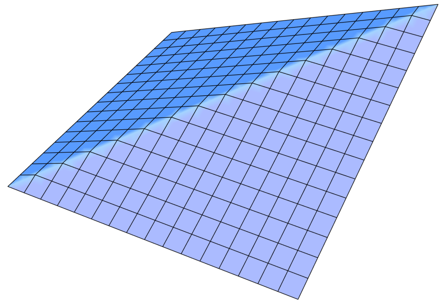

ECON 201: Microeconomic Theory I

Prof. Mohamad Ghaziaskar • Fall 2015 • University of Waterloo
TR 8:30–9:50AM in AL 211

Transcribed, LaTeXed, and edited by Rob Zimmerman

Note: These are personal lecture notes, not official course materials. They may contain errors (which by default you should attribute to me, Rob Zimmerman, rather than the course instructor). The permanent home of these — and all of my other lecture notes — is `https://rob-zimmerman.github.io/coursenotes`.

Contents

Lecture 1 — Tuesday, September 15, 2015

1. Market Models and Consumer Choice

Market Models

Microeconomic theory begins by building models, which are deliberate simplifications of reality. Consider a college town with two types of apartments. Some apartments are close to the university and form the inner ring, while the rest are farther away and form the outer ring. The purpose of the model is to isolate how scarcity and willingness-to-pay determine the rental price of the desirable apartments.

An **endogenous** variable is determined within the model, while an **exogenous** variable is taken as given from outside the model. In the apartment market model, the price of inner ring apartments is endogenous, whereas the price of outer ring apartments is exogenous.

The model also assumes that apartments are identical in every respect except location. That assumption is not literally true in reality, but it allows us to focus on the single force that matters here: people are willing to pay more to live close to campus.

Two organizing ideas are central. The first is the **optimization principle**: people choose the consumption pattern that best serves their own interest, subject to the opportunities available to them. The second is the **equilibrium principle**: prices adjust until the quantity people want equals the quantity available.

Intuition: Optimization explains individual behavior, while equilibrium explains how all individual choices fit together in the market. In this model, renters choose the apartment location they prefer, and market rents move until that choice pattern is consistent with the fixed number of close apartments.

Demand and Supply Curves

To construct demand, line up potential renters from the highest willingness-to-pay for an inner ring apartment to the lowest. The key number for each renter is the **reservation price**.

Definition 1.1 The **reservation price** of a renter is the maximum amount that renter is willing to pay for an inner ring apartment. Equivalently, it is the price at which the renter is indifferent between renting the apartment and not renting it.

Suppose the highest reservation prices are \$500, \$490, \$480, and so on. Then the quantity demanded at any price is the number of renters whose reservation prices are at least that price. Writing out the cases explicitly gives

$$Q_D(p) = \begin{cases} 0, & p > 500, \\ 1, & 490 < p \leq 500, \\ 2, & 480 < p \leq 490, \\ \vdots & \vdots \end{cases}$$

because a renter demands an apartment whenever the market price does not exceed that renter's reservation price. The resulting step demand curve appears in Figure 1.

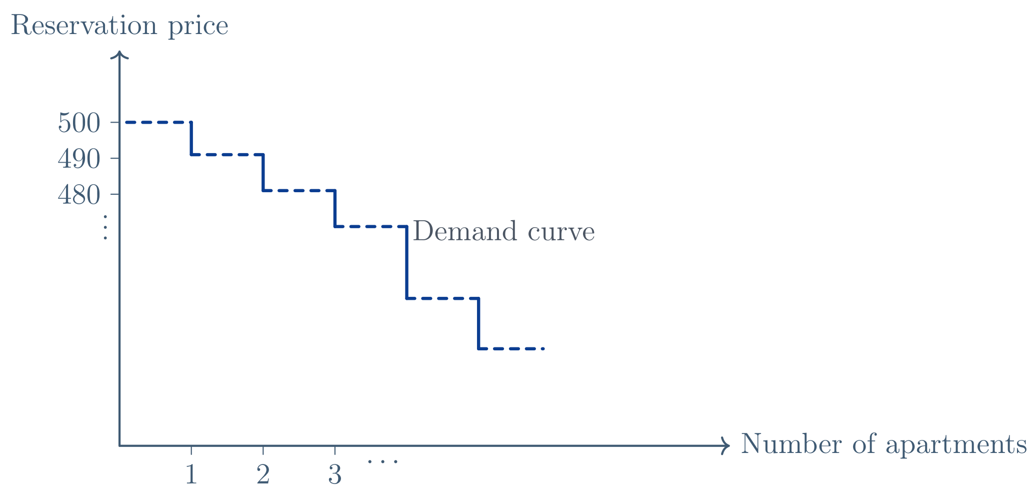

FIGURE 1

Definition 1.2 The **demand curve** records the relationship between the market price p and the quantity Q_D demanded at that price.

When there are many renters, reservation prices differ across people so finely that the step pattern is well approximated by the smooth downward-sloping curve in Figure 2. The curve slopes downward because a lower price makes inner ring apartments attractive to additional renters whose reservation prices were previously too low.

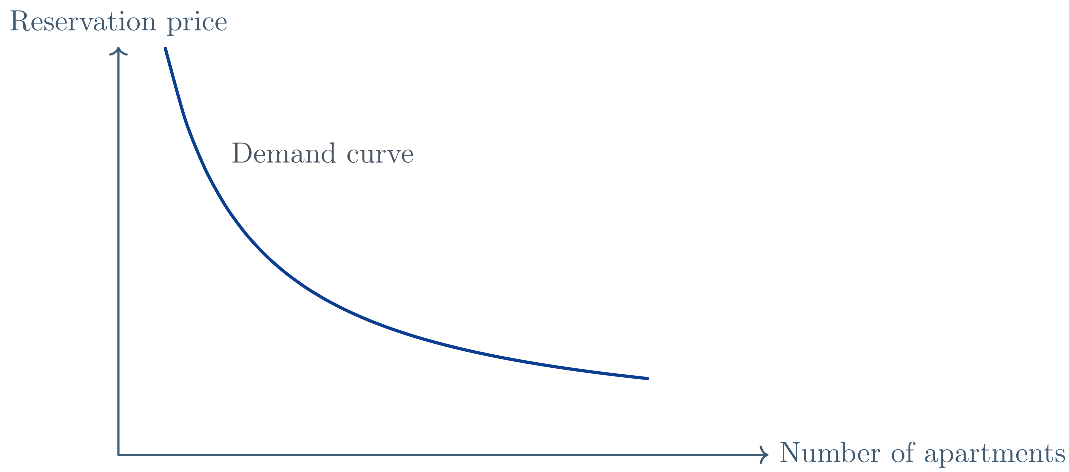

Figure 2

The supply side is modeled as a competitive market with many independent landlords, each trying to rent out an apartment at the highest possible price. Because everyone observes market conditions, all inner ring apartments must rent for the same price in equilibrium. The logic is an arbitrage argument. If one landlord charges a high price p_h and another charges a low price p_ℓ, then any renter facing p_h can approach the cheaper landlord and offer some intermediate price p satisfying

$$p_\ell < p < p_h.$$

That offer makes the renter strictly better off than paying p_h, and it makes the cheaper landlord strictly better off than accepting p_ℓ. Therefore an outcome with two persistent prices cannot survive in a competitive market.

To draw the supply curve, fix a price and ask how many close apartments landlords will supply at that price. The answer depends on the time horizon. In the long run, new apartments can be built. In the short run, however, the stock of inner ring apartments is fixed, so quantity supplied does not depend on price, giving the vertical curve in Figure 3.

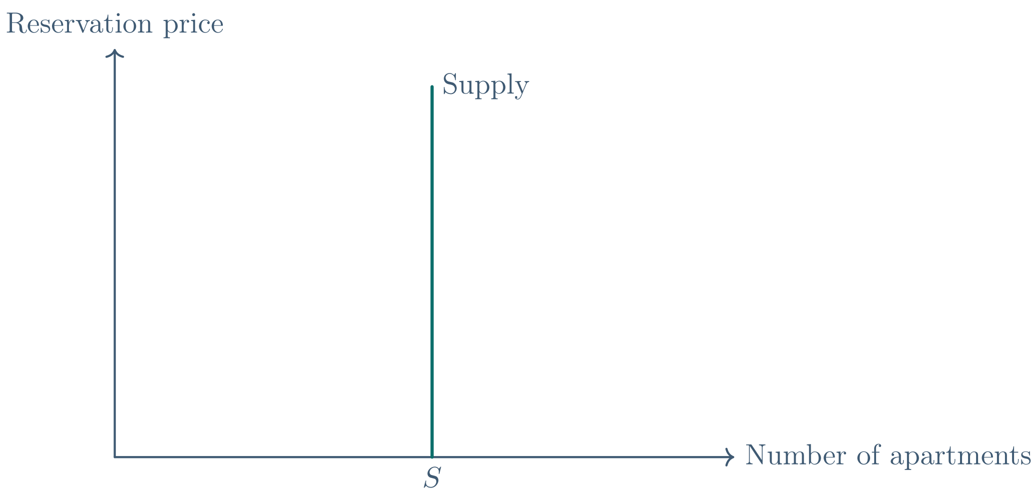

Figure 3

Remark 1.3 The vertical supply curve does not mean price is unimportant. It means that in the short run price adjusts to clear the market, while quantity is pinned down by the number of available close apartments.

Market equilibrium occurs where quantity demanded equals quantity supplied. If the market is at a price p^* such that

$$Q_D(p^*) = Q_S,$$

then no force pushes the price either upward or downward.

Definition 1.4 A **competitive equilibrium** is a price p^* at which quantity demanded equals quantity supplied. In the apartment model, p^* is the rent at which the number of people who want inner ring apartments is exactly the number of such apartments available.

The adjustment process in Figure 4 is straightforward. If $p < p^*$, then demand exceeds supply, lines form, and landlords have an incentive to raise the rent. If $p > p^*$, then supply exceeds demand, vacancies appear, and landlords have an incentive to lower the rent. Only at p^* is there no reason for further price movement.

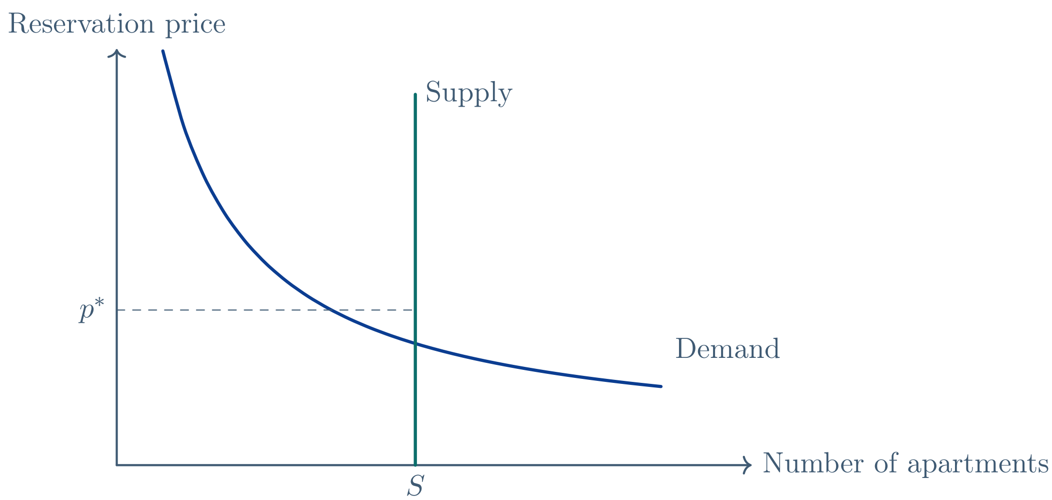

Figure 4

The equilibrium price also determines who gets the scarce close apartments. Renters with reservation prices above p^* are willing to pay enough to obtain an inner ring apartment. Renters with reservation prices below p^* choose the outer ring instead. Renters with reservation price exactly equal to p^* are indifferent between the two locations. Thus the competitive allocation assigns close apartments by willingness-to-pay.

Comparative statics studies how the endogenous variables of a model respond when exogenous variables change. In the apartment market model, the exogenous variables are the price of distant apartments, the number of close apartments, and the incomes of potential renters.

It is useful to write the inverse demand curve as $p = D(q)$, where $D'(q) < 0$, and to write short-run supply as the fixed quantity $q = \bar{q}$. Equilibrium therefore satisfies

$$p^* = D(\bar{q}).$$

This one equation makes the comparative statics results immediate.

First suppose the number of close apartments rises from $\bar{q}$ to $\bar{q}'$, where $\bar{q}' > \bar{q}$. Then the new equilibrium price is

$$p' = D(\bar{q}').$$

Because demand slopes downward, $D(\bar{q}') < D(\bar{q})$, so $p' < p^*$. The market price falls and the quantity traded rises from $\bar{q}$ to $\bar{q}'$, as shown in Figure 5.

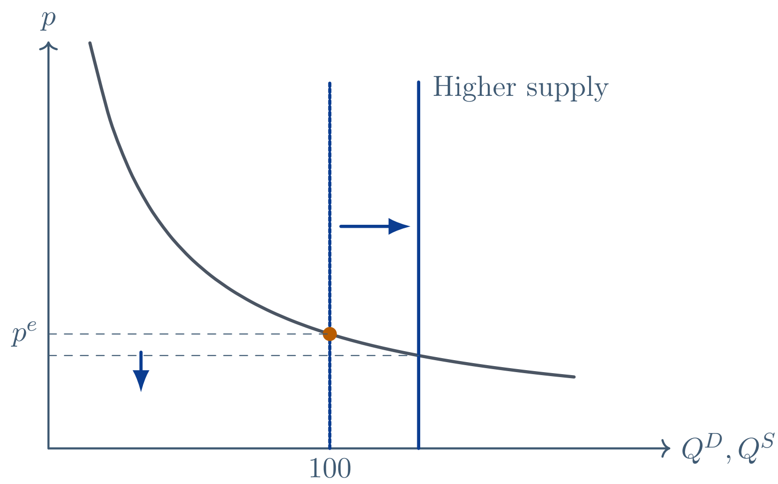

FIGURE 5

Now suppose potential renters' incomes rise and, as a result, their willingness-to-pay for close apartments rises. The demand curve then shifts upward. If the new inverse demand curve is $\widetilde{D}(q)$ with $\widetilde{D}(q) > D(q)$ for each relevant q, then the new equilibrium price is

$$\widetilde{p} = \widetilde{D}(\bar{q}).$$

Since $\widetilde{D}(\bar{q}) > D(\bar{q}) = p^*$, the equilibrium price rises. Because the short-run supply curve is vertical, the quantity traded stays fixed at $\bar{q}$, as shown in Figure 6.

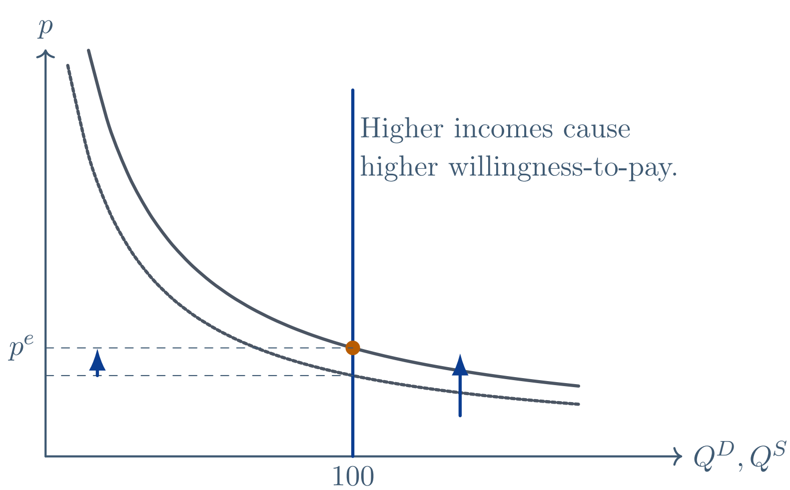

Figure 6

Lecture 2 — Thursday, September 17, 2015

Budget Constraints

A consumption bundle with nonnegative quantities $x_1, x_2, \ldots, x_n$ of n commodities is written as the vector $(x_1, x_2, \ldots, x_n)$. If the corresponding commodity prices $p_1, p_2, \ldots, p_n$ are positive and the consumer's disposable income is $m \geq 0$, then the bundle is affordable exactly when

$$p_1 x_1 + \cdots + p_n x_n \leq m.$$

Consumer theory asks which affordable bundle is best. The first step is to understand the affordability side of the problem.

Definition 2.1 The **budget set** is the collection of all affordable nonnegative bundles. The **budget line** is its expenditure frontier, namely the bundles that cost exactly income:

$$p_1 x_1 + \cdots + p_n x_n = m.$$

For most graphical work we specialize to two goods. A bundle is then (x_1, x_2), and the budget constraint becomes

$$p_1 x_1 + p_2 x_2 \leq m.$$

The left-hand side is total expenditure: $p_1 x_1$ is spending on good 1, $p_2 x_2$ is spending on good 2, and their sum cannot exceed income.

The model with two goods is more general than it first appears. Good 1 can be a specific commodity, while good 2 stands for "everything else." If we measure the second good directly in dollars spent on all other goods, then $p_2 = 1$ and the budget constraint becomes

$$p_1 x_1 + x_2 \leq m.$$

This is a convenient normalization, but the economic ideas are unchanged.

When the consumer spends all income, the budget line satisfies

$$p_1 x_1 + p_2 x_2 = m.$$

Solving for x_2 gives

$$x_2 = \frac{m}{p_2} - \frac{p_1}{p_2} x_1.$$

The vertical intercept is m/p_2, the horizontal intercept is m/p_1, and the slope is $-p_1/p_2$. Figure 7 collects this geometry.

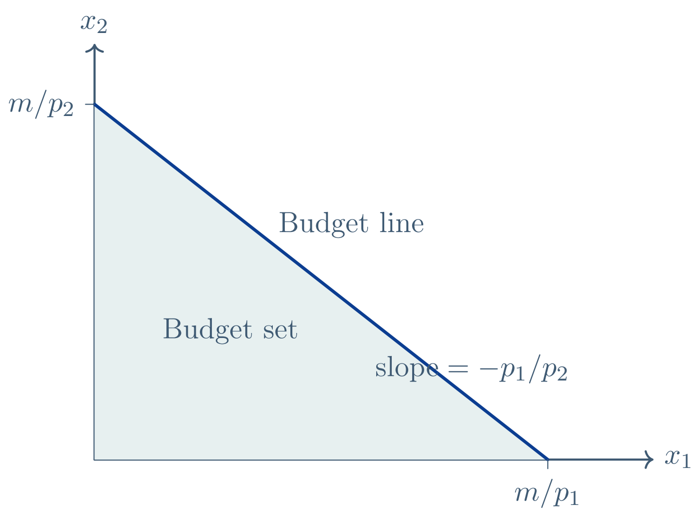

FIGURE 7

The slope has a direct opportunity cost interpretation. If the consumer increases good 1 by a small amount Δx_1 while remaining on the budget line, then expenditure must still balance:

$$p_1 \Delta x_1 + p_2 \Delta x_2 = 0.$$

Therefore

$$\Delta x_2 = -\frac{p_1}{p_2} \Delta x_1.$$

So one more unit of good 1 requires giving up p_1/p_2 units of good 2.

Changes in income and prices move the budget line in different ways.

Fact 2.2 If income rises while prices stay fixed, both intercepts increase proportionally and the budget line shifts outward in parallel.

Fact 2.3 If p_1 rises while m and p_2 stay fixed, the horizontal intercept falls from m/p_1 to m/p_1', the vertical intercept m/p_2 stays fixed, and the budget line becomes steeper.

The income shift and price rotation are compared in Figure 8.

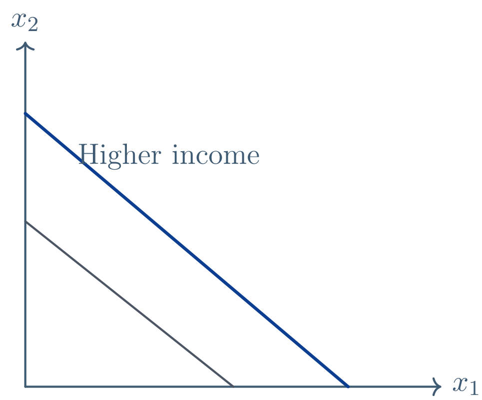

FIGURE 8

If all prices double while income stays fixed, then each intercept is cut in half and the budget line shifts inward in parallel. If both prices and income double, then the budget set is unchanged because all nominal magnitudes have been scaled by the same factor.

It is also useful to consider cases in which the budget line is not a single straight line. If the consumer is rationed and may buy at most $\bar{x}_1$ units of good 1, then the feasible set is truncated at $x_1 = \bar{x}_1$. If instead purchases beyond $\bar{x}_1$ face a per-unit tax t, then the budget line bends inward at $\bar{x}_1$: before the kink the slope is $-p_1/p_2$, and after the kink the slope is $-(p_1 + t)/p_2$. Figure 9 compares the two cases.

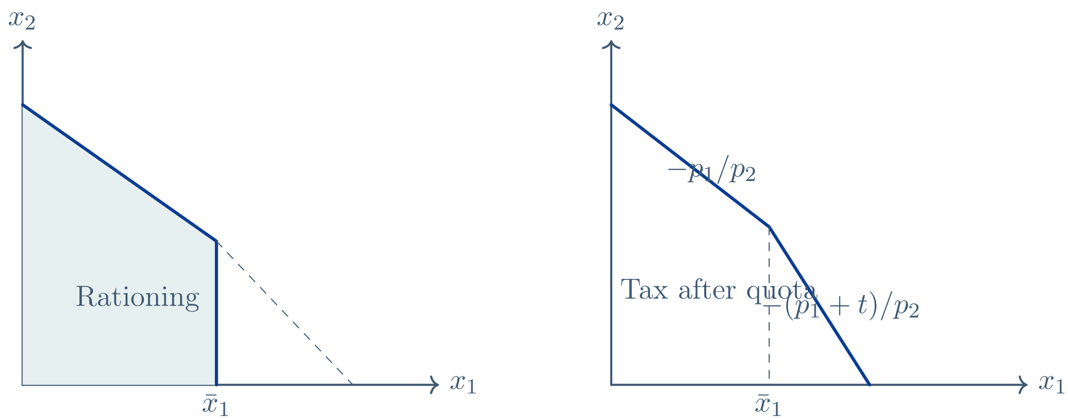

FIGURE 9

Food stamps are coupons that may legally be spent only on food. Suppose income is $m = \$100$, the price of food is $p_F = 1$, and the price of all other goods is $p_G = 1$. Before food stamps, the budget line is

$$F + G = 100.$$

If the household receives \$40 of food stamps that cannot be resold, then the budget set enlarges. The household may spend all cash on other goods and use any desired amount of the stamps for food, allowing unused stamps to expire. Thus, for $0 \le F \le 40$, the highest attainable G remains 100. For $F > 40$, additional food must be purchased with cash, so the upper boundary becomes

$$G = 140 - F.$$

The budget line after adding the stamps is therefore kinked:

$$G = \begin{cases} 100, & 0 \le F \le 40, \\ 140 - F, & 40 \le F \le 140. \end{cases}$$

In the left-hand panel of Figure 10, nontradable food stamps create a horizontal segment. In the right-hand panel, if stamps can be sold at \$0.50 each, the budget set expands again.

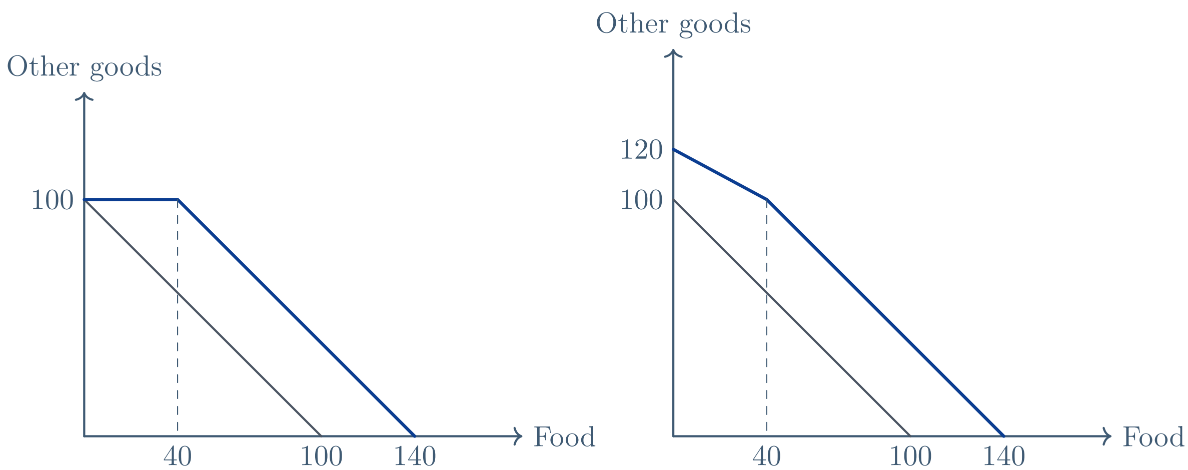

FIGURE 10

If the stamps can be sold informally for \$0.50 each, then selling all 40 stamps yields \$20 of extra cash. More generally, the upper frontier becomes

$$G = \begin{cases} 120 - \frac{1}{2}F, & 0 \le F \le 40, \\ 140 - F, & 40 \le F \le 140. \end{cases}$$

For the first segment, every dollar of stamps used for food forgoes fifty cents of resale proceeds; after all stamps have been used, additional food costs one dollar per unit. The key lesson is that in-kind transfers affect choice differently from unrestricted cash transfers, precisely because they change the shape of the feasible set.

Lecture 3 — Tuesday, September 22, 2015

Preferences and Indifference Curves

The central question in this chapter is what economists mean by "the best" bundle. Preferences compare entire consumption bundles, not isolated goods. If the consumer would choose (x_1, x_2) when (y_1, y_2) is also available, then (x_1, x_2) is preferred to (y_1, y_2).

Definition 3.1 We write $(x_1, x_2) \succ (y_1, y_2)$ when the first bundle is **strictly preferred**, $(x_1, x_2) \sim (y_1, y_2)$ when the consumer is **indifferent**, and $(x_1, x_2) \succeq (y_1, y_2)$ when the first bundle is **weakly preferred**, meaning preferred or indifferent.

If $(x_1, x_2) \succeq (y_1, y_2)$ and $(y_1, y_2) \succeq (x_1, x_2)$, then the consumer must be indifferent between them.

Three standard axioms are the following.

Definition 3.2 Preferences are **complete** if any two bundles can be compared; **reflexive** if every bundle is at least as good as itself; and **transitive** if $X \succeq Y$ and $Y \succeq Z$ imply $X \succeq Z$.

Transitivity is what allows choices to be coherent. Without it, a consumer could rank X above Y, Y above Z, and yet Z above X, making purposeful choice impossible.

Fix a bundle $x = (x_1, x_2)$. The **weakly preferred set** is the set of bundles the consumer likes at least as much as x. Under continuous, locally nonsatiated preferences, its boundary is the indifference curve through x, as shown in Figure 11.

Definition 3.3 The **indifference curve** through x is the set of all bundles y such that $y \sim x$.

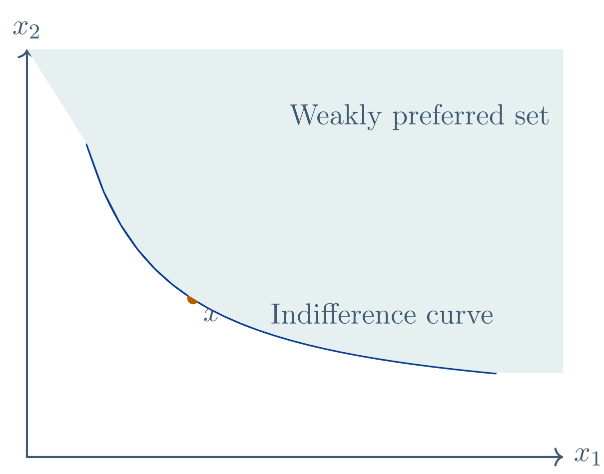

FIGURE 11

Indifference curves cannot cross.

Proposition 3.4 If preferences are transitive, two distinct indifference curves cannot cross.

Proof. Suppose two indifference curves crossed at Z, and suppose X lies on the first curve while Y lies on the second. Then $X \sim Z$ and $Z \sim Y$. By transitivity, $X \sim Y$, so X and Y must lie on the same indifference curve. That contradicts the picture of two distinct curves crossing. □

Several important special cases are collected in Figure 12.

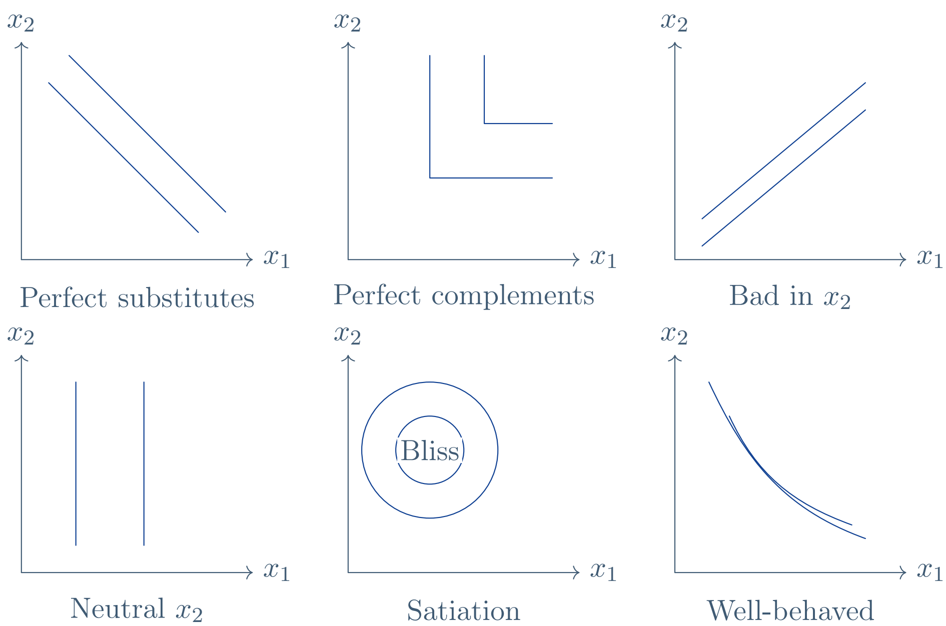

FIGURE 12

For perfect substitutes, the consumer is willing to trade the two goods at a constant rate, so indifference curves are parallel straight lines. A standard example is Pepsi and Coke, where bundles such as $(10, 10)$, $(12, 8)$, and $(3, 17)$ are all indifferent when only the total amount matters.

For perfect complements, goods are consumed together in fixed proportions. Left and right shoes are the classic example, and the indifference curves are L-shaped. More of one shoe without more of the other does not improve well-being.

If one commodity is a bad, then more of it lowers utility, so indifference curves slope upward. If one commodity is neutral, then changing its quantity leaves the consumer indifferent, so the curves are vertical or horizontal. Finally, with satiation there is a **bliss point**, and bundles closer to that point are preferred.

Usually, two additional assumptions are imposed.

Definition 3.5 Preferences are **strictly monotone** if $x_i \geq y_i$ for every good i, with a strict inequality for at least one good, implies $x \succ y$. They are **convex** if $x \succeq y$ implies

$$\lambda x + (1-\lambda)y \succeq y$$

for every $\lambda \in [0,1]$.

Strict monotonicity rules out satiation and bads. Convexity says that averages are no worse than extremes: every bundle between x and y is weakly preferred to y whenever x is. Geometrically, upper contour sets are convex; smooth indifference curves therefore bow weakly toward the northeast, with straight segments still allowed.

The slope of an indifference curve is the **marginal rate of substitution**.

Definition 3.6 At a smooth point, we use the **marginal rate of substitution** for the signed slope of an indifference curve:

$$\mathrm{MRS} = \frac{\mathrm{d}x_2}{\mathrm{d}x_1}.$$

Under strictly monotone preferences the indifference curve slopes downward, so the marginal rate of substitution is negative. Its magnitude, $-\mathrm{MRS}$, is the consumer's marginal willingness to pay for an extra unit of good 1, measured in units of good 2.

Note: The marginal rate of substitution behaves differently across the standard preference classes. For perfect substitutes it is constant. For perfect complements it is zero on the horizontal arm and undefined on the vertical arm, where its magnitude is informally described as infinite. For bads it is positive. For smooth, strictly convex preferences it becomes flatter as x_1 rises, which is the usual diminishing marginal rate of substitution property.

Lecture 4 — Thursday, September 24, 2015

Utility Representation

Utility is a numerical device used to represent preferences, but it is not meant to measure happiness in a literal cardinal sense.

Definition 4.1 A **utility function** assigns a number $u(x_1, x_2)$ to each consumption bundle in such a way that bundles ranked higher receive larger numbers:

$$(x_1, x_2) \succ (y_1, y_2) \qquad \text{if and only if} \qquad u(x_1, x_2) > u(y_1, y_2).$$

The ordering matters, not the magnitude. If $u(x) = 6$ and $u(y) = 2$, then x is preferred to y, but x is not "three times as good" as y.

Fact 4.2 On the standard consumption set $\mathbb{R}^n_+$, every complete, transitive, and continuous preference relation has a continuous utility representation.

Continuity means that small changes in a bundle do not cause abrupt jumps in ranking.

Bundles on the same indifference curve must receive the same utility level. Under monotone preferences, indifference curves lying northeast of one another correspond to higher utility levels. Thus a utility function and an indifference map are two ways of encoding the same preference information, as Figure 13 illustrates.

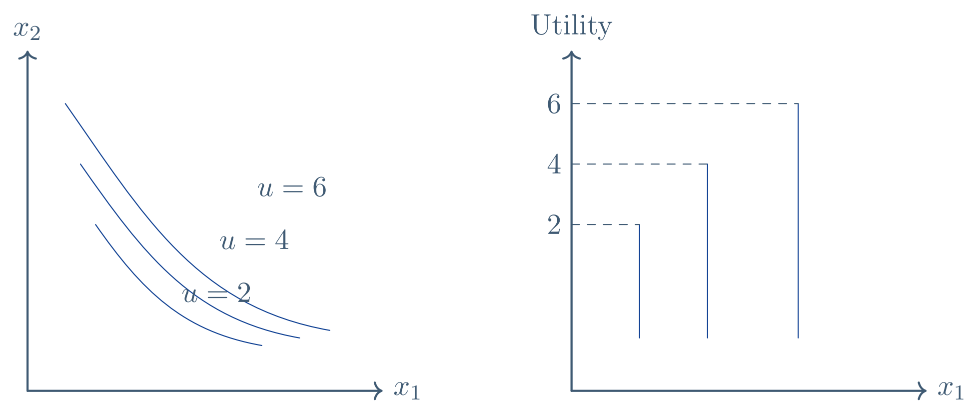

FIGURE 13

There is no unique utility representation of a preference relation. If u represents preferences and f is strictly increasing, then $v = f(u)$ represents the same preferences.

Proposition 4.3 If $u(x)$ is a utility function and f is strictly increasing, then $v(x) = f(u(x))$ represents exactly the same preference ordering as u.

Proof. Because f is strictly increasing,

$$u(x) > u(y) \qquad \text{if and only if} \qquad f(u(x)) > f(u(y)).$$

So x is ranked above y by u exactly when it is ranked above y by v. □

A standard example takes $u(x_1, x_2) = x_1 x_2$ and $v = u^2 = x_1^2 x_2^2$. Both functions rank $(2, 3)$ above $(4, 1) \sim (2, 2)$, so they represent the same preferences.

Standard utility functions for the preference patterns from Chapter 3 are

$$u(x_1, x_2) = ax_1 + bx_2 \quad \text{for perfect substitutes,}$$

$$u(x_1, x_2) = \min(ax_1, bx_2) \quad \text{for perfect complements,}$$

$$u(x_1, x_2) = v(x_1) + x_2 \quad \text{for quasilinear preferences,}$$

and

$$u(x_1, x_2) = x_1^c x_2^d \quad \text{for Cobb–Douglas preferences, where } c, d > 0.$$

The shapes of their indifference curves are compared in Figure 14.

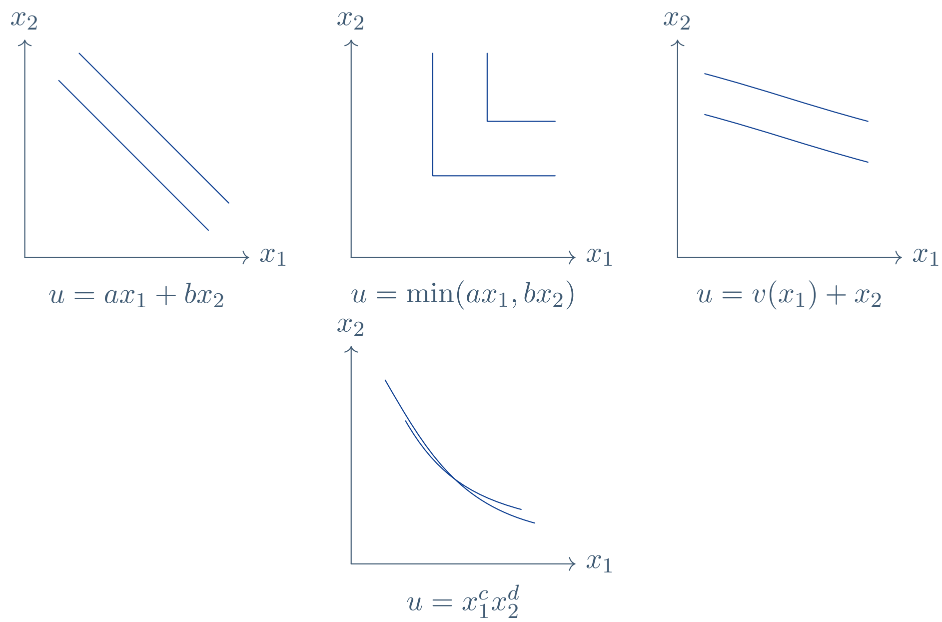

FIGURE 14

The diagram shows utility representations for the standard preference classes used below.

For perfect substitutes, any strictly increasing transformation of $x_1 + x_2$ also works. For perfect complements, the minimum operator is natural because extra units of only one good do not raise the number of usable pairs. Quasilinear preferences are linear in good 2, so indifference curves are vertical shifts of one another. Cobb–Douglas preferences provide the standard smooth, monotone, strictly convex benchmark.

The marginal utility of good i is the partial derivative of utility with respect to x_i:

$$MU_i = \frac{\partial U}{\partial x_i}.$$

If an indifference curve is given by $U(x_1, x_2) = k$, then total differentiation yields

$$\frac{\partial U}{\partial x_1}\,\mathrm{d}x_1 + \frac{\partial U}{\partial x_2}\,\mathrm{d}x_2 = 0.$$

Solving for the slope gives

$$\frac{\mathrm{d}x_2}{\mathrm{d}x_1} = -\frac{\partial U/\partial x_1}{\partial U/\partial x_2} = -\frac{MU_1}{MU_2}.$$

Whenever the utility function is differentiable and $MU_2 \neq 0$, the implicit-function theorem gives a smooth local indifference curve, and our signed slope convention yields

$$\text{MRS} = -\frac{MU_1}{MU_2}.$$

This formula explains why multiplying a utility function by a positive constant does not change behavior: both marginal utilities scale by the same factor, so their ratio, and therefore the marginal rate of substitution, stays unchanged.

Lecture 5 — Tuesday, September 29, 2015

2. Demand, Comparative Statics, and Market Demand

Constrained Choice

The main behavioral postulate is that a consumer chooses the most preferred bundle from the set of affordable bundles. The choice set is the budget set, and the chosen bundle is the most preferred affordable bundle, as depicted in Figure 15.

Definition 5.1 When the utility-maximizing affordable bundle is unique, the **ordinary demand** at prices (p_1, p_2) and income m is denoted

$$x_1^*(p_1, p_2, m) \qquad \text{and} \qquad x_2^*(p_1, p_2, m).$$

If several bundles maximize utility, ordinary demand is instead a set-valued demand correspondence.

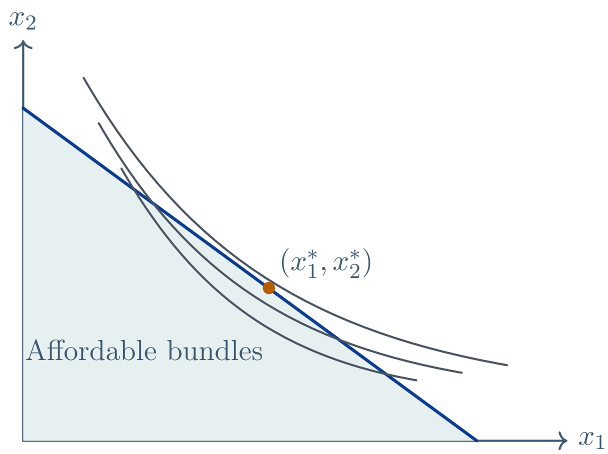

Figure 15

If both demanded quantities are strictly positive, the solution is interior. With locally nonsatiated preferences, the consumer exhausts the budget; when the indifference curve is also smooth at the optimum, it is tangent to the budget line.

Proposition 5.2 If preferences are locally nonsatiated, (x_1^*, x_2^*) is an interior optimum, and the indifference curve is smooth there, then

$$p_1 x_1^* + p_2 x_2^* = m$$

and

$$\text{MRS}(x_1^*, x_2^*) = -\frac{p_1}{p_2}.$$

Take

$$U(x_1, x_2) = x_1^a x_2^b,$$

with $a, b > 0$. Then

$$MU_1 = a x_1^{a-1} x_2^b \qquad \text{and} \qquad MU_2 = b x_1^a x_2^{b-1},$$

so

$$\text{MRS} = -\frac{MU_1}{MU_2} = -\frac{a x_1^{a-1} x_2^b}{b x_1^a x_2^{b-1}} = -\frac{a x_2}{b x_1}.$$

At the optimum,

$$-\frac{a x_2^*}{b x_1^*} = -\frac{p_1}{p_2},$$

which implies

$$x_2^* = \frac{bp_1}{ap_2} x_1^*.$$

Substituting this relation into the binding budget condition gives

$$p_1 x_1^* + p_2 \left(\frac{bp_1}{ap_2} x_1^* \right) = m,$$

so

$$p_1 x_1^* \left(1 + \frac{b}{a} \right) = m.$$

Since $1 + b/a = (a+b)/a$, we obtain

$$x_1^* = \frac{am}{(a+b)p_1}.$$

Plugging this back in yields

$$x_2^* = \frac{bm}{(a+b)p_2}.$$

Example 5.3 For Cobb–Douglas preferences, the consumer spends the fraction $a/(a+b)$ of income on good 1 and the fraction $b/(a+b)$ on good 2.

Tangency is not the only possible optimum. If one of the demanded quantities is zero, the solution is a **corner solution**; Figure 16 gives two standard examples.

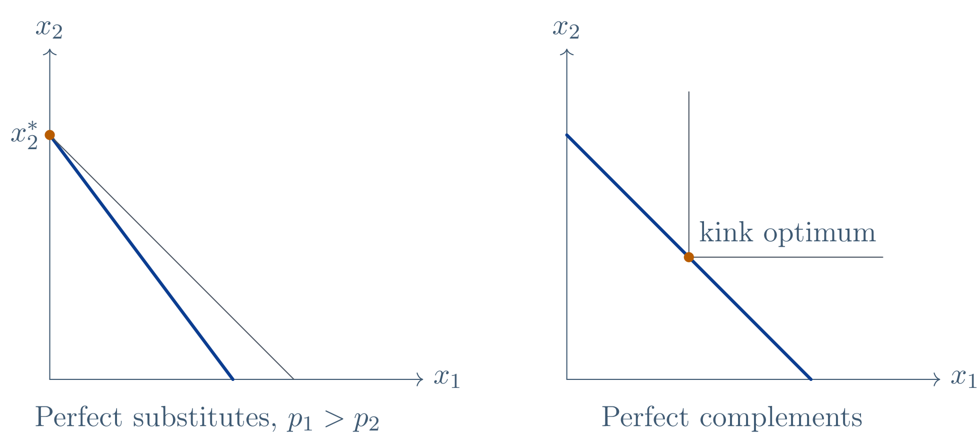

FIGURE 16

For perfect substitutes with $U(x_1, x_2) = x_1 + x_2$, the marginal rate of substitution is constantly -1. If $p_1 < p_2$, good 1 delivers one unit of utility at lower cost, so the consumer spends all income on good 1:

$$(x_1^*, x_2^*) = \left(\frac{m}{p_1}, 0\right).$$

If $p_1 > p_2$, the consumer buys only good 2:

$$(x_1^*, x_2^*) = \left(0, \frac{m}{p_2}\right).$$

If $p_1 = p_2$, every bundle on the budget line is optimal.

For perfect complements with $U(x_1, x_2) = \min(ax_1, x_2)$, the optimum lies at the kink $x_2^* = ax_1^*$. Combining that with the binding budget equation gives

$$p_1 x_1^* + p_2(ax_1^*) = m,$$

so

$$x_1^* = \frac{m}{p_1 + ap_2} \qquad \text{and} \qquad x_2^* = \frac{am}{p_1 + ap_2}.$$

Warning The tangency condition is not sufficient when preferences are nonconvex. A tangency can exist even when the truly preferred affordable bundle lies elsewhere on the boundary.

Lecture 6 — Thursday, October 01, 2015

Price and Income Effects

Comparative statics studies how the ordinary demand functions $x_1^*(p_1, p_2, m)$ and $x_2^*(p_1, p_2, m)$ change when prices or income change. It is natural to begin with changes in the consumer's own price.

Holding p_2 and m fixed, let p_1 vary. Each value of p_1 gives a different optimal bundle. The locus of these utility-maximizing bundles is the $\boldsymbol{p_1}$**-price offer curve**. Plotting the x_1-coordinate of that locus against p_1 gives the ordinary demand curve for good 1, as illustrated in Figure 17.

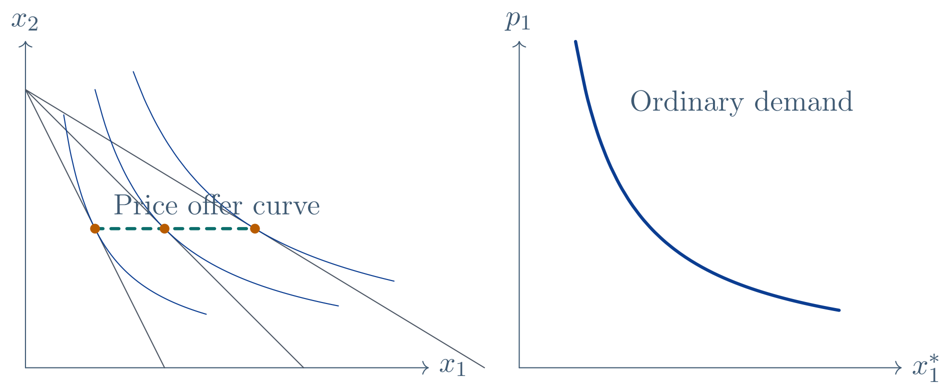

Figure 17

For Cobb–Douglas preferences,

$$x_1^*(p_1, p_2, m) = \frac{a}{a+b}\frac{m}{p_1} \qquad \text{and} \qquad x_2^*(p_1, p_2, m) = \frac{b}{a+b}\frac{m}{p_2}.$$

As p_1 varies, x_2^* stays fixed, so the price offer curve is horizontal and the demand curve is a **rectangular hyperbola**:

$$x_1^* = \frac{am}{(a+b)p_1}.$$

For the normalized perfect-complements utility $U(x_1, x_2) = \min(x_1, x_2)$,

$$x_1^* = x_2^* = \frac{m}{p_1 + p_2}.$$

An increase in p_1 reduces the demand for both goods. For perfect substitutes with $U = x_1 + x_2$, demand is piecewise:

$$x_1^* = \begin{cases} m/p_1, & p_1 < p_2, \\ 0, & p_1 > p_2, \end{cases} \qquad \text{and} \qquad x_2^* = \begin{cases} 0, & p_1 < p_2, \\ m/p_2, & p_1 > p_2. \end{cases}$$

At $p_1 = p_2$, any bundle on the budget line is optimal.

Instead of asking what quantity is demanded at a given price, we can ask which price would support a given demanded quantity. That relation is the **inverse demand function**.

Definition 6.1 An **inverse demand function** expresses price as a function of quantity demanded.

For the Cobb–Douglas case,

$$x_1^* = \frac{am}{(a+b)p_1}$$

implies

$$p_1 = \frac{am}{(a+b)x_1^*}.$$

For the normalized perfect-complements utility $U(x_1, x_2) = \min(x_1, x_2)$,

$$x_1^* = \frac{m}{p_1 + p_2}$$

rearranges to

$$p_1 = \frac{m}{x_1^*} - p_2.$$

Holding prices fixed and varying income traces out the **income offer curve**. Plotting demanded quantity against income gives an **Engel curve**; Figure 18 shows both constructions.

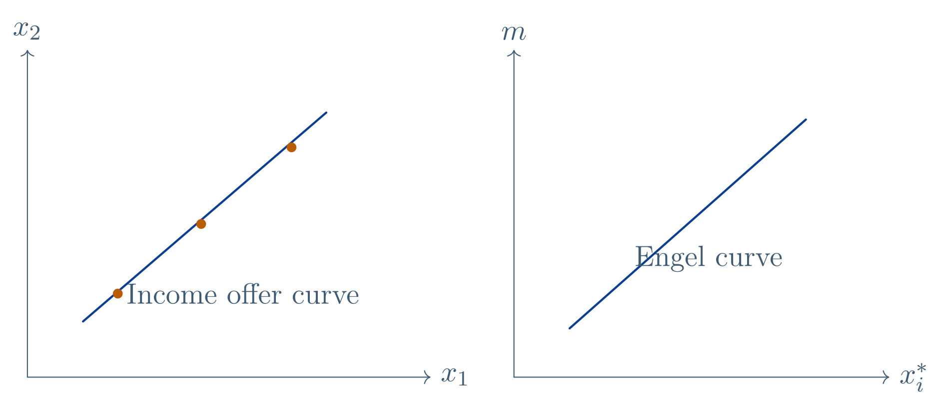

FIGURE 18

Lecture 7 — Tuesday, October 06, 2015

For Cobb–Douglas preferences,

$$x_1^* = \frac{a}{a+b}\frac{m}{p_1} \qquad \text{and} \qquad x_2^* = \frac{b}{a+b}\frac{m}{p_2}.$$

Solving each for income gives the Engel curves

$$m = \frac{(a+b)p_1}{a}x_1^* \qquad \text{and} \qquad m = \frac{(a+b)p_2}{b}x_2^*.$$

Both are straight lines through the origin.

For perfect complements,

$$x_1^* = x_2^* = \frac{m}{p_1+p_2},$$

so the Engel curves are

$$m = (p_1+p_2)x_1^* \qquad \text{and} \qquad m = (p_1+p_2)x_2^*.$$

For perfect substitutes, if $p_1 < p_2$ then

$$x_1^* = \frac{m}{p_1} \qquad \text{and} \qquad x_2^* = 0,$$

so

$$m = p_1 x_1^*$$

is the Engel curve for good 1, while good 2 stays at zero.

Engel curves are not always straight lines.

Definition 7.1 Preferences are **homothetic** if

$$(x_1, x_2) \succeq (y_1, y_2) \qquad \text{if and only if} \qquad (kx_1, kx_2) \succeq (ky_1, ky_2)$$

for every $k > 0$.

For differentiable homothetic preferences, the marginal rate of substitution is constant along any ray from the origin. At fixed prices, homothetic preferences generate linear Engel curves through the origin whenever demand is unique. Quasilinear preferences, such as

$$U(x_1, x_2) = f(x_1) + x_2,$$

are not homothetic, so their Engel curves need not be straight lines.

Definition 7.2 A good is **normal** if quantity demanded rises with income, and **inferior** if quantity demanded falls as income rises.

The Engel curve of a normal good slopes upward. The Engel curve of an inferior good slopes downward.

Definition 7.3 A good is **ordinary** if demand falls whenever its own price rises. A good is **Giffen** if, for some price range, demand rises when its own price rises.

Figure 19 contrasts the corresponding demand curve shapes.

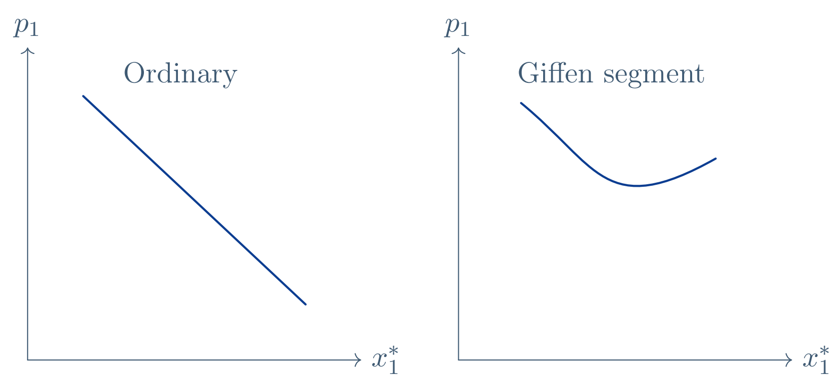

FIGURE 19

Cross-price effects classify pairs of goods.

Definition 7.4 If an increase in p_2 raises demand for good 1, then good 1 is a **gross substitute** for good 2. If an increase in p_2 lowers demand for good 1, then good 1 is a **gross complement** for good 2.

For perfect complements,

$$x_1^* = \frac{m}{p_1 + p_2},$$

so

$$\frac{\partial x_1^*}{\partial p_2} = -\frac{m}{(p_1 + p_2)^2} < 0.$$

Thus the goods are gross complements.

For Cobb–Douglas preferences,

$$x_1^* = \frac{am}{(a+b)p_1},$$

which does not depend on p_2, so

$$\frac{\partial x_1^*}{\partial p_2} = 0.$$

In that sense, neither commodity is a gross complement nor a gross substitute for the other.

Lecture 8 — Thursday, October 08, 2015

Substitution and Income Effects

Chapter 8 explains why a consumer's demand changes when a price changes. The standard decomposition separates that response into two parts, the **substitution effect** and the **income effect**.

Suppose the price of good 1 falls from p_1 to p_1', with $p_1' < p_1$. Let the original chosen bundle be $x^0 = (x_1^0, x_2^0)$. Slutsky's idea is to adjust income so that, at the new prices, the consumer can just afford the original bundle:

$$m^S = p_1' x_1^0 + p_2 x_2^0.$$

Since the original bundle satisfied $p_1 x_1^0 + p_2 x_2^0 = m$, this compensated income can also be written as

$$m^S = m + (p_1' - p_1) x_1^0.$$

Because $p_1' - p_1 < 0$, the compensated income is lower than original income after a price fall. The compensated budget line pivots through the original choice, producing the decomposition shown in Figure 20.

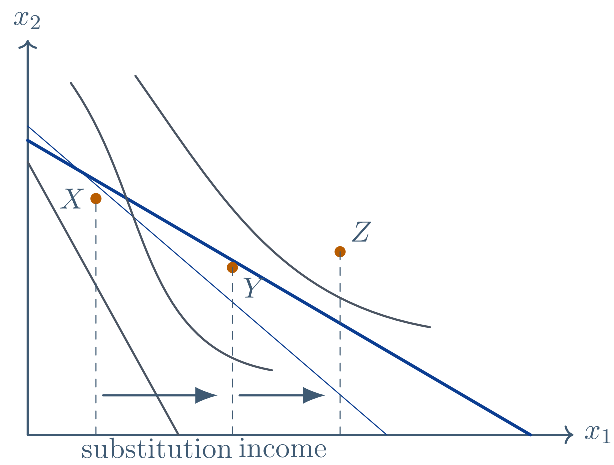

Figure 20

With this notation, the total change in demand for good 1 is

$$x_1(p_1', p_2, m) - x_1(p_1, p_2, m).$$

Adding and subtracting the compensated demand $x_1(p_1', p_2, m^S)$ gives the Slutsky decomposition:

$$\begin{aligned} & x_1(p_1', p_2, m) - x_1(p_1, p_2, m) \\ &= \underbrace{\left[x_1(p_1', p_2, m^S) - x_1(p_1, p_2, m)\right]}_{\text{substitution effect}} + \underbrace{\left[x_1(p_1', p_2, m) - x_1(p_1', p_2, m^S)\right]}_{\text{income effect}}. \end{aligned}$$

The first bracketed term is the substitution effect, while the second bracketed term is the income effect.

For an own-price decrease, the substitution effect is always nonnegative for the demanded good: when good 1 becomes relatively cheaper, the consumer substitutes toward it. The income effect can be positive or negative.

If higher income raises demand for the good, then the good is normal and the income effect from a price decrease is positive. If higher income lowers demand, the good is inferior and the income effect from a price decrease is negative.

So for a normal good, substitution and income effects reinforce one another. For an inferior good, they work in opposite directions. Figure 21 contrasts a Giffen case with an inferior but non-Giffen case.

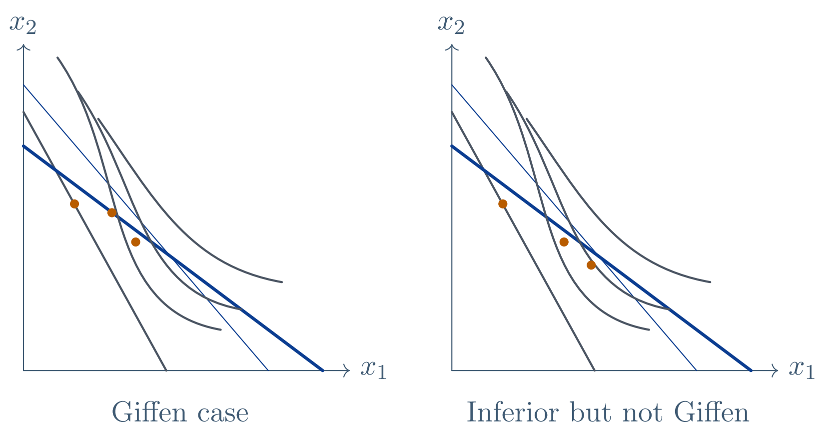

FIGURE 21

This leads directly to the weak law of demand.

Proposition 8.1 If a good is normal, then demand for that good cannot increase when its own price increases. It decreases strictly whenever at least one component of the Slutsky decomposition is strict.

Proof. When price rises, the substitution effect weakly reduces demand for the now relatively more expensive good. If the good is normal, the income effect from the price rise also weakly reduces demand because the consumer has effectively become poorer. Since both components point in the same direction, total demand cannot rise; it falls when either effect is strict. □

This proposition is not a statement about all goods whatsoever. Inferior goods are the important exception, and Giffen goods are the extreme case in which the negative income effect from a price fall overwhelms the positive substitution effect.

Several standard examples illustrate how the decomposition behaves in special cases.

Fact 8.2 For perfect complements, the substitution effect is zero. When the price of one shoe falls, the consumer still wants to consume the goods in fixed proportions, so all of the change in demand comes from the income effect.

Fact 8.3 For perfect substitutes, a price change that switches which good is cheaper can produce a discrete substitution toward the newly cheaper good. In that case the income effect may be zero. If the same good remains uniquely cheaper, however, its own-price response need not be a substitution effect: the consumer continues to spend all income on that good, and the Slutsky income effect generally remains.

Note: The graphical Slutsky decomposition clarifies the sign of substitution and income effects, the law of demand for normal goods, and the standard examples involving inferior goods, perfect complements, and perfect substitutes.

Lecture 9 — Tuesday, October 13, 2015

Market Demand

The **market demand** for a good is obtained by adding individual demands. If there are n consumers and consumer i demands $x_i^1(p_1, p_2, m_i)$ units of good 1, then aggregate demand for good 1 is

$$X_1(p_1, p_2, m_1, \ldots, m_n) = \sum_{i=1}^{n} x_i^1(p_1, p_2, m_i).$$

Market demand therefore depends not only on prices, but also on the distribution of income across consumers. In general, two economies with the same total income

$$M = \sum_i m_i$$

can have different market demand if that total income is distributed differently.

Warning A **representative consumer** with total income M is a convenient shortcut, but it is not always valid. It works only under additional assumptions on preferences and aggregation, so it should be treated as a modeling convenience rather than a general theorem.

If the price of good 2 and all incomes are held fixed, then market demand for good 1 is found by horizontally summing individual demand curves, as in Figure 22.

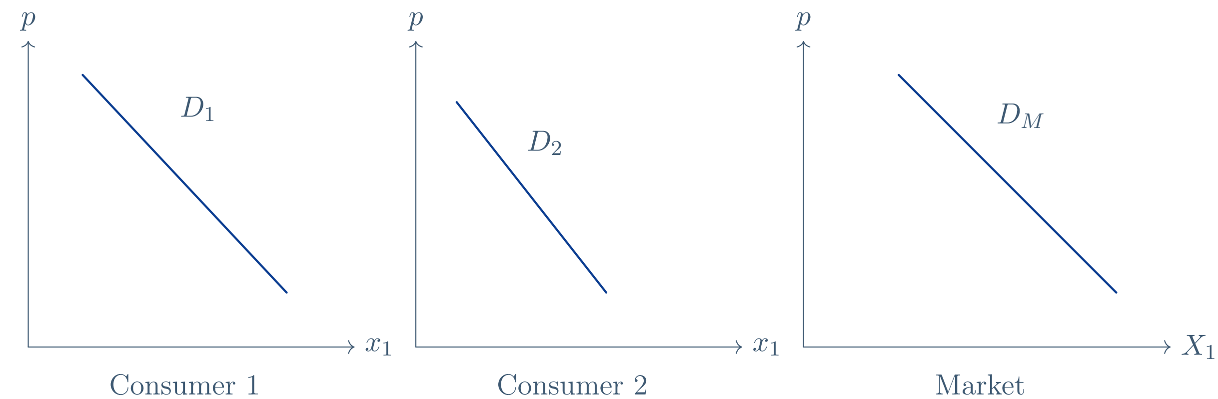

Figure 22

The inverse demand function asks what price must prevail for a given quantity to be demanded. To illustrate, consider two individuals whose demands, over a price range in which both are positive, are

$$x_1 = D_1(p) = 20 - p \qquad \text{and} \qquad x_2 = D_2(p) = 10 - 2p.$$

Within that price range, the market quantity is

$$X(p) = D_1(p) + D_2(p) = (20 - p) + (10 - 2p) = 30 - 3p.$$

Solving for price gives the inverse demand function

$$p(X) = 10 - \frac{X}{3}.$$

Outside that range, a complete market demand function would impose nonnegativity and become piecewise when either consumer reaches a choke price. More generally, when ordinary demand is one-to-one over the relevant range, inverse demand contains the same information with price written as a function of quantity. Without one-to-one demand, the inverse relation may be set-valued or must be restricted to a monotone branch.

Elasticity

The derivative $\mathrm{d}q/\,\mathrm{d}p$ measures local responsiveness, but it depends on the units in which price and quantity are measured. A unit-free measure is the **price elasticity of demand**.

Definition 9.1 At a point where demand $q(p) > 0$ is differentiable, the **price elasticity of demand** is

$$\varepsilon(p) = \frac{p}{q(p)} \frac{\mathrm{d}q(p)}{\mathrm{d}p}.$$

For small finite changes this is approximated by the percentage change in quantity divided by the percentage change in price; for large changes, an arc-elasticity convention must be specified.

For downward-sloping demand, elasticity is negative. In ordinary discussion economists often refer to its absolute value instead. Consider the inverse linear demand

$$p = 6 - q.$$

Writing quantity as a function of price gives $q = 6 - p$, so

$$\frac{\mathrm{d}q}{\mathrm{d}p} = -1.$$

Hence elasticity at a point is

$$\varepsilon = \frac{p}{q}(-1) = -\frac{p}{q}.$$

Along a linear demand curve, elasticity varies from nearly zero near the quantity intercept to arbitrarily large magnitude near the price intercept; Figure 23 marks a representative point.

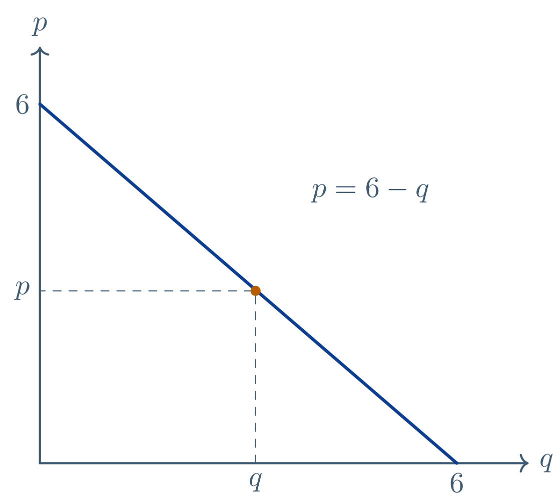

FIGURE 23

Lecture 10 — Thursday, October 15, 2015

For differentiable demand, revenue is

$$R = pq.$$

If price rises, quantity falls, so the effect on revenue depends on how responsive demand is. Differentiating with respect to price gives the exact local relation

$$\frac{\mathrm{d}R}{\mathrm{d}p} = q + p\frac{\mathrm{d}q}{\mathrm{d}p}.$$

Using elasticity,

$$\varepsilon = \frac{p}{q}\frac{\mathrm{d}q}{\mathrm{d}p} \Longrightarrow p\frac{\mathrm{d}q}{\mathrm{d}p} = q\varepsilon,$$

so

$$\frac{\mathrm{d}R}{\mathrm{d}p} = q(1+\varepsilon).$$

Since $\varepsilon < 0$, this is often written as

$$\frac{\mathrm{d}R}{\mathrm{d}p} = q(1-|\varepsilon|).$$

Fact 10.1 If $|\varepsilon| > 1$, demand is elastic and a small price increase lowers revenue. If $|\varepsilon| < 1$, demand is inelastic and a small price increase raises revenue. If $|\varepsilon| = 1$, the local derivative of revenue with respect to price is zero.

Linear demand does not have constant elasticity, so the chapter asks what kind of curve does. Wherever positive differentiable demand has constant revenue, its elasticity is -1. Since $R = pq$, holding revenue fixed at some positive constant R gives

$$pq = R \Longrightarrow q = \frac{R}{p}.$$

This rectangular hyperbola is shown in Figure 24.

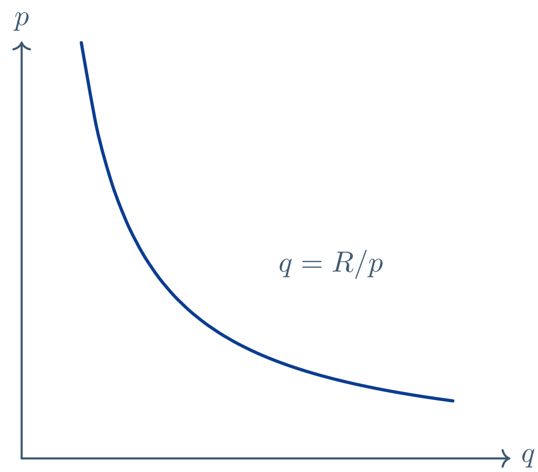

Figure 24

The same revenue identity can be used to compute marginal revenue. Writing inverse demand as $p = p(q)$ and differentiating $R(q) = p(q)q$ gives

$$MR = \frac{\mathrm{d}R}{\mathrm{d}q} = p + q\frac{\mathrm{d}p}{\mathrm{d}q}.$$

To connect this with elasticity, start from

$$\varepsilon = \frac{p}{q}\frac{\mathrm{d}q}{\mathrm{d}p}.$$

Taking reciprocals,

$$\frac{1}{\varepsilon} = \frac{q}{p}\frac{\mathrm{d}p}{\mathrm{d}q}.$$

Hence

$$q\frac{\mathrm{d}p}{\mathrm{d}q} = \frac{p}{\varepsilon},$$

and therefore

$$MR = p + \frac{p}{\varepsilon} = p\left(1 + \frac{1}{\varepsilon}\right).$$

Because demand elasticity is negative, this is commonly written as

$$MR = p\left(1 - \frac{1}{|\varepsilon|}\right).$$

Fact 10.2 If demand is inelastic, then $MR < 0$. If demand is unit elastic, then $MR = 0$. If demand is elastic, then $MR > 0$.

For a linear inverse demand curve

$$p(q) = a - bq,$$

marginal revenue is

$$R(q) = p(q)q = (a - bq)q = aq - bq^2,$$

so

$$MR(q) = \frac{\mathrm{d}R}{\mathrm{d}q} = a - 2bq.$$

Thus the marginal revenue curve has the same intercept as inverse demand, but twice the slope, as Figure 25 makes visually explicit.

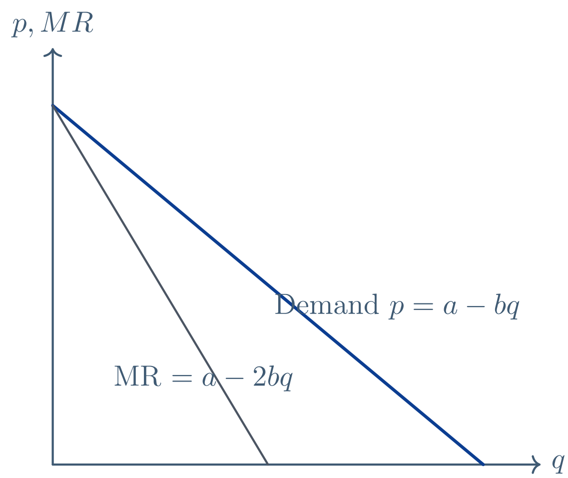

FIGURE 25

Lecture 11 — Tuesday, October 20, 2015

3. Market Equilibrium, Production, and Costs

Market Equilibrium

A competitive market is in equilibrium when the quantity buyers wish to purchase equals the quantity sellers wish to sell at the prevailing price. If demand is written as $q = D(p)$ and supply is written as $q = S(p)$, then an equilibrium price p^* satisfies

$$D(p^*) = S(p^*).$$

The corresponding equilibrium quantity is

$$q^* = D(p^*) = S(p^*).$$

Definition 11.1 A competitive equilibrium is a price–quantity pair (p^*, q^*) such that $q^* = D(p^*) = S(p^*)$ and buyers and sellers are simultaneously optimizing at the price p^*.

If price is above p^*, then $S(p) > D(p)$ and there is excess supply. Unsold output puts downward pressure on price. If price is below p^*, then $D(p) > S(p)$ and there is excess demand. Shortages then put upward pressure on price. Only at the intersection of demand and supply, marked in Figure 26, is there no force pushing price away from its current level.

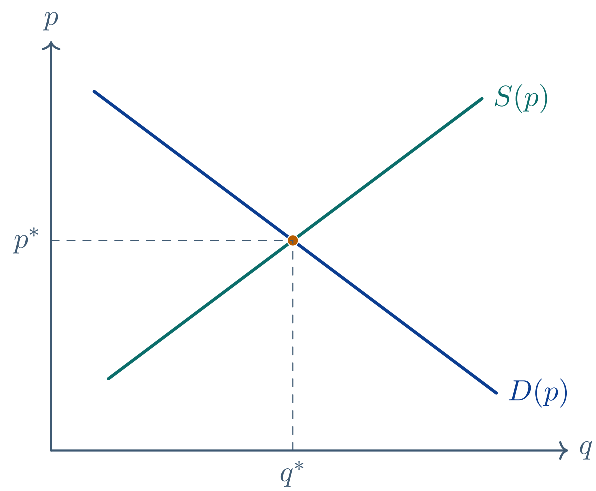

FIGURE 26

Consider the standard linear demand–supply system

$$D(p) = a - bp \qquad \text{and} \qquad S(p) = c + dp,$$

where a, b, c, d are constants with $b > 0$ and $d > 0$. We restrict attention to parameters for which the resulting price and quantity are nonnegative. To solve for an interior equilibrium, set demand equal to supply:

$$a - bp^* = c + dp^*.$$

Collecting the price terms on one side gives

$$a - c = (b + d)p^*,$$

so the equilibrium price is

$$p^* = \frac{a - c}{b + d}.$$

Substituting this into either demand or supply gives equilibrium quantity:

$$q^* = a - b\left(\frac{a - c}{b + d}\right) = \frac{a(b + d) - b(a - c)}{b + d} = \frac{ad + bc}{b + d}.$$

Example 11.2 If $a = 20$, $b = 2$, $c = 2$, and $d = 1$, then

$$p^* = \frac{20 - 2}{2 + 1} = 6 \qquad \text{and} \qquad q^* = \frac{20(1) + 2(2)}{3} = 8.$$

Checking directly,

$$D(6) = 20 - 12 = 8 \qquad \text{and} \qquad S(6) = 2 + 6 = 8.$$

The first special case is fixed short-run supply. If supply does not respond to price, then $d = 0$ and

$$S(p) \equiv c.$$

The supply curve is vertical at $q = c$. Equilibrium quantity is therefore fixed at $q^* = c$, and equilibrium price is whatever price clears demand:

$$c = D(p^*) = a - bp^* \Longrightarrow p^* = \frac{a - c}{b}.$$

The second special case is perfectly elastic supply. In that case sellers are willing to supply any amount at some given price $\bar{p}$, so the inverse supply curve is horizontal:

$$S^{-1}(q) = \bar{p}.$$

Equilibrium price is pinned down immediately by supply:

$$p^* = \bar{p},$$

and equilibrium quantity is then demand at that price:

$$q^* = D(\bar{p}).$$

Lecture 12 — Thursday, October 22, 2015

Figure 27 compares the cases with fixed supply and perfectly elastic supply.

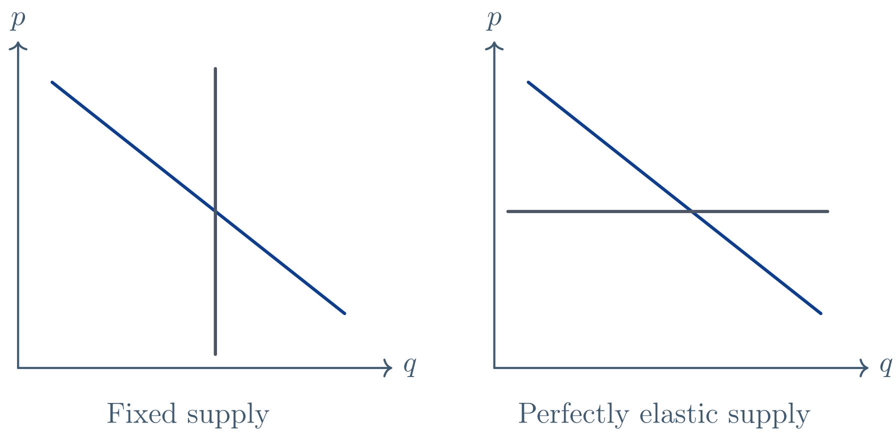

FIGURE 27

A quantity tax of t dollars per unit creates a wedge between the price paid by buyers and the price received by sellers. If p_b denotes the buyers' price and p_s denotes the sellers' price, then the tax wedge is

$$p_b - p_s = t.$$

Equilibrium in the taxed market is described by two conditions:

$$D(p_b) = S(p_s) \qquad \text{and} \qquad p_b - p_s = t.$$

These equations are the same whether the tax is legally imposed on buyers or on sellers. The statutory side of the market need not be the side that ultimately bears the burden.

Using the linear demand and supply functions, the taxed equilibrium solves

$$a - bp_b = c + dp_s \qquad \text{and} \qquad p_b = p_s + t.$$

Substitute the second equation into the first:

$$a - b(p_s + t) = c + dp_s.$$

Expanding and collecting terms yields

$$a - c - bt = (b + d)p_s,$$

so the sellers' price is

$$p_s = \frac{a - c - bt}{b + d}.$$

Then the buyers' price is

$$p_b = p_s + t = \frac{a - c - bt}{b + d} + t = \frac{a - c + dt}{b + d}.$$

The quantity traded is

$$q_t = S(p_s) = c + d\left(\frac{a - c - bt}{b + d}\right) = \frac{ad + bc - bdt}{b + d}.$$

As long as the tax is small enough that the displayed interior equilibrium has positive quantity, the quantity traded falls. Buyers pay more, sellers receive less, and the difference between the two prices is exactly the tax. A larger tax can instead shut the market down, in which case the linear interior formulas no longer apply.

Fact 12.1 The burden of a unit tax is shared by buyers and sellers through the change in the equilibrium prices, and the quantity traded falls because some mutually beneficial trades are no longer carried out.

The resulting price wedge is shown in Figure 28.

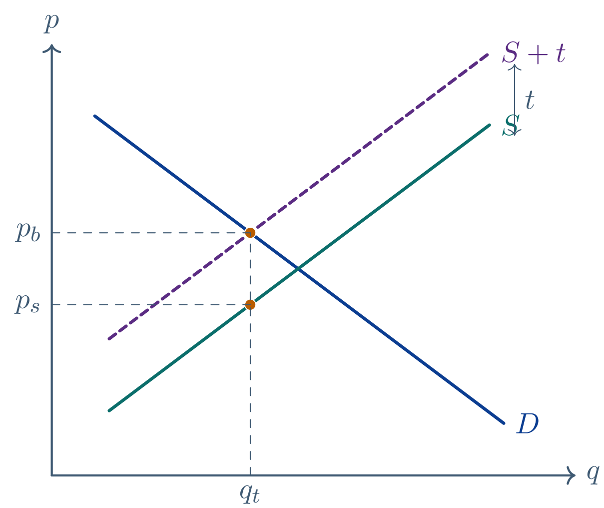

FIGURE 28

Note: The deeper incidence question is not "who remits the payment to the government?" but rather "whose net price changes more?" In general, the less elastic side of the market bears more of the tax burden.

Lecture 13 — Tuesday, October 27, 2015

Production Functions

A technology is a method of converting inputs into output. If the inputs are $x_1, \ldots, x_n$, then an input bundle is the vector

$$(x_1, x_2, \ldots, x_n).$$

The production function records the maximum output attainable from each bundle:

$$y = f(x_1, \ldots, x_n).$$

The word "maximum" matters. A firm may choose to produce less than what is technologically feasible, but the production function describes the frontier of what can be produced.

Definition 13.1 A **production function** assigns to each input bundle the largest output level that can be produced from that bundle.

A useful running example is

$$y = f(x_1, x_2) = 2x_1^{1/3}x_2^{1/3}.$$

This technology has diminishing marginal products in each input, but both inputs contribute positively to output.

For the bundle $(x_1, x_2) = (1, 8)$,

$$y = 2(1)^{1/3}(8)^{1/3} = 2(1)(2) = 4.$$

For the bundle $(x_1, x_2) = (8, 8)$,

$$y = 2(8)^{1/3}(8)^{1/3} = 2(2)(2) = 8.$$

So increasing input x_1 eightfold, from 1 to 8, while keeping $x_2 = 8$ doubles output from 4 to 8 in this example.

The production analogue of an indifference curve is the **isoquant**. It is the set of input bundles that generate the same output level. Starting from

$$y = 2x_1^{1/3}x_2^{1/3},$$

the isoquant for output level y is obtained by solving for x_2:

$$\frac{y}{2} = x_1^{1/3}x_2^{1/3} \Longrightarrow \left(\frac{y}{2}\right)^3 = x_1x_2 \Longrightarrow x_2 = \frac{y^3}{8x_1}.$$

Hence the isoquants are downward sloping hyperbolas.

Definition 13.2 An **isoquant** is the set of input bundles that produce a fixed output level y.

The Cobb–Douglas isoquants in Figure 29 correspond to several output levels.

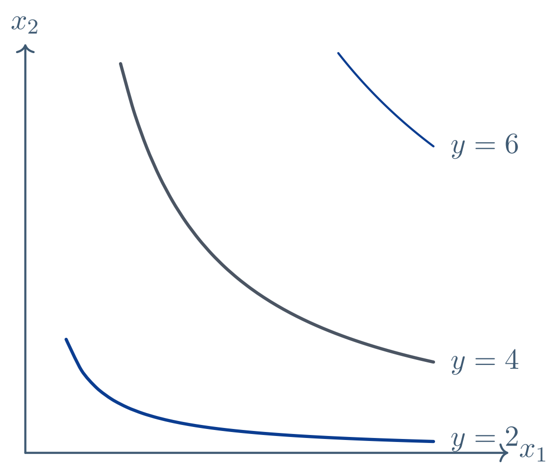

FIGURE 29

Three benchmark technologies are especially important. A **Cobb–Douglas technology** with $A > 0$ and positive exponents has the form

$$y = Ax_1^{a_1}x_2^{a_2}\cdots x_n^{a_n}.$$

Its isoquants are smooth and convex, and they asymptote toward the axes without touching them.

A **fixed-proportions technology** has the form

$$y = \min(a_1x_1, a_2x_2, \ldots, a_nx_n).$$

Inputs must be used in rigid proportions, so the isoquants are L-shaped. For example, if

$$y = \min(x_1, 2x_2),$$

then one extra unit of x_1 is useless whenever $2x_2 < x_1$, because input 2 is then the bottleneck.

A **perfect-substitutes technology** has the form

$$y = a_1x_1 + a_2x_2 + \cdots + a_nx_n.$$

Its isoquants are straight lines because one input can replace another at a constant rate. Figure 30 compares all three benchmark shapes.

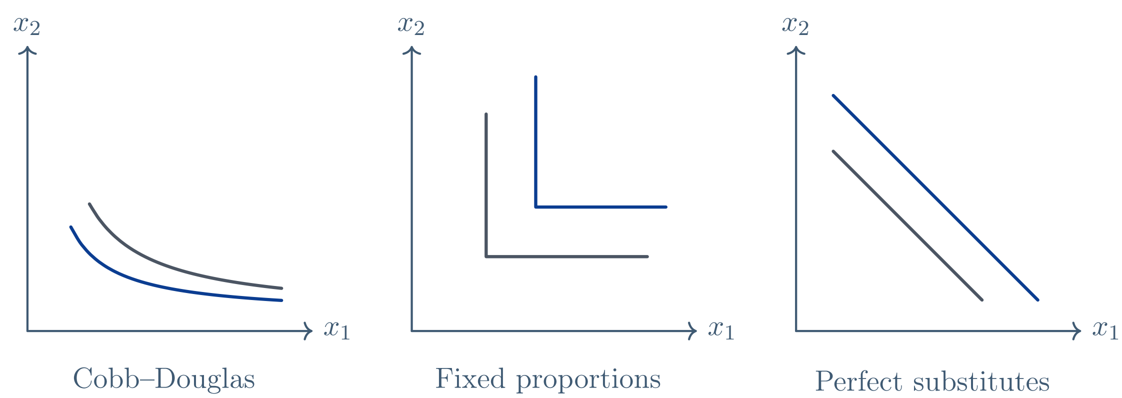

FIGURE 30

Marginal Products and Technical Rate of Substitution

The **marginal product** of input i is

$$MP_i = \frac{\partial f}{\partial x_i}.$$

For the Cobb–Douglas example

$$f(x_1, x_2) = x_1^{1/3}x_2^{2/3},$$

the marginal products are

$$MP_1 = \frac{\partial y}{\partial x_1} = \frac{1}{3}x_1^{-2/3}x_2^{2/3} \qquad \text{and} \qquad MP_2 = \frac{\partial y}{\partial x_2} = \frac{2}{3}x_1^{1/3}x_2^{-1/3}.$$

The ratio of marginal products determines the slope of an isoquant. Along an isoquant, total differentiation gives

$$0 = \mathrm{d}y = MP_1\ \mathrm{d}x_1 + MP_2\ \mathrm{d}x_2,$$

so

$$\frac{\mathrm{d}x_2}{\mathrm{d}x_1} = -\frac{MP_1}{MP_2}.$$

This is the **technical rate of substitution**.

Definition 13.3 The **technical rate of substitution** of input 1 for input 2 is

$$TRS = -\frac{\mathrm{d}x_2}{\mathrm{d}x_1} = \frac{MP_1}{MP_2}$$

along an isoquant. It measures how much of input 2 can be reduced when one more unit of input 1 is used and output is held fixed.

For the Cobb–Douglas example above,

$$\frac{MP_1}{MP_2} = \frac{\frac{1}{3}x_1^{-2/3}x_2^{2/3}}{\frac{2}{3}x_1^{1/3}x_2^{-1/3}} = \frac{1}{2}\frac{x_2}{x_1}.$$

Hence the technical rate of substitution falls as x_1 rises relative to x_2, which is the familiar property of diminishing substitutability behind convex isoquants.

Fact 13.4 Well-behaved production functions are usually assumed to be monotonic, meaning that more of any input cannot reduce output, and quasiconcave, meaning that a mixture of two input bundles producing at least y units can also produce at least y units.

Lecture 14 — Thursday, October 29, 2015

Profit Maximization

A firm's production plan specifies the quantities of outputs it produces and the inputs it hires. If outputs are $y_1, \ldots, y_n$ with prices $p_1, \ldots, p_n$, and inputs are $x_1, \ldots, x_m$ with prices $w_1, \ldots, w_m$, then economic profit is

$$\Pi = p_1 y_1 + \cdots + p_n y_n - w_1 x_1 - \cdots - w_m x_m.$$

For a competitive firm all output prices and all input prices are taken as given. The firm chooses only the quantities.

Consider a firm with one output and two inputs in the short run, where input x_2 is fixed at $\tilde{x}_2$. The short-run production function is therefore

$$y = f(x_1, \tilde{x}_2),$$

and profit is

$$\Pi = py - w_1 x_1 - w_2 \tilde{x}_2.$$

Since $w_2 \tilde{x}_2$ does not change with x_1, it is a **fixed cost**. The only adjustable decision is how much of input 1 to employ.

Holding Π fixed gives the **isoprofit line**

$$\Pi = py - w_1 x_1 - w_2 \tilde{x}_2.$$

Solving for output,

$$y = \frac{w_1}{p} x_1 + \frac{\Pi + w_2 \tilde{x}_2}{p}.$$

Thus an isoprofit line in (x_1, y)-space has slope w_1/p. Larger profits correspond to parallel lines lying higher up, as in Figure 31.

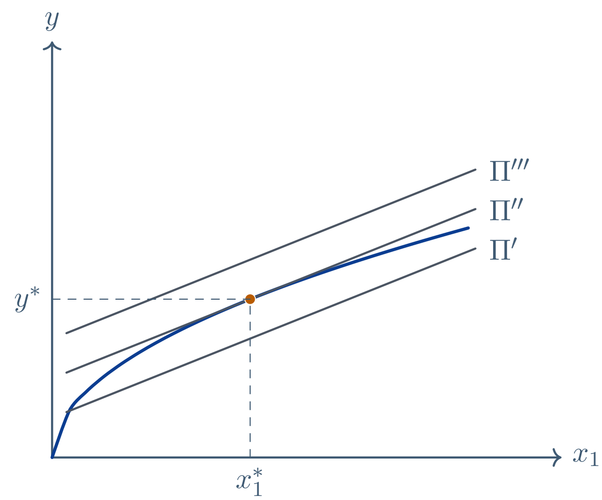

FIGURE 31

At an interior optimum, the highest attainable isoprofit line is tangent to the production function. The slope of the production function in (x_1, y)-space is the marginal product of input 1:

$$MP_1 = \frac{\partial y}{\partial x_1}.$$

Tangency therefore requires

$$MP_1 = \frac{w_1}{p}.$$

Multiplying through by p gives the economically transparent form

$$p\,MP_1 = w_1.$$

The left-hand side is the marginal revenue product of input 1. The firm hires the input up to the point where the extra revenue generated by one more unit equals the extra cost of hiring that unit.

Intuition: If $p\,MP_1 > w_1$, then one more unit of the input adds more to revenue than to cost, so profit rises when the firm increases x_1. If $p\,MP_1 < w_1$, profit rises when the firm reduces x_1.

Suppose the short-run production function is

$$y = x_1^{1/3}\tilde{x}_2^{1/3}.$$

Then
$$MP_1 = \frac{\partial y}{\partial x_1} = \frac{1}{3} x_1^{-2/3} \tilde{x}_2^{1/3}.$$
The profit-maximizing condition $p\,MP_1 = w_1$ becomes
$$\frac{p}{3} x_1^{-2/3} \tilde{x}_2^{1/3} = w_1.$$
Solve this step by step:
$$x_1^{-2/3} = \frac{3w_1}{p\tilde{x}_2^{1/3}},$$
so
$$x_1^{2/3} = \frac{p\tilde{x}_2^{1/3}}{3w_1}.$$
Raising both sides to the power 3/2 gives
$$x_1^* = \left(\frac{p}{3w_1}\right)^{3/2} \tilde{x}_2^{1/2}.$$
Substituting this into the production function gives optimal output:
$$y^* = (x_1^*)^{1/3} \tilde{x}_2^{1/3} = \left(\frac{p}{3w_1}\right)^{1/2} \tilde{x}_2^{1/2}.$$

Fact 14.1 In this example the firm's demand for the variable input rises with output price p, falls with the wage w_1, and rises with the amount of the fixed input $\tilde{x}_2$.

Lecture 15 — Tuesday, November 03, 2015

Cost Minimization

For a given output target y, a firm is a **cost minimizer** if it chooses the cheapest input bundle capable of producing at least that output. With two variable inputs, the problem is
$$\min_{x_1, x_2 \geq 0} w_1 x_1 + w_2 x_2 \qquad \text{subject to} \qquad f(x_1, x_2) \geq y.$$
When input prices are positive and the technology is monotone with free disposal, the output constraint binds at a cost-minimizing bundle, so the subsequent calculations may use $f(x_1, x_2) = y$.

The solution gives the firm's conditional factor demands:

$$x_1^* = x_1^*(w_1, w_2, y) \qquad \text{and} \qquad x_2^* = x_2^*(w_1, w_2, y).$$

Substituting those demands into total expenditure gives the cost function

$$c(w_1, w_2, y) = w_1 x_1^* + w_2 x_2^*.$$

Definition 15.1 A **conditional factor demand** gives the cost-minimizing quantity of an input as a function of input prices and required output.

An **isocost line** contains all bundles with the same total cost c:

$$w_1 x_1 + w_2 x_2 = c.$$

Solving for x_2,

$$x_2 = -\frac{w_1}{w_2} x_1 + \frac{c}{w_2}.$$

Its slope is therefore $-w_1/w_2$. At an interior cost-minimizing bundle, the isoquant for output y is tangent to the lowest attainable isocost line. Since the slope of the isoquant is $-MP_1/MP_2$, tangency requires

$$-\frac{w_1}{w_2} = -\frac{MP_1}{MP_2}.$$

Equivalently,

$$\frac{MP_1}{w_1} = \frac{MP_2}{w_2}.$$

The last equation says the firm equalizes marginal product per dollar across inputs. Figure 32 shows the associated tangency.

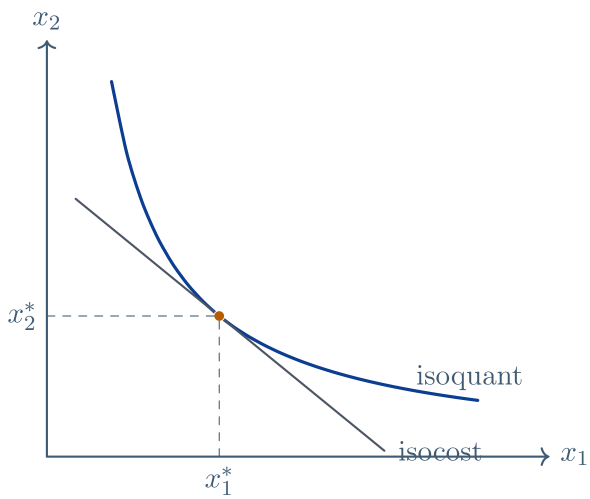

FIGURE 32

Let

$$y = f(x_1, x_2) = x_1^{1/3} x_2^{2/3}.$$

The marginal products are

$$MP_1 = \frac{1}{3} x_1^{-2/3} x_2^{2/3} \qquad \text{and} \qquad MP_2 = \frac{2}{3} x_1^{1/3} x_2^{-1/3}.$$

The tangency condition is

$$\frac{w_1}{w_2} = \frac{MP_1}{MP_2} = \frac{\frac{1}{3} x_1^{-2/3} x_2^{2/3}}{\frac{2}{3} x_1^{1/3} x_2^{-1/3}} = \frac{x_2}{2x_1}.$$

Hence

$$x_2^* = \frac{2w_1}{w_2} x_1^*.$$

Substitute this into the production constraint:

$$y = (x_1^*)^{1/3} \left(\frac{2w_1}{w_2} x_1^* \right)^{2/3} = x_1^* \left(\frac{2w_1}{w_2} \right)^{2/3}.$$

Therefore

$$x_1^* = \left(\frac{w_2}{2w_1} \right)^{2/3} y.$$

Using the earlier relation,

$$x_2^* = \frac{2w_1}{w_2}x_1^* = \left(\frac{2w_1}{w_2}\right)^{1/3} y.$$

The resulting cost function is

$$c(w_1, w_2, y) = w_1x_1^* + w_2x_2^*.$$

Substituting the two conditional demands gives

$$c(w_1, w_2, y) = w_1\left(\frac{w_2}{2w_1}\right)^{2/3} y + w_2\left(\frac{2w_1}{w_2}\right)^{1/3} y = 3\left(\frac{w_1w_2^2}{4}\right)^{1/3} y.$$

If

$$y = \min(4x_1, x_2),$$

then producing y units requires

$$4x_1 = y \qquad \text{and} \qquad x_2 = y,$$

because cost minimization never leaves one input in excess when the technology uses strict proportions. Thus

$$x_1^* = \frac{y}{4} \qquad \text{and} \qquad x_2^* = y,$$

and the cost function is

$$c(w_1, w_2, y) = w_1\frac{y}{4} + w_2y = \left(\frac{w_1}{4} + w_2\right) y.$$

As required output varies, the collection of cost-minimizing bundles traces out the **output expansion path**. For a Cobb–Douglas technology, the expansion path is a ray through the origin because both conditional input demands are proportional to output. For a fixed-proportions technology, the expansion path lies along the kink line where the required input ratio is maintained.

Fact 15.2 At fixed positive input prices, a homogeneous production function with constant returns to scale has constant long-run average cost. Under the corresponding global decreasing- or increasing-returns assumption, long-run average cost rises or falls with output, respectively.

The long-run cost of producing a given output is never greater than the short-run cost of producing that same output. The reason is simple: the long-run problem allows the firm to choose all inputs freely, while the short-run problem imposes extra constraints by fixing at least one input. Adding constraints cannot lower the minimum cost.

Lecture 16 — Thursday, November 05, 2015

Cost Curves

In the short run the firm's cost information can be organized into a family of related curves. Total cost can be decomposed as

$$TC(y) = TFC + TVC(y),$$

where TFC is **total fixed cost** and $TVC(y)$ is **total variable cost**. For $y > 0$, this decomposition gives **average fixed cost**,

$$AFC(y) = \frac{TFC}{y},$$

average variable cost,

$$AVC(y) = \frac{TVC(y)}{y},$$

and **average total cost**,

$$ATC(y) = \frac{TC(y)}{y} = AFC(y) + AVC(y).$$

Average fixed cost is a rectangular hyperbola: as output rises, fixed cost is spread over more units, so $AFC(y)$ falls toward zero.

Definition 16.1 **Marginal cost** is the rate at which total cost changes with output:

$$MC(y) = \frac{\mathrm{d}TC(y)}{\mathrm{d}y} = \frac{\mathrm{d}TVC(y)}{\mathrm{d}y}.$$

Because fixed cost does not vary with output, the total cost curve and total variable cost curve have the same slope everywhere. The only difference is that the total cost curve is shifted upward by the amount of fixed cost. In the short run, diminishing returns imply that variable cost eventually rises more than proportionally with output, so both AVC and ATC eventually rise as output expands. These relationships are collected in Figure 33.

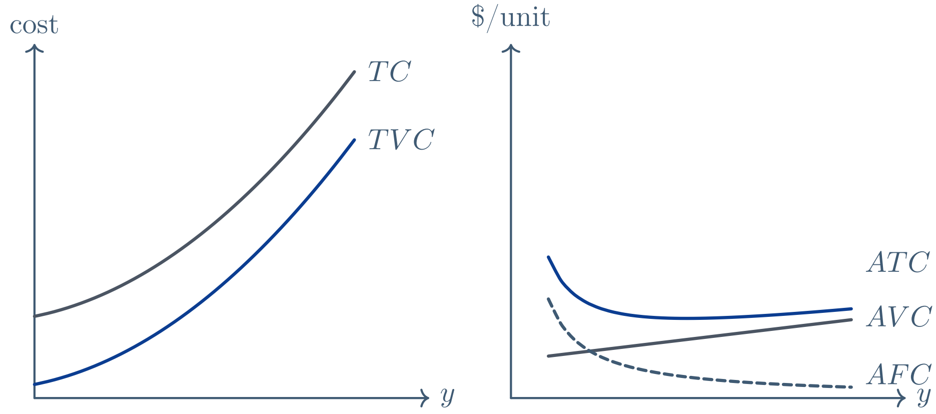

FIGURE 33

The standard intersection results follow from a short calculation. Since

$$AVC(y) = \frac{TVC(y)}{y},$$

the derivative is

$$\frac{\mathrm{d}AVC}{\mathrm{d}y} = \frac{y\,MC(y) - TVC(y)}{y^2}.$$

Because $TVC(y)/y = AVC(y)$, this becomes

$$\frac{\mathrm{d}AVC}{\mathrm{d}y} = \frac{MC(y) - AVC(y)}{y}.$$

Therefore,

$$\frac{\mathrm{d}AVC}{\mathrm{d}y} \begin{cases} < 0, & MC(y) < AVC(y), \\ = 0, & MC(y) = AVC(y), \\ > 0, & MC(y) > AVC(y). \end{cases}$$

Thus $MC = AVC$ at every interior stationary point of AVC. Under the standard U-shaped cost assumptions, marginal cost crosses average variable cost from below at its minimum.

The same calculation works for average total cost:

$$ATC(y) = \frac{TC(y)}{y} \Longrightarrow \frac{\mathrm{d}ATC}{\mathrm{d}y} = \frac{y\,MC(y) - TC(y)}{y^2} = \frac{MC(y) - ATC(y)}{y}.$$

Hence $MC = ATC$ at every interior stationary point of ATC; under the standard U-shaped cost assumptions, the crossing from below occurs at the minimum of ATC.

Fact 16.2 When marginal cost lies below an average, it pulls that average down. When it lies above an average, it pushes that average up.

The resulting intersections are shown in Figure 34.

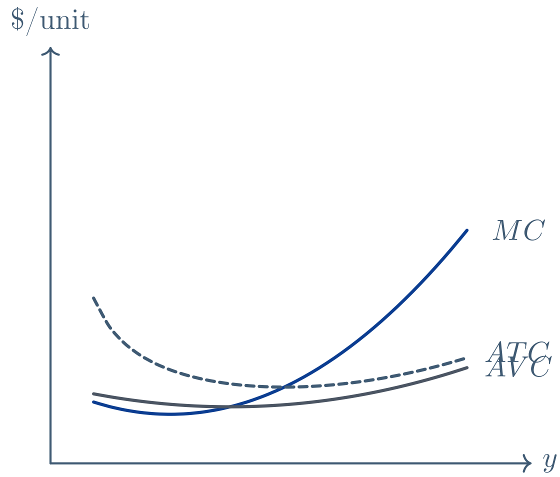

Figure 34

Every possible level of the fixed input generates its own short-run total cost curve. The long-run total cost curve is the lower envelope of all those short-run curves, because in the long run the firm is free to choose whichever level of the fixed input is cheapest for the output it wants to produce. The same envelope logic carries over to average cost: long-run average cost is the lower envelope of short-run average cost curves, as illustrated in Figure 35.

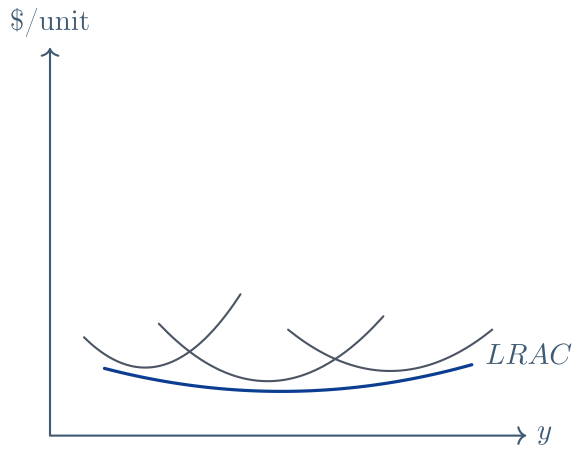

FIGURE 35

Lecture 17 — Tuesday, November 10, 2015

4. Competitive Supply and Monopoly

Competitive Firm Supply

How much a firm supplies depends on its technology, its objectives, and the market structure in which it operates. In a **purely competitive market**, each individual firm is small relative to the entire market and takes the market price as given. Such a firm is called a **price-taking firm**. Monopoly, oligopoly, markets with a dominant firm, and monopolistic competition differ precisely because firms have more scope to influence price.

If the firm tried to charge a price above the market price, buyers would switch to competitors. Charging less would simply give up revenue on units the firm could sell at the market price. Hence the firm's demand curve is horizontal at the market price, as Figure 36 shows alongside market demand and supply.

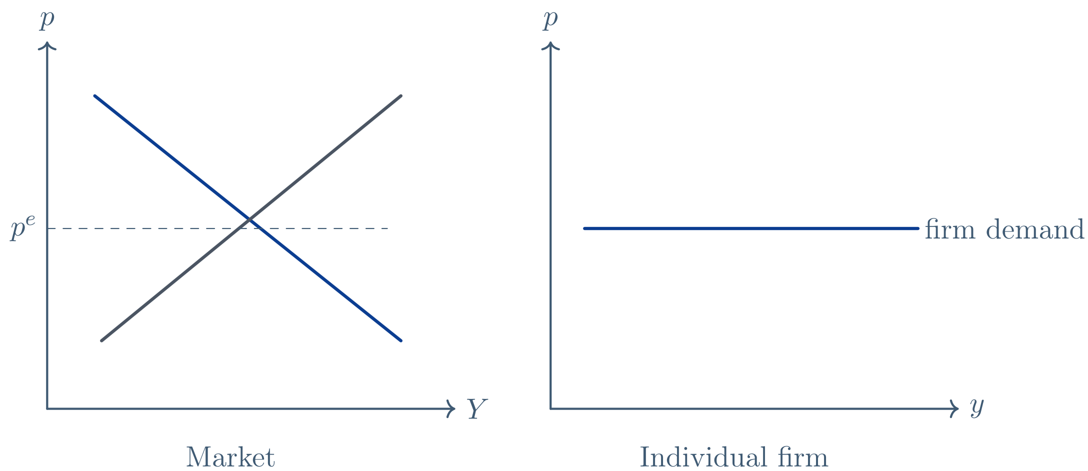

Figure 36

In the short run the firm's problem is

$$\max_{y \geq 0} \Pi^s(y) = py - TC(y).$$

If the optimum has $y^* > 0$, then the first-order condition is

$$\frac{\mathrm{d}\Pi^s}{\mathrm{d}y} = p - MC_s(y) = 0,$$

so

$$p = MC_s(y^*).$$

The second-order condition requires the marginal-cost curve to be upward sloping at that output level.

However, the firm does not produce at every point on the upward-sloping part of marginal cost. It can always choose $y = 0$, in which case its short-run profit is

$$\Pi^s(0) = -TFC.$$

Therefore it produces a positive quantity only if

$$py - TVC(y) \geq 0.$$

Dividing by $y > 0$ gives the **shutdown condition**

$$p \geq \frac{TVC(y)}{y} = AVC(y).$$

With the standard convex variable-cost shape, the firm supplies zero when $p < \min AVC$ and chooses the output on the rising marginal-cost branch satisfying $p = MC(y)$ when $p > \min AVC$. At $p = \min AVC$, it is indifferent between shutting down and producing at an output that minimizes AVC, so supply is set-valued. Thus the familiar positive-output part of short-run supply is the rising portion of marginal cost at or above the shutdown price; Figure 37 highlights that segment. More generally, supply is the set of global maximizers of $py - TC(y)$, so a first-order condition alone is not sufficient.

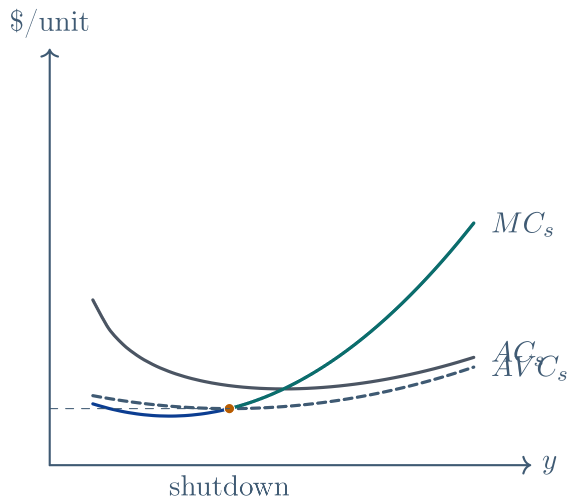

Figure 37

Warning **Shutdown** is not the same as **exit**. A shutdown firm produces zero output in the short run but still incurs fixed cost. Exit is a long-run decision in which the firm leaves the industry.

In the long run all inputs are variable, so the firm solves

$$\max_{y \geq 0} \ \Pi(y) = py - LTC(y).$$

If $y^* > 0$, the first-order condition is

$$p = LMC(y^*).$$

But long-run profit must also be nonnegative, otherwise the firm exits. Thus

$$py - LTC(y) \geq 0 \Longrightarrow p \geq \frac{LTC(y)}{y} = LAC(y).$$

Under the analogous convex-cost assumptions, the positive-output part of the competitive firm's long-run supply is the rising portion of long-run marginal cost at or above $\min LAC$. At the minimum-average-cost price, producing at efficient scale and exiting both yield zero economic profit, so long-run supply is again set-valued.

Fact 17.1 If price falls below the minimum of LAC, the firm exits in the long run because no positive output can be produced without negative economic profit.

Lecture 18 — Thursday, November 12, 2015

Market Supply and Equilibrium

Once each competitive firm has a supply curve, industry supply is obtained by adding those curves horizontally. If firm i supplies $S_i(p)$ units at price p, then short-run industry supply is

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

where n is the fixed number of firms currently operating in the industry.

Definition 18.1 The **industry supply curve** gives the total quantity supplied by all firms in the market at each price.

If the firms are identical and each has supply curve $s(p)$, then aggregation simplifies to

$$S(p) = ns(p).$$

This means that adding firms shifts industry supply outward even when individual firms' technologies do not change.

In the short run the number of firms is fixed, so equilibrium is found where market demand intersects the short-run industry supply curve, as in Figure 38. At that equilibrium price some firms may earn positive economic profit, some may earn losses, and some may break even. Those outcomes do not immediately change the number of firms because entry and exit are long-run adjustments.

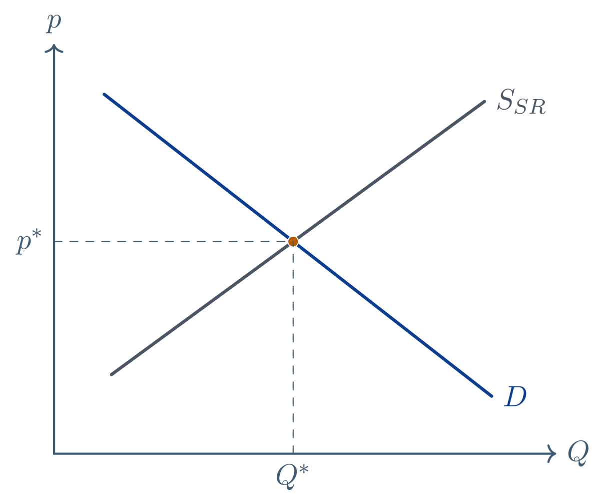

FIGURE 38

With free entry, identical firms, and no entry cost, a market price above minimum long-run average cost allows an entrant to earn positive profit at the efficient scale. Entry then shifts industry supply outward and drives price downward. Conversely, a price below minimum long-run average cost cannot support active firms indefinitely, so exit shifts industry supply inward and pushes price upward.

The long-run equilibrium is reached when no further entry or exit is profitable. In a constant-cost competitive industry, this means

$$p = \min LAC.$$

At that price each active firm earns zero economic profit, even though it still covers all of its explicit and opportunity costs. Figure 39 places the representative firm and industry views side by side.

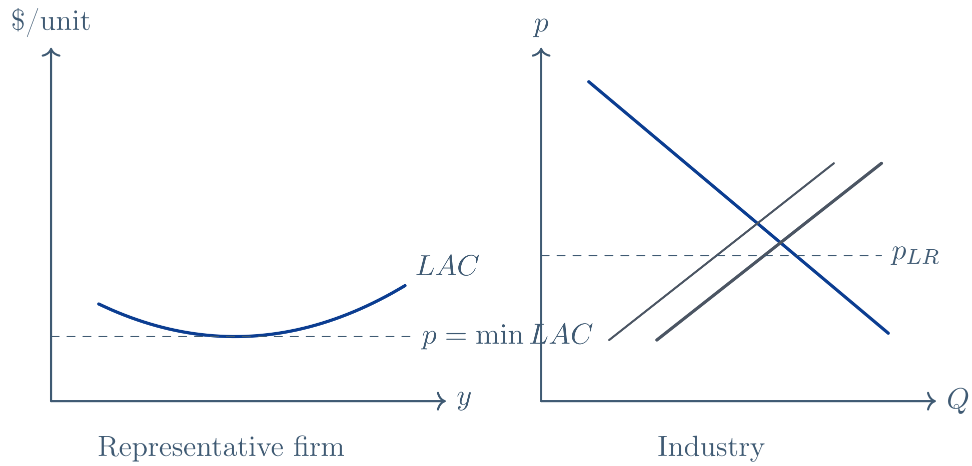

FIGURE 39

Note: The long-run number of firms is not chosen directly by a planner. It emerges through entry and exit: active firms must earn nonnegative economic profit, while another entrant must not be able to earn a positive profit.

Lecture 19 — Tuesday, November 17, 2015

Monopoly

A monopolized market has a single seller. Because the monopolist is the entire supplying side of the market, the firm's demand curve is the market demand curve itself. The firm can therefore influence price by choosing output. Such a market structure is called a **monopoly**.

Common sources of monopoly power include legal protection, patents, sole ownership of a key resource, cartel organization, and large economies of scale.

If inverse demand is $p(y)$ and total cost is $c(y)$, then profit is

$$\Pi(y) = p(y)y - c(y).$$

Differentiating with respect to output gives

$$\Pi'(y) = \frac{\mathrm{d}}{\mathrm{d}y}[p(y)y] - c'(y) = MR(y) - MC(y).$$

At an interior optimum,

$$MR(y^*) = MC(y^*).$$

This is the monopoly analogue of the competitive firm's $p = MC$ rule.

Because the monopolist must lower price to sell additional units, marginal revenue lies below demand. Differentiating revenue directly,

$$R(y) = p(y)y \qquad \text{and} \qquad MR(y) = p(y) + y\,p'(y).$$

Since $p'(y) < 0$, it follows that $MR(y) < p(y)$ whenever $y > 0$.

Example 19.1 If inverse demand is linear,

$$p(y) = a - by,$$

then

$$R(y) = (a - by)y = ay - by^2,$$

so

$$MR(y) = a - 2by.$$

Marginal revenue has the same intercept as demand but twice the slope, as shown in Figure 40.

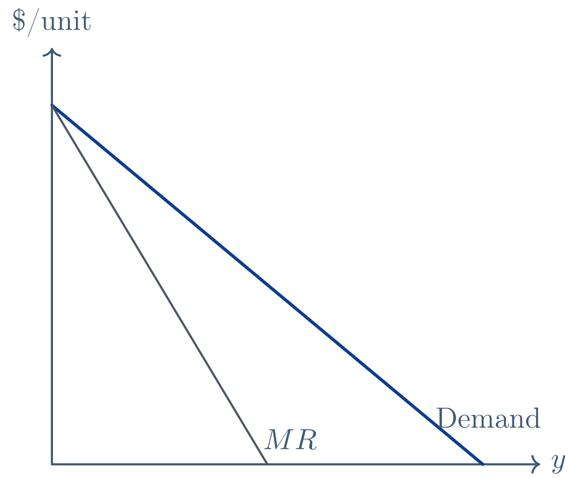

FIGURE 40

To make the monopoly calculation explicit, suppose $A > k$, $B > 0$, and

$$p(y) = A - By \qquad \text{and} \qquad c(y) = F + ky,$$

where k is constant marginal cost. Then

$$MR(y) = A - 2By \qquad \text{and} \qquad MC(y) = k.$$

The profit-maximizing output solves

$$A - 2By^* = k,$$

so

$$y^* = \frac{A-k}{2B}.$$

The monopoly price is

$$p^* = A - B\left(\frac{A-k}{2B}\right) = \frac{A+k}{2}.$$

The fixed cost F does not affect this interior first-order condition, but the firm serves the market only if the resulting operating profit is large enough to cover any avoidable fixed cost. Conditional on operation, the corresponding competitive output is $(A-k)/B$, twice the monopoly quantity. The monopolist therefore restricts output below the competitive level in order to keep price above marginal cost.

Monopoly Pricing and Elasticity

The elasticity form of marginal revenue is especially useful. If demand elasticity is

$$\varepsilon = \frac{p}{y}\frac{\mathrm{d}y}{\mathrm{d}p},$$

then

$$MR = p\left(1 + \frac{1}{\varepsilon}\right).$$

Since $\varepsilon < 0$ for a downward-sloping demand curve, marginal revenue is positive only when demand is elastic. Thus an interior profit maximum with strictly positive marginal cost lies on the elastic portion of the demand curve.

If marginal cost is constant at k, the condition $MR = MC$ becomes

$$p\left(1 + \frac{1}{\varepsilon}\right) = k.$$

Solving for price gives

$$p = \frac{k}{1 + \frac{1}{\varepsilon}} = \frac{k}{1 - \frac{1}{|\varepsilon|}}.$$

As demand becomes less elastic, the denominator gets smaller and the monopolist charges a larger markup over marginal cost.

Fact 19.2 If marginal cost is nonnegative, a monopoly never chooses a positive interior output for which demand is inelastic, because then $MR < 0$ and a small reduction in output would both raise revenue and weakly lower cost.

The efficient output level is the one at which the last unit produced still creates nonnegative gains from trade:

$$p(y^e) = MC(y^e).$$

The monopolist instead chooses y^* from $MR(y^*) = MC(y^*)$. Because $MR(y) < p(y)$, monopoly output is lower than the efficient output. Some mutually beneficial trades are not made, and the resulting lost gains from trade are **deadweight loss**, represented by the shaded wedge in Figure 41.

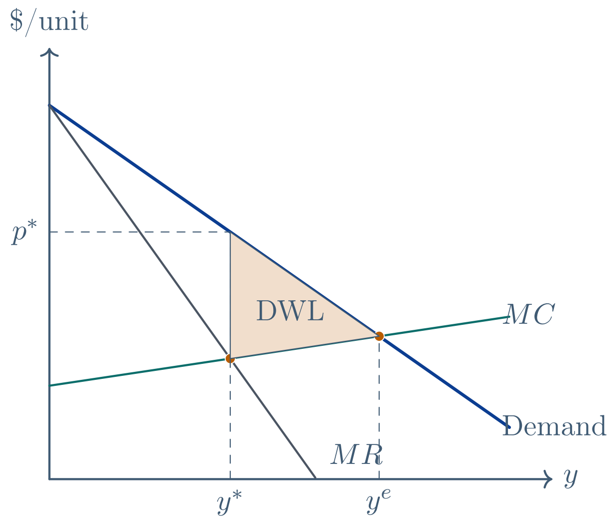

Figure 41

A **natural monopoly** arises when economies of scale are so strong that one firm can serve the entire market at lower average cost than multiple firms could. In that case average cost declines over the relevant range of output. A profit-maximizing natural monopolist still chooses output by $MR = MC$, so the usual monopoly inefficiency remains. Pricing at marginal cost would restore efficiency, but if average cost lies above marginal cost at the efficient output, then marginal-cost pricing would force the firm to operate at a loss. This is the basic regulatory problem created by natural monopoly.

Lecture 20 — Thursday, November 19, 2015

Price Discrimination

Uniform pricing means the monopolist charges the same price for every unit sold to every buyer. **Price discrimination** means the monopolist charges different effective prices across units or across buyers. **First-degree price discrimination** charges each unit at the buyer's exact willingness to pay. **Second-degree price discrimination** uses a pricing schedule, such as bulk discounts, that depends on quantity purchased. **Third-degree price discrimination** charges different uniform prices to different groups of buyers.

Under ideal first-degree price discrimination with divisible output, complete information, and take-it-or-leave-it offers, the monopolist sells each unit at exactly the height of the demand curve for that unit. The final unit sold is the one for which willingness to pay just equals marginal cost. Hence output is efficient:

$$p(y) = MC(y)$$

at the last unit traded. Under these ideal assumptions, the difference from competition is distributional rather than allocative: the monopolist can capture all the gains from trade, leaving consumers at their participation constraints, as the shaded region in Figure 42 indicates.

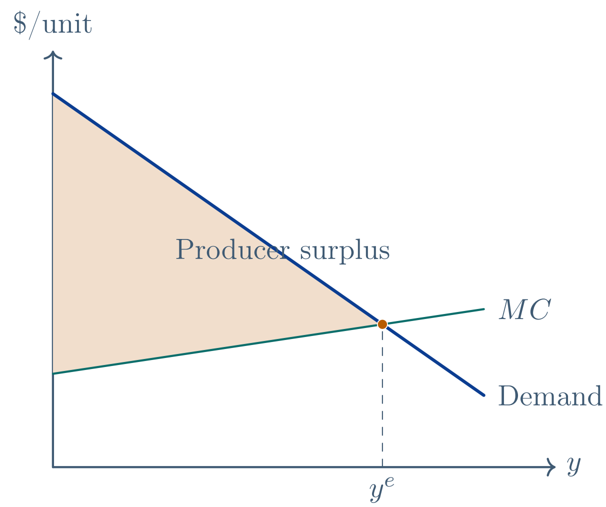

FIGURE 42

Suppose the monopolist serves two markets. Let y_1 and y_2 denote the quantities sold in those markets, and let inverse demand be $p_1(y_1)$ and $p_2(y_2)$. Profit is

$$\Pi(y_1, y_2) = p_1(y_1)y_1 + p_2(y_2)y_2 - c(y_1 + y_2).$$

Differentiating with respect to each market quantity yields

$$MR_1(y_1) = MC(y_1 + y_2) \qquad \text{and} \qquad MR_2(y_2) = MC(y_1 + y_2).$$

Thus the optimal allocation satisfies

$$MR_1(y_1^*) = MR_2(y_2^*) = MC(y_1^* + y_2^*).$$

The interpretation is that the monopolist first chooses total output and then allocates it across markets so that the marginal revenue from the last unit sold is the same in every market.

Using the elasticity formula

$$MR_i = p_i \left(1 + \frac{1}{\varepsilon_i}\right),$$

we see that, at an interior solution facing a common marginal cost, the higher price is charged in the market with the less elastic demand.

A **two-part tariff** charges an entry fee plus a per-unit price. If there is a single consumer and the monopolist knows the consumer's demand curve, then the efficient design is to set the per-unit price equal to marginal cost and then charge an entry fee equal to the consumer's resulting surplus. The per-unit charge ensures efficient usage, while the fixed fee transfers surplus to the monopolist.

Product Differentiation and Monopolistic Competition

Under **monopolistic competition**, each firm sells a differentiated product and therefore faces a downward-sloping demand curve. Profit maximization still requires

$$MR = MC.$$

But free entry drives profit to zero in the long run, so equilibrium also satisfies

$$p = AC.$$

These two conditions can hold simultaneously because $MR < p$. Under the standard model without offsetting externalities or other distortions, the result is zero profit with price still above marginal cost and therefore allocative inefficiency. Figure 43 shows the long-run tangency together with the gap between price and marginal cost.

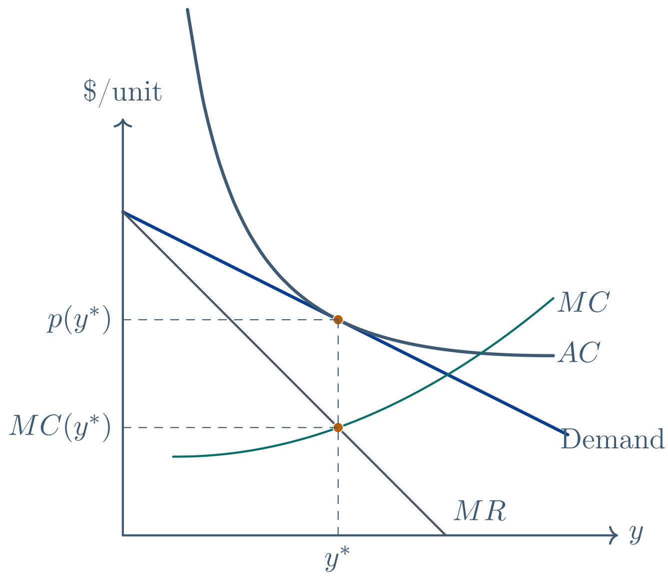

FIGURE 43

5. Oligopoly and Game Theory

Lecture 21 — Tuesday, November 24, 2015

Cournot Competition

An oligopoly contains only a few firms, and each firm's decisions affect the others' profits. In the Cournot model firms choose output simultaneously. If firm 1 produces y_1 and firm 2 produces y_2, market price is $p(y_1 + y_2)$, and profits are

$$\Pi_1(y_1; y_2) = p(y_1 + y_2)y_1 - c_1(y_1) \qquad \text{and} \qquad \Pi_2(y_2; y_1) = p(y_1 + y_2)y_2 - c_2(y_2).$$

Each firm chooses its best response to the other firm's output.

Suppose

$$p(y_T) = 60 - y_T, \qquad c_1(y_1) = y_1^2, \qquad \text{and} \qquad c_2(y_2) = 15y_2 + y_2^2.$$

Firm 1's profit is

$$\Pi_1(y_1; y_2) = (60 - y_1 - y_2)y_1 - y_1^2.$$

Differentiate with respect to y_1:

$$\frac{\partial \Pi_1}{\partial y_1} = 60 - 2y_1 - y_2 - 2y_1 = 60 - 4y_1 - y_2.$$

Setting this equal to zero gives firm 1's interior reaction function. Imposing nonnegative output, the full best response is

$$y_1 = R_1(y_2) = \max\left(0, 15 - \frac{1}{4}y_2\right).$$

Similarly, firm 2's profit is

$$\Pi_2(y_2; y_1) = (60 - y_1 - y_2)y_2 - 15y_2 - y_2^2,$$

so

$$\frac{\partial \Pi_2}{\partial y_2} = 60 - y_1 - 2y_2 - 15 - 2y_2 = 45 - y_1 - 4y_2.$$

Thus the full best response is

$$y_2 = R_2(y_1) = \max\left(0, \frac{45 - y_1}{4}\right).$$

The Cournot–Nash equilibrium is interior, so it solves the untruncated equations simultaneously:

$$y_1 = 15 - \frac{1}{4}\left(\frac{45 - y_1}{4}\right).$$

Multiplying by 16 gives

$$16y_1 = 240 - (45 - y_1) = 195 + y_1,$$

so

$$15y_1 = 195 \Longrightarrow y_1^* = 13.$$

Then

$$y_2^* = \frac{45 - 13}{4} = 8.$$

The two reaction curves intersect at this Cournot–Nash equilibrium in Figure 44.

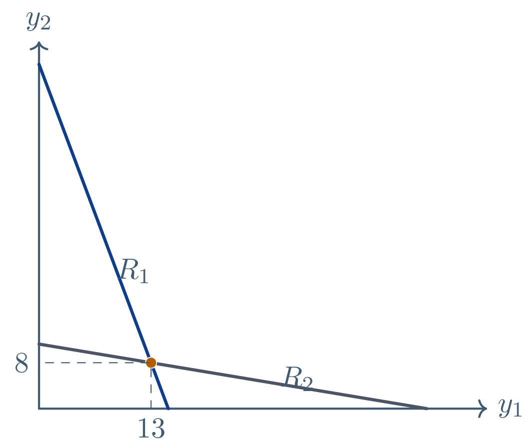

FIGURE 44

Collusion and Cartels

The firms can sometimes do better jointly than they do at the Cournot equilibrium. If they collude, they choose (y_1, y_2) to maximize total cartel profit:

$$\Pi^M(y_1, y_2) = p(y_1 + y_2)(y_1 + y_2) - c_1(y_1) - c_2(y_2).$$

This can make both firms better off than the noncooperative equilibrium. However, a cartel is intrinsically unstable because once the collusive output pair (y_1^M, y_2^M) is chosen, each firm is tempted to move to its own reaction curve and expand output while the other firm continues to restrict production.

Fact 21.1 Cartel agreements are fragile because the output restriction that raises joint profit also gives each member an incentive to cheat.

Lecture 22 — Thursday, November 26, 2015

Repeated play can sometimes sustain collusion. Using the repeated game example,

$$p(y_T) = 24 - y_T, \qquad c_1(y_1) = y_1^2, \qquad \text{and} \qquad c_2(y_2) = y_2^2.$$

Joint profit is

$$\Pi^M(y_1, y_2) = (24 - y_1 - y_2)(y_1 + y_2) - y_1^2 - y_2^2.$$

The first-order conditions are

$$24 - 4y_1 - 2y_2 = 0 \qquad \text{and} \qquad 24 - 2y_1 - 4y_2 = 0,$$

so the cartel optimum is

$$y_1^M = y_2^M = 4.$$

At that point total profit is

$$\Pi^M = (24 - 8)(8) - 4^2 - 4^2 = 128 - 16 - 16 = 96,$$

so an equal sharing rule gives each firm 48 per period.

If the firms instead play the Cournot game, the interior reaction functions are

$$y_1 = \frac{24 - y_2}{4} \qquad \text{and} \qquad y_2 = \frac{24 - y_1}{4},$$

which imply

$$y_1^* = y_2^* = \frac{24}{5} = 4.8.$$

The full best responses take the maximum of each displayed expression and zero, but the equilibrium lies in the interior. The resulting profit per firm is

$$\Pi_i^* = (24 - 9.6)(4.8) - (4.8)^2 = 46.08.$$

If firm 1 deviates while firm 2 keeps producing 4, then firm 1 solves

$$\max_{y_1} \; (24 - y_1 - 4)y_1 - y_1^2 = \max_{y_1} \; (20y_1 - 2y_1^2),$$

so the cheating output is $y_1 = 5$, and the one-period deviation profit is

$$\Pi_1^{ch} = 15 \cdot 5 - 5^2 = 50.$$

If the discount factor is $\delta = 1/(1+r)$ and cheating is punished forever by reversion to Cournot, then collusion is sustainable when

$$\frac{48}{1-\delta} \geq 50 + \frac{\delta \cdot 46.08}{1-\delta}.$$

Solving gives

$$\delta \geq \frac{2}{3.92} \approx 0.5102.$$

Thus firms must care enough about future profits for collusion to survive.

Stackelberg, Bertrand, and Price Leadership

If one firm moves first and the other responds, the game becomes a Stackelberg quantity game. In the earlier example with

$$p = 60 - y_T, \qquad c_1(y_1) = y_1^2, \qquad \text{and} \qquad c_2(y_2) = 15y_2 + y_2^2,$$

suppose firm 1 is the leader and firm 2 is the follower. Over the range relevant to the interior solution, the follower's reaction function is

$$R_2(y_1) = \frac{45 - y_1}{4}.$$

Substituting this into the leader's profit gives

$$\Pi_1^S(y_1) = \left(60 - y_1 - \frac{45 - y_1}{4}\right) y_1 - y_1^2 = \frac{195}{4} y_1 - \frac{7}{4} y_1^2.$$

The first-order condition is

$$\frac{195}{4} - \frac{7}{2} y_1 = 0,$$

so

$$y_1^S = \frac{195}{14} \approx 13.93 \qquad \text{and} \qquad y_2^S = \frac{45 - y_1^S}{4} = \frac{435}{56} \approx 7.77.$$

The leader produces more than in Cournot and the follower produces less.

If firms compete in prices instead of quantities and choose prices simultaneously, the model becomes Bertrand competition. With at least two identical firms, homogeneous output, positive market demand at constant marginal cost c, no capacity constraints, and consumers who split demand among firms tied at the lowest price, the unique Nash equilibrium is

$$p = c,$$

because any common price above c can be profitably undercut by a rival.

If one large firm sets price first and many competitive fringe firms follow, the leader faces the nonnegative residual demand

$$L(p) = \max\left(D(p) - Y_f(p), 0\right),$$

where $Y_f(p)$ is the followers' aggregate supply. The leader then chooses price to maximize profit on that residual demand.

Lecture 23 — Tuesday, December 01, 2015

Games and Strategies

Game theory studies situations in which each agent's payoff depends on what other agents do. A game consists of a set of players, a set of strategies for each player, and a payoff assigned to each possible strategy profile.

Definition 23.1 A **Nash equilibrium** is a strategy profile in which each player's strategy is a best response to the strategies chosen by the others.

Consider the payoff matrix

	L	R
U	$(3, 9)$	$(1, 8)$
D	$(0, 0)$	$(2, 1)$

If player B chooses R, then player A prefers D to U because $2 > 1$. So (U, R) cannot be a Nash equilibrium. If B chooses L, then A prefers U to D because $3 > 0$. Meanwhile, if A chooses U, then B prefers L to R because $9 > 8$, and if A chooses D, then B prefers R to L because $1 > 0$. Therefore the game has two pure-strategy Nash equilibria:

$$(U, L) \qquad \text{and} \qquad (D, R).$$

Prisoner's Dilemma

The prisoner's dilemma matrix is

	S	C
S	$(-5, -5)$	$(-30, -1)$
C	$(-1, -30)$	$(-10, -10)$

For each player, confessing strictly dominates silence. If the other prisoner stays silent, confessing yields -1 instead of -5. If the other prisoner confesses, confessing yields -10 instead of -30. Hence the unique Nash equilibrium is

$$(C, C),$$

even though both players would prefer (S, S). The example shows that Nash equilibrium need not be Pareto efficient.

Sequential Play and Mixed Strategies

When one player moves first and the other observes that move before responding, the game is sequential rather than simultaneous. In the first example, if A moves first, then B's best reply to U is L and B's best reply to D is R. Anticipating this, A compares the resulting payoffs:

$$\text{if } A \text{ chooses } U, \text{ payoff is } 3 \qquad \text{and} \qquad \text{if } A \text{ chooses } D, \text{ payoff is } 2.$$

So backward induction selects (U, L) as the subgame-perfect equilibrium outcome of the sequential game represented in Figure 45.

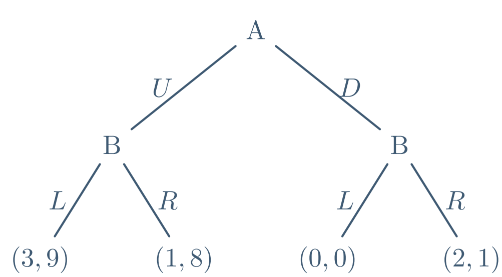

FIGURE 45

Not every game has a pure-strategy Nash equilibrium. Consider the game

	L	R
U	$(1,2)$	$(0,4)$
D	$(0,5)$	$(3,2)$

for which none of the four pure action profiles is stable.

Let player A choose U with probability p_U, and let player B choose L with probability p_L. For A to be willing to mix, B must choose p_L so that A is indifferent between U and D. A's expected payoff from U is

$$1 \cdot p_L + 0 \cdot (1 - p_L) = p_L,$$

and A's expected payoff from D is

$$0 \cdot p_L + 3 \cdot (1 - p_L) = 3(1 - p_L).$$

Setting these equal gives

$$p_L = 3(1 - p_L) \Longrightarrow p_L = \frac{3}{4}.$$

Similarly, for B to be willing to mix, A must choose p_U so that B is indifferent between L and R. B's expected payoff from L is

$$2p_U + 5(1 - p_U),$$

while B's expected payoff from R is

$$4p_U + 2(1 - p_U).$$

Setting them equal yields

$$2p_U + 5(1 - p_U) = 4p_U + 2(1 - p_U) \Longrightarrow p_U = \frac{3}{5}.$$

Therefore the mixed-strategy Nash equilibrium is

$$\text{A mixes } \left(\frac{3}{5}, \frac{2}{5}\right) \qquad \text{and} \qquad \text{B mixes } \left(\frac{3}{4}, \frac{1}{4}\right).$$

The resulting outcome probabilities are

$$\mathbb{P}(U, L) = \frac{3}{5} \cdot \frac{3}{4} = \frac{9}{20} \qquad \text{and} \qquad \mathbb{P}(U, R) = \frac{3}{5} \cdot \frac{1}{4} = \frac{3}{20},$$

$$\mathbb{P}(D, L) = \frac{2}{5} \cdot \frac{3}{4} = \frac{6}{20} \qquad \text{and} \qquad \mathbb{P}(D, R) = \frac{2}{5} \cdot \frac{1}{4} = \frac{2}{20}.$$

Fact 23.2 Every finite game has at least one Nash equilibrium in mixed strategies. Thus if there is no equilibrium in pure strategies, there is still at least one equilibrium involving randomization.

Appendix: Lecture and Tutorial Index