I. The Fundamental Geometry of Net Benefit Maximizing Choice

A. Nearly all economic models can be developed from a fairly simple model of rational decision making that assume that individuals **maximize** their private **net benefits**.

Consumers maximize consumer surplus: the difference between what a thing is worth to them and what they have to pay for it. CS(Q) = TB(Q) - TC(Q)

Firms maximize their profit:, the difference in what they receive in revenue from selling a product and its cost of production: Π = TR(Q) - TC(Q)

B. The change in benefits, costs, etc. with respect to quantity consumed or produced is generally called Marginal benefit, or Marginal cost.

DEF: **Marginal "X"** is the change in Total "X" caused by a one unit change in quantity. It is the slope of the Total "X" curve. "X" $\in$ {cost, benefit, profit, product, utility, revenue, etc.}

Important Geometric Property: Total "X" can be calculated from a Marginal "X" curve by finding the area under the Marginal l "X" curve over the range of interest (often from 0 to some quantity Q).

This property allows us to determine consumer surplus and/or profit from a diagram of marginal cost and marginal revenue curves.

It also will be used through out the course to calculate similar net beneftis for industries, consumers, or society as a whole, and helps to make the logic of (normative) welfare economics clear.

C. Examples:

Given the marginal cost and marginal benefit curves in Figure 1, it is possible to calculate the total cost of Q' and the total benefit of Q' .

The first step is normally to label the areas so that one can do some "area arithmetic." The relevent areas are those under the MB and MC curves in the range of interest.

In Figure 1, total cost is the area under the relevant actor's (consumer or firm) marginal cost curve, here TC(Q') = II ; and total benefit is the area under the relevant actor's (consumer or firm) marginal benefit curve, here TB(Q') = I + II .

Because net benefits are total benefits minus (less) total costs, one **can calculate the net benefits** by finding the actor's total benefit and total cost for the quantity or activity level of interest, and subtracting them.

Thus the net benefit of output Q' is TB(Q') - TC(Q') = [I + II] - [II] = I.

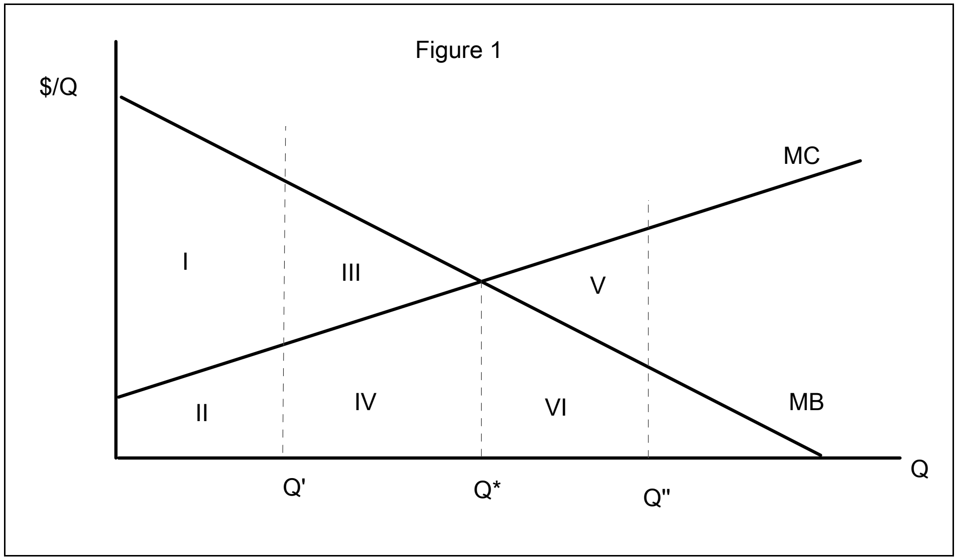

Use Figure 1 to determine the areas that correspond to the total benefit, cost and net benefit at output Q* and Q".

Answers:

TB(Q*) = I + II + III + IV , TC(Q*) = II + IV , NB(Q*) = I + III

TB(Q') = I + II + III + IV + VI , TC (Q") = II + IV + V + VI , NB(Q") = I + III - V

D. If one attempts to maximize net benefits, it turns out that generally he or she will want to consume or produce at the point where marginal cost equals marginal benefit (at least in cases where Q is very divisible).

There is a nice geometric proof of this. (The example above, C, nearly proves this. Note that NB(Q*) > NB(Q') and NB(Q*) > NB(Q").)

In the usual chase, a net-benefit maximizing decision maker chooses consumption levels (Q) such that their own marginal costs equal their own marginal benefits. They do this not because they care about "margins" but because **this is how they maximizes net benefits** in most common choice settings of interest to economists. (Another common choice that maximizes net benefits is Q* = 0. Why?)

This characterization of net benefit maximizing decisions is quite general, and can be used to model the behavior of both firms and consumers in a wide range of circumstances.

Moreover, the same geometry can be used to characterize ideal policies if "all" relevant costs and benefits can be computed, and one wants to maximize *Social Net Benefits.*

E. That each person maximizes their own net benefits does not imply that every person will agree about what the ideal level or output of a particular good or service might be.

Most individuals will have different marginal benefit or marginal cost curves, and so will differ about ideal service levels.

To the extent that these differences can be predicted, they can be used to model both private and political behavior:

(What types of persons will be most likely to lobby for subsidies for higher education?

What types of persons will prefer progressive taxation to regressive taxation?

What industries will prefer a carbon tax to a corporate income tax?)

F. One can **use the consumer-surplus maximizing model to derive a consumer's demand curve** for any good or service (given their marginal benefit curves) by: (i) choosing a price, (ii) finding the implied marginal cost curve for a consumer, (iii) use MC and MB to find the CS maximizing quantity of the good or service, (iv) plot the price and the CS maximizing Q*, and (v) repeat with other prices to trace out the individual's demand curve.

G. Similarly, one can use a profit maximizing model (another measure of net benefit) to derive a competitive firm's **short run** supply curve, given its marginal cost curve. Again, one (i) chooses a price (which is a price taking firm's MR curve), (ii) finds the profit maximizing output, (iii) plot P and Q*, (iv) repeat to trace out a supply curve.

II. Deriving Consumer's Demand Curve from MB curves using the Net Benefit Maximizing Model

A. If consumers are rational net benefit maximizers, it turns out that both individual and market demand curves are downward sloping.

B. To demonstrate this, we will first determine what net benefit maximizing behavior implies about individual demand curves and then what individual demand curves imply about market demand curves.

To derive a demand curve from a consumer's MB curve:

(i) pick a price, P

(ii) find the associated MC curve

(iii) find the NB maximizing quantity, Q (the amount bought by a rational consumer).

(iv) plot P and Q

(v) repeat with another price

C. The next figure illustrates the geometry of this process. It begins with a marginal benefit curve (MB) and uses two prices, P1 and P2, to find two points on this person's demand curve.

Note that the **price**(s) characterizes a consumer's **marginal cost curve**(s), since he or she can has to pay that amount to get each successive unit of this good. The first MC curve is simply a horizontal line at P1.

Given that MC curve, this person will chooses the quantity that maximizes his or her consumer surplus, which is labeled Q1.

This price-quantity combination is plotted on the Demand diagram as the point P1, Q1. (At price P1, the consumer will buy quantity Q1.)

Now we repeat this process with another price, P2. The implied new (second) MC curve is a horizontal line passing at P2.

Given that MC curve, a net benefit maximizing (rational) consumer chooses the quantity that maximizes his or her consumer surplus, which is found at the quantity where MC=MB, and is labeled Q2.

This price-quantity combination is plotted on the Demand diagram as the point P2, Q2. (At price P2, the consumer will buy quantity Q2.)

In principle this process is continued forever, choosing prices, finding the associated CS-max quantities, and plotting the price and quantity.

But in most cases three or four points are enough to get the basic geometry of a consumer's demand curve.

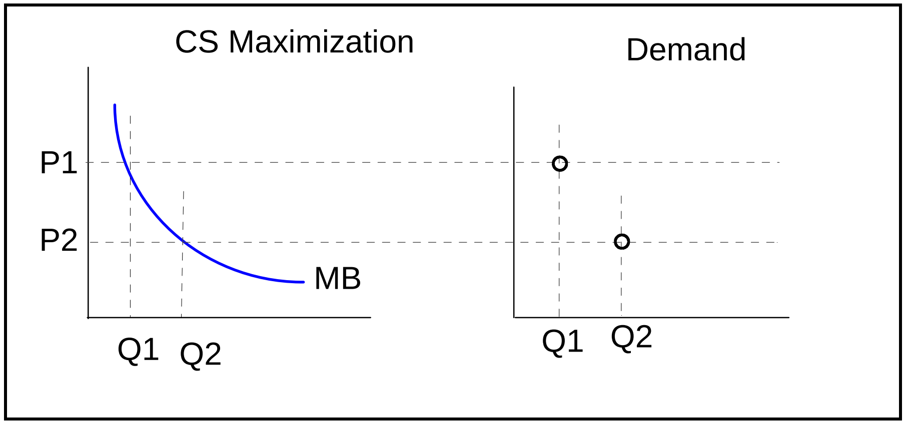

THE POINT HERE IS NOT TO MEMORIZE DIAGRAMS but to understand how to prices affect purchases by rational consumers.

It is the logic that is general, rather than the particular diagrams.

Given a consumer's MB curve, we can deduce his or her demand curve!

The rational person or organization will do X in circumstance Y, because X maximizes his or her net benefits in those circumstances.

D. This basic process can be used to derive a consumer's Demand curve from any sort of MB curve, although they are not always as easy as this one.

[Several less straightforward derivations will be done in class!]

Note that whenever the marginal benefit curve is downward sloping over its whole range, the consumer's demand curve goes through exactly the same points as the MB curve and is also downward sloping.

In such cases, one can use estimates of individual demand curves as estimates of MB curves.

However, although they go through the same points, marginal benefits and demand are not the same function, because they **have different meanings** and the **functions go in opposite directions**.

MB functions go from Q into $/units, whereas demand functions go from P into Q. They are inverse functions of each other.

E. **IN CLASS, A VARIETY OF SIMILAR DERIVATIONS WILL BE DONE ON FOR VARIOUS MB CURVES, ONE OF WHICH WILL BE ASSIGNED AS A HOMEWORK PUZZLE. (So do show up for class!!!)**

In some cases, the demand curves include only a subset of the point on an individual's MB curve, but the points in common are always from the downward sloping portions of the MB curve.

Thus, when a demand curve is derived in this way from MB curves, it turns out that **every** individual demand curve slopes downward.

[Later in the course, we will see that this is **not necessarily the case when a demand curve is derived from a utility** function.]

F. Notice that we **already have some predictions** about behavior that follows simply from the net-benefit maximizing model of consumer choice:

Demand curves generally slope downward.

As price rise, consumers tend to purchase fewer and firms tend to produce more.

A change in MB will cause demand curves to shift up or down according the shift in MB.

Consumer net benefits tend to fall as prices rise, other things being equal.

Firm profits tend to rise as prices rise, other things being equal.

III. Deriving a Firm's Supply Curve from MC Curves

A. In a similar fashion, one can use the **profit maximizing model** (another measure of net benefit) to derive a competitive firm's short run supply curve from its short run marginal cost (MC) curve.

B. The method used to derive a firm's supply curve is very similar to that used to derive a consumer's demand curve:

(i) Choose a price

(ii) Find the associated marginal revenue curve (MR)

(iii) Find the profit maximizing quantity of the good or service, given that MC curve

(iv) Plot the price and the profit-maximizing quantity on another diagram.

(v) Repeat with several other prices to trace out a supply curve.

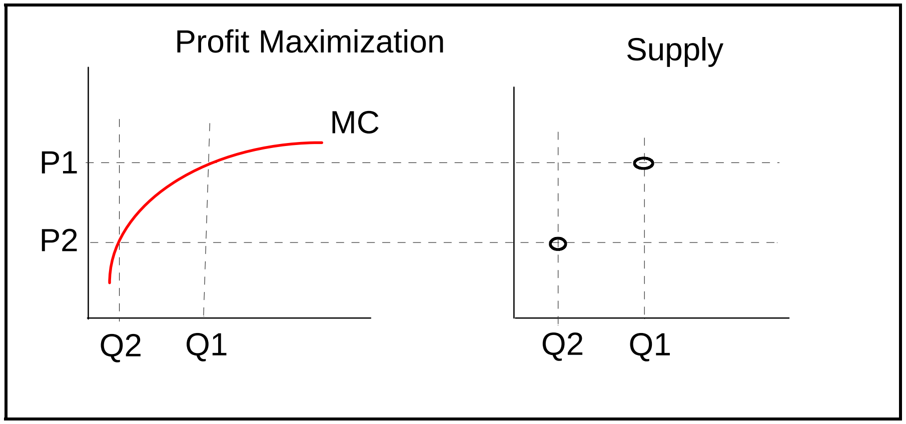

C. The figure above illustrates these steps for a given marginal cost curve (MC) and two prices, P1 and P2.

(i) In a competitive market, the market price is every firm's marginal revenue curve, since the firms gets "P" new dollars of revenue every time it sells an additional unit at that price.

So the first MR curve is a horizontal line through P1.

Given that MR curve, the firm will chooses the quantity that maximizes his or her profit (the firm's net benefit), which occurs where MR=MC, and is labeled Q1.

This price and output combination is plotted on the Supply diagram as the point P1, Q1. (At price P1, the firm will produce and sell quantity Q1.)

(ii) Next we try another price, P2. A horizontal line through P2 is the firm's marginal revenue curve, since its revenue now increases by that P2 dollars every time another unit of the good is sold.

Given that new MR curve, this firm chooses the quantity that maximizes profits, which is again found where MC=MR and is labeled Q2.

This price output combination is plotted on the Supply diagram as the point P2, Q2. (At price P2, the firm will produce and sell quantity Q2 if the price is P2.)

(iii) In principle this process is continued forever, choosing prices, finding the associated profit-maximizing quantity, and plotting the price and quantity.

But in most cases three or four points are enough to get the basic geometry of a firm's supply curve, but more will be needed when MC has a complex shape.

(iv) Note that when the MC is upward sloping (and MC>0) that a firm's supply curve goes through exactly the same points as its MC curve.

This implies that one can use estimates of a firm's supply curve as estimates of its MC curve.

(As in the case of demand curves, a firm's supply curve is the inverse of its MC curve. Supply goes from prices ($/unit) into quantities; whereas MC goes from quantity into $/unit measures of how costs increase as output increases.)

However, if MC has a more complex shape, only the points on MC that can represent profit maximizing outputs are on both the MC and firm supply curve.

Since it turns out that those points in common are all from the upward sloping parts of the firm's MC curve, a firm's supply curve slopes upward.

D. Again geometry and deduction generate a series of predictions about a firm's behavior

Firm supply curves tend to be upward sloping.

As prices rise, other things being equal, firms produce and sell more units of their products.

Any change in market conditions that affects a firm's marginal costs will shift the firm's supply curve.

For example, an increase in input prices tends to cause the marginal cost of production to increase.

If MC rises for a firm, its supply curve shifts back to the left.

If MC rises for a , profits tend to fall for that firm, other things being equal.

If market prices rise, MR shift upward, and profits tend to increase.

IV. Markets Demand and Supply

A. A market demand curve tells us how much of the product or service will be purchased by all consumers at a given price. Each consumer will be purchasing the quantities that he or she believes maximizes his or her net benefits in the above model. This implies that each person will be purchasing goods in the manner described by their individual demand curves.

B. As a consequence, Market Demand can be determined by varying price and adding up the amounts that consumers want to buy at each price.

The process is similar to that used to derive individual demand curves.

One picks a price, P',

Find the relevant quantities demanded by each consumer,

Add up those quantities. (Q' = sum of the individual consumer purchases at P'.)

Plot P' and Q'

Repeat with other prices to trace out the market demand curve.

[The geometry of this process was illustrated in class.]

C. **Market Demand** curves for ordinary private goods are thus "horizontal" sums of individual demand curves

Note that market demand can also be considered to be horizontal sums of individual MB curves, because all the points on individual demand curves are also points on the individual MB curves.

This allows us to use market demand curves as the MB curve for all consumers in the market.

This as, will be seen in later lectures, lets us understand how public policies like taxes affect consumer interests.

D. Similarly, **Market Supply** (for an industry with a fixed number of firms) can be derived by varying price and adding up the amounts that each firm in the industry is willing to sell at each price.

The process is similar to that used to derive individual supply curves.

One picks a price, P',

Find the relevant quantities supplied by each firm,

Add up those quantities. (Q' = sum of the individual firm production at P'.)

Plots P' and Q'

And repeats with other prices to trace out the market demand curve.

[The geometry of this process was illustrated in class., see your class notes for more details.]

E. **Market Supply** curves for ordinary private goods are "horizontal" sums of individual firm supply curves.

Note that market supply can also be regarded to be a horizontal sum of individual firm MC curves.

For a short run supply curve, this will be the sum of short-run firm MC curves.

For a medium run supply curve this will be the sum of medium-run MC curves.

For a long run supply curve this will be the sum long-run MC curves.

[For the long run, we will always use this "Ricardian" long run supply, to simplify our geometry and because the Ricardian model is often the best one for the markets of interest.]

V. Competitive Market Equilibrium and the Gains to Trade

A. In competitive markets, prices tend to move to "market clearing levels," that is to prices that set the total quantity supplied by all firms equal to the total amount demanded by consumers.

In competitive markets, this occurs where the supply and demand curve cross.

At any other price, there will either be surpluses (which tend to cause prices to fall) or shortages (which tend to cause prices to rise).

Note that this is, in principle, an entirely decentralized process requiring governments to do nothing more than enforce property rights and contracts.

The market clearing price is often denoted P* and the market clearing quantity as Q* where $Q^* = Q^D(P^*) = Q^S(P^*)$.

B. One can use our previous geometry to characterize the net benefits realized by consumers and firms at the market clearing prices

Finding Consumer Surplus:

The area under the demand curve from 0 to Q* is the total benefit realized by all consumers.

The area under P* from 0 to Q* represents the total cost (amount paid) for the Q* units.

Subtracting the TC from TB gives you consumer surplus.

[The geometry was illustrated in class and is also illustrated below.]

Finding Industry Profit:

The area under the supply curve from 0 to Q* is the total production cost for all firms in the industry.

The area under P* from 0 to Q* represents the total revenue (or benefit, the amount received) for the Q* units.

Subtracting the TC from TR gives you industry profit.

[The geometry was illustrated in class and is also illustrated below.]

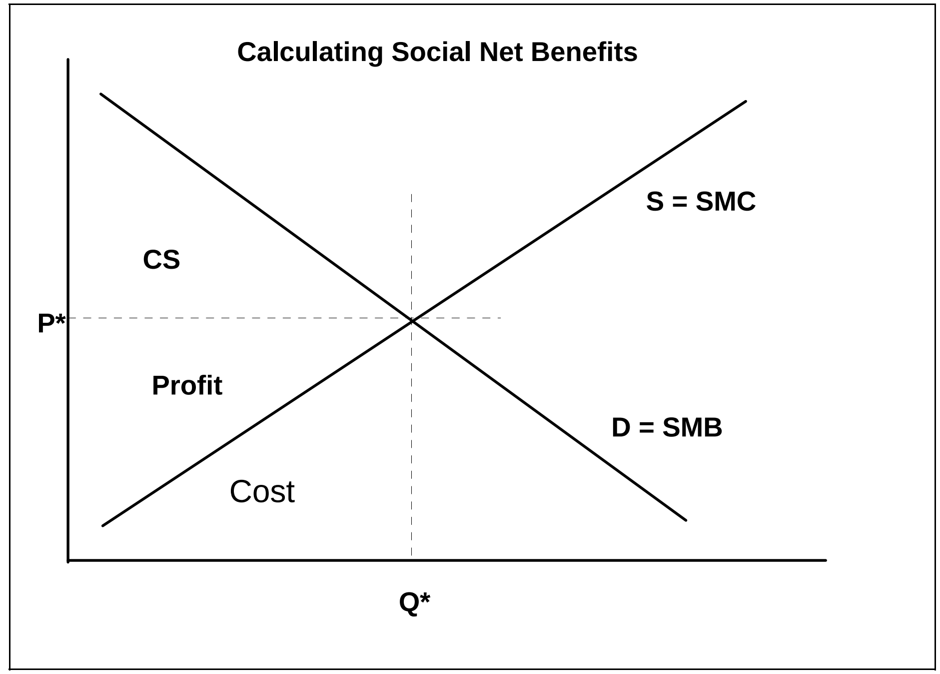

C. **A digression on the difference between Ricardian and Marshallian long run supply.**

In the short and medium run, the number of firms in the industry can be taken as fixed.

As a first approximation, this is also often true in the long run.

On this point, however, the Ricardian and Marshallian conceptions of long run supply are quite different and imply different things about the nature of the long run supply curve.

Long run supply in the "Marshallian" long run mostly reflects entry and exit of **identical** efficiently-sized firms into and out of the industry. This process takes place until profits fall to zero. As long as input prices do not change, **Marshallian long run supply curves are always horizontal** and reflect the industries average cost (the minimum average cost for each identical firm).

Long run supply in the Ricardian framework reflects “entry and exit” of firms with different cost functions. So, the low cost producers enter first, other things being equal, and the high cost firms last. **Ricardian long run supply thus is always upward sloping** and reflect the marginal production costs of firms in the industry. In general, relatively efficient firms (those with better technologies, locations, or natural resources) will continue to earn profits greater than zero in the long run, because they have lower costs than "marginal" firms.

Note the derivation of long run supply in the Ricardian sense, is exactly like that used above for both short and medium run analysis. The number of potential suppliers is essentially fixed, although the individual firms will alter their production as prices vary.

It bears noting that the Marshallian model of long run supply works well for some industries (say auto repair shops), and the Ricardian model works well for others (oil, minerals, sports, farming, etc.).

We will use the Ricardian model of long run supply in this class, but it should be kept in mind that the Marshallian entry and exit may also be important in some cases.

VI. Summary: Markets, Net Benefits, and "Social Welfare"

A. Our derivations of demand and supply make it clear that:

Every market **demand curve is the sum of the marginal benefit** curves of the individual consumers, because each consumer's demand curve is essentially his or her MB curve.

Every short (or middle) run **market supply curve is the sum of the short (or middle) run marginal cost** curves of the individual firms in the market. Short and middle run supply curves can thus be used to represent an industry's short or middle run marginal cost of production.

Every Ricardian long run supply curve is the sum of the long run marginal cost curves of potential suppliers of the goods. Consequently, a Ricardian long run supply curve can be used as the industry's long run marginal cost curve.

B. In cases in which there are **no externalities** and little monopoly power, firms and consumers are the only ones affected by market transactions.

In such cases, the demand curve can be used as the Social Marginal Benefit Curve and the Supply curve as the Social Marginal Cost curve.

This allows **social net benefits** to be calculated for individual markets and for policies that affect markets--as demonstrated below and in subsequent lectures.

In the absence of externalities or market concentration, **markets tend to produce social net benefit maximizing** levels of goods and services!

From our earlier work, we know that given SMB and SMC curves, social net benefits are maximized at the quantity where SMB crosses SMC (assuming that it does so only once, and that SMB crosses SMC from above).

Because D is approximately SMB (at least for all consumers) and S is approximately SMC (at least for firms in the industry), the quantity where demand and supply cross is exactly the output required to maximize social net benefits.

Consequently, the "market clearing" price, P*, causes markets to produce the **social net benefit maximizing** level of output (in cases where there are not externalities, e.g. relevant costs or benefits).

The market clearing price, P*, causes social marginal benefit (the demand curve) equal to social marginal cost (the supply curve). Both are equal to price in this case!

Social net benefits turn out to be the sum of profits and consumer surplus in this case. (Use diagrams like that above to prove this.)

Many economists use the **maximize social net benefit norm.**, and the above calculations allow them to rank states of the world and public policies.

Note that competitive market outcomes are very desirable under this norm, because they tend to maximize social net benefits.

This is one very widely used **normative** argument favoring markets as a method of social organization.

[In cases where there are additional costs or benefits because of externalities, those additional marginal benefits or marginal costs will need to be taken account of.]

C. In cases where **external costs** exist, however, market outcomes will (often) fail to maximize social net benefits, because those costs will not be fully accounted for in the calculations of economic agents.

We will take this case up later in the course.

When there are external benefits or costs, demand and/or supply will not include all the relevant marginal benefits and/or all marginal costs.

The existence of externality problems provides a **normative** basis for government policy (if one wants to maximize social net benefits).

In cases where significant external costs exist at the margin (at Q*), markets will tend to **over produce** the output of interest relative to that which maximizes social net benefits.

In cases where significant external benefits exist at the margin (at Q*), markets will **under produce** the service of interest relative to that which maximizes social net benefits.

Governments might adopt policies to discourage production in the first case (perhaps with taxes) or encourage it (perhaps with subsidies) in the second case.

(Tools for analyzing externality and public goods problems will be developed after the midterm.)