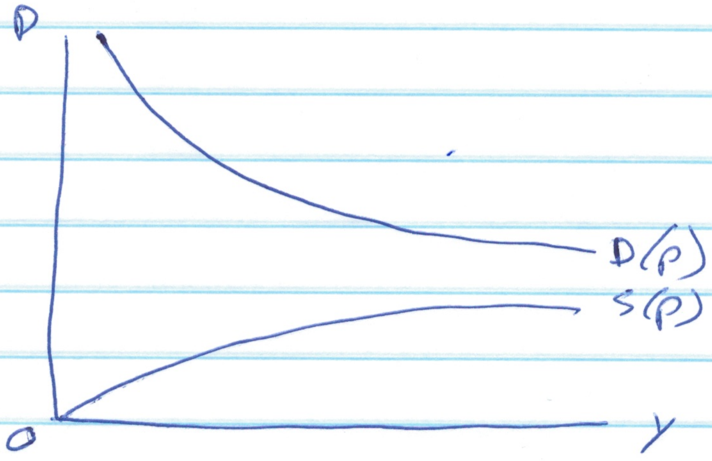

Suppose instead we have a case like this, where D and S do not intersect. This does not make much economic sense. First, it says that firms are willing to produce unlimited output at a finite price. Second, it says that consumers are willing to buy an unlimited amount at a positive price, which will violate their budget constraints. So we can ignore this possibility.

What if we have a case like This?

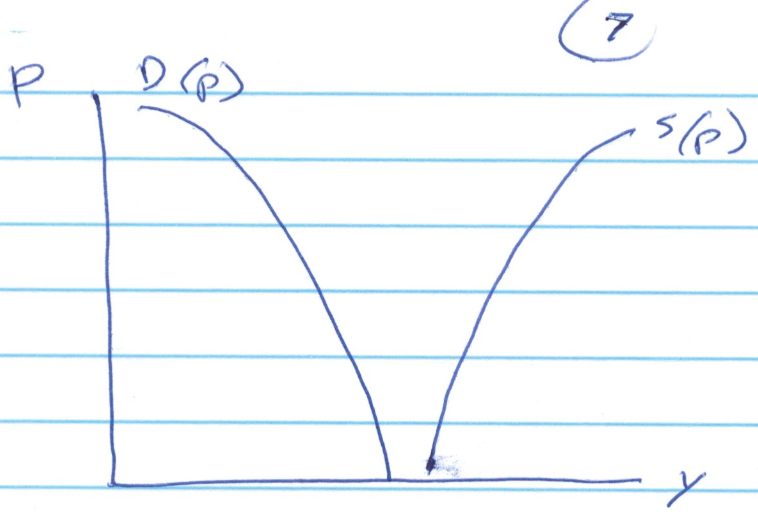

Again This does not make much sense; it says That firms are willing to produce a positive output at a zero price, and it says consumers want a finite amount even if The price of the good is zero. Neither idea is very plausible so we can ignore This possibility too.

Conclusion: in any sensible economic model, at least one equilibrium price p_e should exist.

Once we have identified p_e, we may be interested in The output of an individual firm and The resulting profit level. For example, if The cost curves look like This, we have

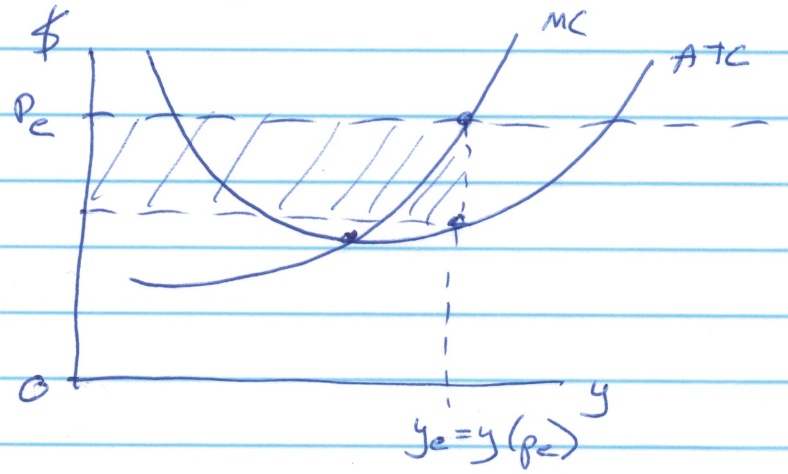

where y_e is The profit-maximizing output at p_e and profit is The shaded area: $y_e\left(p_e - \frac{c(y_e)}{y_e}\right)$

$= y_e(p_e - ATC(y_e))$

Note: it is important not to ~~consider~~ confuse The optimality condition $P = MC$ with The zero profit condition $P = AC$. We sometimes have $P = AC$ (for example in certain long run equilibrium situations) but it is not implied by profit max.

Next let's think about the difference between the short run and the long run.

In the short run market equilibrium is determined using two assumptions.

1. There are some fixed inputs. Therefore we need to use SMC, SAC, etc for each firm. We also need to consider whether $P \geq AVC$ in order to find out whether it is optimal to shut down. In general, profit could be positive, zero or negative in the SR.
2. There is a fixed set of firms (no entry or exit). The idea is that the same set of inputs that are fixed for the existing firms are also fixed for potential firms, so positive profit will not attract entry. With a fixed set of firms, we can sum up the outputs of the individual firms to get market supply (graphically, the horizontal sum of the individual ~~dem~~ supply curves gives the market supply curve).

In the long run,

1. There are no fixed inputs so we use LMC and LAC. There is no fixed cost, and profit is generally non-negative (a firm can always produce zero output and have both zero cost and zero revenue).
2. The set of firms may not be fixed (entry or exit could occur). However, this depends on the situation (for example there might be a fixed number of firms if it is necessary to get a license from the government in order to operate).

I have already discussed the derivation of SR supply curves, so now I will run through a series of special cases that differ according to the shape of the LAC curve. In each case I distinguish between the LR supply curve that would arise with a fixed number of firms versus free entry and exit.

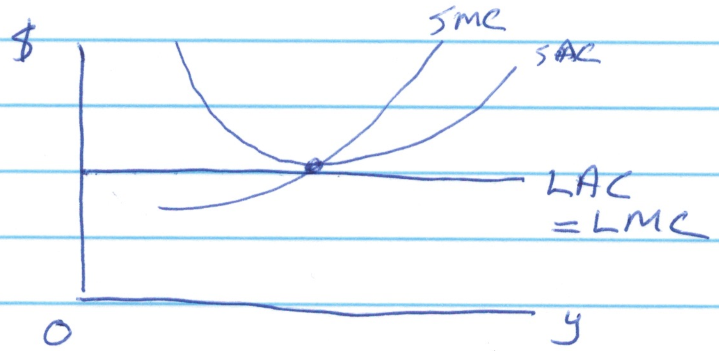

① constant LAC

In this case, the LR market supply curve is horizontal, regardless of whether entry is allowed or not.

For an individual firm, profit maximization is only well defined if $p \leq LAC$. If $p < LAC$, the firm sets $y = 0$. If $p = LAC$, any output level is optimal because profit is always zero. So we can think of the firm's supply function as

$y^* = 0$ for $p < LAC$

y indeterminate for $p = LAC$

no solution for $p > LAC$.

At the market level, the same thing is true with a fixed number of identical firms: the supply function is zero for $p < LAC$ and horizontal for $p = LAC$; the size of the individual firms is indeterminate. If we have free entry and exit the shape of the market supply curve is unchanged (no one enters for $p < p_{min}$ and firms are indifferent about entry for $p = p_{min}$)

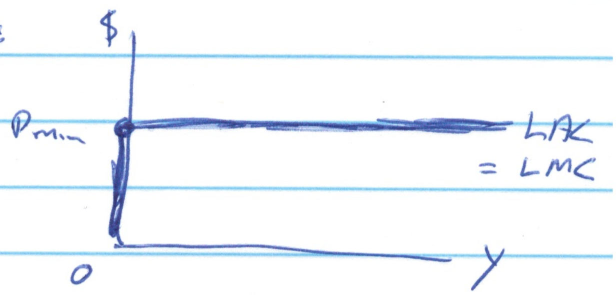

To solve for a market equilibrium, we would think about a graph like this:

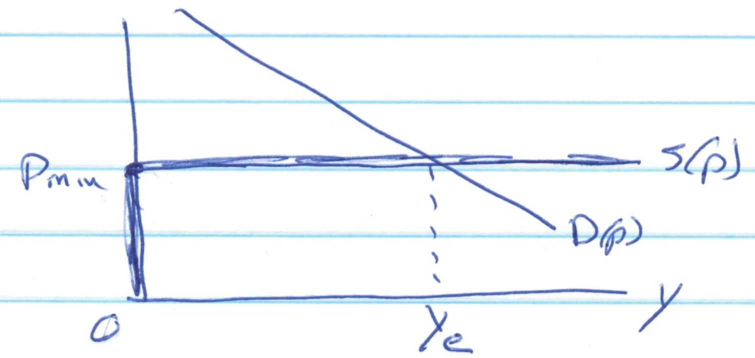

The equilibrium price p_e is obtained directly from $p_e = p_{min} = LAC$, and the equilibrium quantity is obtained by plugging this into demand: $y_e = D(p_e)$. The only exception is when the demand curve has a vertical intercept below p_{min}. In that case there will be an interval of equilibrium prices that all give $y_e = 0$.

Note that in this framework if the firms have different levels of LAC, it will only be the firms with the lowest level of LAC that produce positive output. These firms will drive price down to a level where the ~~[illegible]~~ higher-cost firms would have negative profit and prefer to exit.

→ Increasing LAC

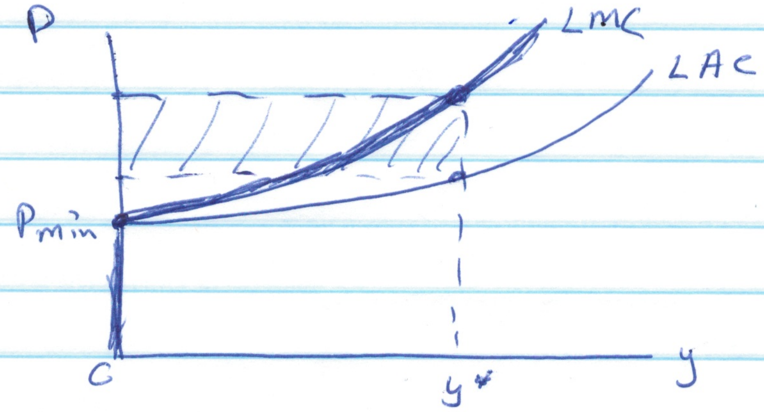

Remember that when LAC is rising, LMC must be above it. In general, these curves do not need to pass through the origin, so we can have $p_{min} > 0$. Whenever $p > p_{min}$, the firm produces some positive output $y^* > 0$ and has positive profit (see shaded area). The firm's supply curve is shown by the heavy line and curve.

(a) restricted entry. We often assume that when profit is positive, in the long run new firms will enter. But suppose there is some barrier to entry so the number of firms is fixed. It should be clear from the graph that each firm has a well defined output supply function, so we can sum over the firms to get

$$S(p) = \sum_{i=1}^{n} y_i(p)$$

where here the $y_i(p)$ refer to the long run supply functions of the individual firms.

Note: In this situation, Varian says there is positive "accounting" profit but zero "economic" profit. What he means by this is that there is some non-marketed fixed factor that is causing LAC to rise, and the shaded area in my graph is a return to this fixed factor. He calls this return a "rent" and treats it as a cost from an economic point of view, so profit is zero.

I think this is overly complicated and potentially misleading. Yes, you can define profit this way if you want to. But this eliminates the usefulness of the profit concept. I think it is better to allow the possibility of positive profit, because this doesn't have to be caused by a fixed input; it could be caused by a government restriction on the number of firms, for example. It is also possible that LAC rises due to inputs from nature like sunshine or rain, and it doesn't make much sense to say that sunshine gets a rent.

(b) free entry. If new firms can enter in response to positive profit, the market supply curve shifts out like this:

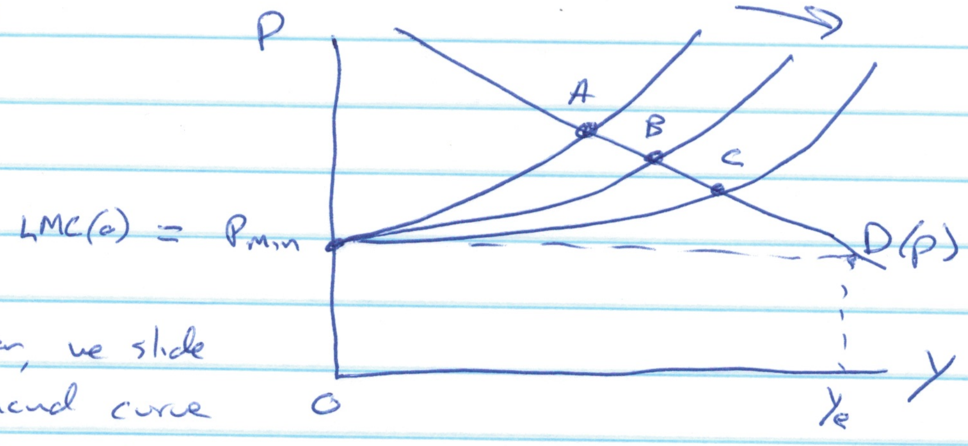

As firms enter, we slide down the demand curve from A to B to C. As price drops each individual firm produces less (see earlier graph of LMC), but aggregate output rises because there are more firms. As long as $P > P_{min}$, profit is still positive for each firm and entry continues. In the limit we have $P \to P_{min}$, $y(p) \to 0$ for each firm, $S(p) \to Y_e$ at the aggregate level, and profit $\to 0$ for each firm. Economists normally treat this limiting result as the equilibrium in the LR.

(3) Decreasing LAC.

(a) restricted entry:

If we try to construct a supply function for an individual firm, we run into problems.

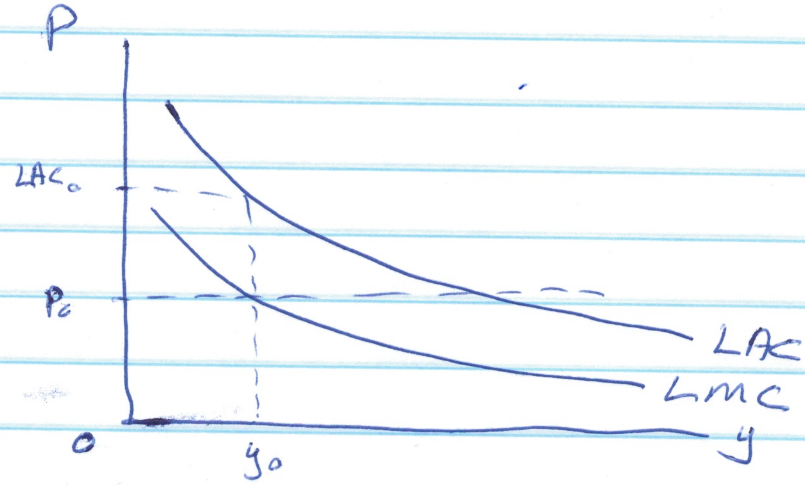

Suppose price is p_0 and we use $p_0 = LMC(y_0)$. The resulting profit is negative ($p < LAC_0$). The issue is that the necessary SOC is violated (LMC is falling) so the FOC is giving a local minimum, not a maximum.

As we learned earlier in the course, falling LAC occurs when the firm has increasing returns. In such cases, there is generally no solution to the profit max problem. Thus, the firm's output supply function is undefined, and even with a fixed set of firms, we cannot sum over the firms to get a market supply function.

If there is only one firm, we would use a monopoly model. With more than one firm, we would use an oligopoly model.

(b) free entry. If we have falling LAC and free entry, we need to take a course in industrial organization to understand what might happen.

④ U-shaped LAC

(a) restricted entry

This is like the SR situation where AVC is U-shaped. All inputs are variable, so we replace AVC by LAC.

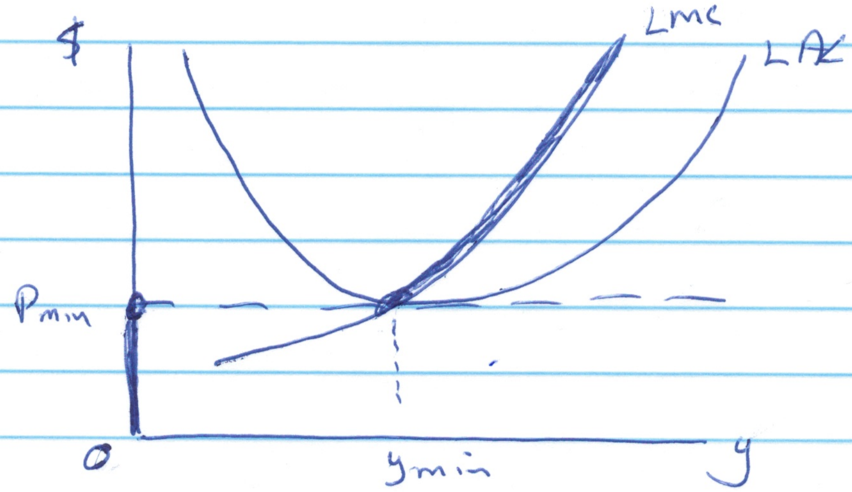

can sum over firms to get $S(p)$

For $p < p_{min}$, the firm has negative profit at all positive output levels and prefers to do nothing ($y = 0$).
At $p = p_{min}$ the firm is indifferent between $y = 0$ and $y = y_{min}$ (both give zero profit). However, it will not produce any intermediate output because profit would be < 0.
For $p > p_{min}$, we use the standard FOC $p = LMC(y^*)$.
The firm's supply curve is indicated by the heavy lined curve, with a discontinuity at p_{min}.

(b) free entry

This case is complicated. If $p > p_{min}$, profit is positive, so new firms enter. This shifts out the market supply curve. If there is some integer number of firms n, where $p_{min} = D(n y_{min})$, then we can have a LR equilibrium where $p_e = p_{min}$, there are n firms, and each firm produces y_{min} units of output. Although the firms are indifferent between y_{min} and zero, y_{min} does maximize profit (it just turns out that the maximum profit is zero at p_e). So this will work.

More generally, there may not be any such integer. To see what happens in this case, it is simplest to think about LMC as a straight line:

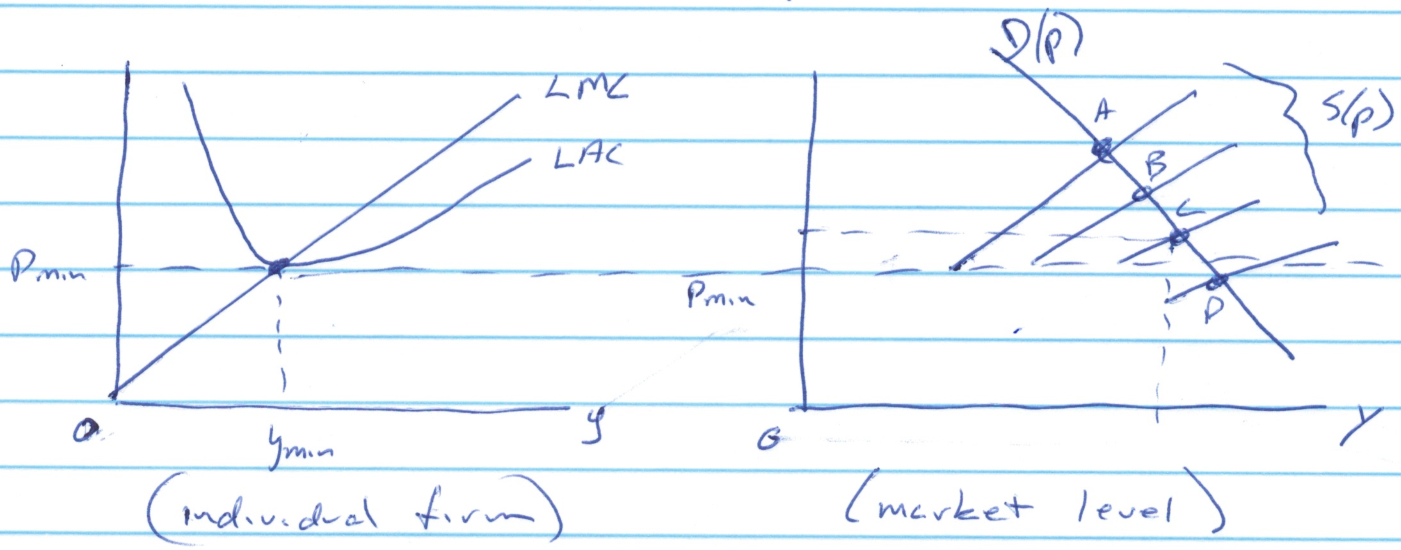

As entry occurs, we shift out $S(p)$ and price falls from A to B to C. Suppose A corresponds to 1 firm, B corresponds to 2 firms and C corresponds to 3. Will a fourth firm enter? Probably not, because this leads to $P < p_{min}$ and negative profit. So it might make sense to say that entry stops at C, and this is the LR equilibrium.

Following this logic, LR equilibrium occurs when we have the largest possible number of firms n such that the supply/demand equilibrium satisfies $P \geq LAC$.

The reason I say this "might make sense" is that technically we are violating the assumption of price-taking behavior. The argument on p. 14 assumes that the fourth firm knows it would make profit negative, and thus it does not enter. But this implies that the firm knows it has some effect on the market price. We assumed earlier that firms do not believe they have any effect on price in a competitive market.

If we were being careful, we would have to admit that this is a contradiction. However, most economists don't worry about it.

Note: even if we don't worry about the violation of price-taking (with free entry), there is still an issue with non-existence of equilibrium with a fixed number of firms. Suppose we have this (there are n firms):

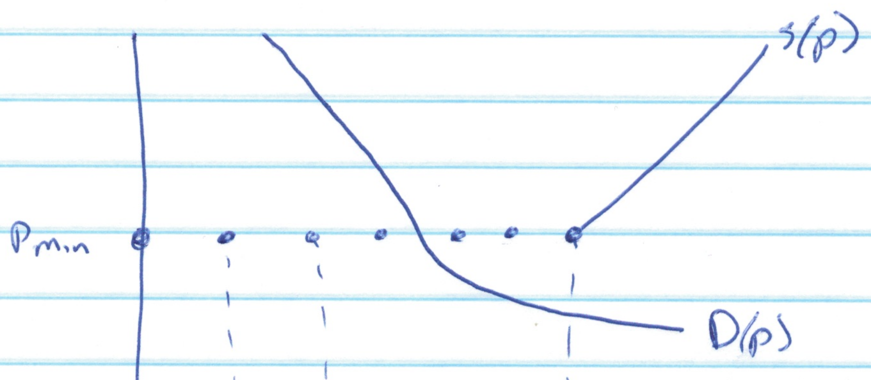

Each firm is indifferent between 0 and y_{min}, and is unwilling to produce an intermediate output. If the demand curve gives an output at P_{min} that is not an integer multiple of y_{min} we have no equilibrium: $P < P_{min}$ and $P > P_{min}$ won't work and we can't get $S(p) = D(p)$ at P_{min} (this would require that some firm must have negative profit).

Taxation

Varian postpones this topic until the end of Chapter 13 but I will deal with it here in order to illustrate how the distinction between short run and long run supply curves can be important.

Suppose we have CRS technology and the government imposes a tax on firms of t per unit produced.

Short run analysis:

Suppose we start from the price p_0 where the firm has zero profit. For $p > p_0$, its SR supply curve is SMC before the tax. When the government imposes a tax of t per unit, this shifts up marginal cost vertically by t, so the new supply curve with the tax is SMC′

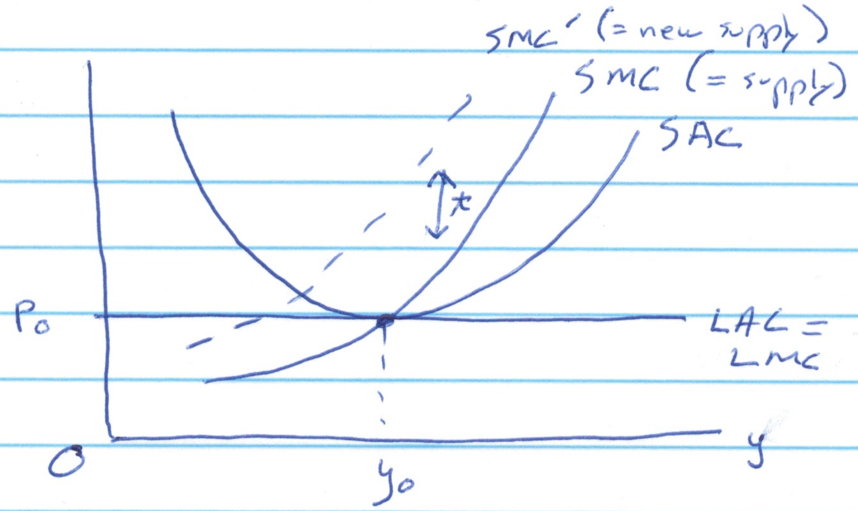

Now go to the market level of analysis. Shifting each firm's SMC up by t will shift the SR market supply curve up by t.

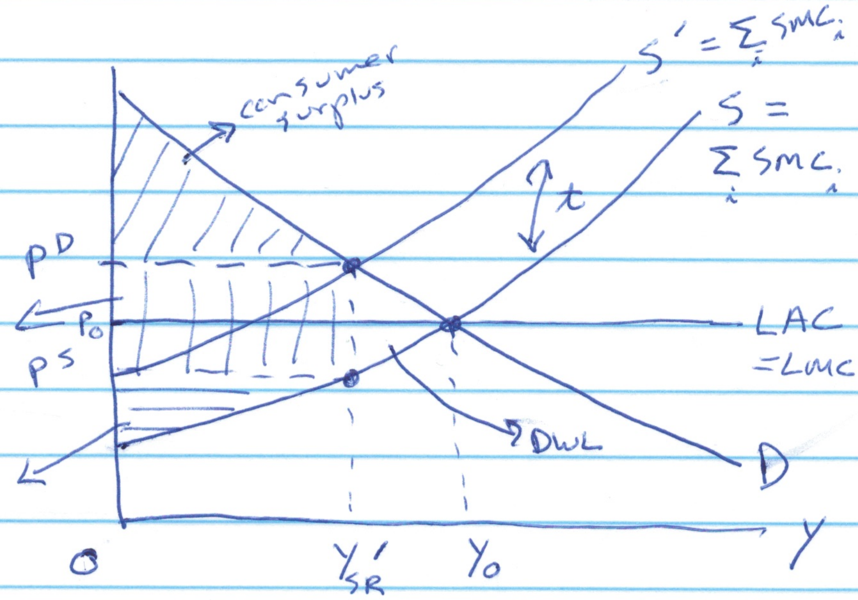

In the new equilibrium with the tax, consumers pay P^D, which is called the "demand price," and total output falls from Y_0 to Y_{SR}'. The price firms actually receive (after paying the tax) is P^S, which is called the "supply price." The difference between the two is the amount of the tax: $t = P^D - P^S$.

Consumer surplus is defined to be the area below the demand curve but above the price consumers actually pay (which here is P^D).

Producer surplus (in the SR) is defined to be the area above the supply curve but below the price firms receive (which here is P^S).

Revenue to the government from the tax is the area $(P^D - P^S) \cdot Y_{SR}' = t\, Y_{SR}'$ because it collects t per unit.

The triangular area between Y_{SR}' and Y^0, below the demand curve and above the supply curve, is called "deadweight loss." Without the tax, the price would have been P_0 and output would have been Y_0, so the DWL triangle would have added something to consumer and producer surplus, but with the tax, this part of surplus disappears.

An important point in this graph is that consumers do not bear the entire burden of the tax. The price increase to consumers is $P^D - P_0$ which is less than t. The rest of the tax is borne by firms through the drop from P_0 to P^S. The amount of the tax borne by consumers and firms depends on the elasticities of the supply and demand curves.

Also note that in the short run with the tax, firms must have negative profit, because $P^S < P_o = LAC \leq SAC$.

Long run analysis.

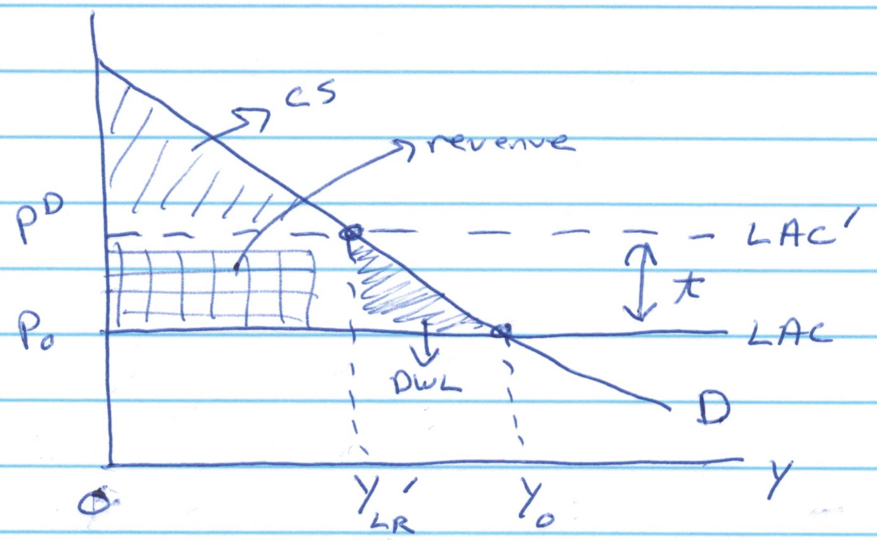

With CRS, the long run market supply curve in the absence of a tax is LAC. When the tax is imposed, this shifts up to LAC', where $LAC' - LAC = t$.

The new LR equilibrium is at P^D and Y'_{LR}. Notice that in this situation $P^D - P_o = t$ so consumers bear the entire burden of the tax. Firms face the original price P^o, which is what they keep after collecting P^D from consumers and then paying the tax.

Firms have zero profit in the new LR equilibrium because $P^S = P_o = LAC$. However, total output is lower than Y_o. Note that there is no producer surplus here because $LAC = LMC = P_o = P^S$, so there is no area between the LR supply curve and the price firms receive.

The reason why consumers bear the entire tax in the LR is that we have assumed CRS, which implies a horizontal LR supply curve, which is infinitely elastic.

A few more comments on tax incidence.

With CRS, firms do not "bear the tax."

In fact in the long run, they are indifferent; they had zero profit in the original equilibrium and they end up with zero profit again in the new LR equilibrium.

So why do firms complain about taxes on the goods they produce?

1. in real markets, there may not be CRS
2. in real markets, firms may not be perfectly competitive
3. because aggregate output drops with a tax, there is less demand for inputs, and some of those inputs may be supplied by the owners of the firm
4. firms may be worried about short run adjustments (recall that the SR effect of the tax was to give firms negative profit).

Welfare Analysis

The last thing I want to cover in chapter 13 is how we can think about the welfare generated by competitive markets in a partial equilibrium framework. This will link back to the concepts of consumer and producer surplus.

Consider the simplest possible model: one consumer, one firm, and two goods.

The consumer's utility function is $u(x, y) = b(x) + y$ where $b' > 0$, $b'' < 0$, and $b(0) = 0$.

This is a quasi-linear utility function. You can think of x as a specific good and y as "expenditures on everything else." Or you can just think of y as a second good.

Note: Varian writes utility as $u(x) + y$, which is a bit misleading because $u(x)$ is not the entire utility function.

Throughout the analysis I set $p_y \equiv 1$ so the price of the y good does not change (all that matters is the price of x relative to this).

The reason we are using quasi-linearity is that this ensures there will be no income effects for the x good (the Marshallian and Hicksian demands will be identical). This is a big simplification.

Assume the consumer has an endowment of ω units of the y good, but no endowment of the x good.

However, there is a firm that can convert y into x, and this firm is controlled by the consumer.

The firm's technology is described by a cost function $c(x)$, which is the physical amount of the y good needed to produce x units of the x good. Because we set $p_y \equiv 1$, you can also think of $c(x)$ as the cost in dollars of producing x.

We assume $c' > 0$, $c'' > 0$, and $c(0) = 0$.

What would the consumer like to do? The answer is: maximize utility subject to the technology constraint. That is,

$$\max\ b(x) + y \quad \text{subject to } y = \omega - c(x).$$

(The amount of the y good available for consumption will be the endowment ω minus whatever is used up to produce the x good).

Substituting the constraint into the objective function, the consumer chooses x to solve

$$\max_{x \geq 0}\ b(x) - c(x) + \omega$$

FOC: $b'(x^*) = c'(x^*)$ Note: ω is a constant and therefore irrelevant.

SOC: $b''(x^*) - c''(x^*) < 0$ (true under our assumptions)

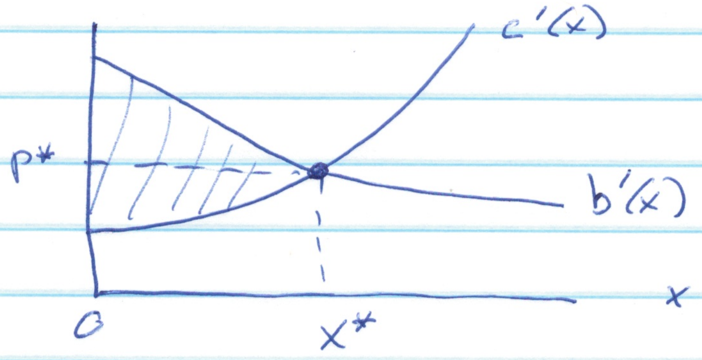

Although we did not use prices or a budget constraint, this looks a lot like a graph of supply and demand!

Let's do the same thing using prices.

Demand side: $\max b(x) + y$

subject to $px + y = \omega$

Setting this up as a normal utility max problem, we have the budget constraint $px + y = \omega$ because p is the price of x, we fixed the price of y at $p_y = 1$, and ω is the value of the consumer's endowment ($p_y \omega$ = income, or simply ω because $p_y = 1$; there is no endowment of the x good.)

Substituting the budget constraint into the utility function gives $\max b(x) + \omega - px$

FOC: $b'(x^*) = p$ (SOC holds)

Supply side: $\max\limits_{x \geq 0} px - c(x)$ because the firm maxes profit

FOC: $p = c'(x^*)$ (again SOC holds)

So we get a competitive equilibrium where the amount the consumer wants to buy is equal to the amount the firm wants to sell when $p = p^*$ (see graph on p. 21)
The marginal utility $b'(x)$ functions as the demand curve and the marginal cost $c'(x)$ functions as the supply curve. They intersect where $p = p^*$ and $x = x^*$.

Note: don't worry about the price taking assumption here. Yes, it looks silly with just one consumer and one firm, but we are trying to develop concepts in a simple way.

Now let's introduce the concept of total surplus. First notice that on the graph on p. 21, x^* maximizes the area between the demand curve $b'(x)$ and the supply curve $c'(x)$. This is the sum of consumer + producer surplus.

Mathematically: The area under the demand curve up to x is

$$\int_0^x b'(s)\,ds = b(x)$$

where we use $b(0) = 0$.

The area under the supply curve is $\int_0^x c'(t)\,dt = c(x)$

where we use $c(0) = 0$.

So maximizing $b(x) - c(x)$, as we did on p. 21, is the same as maximizing the size of the area under the demand curve and above the supply curve, which is total surplus.

We can write this equivalently as

$$\underbrace{[b(x) - px]}_{\text{consumer surplus}} + \underbrace{[px - c(x)]}_{\text{producer surplus}}$$

This gives a basic intuition for why economists tend to like competitive markets: price taking behavior by consumers and firms leads to the maximization of total surplus and this seems like a good thing.

However, we need to show that this general argument carries over to many consumers and many firms, where price-taking would be more reasonable.

Assume we have many consumers $i = 1..n$
with utilities $u_i(x_i, y_i) = b_i(x_i) + y_i$.
Each has an endowment ω_i of the y good.

Also assume we have many firms $j = 1..m$
with cost functions $c_j(z_j)$ where z_j is firm j's output of the x good and $c_j(z_j)$ is the amount of the y good it uses up as an input.

Define an allocation to be (x_i, y_i) for $i = 1..n$
and z_j for $j = 1..m$.

Imagine we have a benevolent social planner who wants to choose an allocation that maximizes the sum of the utilities of all consumers, subject to feasibility constraints:

$$\max \sum_{i=1}^{n} b_i(x_i) + \sum_{i=1}^{n} y_i$$

$$\text{subject to} \quad \sum_{i=1}^{n} y_i = \sum_{i=1}^{n} \omega_i - \sum_{j=1}^{m} c_j(z_j)$$

$$\text{and} \quad \sum_{i=1}^{n} x_i = \sum_{j=1}^{m} z_j$$

Note: There are no prices here just physical quantities.

The constraints ensure that we are not giving a higher aggregate consumption level for either good than the total amount available.

Set up a Lagrangean where we eliminate the $\sum_i y_i$ constraint by substitution and use a multiplier λ for the $\sum_i x_i$ constraint:

$$L = \sum_i b_i(x_i) - \sum_j c_j(z_j) + \sum_i \omega_i - \lambda \left[\sum_i x_i - \sum_j z_j \right]$$

$$\text{FOC:} \quad \left.\begin{aligned} b_i'(x_i^*) - \lambda &= 0 \qquad i = 1..n \\ -c_j'(z_j^*) + \lambda &= 0 \qquad j = 1..m \end{aligned}\right\}$$

Note: FOC are sufficient due to strict concavity of the objective.

Interpretation: marginal utility of x_i is equal for all i

marginal cost of z_j is equal for all j

Because we have quasi-linear utility, the marginal utility of the y good is always 1, so another way to say this is that the marginal rate of substitution (MRS) is equal for all consumers: $MRS_i = \dfrac{MU_i^x}{MU_i^y} = \lambda$ for all i.

Similarly, we can interpret c_j' as the marginal rate of transformation (MRT) for firm j; that is, the rate at which it can convert the y good into the x good. FOC says $MRT_j = \lambda$ for all j.

Furthermore, we must have MRS = MRT because the same multiplier λ is used in each case.

These are standard necessary conditions for a Pareto efficient allocation of resources.

What if we get rid of the social planner and use market prices instead?

Each consumer solves $\max\ b_i(x_i) + y_i$

subject to $px_i + y_i = w_i$

substitute for y_i and differentiate with respect to x_i:

FOC: $b_i'(x_i^*) = p$ all $i = 1..n$

Each firm solves $\max \{pz_j - c_j(z_j)\}$

FOC: $p = c_j'(z_j^*)$ all $j = 1..m$

It should be clear that the resulting allocation is exactly the same as with the social planner; we have just replaced the multiplier λ by the price p.

The condition $\sum_i x_i^* = \sum_j y_j^*$ is the same as the feasibility constraint in the planner's problem but now we interpret it as demand = supply for the x good

Likewise, the condition $\sum_i y_i^* = \sum_i w_i - \sum_j c_j(z_j^*)$ was a feasibility constraint in the planner's problem, but now it is demand = supply for the y good.

Conclusion: a system of competitive markets will maximize the sum of the consumers' utilities, subject to physical feasibility constraints (assuming we have quasi-linear utility).

→ Note: we will say a lot more about Pareto efficiency in Chapter 17.

Last question: why is it desirable to max the sum of the utilities? In the quasi-linear framework, this is necessary for Pareto efficiency. We can always take away a unit of the y good from one consumer and give it to another, which shifts a unit of utility from one person to another. So there are no restrictions on redistribution of utility; for a given allocation of the x good, we can achieve any desired distribution of utility by reallocating the y good. But if the allocation of x does not maximize total utility, we can make everyone better off simultaneously.

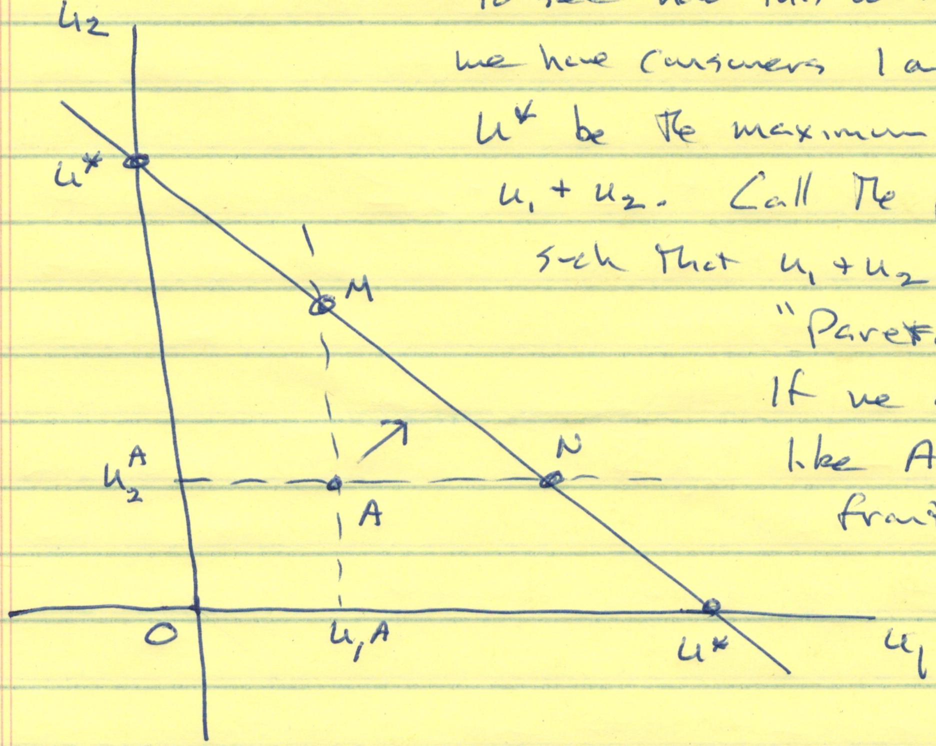

To see how this works, suppose we have consumers 1 and 2. Let u^* be the maximum possible sum $u_1 + u_2$. Call the points (u_1, u_2) such that $u_1 + u_2 = u^*$ the "Pareto frontier". If we are at a point like A below the frontier, we can reallocate x in a way that increases total utility and achieves u^*. Then we can distribute the y good in a way that gives a point along the frontier between M and N, which makes both people better off simultaneously.