Chapter 16: General Equilibrium and Efficiency

So far, we have considered the microeconomics of one particular market or one particular firm or one particular consumer. Does microeconomic theory have anything to say about the economy as a whole? Yes. In this chapter, we will study the concept of general equilibrium, which occurs when all markets in the economy are in equilibrium. We will also introduce some basic concepts of economic efficiency.

To keep things simple, we will consider an economy with only two consumers and two markets, although the results presented in this chapter hold for any number of markets, firms and consumers. In our imaginary economy, Bob and Teresa each wish to buy coffee and donuts.

Efficiency in Exchange

An Edgeworth Box depicts the distribution of the two goods between the two consumers.

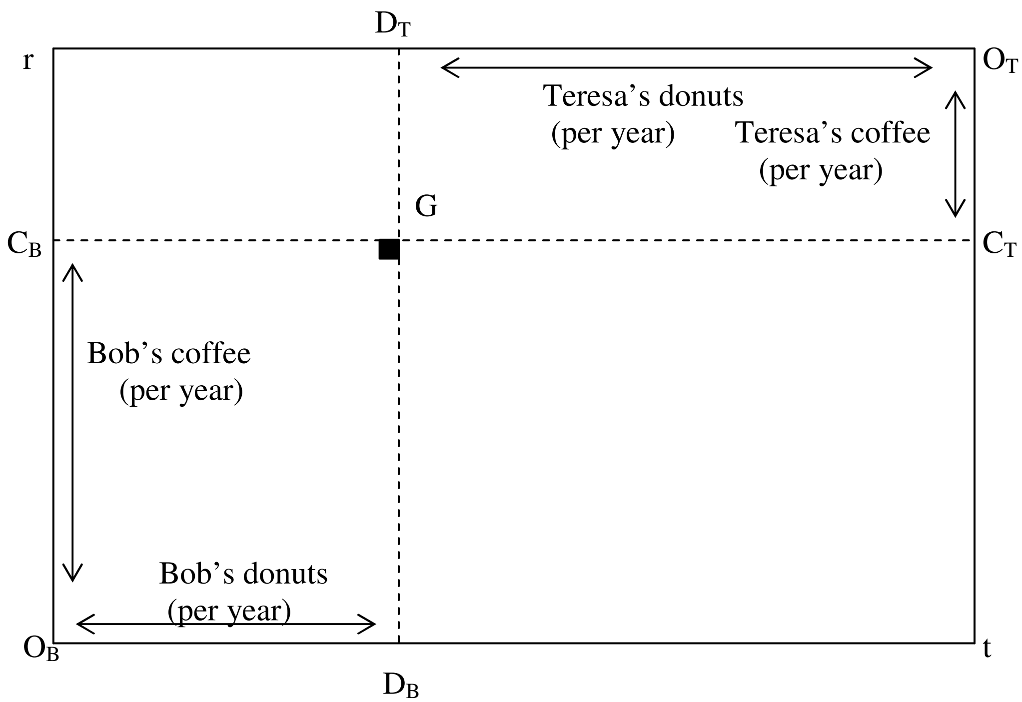

r: total amount of coffee available
t: total amount of donuts available
Bottom axis: Bob's consumption of donuts
Left axis: Bob's consumption of coffee
Top axis: Teresa's consumption of donuts
Right axis: Teresa's consumption of coffee

Point G: Bob consumes D_B donuts and C_B coffees while Teresa consumes $(t – D_B) = D_T$ donuts and $(r – C_B) = C_T$ coffees.

We can also graph indifference curves for Bob and Teresa in the Edgeworth Box.

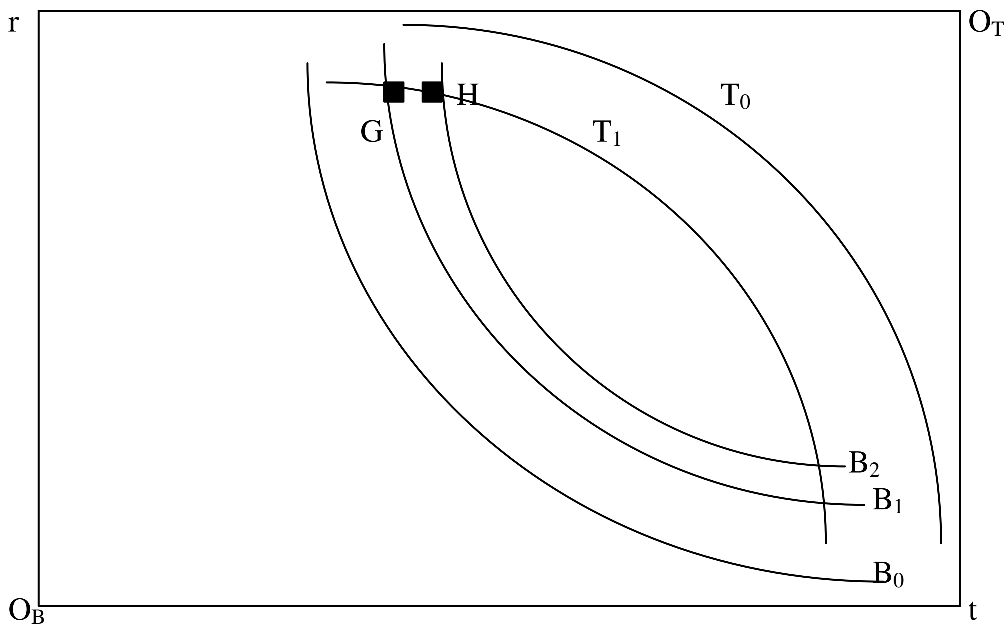

B_0, B_1 and B_2 denote indifference curves for Bob. As usual, they are convex, and Bob's utility increases as his position moves to the northeast. T_0 and T_1 denote indifference curves for Teresa. She becomes better off as her position moves toward the southwest, and her indifference curves are also convex with respect to her origin.

An allocation is <u>Pareto efficient</u> if no one can be made better off without making someone else worse off. Can point G be a Pareto efficient allocation of goods? No, since Bob can be made better off by moving to point H without making Teresa any worse off.

In fact, if we start at point G, all the possible allocations lying in or on the lens formed by the intersection of B_1 and T_1 represent Pareto improvements over point G (including point H). However, H is not Pareto efficient either.

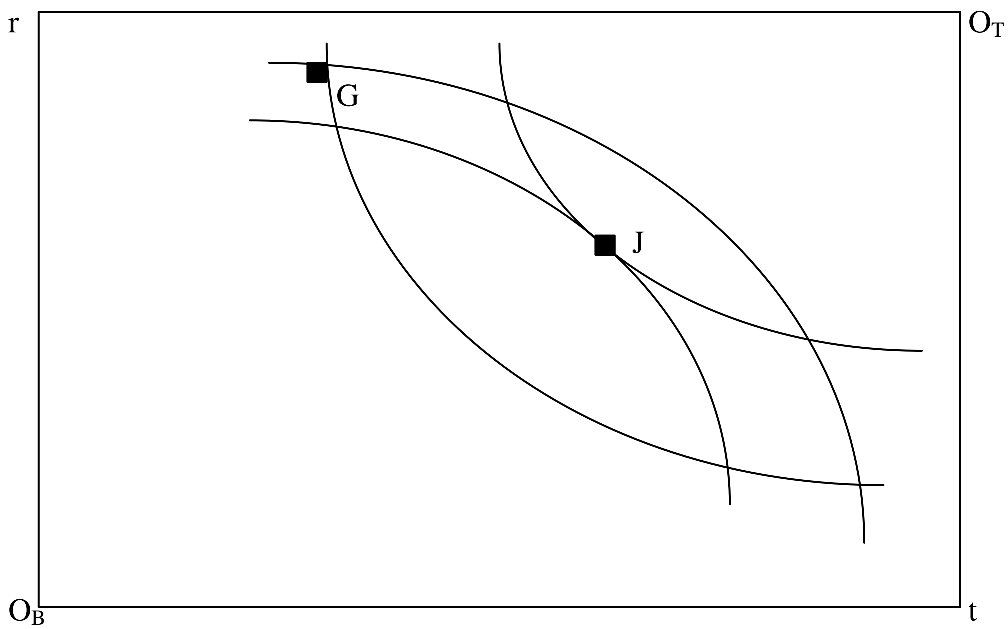

Starting at point G, we observe a Pareto efficient allocation at point J. Here, the allocation of both goods has been rearranged to make both Bob and Teresa better off. Note that their indifference curves are tangent at this point, and it is impossible to move from point J without making at least one party worse off.

Recall that the (absolute value of the) slope of the indifference curve is simply the <u>marginal rate of substitution of donuts for coffee</u> ($MRS_{D,C}$): the number of coffees the consumer would be willing to give up to get one more donut.

$$MRS_{D,C} = (\Delta C/ \Delta D) \text{ or } - (\partial U/\partial D)/(\partial U\partial /C)$$

Since the slopes of the indifference curves passing through a Pareto efficient point must be equal, we see that one necessary condition for Pareto efficiency is that (the <u>allocative efficiency condition</u>)

$$MRS_{D,C}^{B} = MRS_{D,C}^{T}$$

Alternative explanation: suppose the this condition does not hold. Let $MRS_{D,C}^{B} = 3$ and $MRS_{D,C}^{T} = 1$. This means that if Bob gives up one donut, he must receive three coffees to be just as well off, while if Teresa gives up one donut, she needs only one coffee to be kept as well off. If this is really the case, there is a mutually beneficial trade to be made! Teresa trades one donut to Bob in exchange for two coffees: Teresa gets more coffee than she needed to remain indifferent, so she is better off, and Bob gives up fewer coffees than he would have had to if he were to remain indifferent, so he is made better off. Since both parties are gaining from this exchange, the initial allocation could not have been Pareto efficient.

In fact, there is a continuum of Pareto efficient allocations of goods., as shown below:

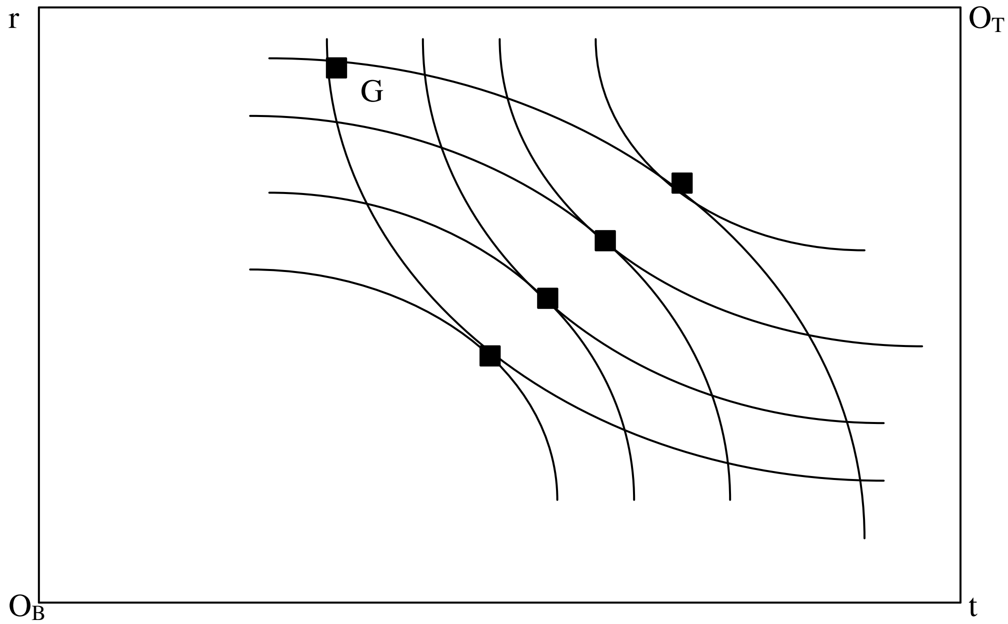

It is easy to verify that each of the points other than the initial allocation G in the graph above is Pareto efficient. We can derive the contract curve by repeating this procedure for all starting points.

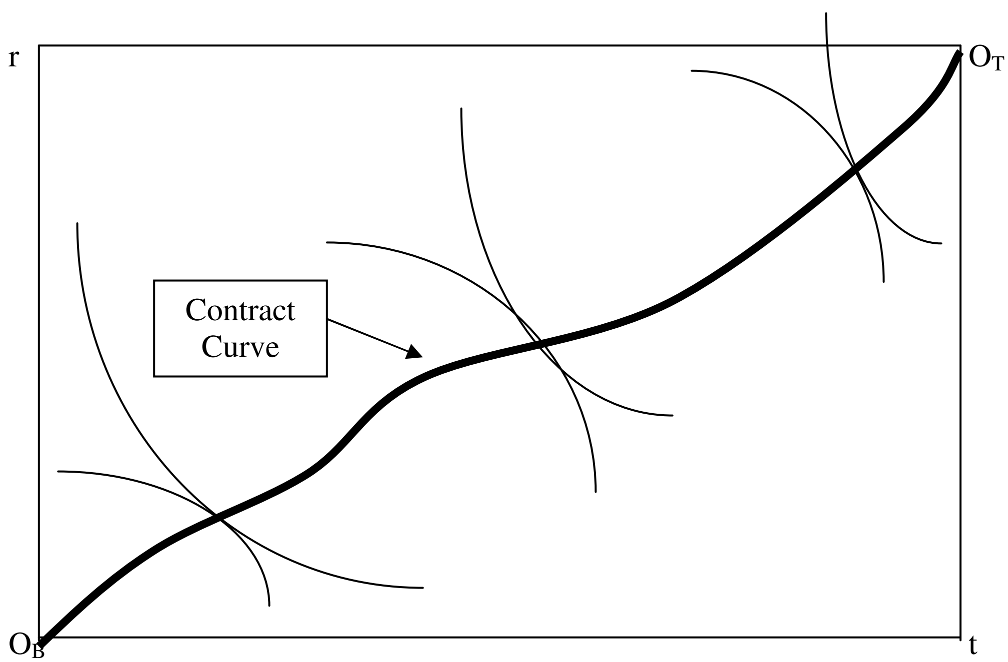

The contract curve shows the locus of all Pareto efficient allocations.

Efficiency in Production

Our analysis of allocative efficiency assumed that the number of donuts and the amount of coffee were both fixed. In fact, society can choose how many donuts and how much coffee to produce, given existing productive resources. The maximum amount of coffee that can be produced given a particular level of donut production is given by the production possibilities frontier (PPF).

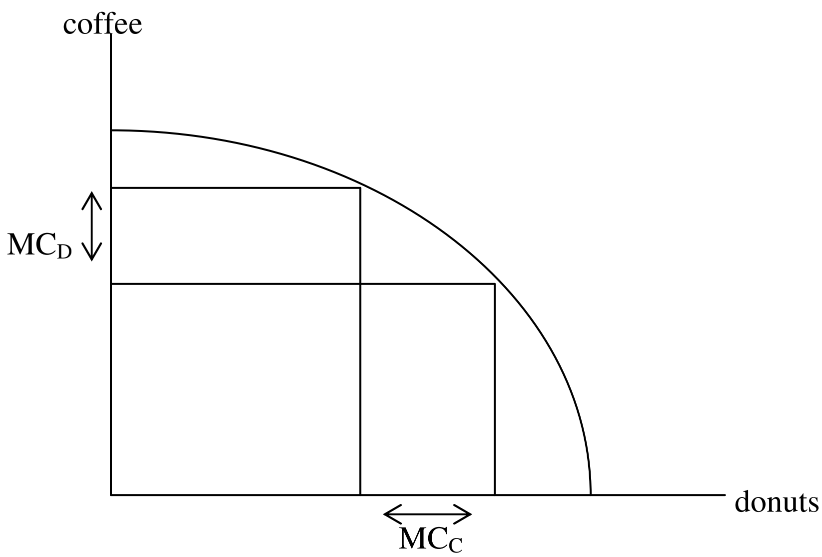

The absolute value of the slope of the PPF is the marginal rate of transformation of donuts for coffee, $MRT_{D,C}$ (it indicates how much coffee has to be given up to obtain an additional donut). Notice that the $MRT_{D,C}$ can be written in terms of the relevant marginal costs as

$$MRT_{D,C} = MC_D / MC_C$$

because the MRT is defined as the (absolute value) slope of the PPF and the equation given is just the definition of slope.

We all know (from Econ 201 or 202) that a Pareto efficient production of goods must lie on the PPF. However, such an allocation must also satisfy the productive efficiency condition that

$$MRT_{D,C} = MRS_{D,C}^B = MRS_{D,C}^T$$

Why? Well, we already know why the last equality, the allocative efficiency condition, holds, so let's focus on the first equality. Specifically, lets suppose that the first equality does *not* hold, and that $MRT_{D,C} = 4/5$ while $MRS_{D,C}^{B} = 2/5$. In other words, society must give up five donuts to produce four more coffees; conversely, it must give up four coffees to produce five more donuts. However, Bob is willing to give up five donuts to get just two more coffees. If society actually transformed five donuts into four coffees, Bob could be made better off, so our initial production level could not possibly have been efficient.

It will be convenient to rewrite the productive efficiency condition as:

$$(MC_D / MC_C) = MRS_{D,C}^{B} = MRS_{D,C}^{T}$$

The economy is producing and allocating its goods in a Pareto efficient manner if and only if this condition holds.

The First Fundamental Theorem of Welfare Economics

This theorem basically states that if consumers and producers are price-takers and have complete information about prices and production, then the market equilibrium will be Pareto efficient.

The intuitive explanation for this proof is really pretty simple. We just need to show that a combination of well-informed, profit maximizing, perfectly competitive firms and well-informed, utility maximizing, price-taking consumers produce a market equilibrium which satisfies the productive efficiency condition above. We can do this in two parts:

1) $MRS_{D,C}^{B} = MRS_{D,C}^{T}$ because utility maximizing consumers will always consume at a point where $MRS_{D,C} = P_D / P_C$. Why? Because if this condition weren't satisfied, the consumer could make herself better off by selling one good and buying another. For instance, suppose that this condition doesn't hold for Teresa; her $MRS_{D,C} = 2$ but $P_D = P_C = \$1$. This means that if she gives up two coffees, she will be just as well off if she gets one donut in return. But at these prices, she can give up (sell) two coffees and get two donuts in return! Hence, she could not possibly have been maximizing her utility.

We learned in chapter 3 that utility-maximizing consumers equate their MRS to the price ratio. Since all consumers consume so that $MRS_{D,C} = P_D / P_C$, and all consumers face the same prices (they're price-takers), it must be the case that $MRS_{D,C}^{B} = MRS_{D,C}^{T}$.

2) $(MC_D / MC_C) = MRS_{D,C}^{B}$ because under perfect competition, $P_D = MC_D$ and $P_C = MC_C$, so the marginal cost ratio is equal to the price ratio which we just saw is equal to the MRS.

The first fundamental theorem is probably the most important result in the entire discipline of economics. Note that its proof is actually quite elegant. It follows directly from the assumptions that $P_D = MC_D$ and $P_C = MC_C$; the assumption of perfect competition is the link that combines efficiency in exchange with efficiency in production.

A perfectly competitive market will automatically produce a Pareto efficient allocation of resources. So I guess there's nothing for the government to do except get out of the way, right? Wrong. Actually, one could justify government intervention in the economy, in spite of the first fundamental theorem. This is because the first fundamental theorem simply guarantees that the economy will reach a Pareto efficient point; it says nothing about . . .

Social Welfare

Lets return to the contract curve for Bob and Teresa.

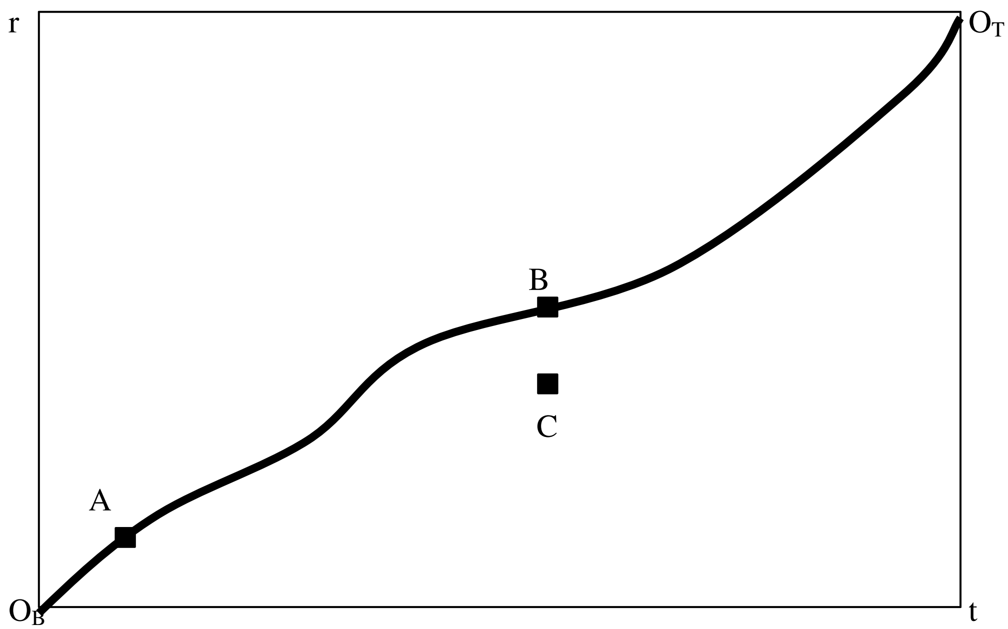

Points A and B are both on the contract curve, so both are Pareto efficient. However, are both points equally desirable? Many people would say no, because at point A Teresa has a large amount of both goods and enjoys a high utility while Bob consumes little of either good and receives little utility. It can be argued, then, that point A is less equitable than point B. (Obviously, such a value judgement is controversial and reasonable people could disagree on whether point A was more or less equitable than point B).

Furthermore, consider point C. Point C isn't even Pareto efficient, yet many would argue that point C is preferable to point A because it seems "fairer" than point A.

Starting at O_B and moving to the northeast along the contract curve, we see that Bob's utility is rising while Teresa's is falling. We use the contract curve to derive a utility possibilities frontier, which is analogous to a production possibilities frontier. The main differences are that the "goods" produced are Bob's and Teresa's utility, and that the UPF need not be strictly "bowed out" since we have no way of knowing whether utility tradeoffs involve diminishing marginal returns. The following graph shows a hypothetical UPF.

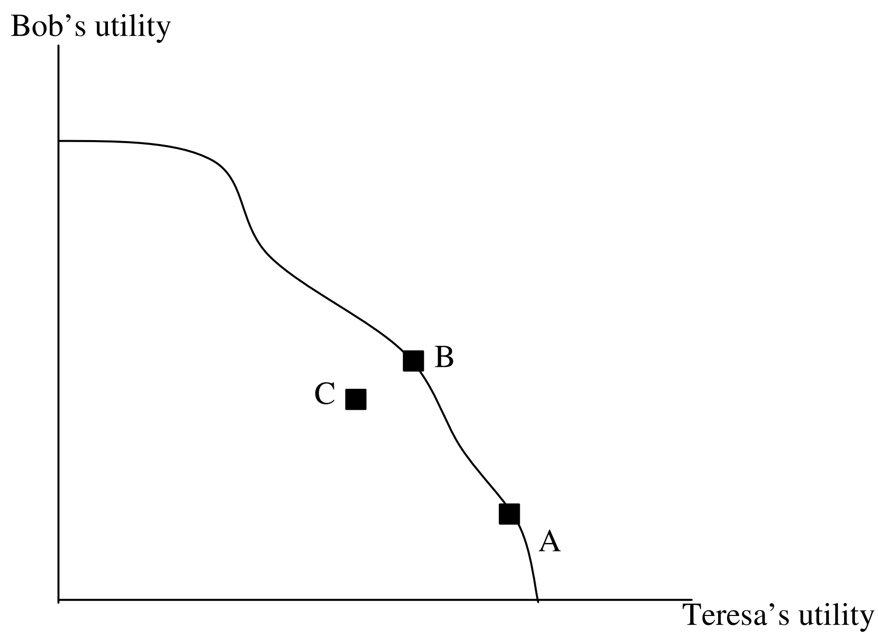

Points A, B and C correspond to the points from the contract curve. Again, notice that C is inefficient since both Bob and Teresa could be made better off by a move to B, for example. However, Bob prefers certainly prefers C to A!

Since all the points on the UPF are efficient, how can we judge which one is "best?" Suppose we can write a <u>social welfare function</u> that represents society's views regarding tradeoffs between Bob's and Teresa's utility. Basically, such a function would behave like an individual's utility function where the "goods" involved were simply the two consumers' utility levels. Thus, we can write the social welfare function as

$$W(U^B, U^T)$$

where both partial derivatives are positive (society wants to make both better off) and indifference curves from this function are convex to the origin. We can graph society's indifference curves on the UPF (treating it as a budget constraint) to determine the optimal allocation of utility, which in turn tells us the optimal point on the contract curve.

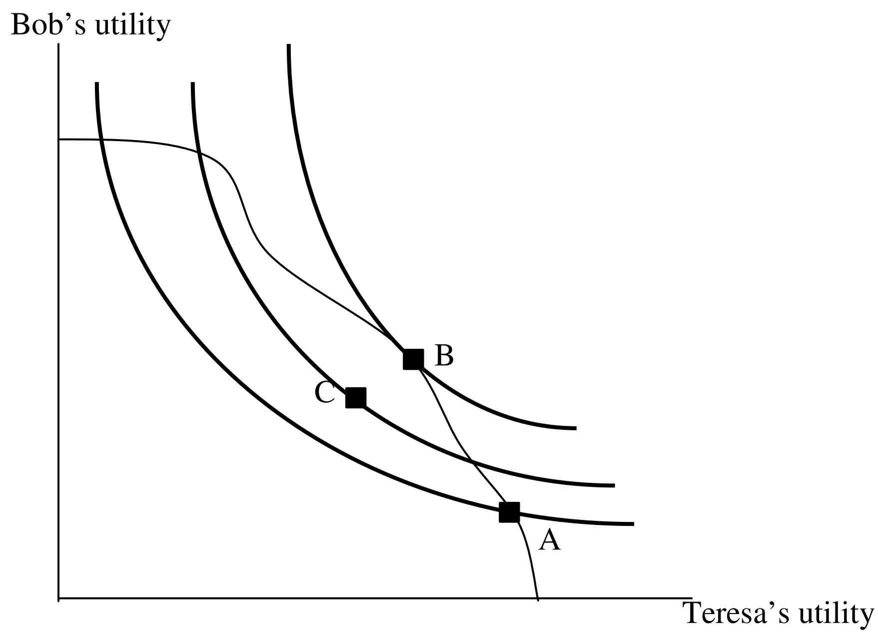

Based on the social welfare function described by these indifference curves, society prefers point C to A, and point B indicates the optimal allocation of goods.

One role for the government, then, is to ensure that society reaches its optimal outcome. Barring that, the government can decide to accept some inefficiency to keep the economy away from point A to reach a "superior" outcome like the inefficient point C.

It's worth noting that a result called, appropriately, the <u>second fundamental theorem of welfare economics</u> states that any point on the UPF can be reached given a certain initial allocation of goods. In other words, the government can, in principle, redistribute goods to enhance social welfare.

This discussion has assumed that we can count on the first fundamental theorem actually to produce a Pareto efficient equilibrium. However, that is not necessarily the case when we are confronted by . . .

Market Failure

If the conditions assumed by the FFT are not met, it would not be surprising if the market failed to achieve Pareto efficiency. In fact, there are a variety of commonly occurring reasons why the market may not produce efficient outcomes on its own.

1) Market Power

The FFT assumed that consumers and firms were both price takers. In real life, firms seldom face perfect competition, meaning that price will often exceed MC. If $P \neq MC$, the argument for point #2 of the FFT falls apart.

2) Imperfect Information

If consumers are not well informed about prices, it would be a lucky coincidence if each of them managed to set his MRS equal to the price ratio (or MB = P). Since all consumers' marginal rates of substitution would probably not be equal, the resulting allocation of goods would not lie on the contract curve and would therefore not be Pareto efficient.

3) Nonexistence of Markets

We can't expect "the market" to produce the socially efficient amount of a good and allocate it efficiently if a market for that good doesn't exist! The lack of a market primarily occurs in two cases:

Public Goods: goods that are not diminished when consumed by an individual.
e.g.) light house: the fact that one ship sees the lighthouse doesn't prevent another ship from seeing it, too ⇒ incentive for individual consumers to hide true preferences (free rider problem) ⇒ too little production.

Externalities: one person's behavior imposes a cost on another.
e.g.) air pollution produced by 3M imposes a cost on surrounding homeowners, but there is no "market" for pollution ⇒ no price ⇒ too much produced

In each of these cases, the government can attempt to improve the market's performance. Antitrust regulation attempts to prevent the buildup or exploitation of market power. A variety of regulations (like the Truth in Lending Act) attempt to keep consumers well informed.