to get the quantities they want, and there are sellers who realize that they can boost prices and still sell all they want to sell. Similarly, in the everyday world excess supply means surpluses. When there are surpluses, prices tend to go down. (This phenomenon might be masked in an economy in which there is persistent inflation which hides price declines. But even if all absolute prices are rising, the goods for which there are surpluses will fall in price relative to other goods. And, only relative prices matter in our exchange economy model, since doubling all prices has no real effect on any consumer's budget equation.) When there are surpluses, sellers have unplanned and unwanted inventories, so they have "special sales." Buyers see extra stocks of merchandise, so they try to bargain with sellers, again pushing prices down.

Let's now incorporate these natural price movements in our analysis. Good 1 is in excess demand, and good 2 is in excess supply. Therefore, p_1 will tend to rise, and p_2 will tend to drop, or, at any rate, p_1 will tend to rise relative to p_2. Therefore, p_1/p_2 will go up, and the budget line will get steeper. This shift will continue until supply equals demand for both goods, that is, until the desired bundles $\hat{x}_1$ and $\hat{x}_2$ of the two traders coincide.

An allocation $\hat{x} = (\hat{x}_1, \hat{x}_2)$, where the traders are consistently maximizing their respective utilities, subject to their budget constraints, is called a competitive equilibrium allocation. It is competitive because both traders are acting as price takers, which is presumably how people act in large competitive markets with lots of traders, and it is an equilibrium because their consumption plans are consistent; there is no excess supply or excess demand for either good; and there is no reason for prices to change further. A competitive equilibrium allocation is illustrated in Figure 3.4.

The figure suggests a very important result. Since the two traders' indifference curves are tangent to the same budget line at $\hat{x}$, they are also tangent to each other. Consequently, $\hat{x}$ is Pareto optimal. It is also clear that each person likes $\hat{x}$ at least as much (in fact more) than he likes ω. Consequently, $\hat{x}$ is in the core.

3. Competitive Equilibrium in an Exchange Economy: Formal Preliminaries

We formally define a competitive equilibrium for an economy of self-interested traders as follows. Suppose there are n people and m goods.

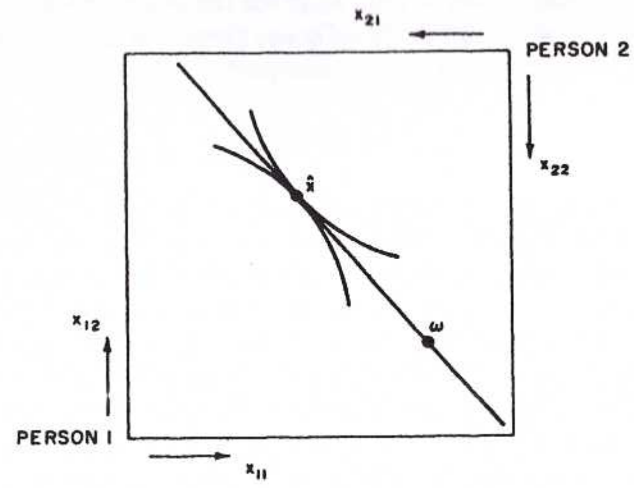

Figure 3.4.

The allocation $\hat{x}$, along with the price vector $p = (p_1, p_2, \ldots, p_m)$, form a *competitive equilibrium* if, for all i,

$\hat{x}_i$ maximizes $u_i(x_i)$ subject to the budget constraint

$$p_1 x_{i1} + p_2 x_{i2} + \ldots + p_m x_{im} \leq p_1 \omega_{i1} + p_2 \omega_{i2} + \ldots + p_m \omega_{im}.$$

Let's make three remarks about the definition. First, $p_1 x_{i1} + p_2 x_{i2} + \ldots + p_m x_{im}$ is the value of the bundle x_i, while $p_1 \omega_{i1} + p_2 \omega_{i2} + \ldots + p_m \omega_{im}$ is the value of the bundle ω_i. The budget inequality then says that, when valued at the prices given by p, the value of what person i consumes must be less than or equal to the value of what he starts with. Second, since $\hat{x}$ is an allocation, we must by definition have supply equal to demand for every good:

$$\sum_{i=1}^{n} \hat{x}_{ij} = \sum_{i=1}^{n} \omega_{ij}, \text{ for all goods } j.$$

Third, a competitive equilibrium depends on ω, the initial allocation. If you change ω, you change the competitive equilibria.

At this point we can introduce an important notational simplification. We define $p \cdot x_i$, the *dot product* of the vector p and the vector x_i as follows:

$$p \cdot x_i = \sum_{j=1}^{m} p_j x_{ij} = p_1 x_{i1} + p_2 x_{i2} + \ldots + p_m x_{im}.$$

Thus $p \cdot x_i$ is the value of the bundle x_i, given the price vector p. Similarly, $p \cdot \omega_i$ is the value of the initial bundle ω_i, given the price vector p, and person i's budget constraint can be compactly rewritten

$$p \cdot x_i \leq p \cdot \omega_i.$$

Now for the theorems below we need a formalization of the idea that a person prefers more goods to less. Suppose person i is self-interested, so his utility depends only on his own consumption. Let x_i be any bundle of goods for him. If i would prefer to x_i any bundle y_i that includes more of every good, then i's utility function is said to be monotonic. That is, u_i is *monotonic* if $y_{ij} > x_{ij}$ for all j implies $u_i(y_i) > u_i(x_i)$.

For one of the theorems below we need two simple preliminary observations, in addition to the definitions and notation developed so far. The first observation says if i prefers some other bundle to the bundle he chooses to buy, that other bundle must be too expensive for him: it must be more valuable than the bundle he starts with. This makes splendid sense, for if the preferred bundle weren't too expensive for him, he would buy it, since his goal, after all, is to maximize his utility. The second observation says that if i likes some other bundle at least as much as the bundle he chooses to buy, that other bundle can't be less valuable than the bundle he starts with. This depends on the prior assumption of monotonicity, but it too makes fine sense. For if the other bundle were actually less valuable than the bundle he starts with, he could afford a bundle that is slightly more expensive than the other bundle, one that contains slightly more of every good. He would prefer this third bundle to the bundle he chooses to buy. So he would buy it instead, since his goal is, again, to maximize his utility.

Let's turn to the formal statements of these observations.

Observation 1. Let $(\hat{x}, p)$ be a competitive equilibrium. If $u_i(y_i) > u_i(\hat{x}_i)$ for some bundles y_i, then

$$p \cdot y_i > p \cdot \omega_i.$$

Proof. If this were not the case, then $\hat{x}_i$ would not be the bundle that maximizes the utility function u_i subject to i's budget constraint, which would contradict the assumption that $(\hat{x}, p)$ is a competitive equilibrium.

Q.E.D.

Observation 2.. Let $(\hat{x}, p)$ be a competitive equilibrium. Assume i has a monotonic utility function. If $u_i(y_i) \geq u_i(\hat{x}_i)$ for some bundle y_i then

$$p \cdot y_i \geq p \cdot \omega_i.$$

Proof. Suppose to the contrary that $p \cdot y_i < p \cdot \omega_i$. Define a bundle z_i by adding a small quantity ϵ of every good to the bundle y_i:

$$z_{ij} = y_{ij} + \epsilon \text{ for all } j.$$

Choose ϵ small enough so that

$$\begin{aligned} p \cdot z_i &= p_1 y_{i1} + p_1\epsilon + p_2 y_{i2} + p_2\epsilon + \ldots + p_m y_{im} + p_m\epsilon \\ &= p \cdot y_i + \epsilon(p_1 + p_2 + \ldots p_m) \\ &\leq p \cdot \omega_i. \end{aligned}$$

This can be done since, by assumption, $p \cdot y_i < p \cdot \omega_i$. By the monotonicity assumption, person i prefers the bundle z_i to the bundle y_i. That is,

$$u_i(z_i) > u_i(y_i) \geq u_i(\hat{x}_i).$$

Therefore, $u_i(z_i) > u_i(\hat{x}_i)$, while $p \cdot z_i \leq p \cdot \omega_i$, which again contradicts the assumption that $(\hat{x}, p)$ is a competitive equilibrium.

Q.E.D.

This completes our formal preliminaries to the first main result.

4. The First Fundamental Theorem of Welfare Economics

Since around the time of Adam Smith (*Wealth of Nations*, 1776), economists have argued that the price system has a special virtue that other allocative systems might not share. The price system induces self-interested individuals, independently maximizing their private well being, to bring the economy to a socially optimal state. The invisible hand of competition transforms private greed into public welfare. Why should this be the case? In the complicated real world, people see prices and adjust their behavior accordingly, in order to maximize private welfare or private profit. High prices steer people away from wasteful uses of resources and technology, low prices attract them to effective production processes, effective technologies, and effective ways to satisfy wants. With prices giving the signals, wants are fulfilled in an economical way.

Moreover, the system works automatically. When supply is greater than demand, prices fall; when demand is greater than supply, prices rise. There is no need for an expensive or cumbersome centralized bureaucracy to tell us what to do; the hand of competition will lead in the right direction, and without force or coercion. Respecting the individual's property rights, the market system is based on the voluntary participation of people. Each person's wants and the resources that he has to satisfy them determine his final position in the system.

That's not all, though. Smith famously argued that the market system maximizes wealth. But 19th and 20th century economists, including Francis Ysidro Edgeworth, Martin Shubik, Herbert Scarf Gerard Debreu and Robert Aumann, proved that the market does more: it produces an outcome that is in the core. As we will see, a market equilibrium has a special quality, that no coalition of traders, large or small, can object to it.

In the simple world of exchange, the market mechanism automatically distributes goods among people in an optimal way. The person who likes coffee will end up with a lot of coffee, while the one who likes tea will end up with a lot of tea. The person who likes bread will have bread, while the one who likes potatoes will have potatoes. Each consumer will get the bundle of goods he likes best (given his budget constraint), and all those diverse desires will be invisibly reconciled through the adjustment of prices.

Let us be more precise. First, for us, a socially optimal state in the exchange economy model is an allocation that is in the core — one with the property that no coalition of traders, large or small, could better itself by an internal redistribution of its own resources. Second, the outcome of a competitive system or the market mechanism is simply a competitive equilibrium: An allocation $\hat{x}$ and a price vector p with the property that, given his initial holdings ω_i and the prices p, each and every trader is maximizing his own (self-interested) utility function. So the fundamental theorem we are after is: "A competitive equilibrium allocation is in the core."

In the last chapter, we saw that it is generally computationally impossible to determine whether or not an allocation is in the core via an examination of all possible blocking coalitions. There are just too many possible coalitions or groups of traders. But the fundamental theorem points the way to arriving at a core allocation: Use the competitive mechanism; use the free market.

Now we can analyze the theorem to see whether or not it is right. Very few meaningful propositions are always true, and the one at hand is not one of the few. In fact, we can construct a clear example, with two self-interested traders and two goods, where the proposition is false.

Consider a two-person, two-goods economy, in which person 1 has the following odd utility function:

$$u_1(x_1) = \begin{cases} 1 \text{ when } x_{11} + x_{12} < 1 \\ x_{11} + x_{12} \text{ when } x_{11} + x_{12} \geq 1. \end{cases}$$

Note that this utility function is not monotonic in the region where $x_{11}+x_{12} < 1$. That is, person 1 is indifferent among all bundles satisfying $x_{11} + x_{12} < 1$; if you start at one such bundle, and give him just a little more of both goods (so that $x_{11} + x_{12} < 1$ remains true), then he is no better off. He is said to have a "fat" indifference curve in this region. Let person 2's utility function be

$$u_2(x_2) = x_{21}x_{22}.$$

Suppose the initial allocation is $\omega_1 = (1, 0), \omega_2 = (0, 1)$.

Now let $p = (1, 1)$, let $\hat{x}_1 = (1/2, 1/2)$, and $\hat{x}_2 = (1/2, 1/2)$. Obviously $\hat{x}$ is an allocation. The totals of the two goods are 1 and 1, as they must be. With the prices $p = (1, 1)$ and the initial bundle $\omega_1 = (1, 0)$, person 1 can afford any bundle that costs no more than $p_1\omega_{11} + p_2\omega_{12} = 1$. That is, he can afford any bundle in his fat region of indifference, as well as any bundle on the line $x_{11} + x_{12} = 1$. But all these bundles give him 1 unit of utility. Consequently, $\hat{x}_1$ maximizes person 1's utility subject to his budget constraint.

With the prices $p = (1, 1)$ and the initial bundle $\omega_2 = (0, 1)$, person 2 can afford any bundle satisfying

$$p_1x_{21} + p_2x_{22} \leq p_1\omega_{21} + p_2\omega_{22} = 1.$$

He finds the best such bundle using the tangency condition:

$$\text{MRS for person } 2 = \frac{p_1}{p_2}$$

or

$$\frac{x_{22}}{x_{21}} = \frac{p_1}{p_2} = 1.$$

The bundle $\hat{x}_2 = (1/2, 1/2)$ solves person 2's problem; so $\hat{x}_2$ maximizes person 2's utility subject to his budget constraint.

In short, $\hat{x}$ is a competitive equilibrium allocation. But it is not Pareto optimal, and therefore it is not in the core. There are allocations that make person 2 better off and person 1 no worse off. For instance, let $\hat{y}_1 = (0, 0)$ and $\hat{y}_2 = (1, 1)$. Then $u_1(\hat{y}_1) = 1 \geq u_1(\hat{x}_1) = 1$, while $u_2(\hat{y}_2) = 1 > u_2(\hat{x}_2) = 1/4$.

The source of the difficulty in this example is the first person's fat indifference curve. But this fat indifference curve is really quite bizarre. People with fat indifference curves in the everyday world are the ones who literally throw their money away. Lots of us claim that we know someone else who throws money away, but we deny that we do it ourselves! (Putting money in a bank for future use is not throwing it away, nor is donating it to a worthwhile charity.) In short, fat indifference curves can be comfortably assumed away. They are in fact ruled out by the assumption of monotonic preferences.

How do things change when monotonicity is assumed? It turns out that in our exchange model, with self-interested monotonic utility functions, the competitive mechanism automatically distributes goods among people in an optimal way. A competitive equilibrium allocation is in the core, and the market does achieve a socially desirable state of affairs.

Let's now turn to a formal statement and proof of this most basic theorem of welfare economics.

First Fundamental Theorem of Welfare Economics. If all traders have monotonic self-interested utility functions, and if $(\hat{x}, p)$ is a competitive equilibrium, then $\hat{x}$ is in the core (and is, therefore, Pareto optimal as well).

Proof. Suppose $(\hat{x}, p)$ is a competitive equilibrium.

Suppose, contrary to the theorem, that $\hat{x}$ is not in the core.

Then some coalition can block $\hat{x}$ from ω. Let us say S can block x, and let us say it can do so with the bundles $\{s_i\}_{i \text{ in } S}$. This means that

$$\sum_{i \text{ in } S} s_{ij} = \sum_{i \text{ in } S} \omega_{ij} \text{ for all goods } j \qquad (i)$$

$$u_i(s_i) \geq u_i(\hat{x}_i) \text{ for all } i \text{ in } S \qquad (ii)$$

$$u_i(s_i) > u_i(\hat{x}_i) \text{ for at least one } i \text{ in } S \qquad (iii)$$

Combining (ii) with Observation 2, we have

$$p \cdot s_i \geq p \cdot \omega_i \text{ for all traders } i \text{ in } S.$$

Combining (iii) with Observation 2, we have

$$p \cdot s_i > p \cdot \omega_i \text{ for at least one } i \text{ in } S.$$

Now let us add these inequalities over all the traders in S, to get

$$\sum_{i \text{ in } S} p \cdot s_i > \sum_{i \text{ in } S} p \cdot \omega_i.$$

This inequality can be rewritten

$$\sum_{i \text{ in } S} (p_1 s_{i1} + p_2 s_{i2} + \ldots p_m s_{im})$$

$$> \sum_{i \text{ in } s} (p_1 \omega_{i1} + p_2 \omega_{i2} + \ldots p_m \omega_{im})$$

Rearranging, we have

$$p_1 \sum_{i \text{ in } S} s_{i1} + p_2 \sum_{i \text{ in } S} s_{i2} + \ldots + p_m \sum_{i \text{ in} S} s_{im}$$

$$> p_1 \sum_{i \text{ in } S} \omega_{i1} + p_2 \sum_{i \text{ in } S} \omega_{i2} + \ldots + p_m \sum_{i \text{ in } S} \omega_{im}$$

or

$$p_1 \left[\sum_{i \text{ in } S} s_{i1} - \sum_{i \text{ in } S} \omega_{i1} \right] + p_2 \left[\sum_{i \text{ in } S} s_{i2} - \sum_{i \text{ in } S} \omega_{i2} \right]$$

$$+ \ldots + p_m \left[\sum_{i \text{ in } S} s_{im} - \sum_{i \text{ in } S} \omega_{im} \right] > 0$$

But by (i), every term in brackets is zero. Therefore, the assumption that $\hat{x}$ is not in the core leads to a contradiction, which proves the theorem. Q.E.D.

5. The Second Fundamental Theorem of Welfare Economics

The first fundamental theorem says that a competitive equilibrium allocation is in the core, and is Pareto optimal. The competitive market, in other words, brings about a distribution of goods that is desirable in the sense that no group of traders could do better on its own. Moreover, it brings about this desirable distribution automatically: prices tend to rise in response to excess demand and tend to fall in response to excess supply; the prices adjust by themselves to solve the distribution problem.

However, the ideal distribution of goods brought about by the competitive mechanism depends heavily on the initial allocation. That is, the competitive allocation and the core are determined by the initial allocation ω as well as by preferences, and if, for example, the initial allocation is very unequal, so will be the competitive allocation.

This is an important objection to complete reliance on the competitive market: it might produce great inequalities. What does this mean in a real economy? An economy with production as well as exchange has people selling (or renting) their assets to firms, as well as buying and consuming goods and services. People sell their labor, or rent their capital goods. Some people have lots of capital to rent, and some people have very valuable labor to sell. Usually we think of industrialists, doctors and lawyers in this regard, but T.V. personalities, baseball players, and movie stars are better examples. And some people have no capital to lend or rent, and very little valuable labor to sell. Some people have few talents, few skills, and maybe not even much muscle power. There are haves, and there are have nots. And the have nots might be have nots through no fault of their own. They might be disabled, afflicted by disease, or just very unlucky. The free market mechanism will produce a distribution of goods that gives Rolls-Royces and homes in Palm Springs to the baseball stars. The distribution will give Fords and suburban tract homes to many of us. But it will give poverty and hunger to the have nots. And the result will very likely be Pareto optimal and in the core.

We can illustrate the inequality that might result from a competitive equilibrium allocation in an Edgeworth box diagram.

In Figure 3.5, ω is an initial allocation that strongly favors person 1, and $\hat{x}$ is the competitive equilibrium based on it. A Pareto optimal allocation like y might seem preferable to $\hat{x}$ on equity grounds. So the question arises: Can the competitive mechanism with some modifications be used to move the economy to an alternative Pareto optimal allocation like y, even given the initial distribution ω?

Do we really need to ask this question? Examination of Figure 3.5 seems to indicate an obvious solution to the inequality problem. Simply make person 1 give to person 2 $\omega_{11} - y_{11}$ units of good 1, and $\omega_{12} - y_{12}$ units of good 2. End of discussion. (Person 1 doesn't want to do this, of course, but reducing inequality almost always means causing someone to do something he doesn't want to do.)

This obvious solution to the inequality problem involves the direct transfer of quantities of various goods from one person to another. Why can't this solution be used in general for cases more complicated than the one drawn in the figure? It cannot be used for the same reason that a proposed allocation cannot be checked to see whether or not it is in the core by examining all possible blocking coalitions. That is, when there are many people and many goods the direct transfer solution is computationally impossible. It is just too burdensome to work.

Consider the economy of the United States, with its nearly 300 million consumers. How many different goods do we have? Obviously, the answer to this question depends on the level of aggregation we use: we might say food is a good, rather than apples, tomatoes, bread, etc. But if we talk about reasonable levels of aggregation, there are surely hundreds of things we want to identify as distinct goods. And if we walk into a large department store, we can count tens of thousands of different items; in a large supermarket there are probably tens of thousands of different food, grocery, and household goods available. So on a finely disaggregated level, there are hundreds of thousands of different goods being produced and consumed in the United States. Now imagine the

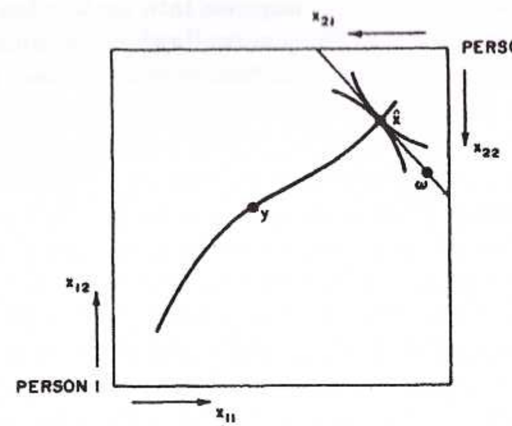

Figure 3.5.

problem of directly transferring either hundreds, or hundreds of thousands, of goods among hundreds of millions of people. Could a Central Authority, say a branch of the U.S. Government, effect such a transfer in a reasonable way?

To answer that question we can look at an effort that took place in the 70's by the U.S. Government to partially direct the reallocation of one good: gasoline. The United States Department of Energy employed some 20 thousand people (of whom, in truth, only a fraction worked on the chore of gas allocation). It attempted, in the spring and summer of 1979, to direct the distribution of gasoline in the United States, on a regional basis. That is, it attempted to dictate how much each state should get and how the gas should be distributed between urban and rural areas within each state. It did not attempt to decide how much each driver should get. Now the Energy Department allocated gasoline by, in effect, short circuiting the market mechanism. Gasoline sellers were not allowed to adjust their prices freely, and they were not allowed to decide by themselves where to sell what they want to sell. Their actions were governed by Department regulations, which were literally thousands of pages in length.

Did the Energy Department, with its thousands of employees and thousands of pages of regulations, succeed in distributing gasoline to U.S. consumers in a reasonable way? The answer to the question is a rather clear No. With the suppression of the price mechanisms a new rationing device appeared: the gas station line. Consumers spent, in the aggregate, millions of man hours waiting in lines. They burned, in the aggregate, millions of gallons of gasoline simply looking for open gas stations or waiting on lines. They ran into each other's cars while jockeying for positions on lines. They actually shot and killed each other over gallons of gas. This was a distribution system that wasted time, gasoline, people's nerves, and even occasionally lives. The attempt to partially direct the distribution of one good was a dramatic failure.

Now, imagine the complexities of having the Central Authority effect a transfer of hundreds or thousands of goods among 300 million people. If an initial allocation is unequal or inequitable in a large economy, like the U.S. economy, and if a more equitable allocation is sought on equity grounds, a movement from the initial allocation to the more equitable allocations might theoretically be brought about by a centralized and purely nonmarket mechanism. The Central Authority might theoretically inform everyone of the precise quantities of the thousands of goods to which he is entitled. But the costs, the wastes, and the information requisites of such nonmarket reallocations are enormous.

Back then, to the question: Can the automatic, decentralized, competitive market mechanism with some modifications be used to move an economy to a more equitable Pareto optimal allocation? In terms of Figure 3.5, can a modified market mechanism be used to get the economy from ω to a point like y?

The answer is generally yes, and the modified mechanism works like this. Instead of transferring quantities of hundreds of thousands of goods among millions of people, the Central Authority transfers cash. That is, generalized purchasing power, or money, is taken from some people, and given to others. After people's bank accounts have been lightened, or enhanced, as the case may be, they are left to their own devices, and the market proceeds to work as usual. Prices adjust to automatically equate supply and demand in each market, and a new, more equitable competitive equilibrium allocation comes about. The new allocation is Pareto optimal. But it is probably not in the core for the original, pretransfer allocation.

The cash transfer system has important advantages over the direct transfer of goods system. First, it is not computationally staggering. A single human mind can grasp the idea that everyone should have a minimum (cash) income of X dollars. A single human mind cannot devise a good distribution of tens of thousands of goods among hundreds of millions of people. Second, it preserves people's freedom of choice. Adjusting a person's bank book and then letting him do his shopping is less onerous than adjusting his consumption bundle directly. And third, adjusting bank balances will not create the disastrous and wasteful market problems that direct transfers have created in, for instance, the market for gasoline.

Now let's concentrate on the meaning of the cash-transfer system in the context of our exchange economy model. This might be confusing, since we have said there is no cash, no money, in the exchange model: goods trade for goods. That position must be modified. To picture what is going on here, imagine that everyone has an account with a central bank. The bank lists, in person i's account, all the goods he has. Initially the accounts lists ω_i. Now suppose there is some list of prices for the goods, or price vector p. At the bottom of person i's account book, the bank evaluates i's goods. Initially, this value is $p_1\omega_{i1}+p_2\omega_{i2}+\ldots+p_m\omega_{im}$, or $p\cdot\omega_i$. If there are no cash transfers and if i wants to trade his initial bundle ω_i for a consumption bundle x_i, i must clear it through the bank, which approves the transaction if

$$p\cdot x_i \le p\cdot\omega_i$$

that is, if $p\cdot x_i$ is less than or equal to the value of i's account. What we have described so far is just a variation of our basic exchange model, made a little complex by the existence of the bank, whose sole function is to keep an eye on people's budget constraints.

When there are cash transfers, the bank is instructed by the Authority running the system to add an amount T_i to person i's bottom line amount $p\cdot\omega_i$. The number T_i could be positive or negative. The right-hand side of i's budget inequality becomes $p\cdot\omega_i+T_i$. And now if i wants to trade his initial bundle ω_i plus his transfer T_i for the consumption bundle x_i, he must clear it through the bank, which approves the transaction if and only if

$$p\cdot x_i \le p\cdot\omega_i + T_i.$$

What does the Central Authority do to effect a more equitable allocation of goods? It assigns positive T_i's to those people who are judged too poor, and negative T_i's to those who are judged too rich. In other words, it transfers bank balances from the rich to the poor. Once the balances are transferred, the individuals buy and sell as usual and the market mechanism is allowed to work by itself.

We have asserted that a modified mechanism can be devised to get the economy to an equitable allocation, like y in Figure 3.5. At this point we must indicate the formal nature of the problem, and be more precise about the assertion. Mathematically, the problem is this. Suppose we are given an initial allocation ω, and a (desired) Pareto optimal allocation y. Does there exist a vector of bank balance transfers $(T_1, T_2, \ldots, T_n)$ and a price vector $p=(p_1,\ldots,p_m)$, such that, for every person i, y_i maximizes u_i subject to $p\cdot y_i \le p\cdot\omega_i + T_i$?

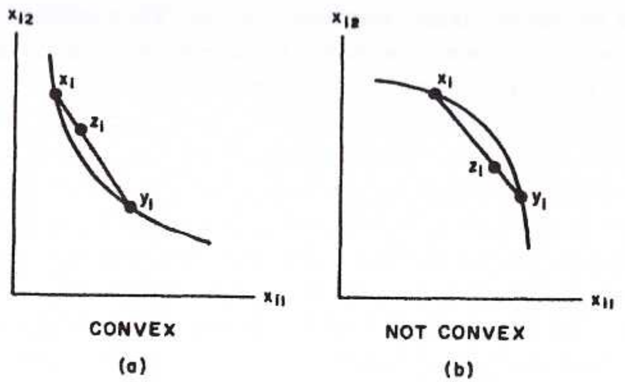

Figure 3.6.

The affirmative answer depends on three assumptions about the utility functions: (1) self-interestedness and (2) monotonicity, both of which have already been defined, and (3) convex indifference curves. Formally, we say u_{i1} has *convex indifference curves* if the following is true: If $u_i(x_i) \geq u_i(y_i)$ for two different bundles x_i and y_i, then $u_i(z_i) \geq u_i(x_i)$ for any bundle z_i that lies on the straight-line segment connecting x_i and y_i. the idea of convexity for indifference curves is illustrated in the figure below. Figure 3.6a shows indifference curves that are convex; while 3.6b shows indifference curves that aren't. Note that the figure illustrates the $u_i(x_i) = u_i(y_i)$ case, and that it assumes monotonicity.

The formal assertion that a modified competitive mechanism can be used to get the economy to almost any desired Pareto optimal allocation is called the Second Fundamental Theorem of Welfare Economics. The proof of this important result is considerably more complex than the proof of the First Fundamental Theorem, so we shall not include it here.

Instead we only state the theorem for the general n-person, m-good case:

Second Fundamental Theorem of Welfare Economics. Suppose all traders have self-interested, monotonic utility functions, and convex indifference curves. Let y be any Pareto optimal allocation that assigns positive quantities of every good to every trader: $y_{ij} > 0$ for all i and j.

Then there exists a vector of bank balance transfers $(T_1, T_2 \ldots, T_n)$ and a price vector $p = (p_1, \ldots, p_m)$ such that y and p are a competitive equilibrium given the transfers. That is, for all i, y_i maximizes u_i subject to $p \cdot y_i \leq p \cdot \omega_i + T_i$.

Let's make some observations about the theorem at this point.

First, the assumption that $y_{ij} > 0$ for all i and j, which seems quite restrictive, is made largely for mathematical convenience, and can be greatly weakened without affecting the conclusion of the theorem.

Second, the T_i's must sum to zero: All transfers to people must be financed by taxing other people. This is so because monotonicity ensures that people will want to spend to the limits of their budgets; therefore, for all i,

$$p \cdot y_i = p \cdot \omega_i + T_i.$$

Adding over all the i and rearranging terms gives

$$\sum_{j=1}^{m} p_j \sum_{i=1}^{n} y_{ij} = \sum_{j=1}^{m} p_j \sum_{i=1}^{n} \omega_{ij} + \sum_{i=1}^{n} T_i$$

or

$$\sum_{i=1}^{n} T_i = \sum_{j=1}^{m} p_j \left[\sum_{i=1}^{n} y_{ij} - \sum_{i=1}^{n} \omega_{ij} \right] = \sum_{j=1}^{m} p_j [0] = 0.$$

There is no way to induce the desired change by simply printing money and giving it to the poor. Some people must be taxed so that others can be subsidized.

Third, the theorem does not indicate how y ought to be chosen! We have argued that y is a more equitable allocation than the original ω, or the competitive allocation that would arise without transfers. But the decision to pick a particular y is made, somehow, by some people, and the process they use to make that decision is left in the dark at this point. In subsequent chapters we will discuss at some length the problem of choosing y.

Fourth, the Second Fundamental Theorem says that the competitive mechanism, with modifications, is even more useful, and more robust, than the First Fundamental Theorem indicates. Even if unmodified competition brings about distributions of goods that are inequitable, the price mechanism with modifications can be used to bring about almost any equitable and optimal allocation.

Fifth and finally, although the second fundamental theorem is an important mathematical result, one should also be aware of its limitations in terms of its practical application. In the real world, the Central Authority in charge of bringing about a desired Pareto efficient allocation y will often not have the necessary information to even determine whether or not y is Pareto optimal whether it is or is not depends on the true preferences of the agents, who may have an incentive to misrepresent them before the Authority.) Thus, the task is hard because one is effectively trying to shoot at a "moving target." We will address these issues in the last chapters of this book on the theory of implementation.

6. An Algebraic Example

Consider first a consumer, person 1, with a linear utility function, $u_1 = x_{11} + 2x_{12}$, and an initial bundle $\omega_1 = (1, 1/2)$. If he is faced with prices $p = (p_1, p_2)$, he will want to buy the best bundle he can afford, that is, the best bundle x_1 satisfying

$$p_1 x_{11} + p_2 x_{12} \leq p_1 \cdot 1 + p_2 \cdot \frac{1}{2}.$$

His indifference curves are straight lines, with slope in absolute value equal to

$$\text{MRS person } 1 = \frac{\text{MU of good 1}}{\text{MU of good 2}} = \frac{1}{2}.$$

Figure 3.7a illustrates indifference curves and a budget line for person 1. The dark line is 1's budget line; the absolute value of its slope, which evidently exceeds 1/2, is equal to p_1/p_2. Given this p_1/p_2, person 1 wants to consume the bundle $\hat{x}_1$. For a different p_1/p_2, he wants to consume a different bundle: As p_1/p_2 rises, the budget line pivots around ω_1 and gets steeper, and the desired bundle $\hat{x}_1$ climbs up the vertical axis. As p_1/p_2 drops, the budget line gets flatter, and the desired bundle $\hat{x}_1$ climbs down the vertical axis, until $p_1/p_2 = 1/2$. When $p_1/p_2 = 1/2$, the budget line coincides with 1's indifference curve through ω, and every bundle on that indifference curve maximizes 1's utility subject to his budget constraint. When $p_1/p_2 < 1/2$, 1's desired bundle $\hat{x}_1$ moves to the horizontal axis, and as the budget lines get flatter, the desired bundle $\hat{x}_1$ moves out the horizontal axis. The locus of all desired bundles, for all possible prices, is the dashed line in Figure 3.7b. This is the path that $\hat{x}_1$ traces as p_1/p_2 goes from zero to infinity, and it is called person 1's *offer curve.* Figure 3.7b also includes the budget line that appears in 3.7a. Note that when we have the offer curve and the budget line we can read off the bundle $\hat{x}_1$ that 1 wants to consume. This is the point (other than ω_1) where the offer curve and the budget line intersect.

Now suppose we have another consumer, say person 2, with the utility function

$$u_2 = x_{21} x_{22}$$

and the initial bundle $\omega_2 = (0, 1/2)$. If person 2 is faced with prices $p = (p_1, p_2)$, he will want to buy the best bundle he can afford, that is, the best bundle x_2 satisfying

$$p_1 x_{21} + p_2 x_{22} \leq p_1 \cdot 0 + p_2 \cdot \frac{1}{2}.$$

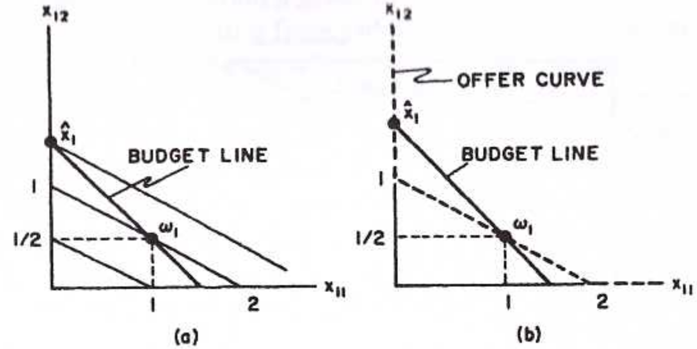

Figure 3.7.

Since his utility function is monotonic he will always want to consume at a point on, rather than below, his budget line; so the bundle he buys will satisfy

$$p_1 x_{21} + p_2 x_{22} = p_1 \cdot 0 + p_2 \cdot \frac{1}{2}.$$

Person 2's indifference curves are hyperbolic, and the absolute value of the slope of an indifference curve for him is given by

$$\text{MRS person } 2 = \frac{\text{MU of good 1}}{\text{MU of good 2}} = \frac{x_{22}}{x_{21}}.$$

With these indifference curves, his utility maximizing bundles will always be at points of tangency of his indifference curves and budget lines. That is, for the bundles person 2 wants to consume we will have

$$\text{MRS person } 2 = \frac{x_{22}}{x_{21}} = \frac{p_1}{p_2}.$$

Consequently, person 2's choices are governed by two equations

$$p_1 x_{21} + p_2 x_{22} = \frac{1}{2} p_2$$

and

$$p_2 x_{22} = p_1 x_{21}.$$

The solution to this pair of equations is

$$x_{21} = \frac{1}{4} \frac{p_2}{p_1} \text{ and } x_{22} = \frac{1}{4}.$$

This is the algebraic form of person 2's offer curve.

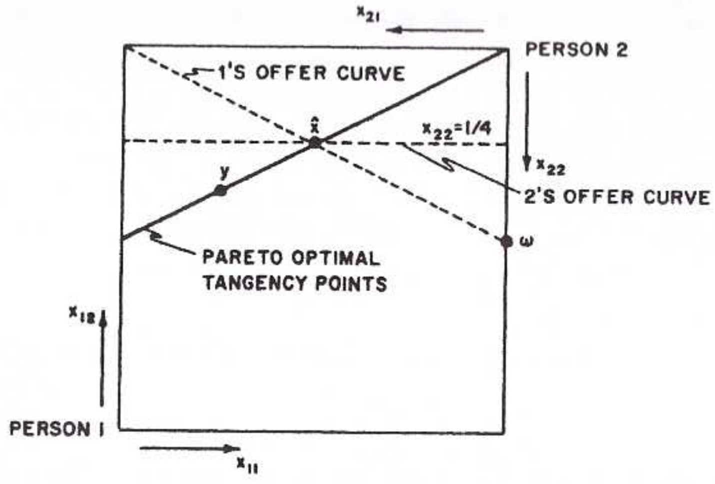

Figure 3.8.

Now consider an economy made up of persons 1 and 2. We have pictured it in an Edgeworth box diagram in Figure 3.8. To minimize complexity, all the figure shows is the offer curve of person 1 (taken from Figure 3.7b), the offer curve of person 2 (the dashed line where $x_{22} = 1/4$), and the locus of Pareto optimal points where the two individuals' indifference curves are tangent. (This is determined by setting MRS for person 1 = MRS for person 2, or $x_{22}/x_{21} = 1/2$.) The initial allocation is the point ω.

The two offer curves intersect at the point $\hat{x}$, or $\hat{x}_1 = (1/2, 3/4)$, $\hat{x}_2 = (1/2, 1/4)$. This means that for the appropriate prices person 1 wants to consume $\hat{x}_1$ and person 2 wants to consume $\hat{x}_2$. What are the prices? The budget line needed to get the two to $\hat{x}$ goes through ω and $\hat{x}$; so its slope in absolute value is 1/2. Consequently, $p_1/p_2 = 1/2$ is required, and, since one of the prices can be chosen arbitrarily, the price vector $p = (p_1, p_2) = (1, 2)$ works fine. In short, $p = (1, 2)$ and $\hat{x}_1 = (1/2, 3/4)$, $\hat{x}_2 = (1/2, 1/4)$ is a competitive equilibrium based on ω.

The First Fundamental Theorem says that $\hat{x}$ is Pareto optimal, and in the core. The figure shows it is Pareto optimal because it lies on the locus of tangencies, and it's in the core because it is Pareto optimal and it makes each individual at least as well off as ω.

Now we turn to the Second Fundamental Theorem. Suppose $\hat{x}$ is judged objectionable because it gives too much to person 1. Consider the Pareto optimal allocation y given by $y_1 = (1/4, 5/8)$, $y_2 = (3/4, 3/8)$. Suppose we want to get to y from ω via a modified competitive mechanism. What bank balance transfers T_1 and T_2 are required?

The first thing to notice about this example is that person 1, whose indifference curves are straight lines with slope 1/2, in absolute value, will only choose y_1 to maximize his utility if his budget line also has slope 1/2 in absolute value. Consequently, the two people will end up at y only if $p_1/p_2 = 1/2$. Again, one price can be chosen arbitrarily, so assume $p = (p_1, p_2) = (1, 2)$ is the competitive equilibrium price vector.

Next note that person 1 will choose point y_1 only if y_1 satisfies the equation

$$p \cdot y_1 = p \cdot \omega_1 + T_1.$$

Similarly, y_2 must satisfy the equation

$$p \cdot y_2 = p \cdot \omega_2 + T_2.$$

Substituting (1, 2) for p, and the given values for y_1, ω_1, y_2 and ω_2, we find that

$$T_1 = -\frac{1}{2} \text{ and } T_2 = +\frac{1}{2}.$$

(Obviously, these magnitudes depended on our setting $p_1 = 1$. If we let $p_1 = 2$, then p_2, T_1 and T_2 would have to be doubled.)

The reader should check that when $p = (1, 2)$, $T_1 = -1/2$ and $T_2 = +1/2$, person 1 will actually maximize his utility at y_1, and person 2 will actually maximize his utility at y_2.

7. Exercises

1. Let $u_1 = 3x_{11} + x_{12}$ and

   $u_2 = x_{21}x_{22}$.

   Let the initial allocation be $\omega_1 = (2, 1)$, $\omega_2 = (1, 2)$. Solve for the competitive equilibrium.

2. An economy is made up of two individuals and two goods. Their utility functions are:

$$u_1 = x_{11} + x_{12}$$
$$u_2 = 5x_{21}x_{22}.$$

   Their initial endowments are:

$$\omega_1 = (100, 0)$$
$$\omega_2 = (0, 50).$$

Chapter 6

UNCERTAINTY IN EXCHANGE

1. Introduction

In this chapter we consider economies that operate in uncertain environments. Therefore, it will be necessary to employ the tools of decision theory under uncertainty. We will use the expected utility preferences, as defined in Chapter 1. Our concern now is how uncertainty may affect the economic activity. What are the implications of uncertainty for the different economic institutions that we have studied, such as pure barter, price-governed exchange, production decision making, etc.? Can one recover versions of the welfare theorems in these contexts? For our purposes, there is no loss of generality in excluding production, so we shall concentrate on exchange economies for simplicity.

As we shall see, what will be important for the extension of the welfare theorems is that, despite there being uncertainty, no person in the economy holds private information with respect to the others. Instances of such private information give something of a "monopoly power" to the person that holds it, and this causes difficulties for the performance of market institutions (think of insider trading, for instance). On the other hand, the conclusions are rather different If uncertainty exists but the information is always held symmetrically among individuals (e.g., today no one knows the true state of the world, while tomorrow it is observable by everyone). For this case, financial markets may work very well to deliver an efficient allocation of risk in the economy. We begin by describing the basic model of an economy with uncertainty, which uses the notion of states of the world.

2. States of the World

Suppose there are two dates in our model. At date 0 there is uncertainty concerning the true state of the world (or simply, state) that will be realized at date 1. At date 1 no uncertainty remains because the true state of the world is observable by all agents.

A state or *state of the world* is therefore a full description of everything relevant to the economy, i.e., the true preferences of all economic agents, their true initial endowments, and the true technology available to each firm (if we also included production in our model, which is not the case here). That is, in our exchange economy, there may be uncertainty at date 0 around the agents' true preferences and endowments. For example, suppose the weather at date 1 is uncertain at date 0: it might be rainy or sunny. We shall say then that there are two states of the world: the rainy state and the sunny state. In the rainy state (i.e., if it rains at date 1), agents' date 1 preferences over sun glasses and umbrellas may be quite different from those in the sunny state (if it turns out to be a beautiful sunny day). In another example, uncertainty could involve endowments. (For example, suppose there are four possible states, which vary according to the proportion of the population with a college degree. In the first state, the proportion is between 0 and 1/4; in the second, between 1/4 and 1/2; in the third, between 1/2 and 3/4; and in the fourth between 3/4 and 1. In each of these four states, there are different aggregate endowments in the population, in terms of skilled labor versus unskilled labor).

For simplicity, we shall assume that there is a finite set of possible states, $\{1,\ldots,k\}$. Each state s consists of an exchange economy with the same n traders and the same m goods. In state s, trader i's preferences over the m goods are represented by the utility function $u_{is}(x_{is})$, where $x_{is} = (x_{i1s},\ldots,x_{ims})$ is the bundle of the m goods that trader i will consume at date 1 if state s is realized. In state s, $\omega_{is} = (\omega_{i1s},\ldots,\omega_{ims})$ is person i's endowment of goods.

At date 0, when agents are uncertain about the state of the world to be realized at date 1, agents have expected utility preferences. That is, person i maximizes $\sum_{s=1}^{k} q_{is}u_{is}(x_{is})$, where q_{is} is the probability that person i assigns to state s.

At this point it will be useful to introduce two alternative assumptions on information that will be used in the chapter. First, we speak of *symmetric uncertainty* if the situation is such that at date 0 all agents think that each state might happen at date 1. They may differ in the probability assignments given to each state, as long as all of these probabilities are positive. (If all agents agree on the probability of a state being 0, this is also symmetric uncertainty, and without losing anything,

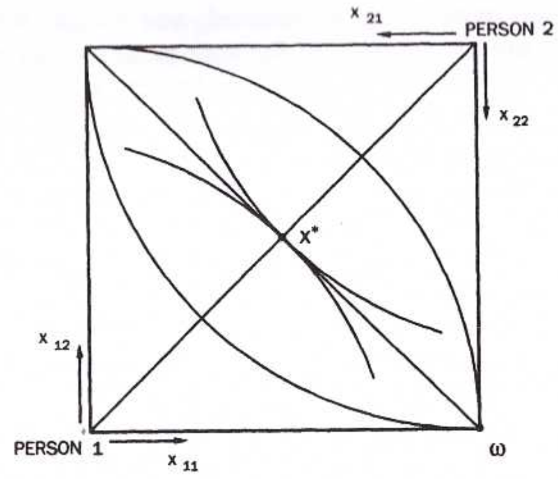

Figure 6.1.

one can remove such states from the analysis). In contrast, we speak of *asymmetric information* if for some individual, but not for all, the probability he assigns to at least one state s is 0. This means that different people have different information, in the sense that some of them know that certain things are just not true, will not happen, while other people still think that they might.

3. Symmetric Uncertainty

Let's begin with a simple example, in which we have $n = 2$ traders, $m = 1$ consumption good in each state, and $k = 2$ states. Suppose that the utility function for each person i and each state s is $u_{is}(x_{is}) = \sqrt{x_{is}}$. There is uncertainty concerning the endowment of the (only) consumption good. In state $s = 1$, agent 1 will hold one unit of the good, and agent 2 will initially have nothing. Exactly the opposite will happen in state $s = 2$. Suppose each person believes that the two states are equally likely at date 1.

First, it is interesting to note that we can still use the Edgeworth box diagram to represent this economy. See Figure 6.1. In it, the endowment is the lower right corner of the Edgeworth box, and the indifference curves depicted correspond to the expected utility function $(1/2)\sqrt{x_{i1}} + (1/2)\sqrt{x_{i2}}$ for each person i.

These individuals have an important problem to solve. If they do nothing about it, tomorrow one of them will hold all the consumption good and the other will have to starve. Furthermore, when they get to date 1 and the uncertainty is resolved, one should see no trade if these individuals are self-interested: for each person i, his *ex-post utility function* is $\sqrt{x_{is}}$, so there is no incentive for the agent who holds the unit of the good to give any to the other.

From the perspective of date 0, this is not a good situation. Each person may well be scared at the prospect of having to starve with probability one half. Note in Figure 6.1 that the endowment point gives each of them an expected or *ex-ante utility* of $1/2$, while other points of the Edgeworth box, such as its center x^*, generate higher levels of ex-ante utility. Could they do something at date 0 to prevent this situation? Well, they could sign a contract to try to provide insurance to each other. Such a contract is a state contingent contract.

A *state contingent commodity* is a title, a piece of paper, that reads "whoever holds this piece of paper is entitled to having x units of good m if and when state s happens." The idea is to create markets where these titles, these assets, will be traded. Note how, for these state contingent commodities to be worth anything and make sense, the assumption of observability of the state at date 1 is important. Otherwise, the execution and enforcement of these contracts would be impossible. In the example of Figure 6.1, a state contingent contract could be the following: if state 1 occurs, person 2 has a claim to receive half a unit of the good to be delivered by person1, and vice versa if state 2 occurs.

We shall assume the existence of *complete markets*, which is to say that there is a market for each such state contingent commodity. There are km such markets because that's the number of state contingent commodities (the product of the number of goods times the number of states of the world). Traders will participate in these markets for financial assets at date 0. Let x_{isj} denote the number of units of contingent commodity sj purchased by agent i, i.e., if agent i purchases x_{isj} units of that state contingent commodity, he is effectively purchasing the right to consume those same units of good j if state s occurs at date 1. Let p_{sj} be the price of such a state contingent commodity.

The income that person i has to participate in these financial markets is simply the market value of his initial endowments, also interpreted as state contingent commodities, i.e., $\sum_{s=1}^k p_s \cdot \omega_{is}$, which can be written out as $\sum_{s=1}^k p_{s1}\omega_{is1} + \ldots + p_{sm}\omega_{ism}$.

A competitive equilibrium in this setting is known as an *Arrow-Debreu equilibrium*, named after economists Kenneth Arrow and Gerard Debreu. It is defined as follows: a vector of prices of state contingent commodities $p = (p_{11}, \ldots, p_{1m}, \ldots, p_{k1}, \ldots, p_{km})$ and an allocation of state contingent commodities $x = (x_1, \ldots, x_n)$, where each $x_i = (x_{i11}, \ldots, x_{i1m}, \ldots, x_{ik1}, \ldots, x_{ikm})$ is an Arrow-Debreu equilibrium if two conditions are satisfied:

Expected utility maximization: for each agent i, the bundle x_i of state contingent commodities maximizes his expected utility $\sum_{s=1}^k q_{is} u_{is}(z_{is})$ subject to the budget constraint $\sum_{s=1}^k p_s \cdot z_{is} \leq \sum_{s=1}^k p_s \cdot \omega_{is}$.

Market clearing: in the market for each state contingent commodity sj, $\sum_{i=1}^n x_{isj} = \sum_{i=1}^n \omega_{isj}$.

Observe that, if each person i's ex-post utility function in state s is monotonic, then the budget constraint will be satisfied with equality in equilibrium, because at date 0 each agent will want to allocate all his income among the different state contingent commodities.

Let's try to understand the logic behind the notion of Arrow-Debreu equilibrium. According to this notion, agents try to eliminate the effects of uncertainty by insuring against it. The insurance that each person buys and sells is formulated in terms of his demand for the different state contingent commodities, which will give him the right to receive the real goods in each state. These demands take into account all relevant factors in his decision making, such as his endowment of the different commodities in each state, his beliefs about which states are more likely, and his preferences in each state. Thus, if many people believe that a particular state is very likely to happen, the prices of the state contingent commodities pertinent to that state will be high because many people will demand them, and vice versa if a state is generally deemed unlikely. Similarly, if almost no one likes to wear sun glasses when it rains, the price of that state contingent commodity (sun glasses if it rains) will be low; or if a particular good is expected to be abundant in some states, the corresponding state contingent commodities will also have a low price. In equilibrium, given all these different factors, the prices of all state contingent commodities will adjust to eliminate any excess demand or excess supply.

What can one say about the welfare properties of the Arrow-Debreu equilibrium? As we shall indicate below, the two fundamental theorems hold for this concept. But before we do that, we should spend some time talking about the definitions of efficiency. When one asks the efficiency question in these settings, the issue is the timing at which the Central Authority wishes to perform the improvement. Should it think of making the improvement at date 0, when the state is uncertain, or at date 1, when the state is known? In what follows, we refer to decisions made at date 0 as *ex-ante* or "before the fact," and decisions made at date 1 as *ex-post* or "after the fact." We have two different concepts of Pareto optimality or efficiency.

We say that an allocation of state contingent commodities x is not ex-ante efficient if there is another allocation of state contingent commodities y such that

$$\sum_{s=1}^{k} q_{is}u_{is}(y_{is}) \geq \sum_{s=1}^{k} q_{is}u_{is}(x_{is}) \text{ for all } i = 1, 2, \ldots, n$$
$$\sum_{s=1}^{k} q_{is}u_{is}(y_{is}) > \sum_{s=1}^{k} q_{is}u_{is}(x_{is}) \text{ for at least one } i.$$

If there is no such alternative allocation, x is *ex-ante efficient.*

That is, at an ex-ante efficient allocation, one cannot make an improvement on the entire economy, even when all the financial markets for state contingent commodities are open, by redistributing assets in some fashion. If one performs any such redistribution and makes some people better off in ex-ante or expected utility terms, it must be at the expense of making someone else's ex-ante utility lower. Since the Arrow-Debreu equilibrium is based on ex-ante trade in markets when uncertainty is present, one should evaluate its efficiency properties by using this notion.

Now we turn to *ex-post* or "after the fact."

We say that an allocation of state contingent commodities x is not ex-post efficient if there exists a state s and another allocation of state contingent commodities y such that

$$u_{is}(y_{is}) \geq u_{is}(x_{is}) \text{ for all } i = 1, 2, \ldots, n$$
$$u_{is}(y_{is}) > u_{is}(x_{is}) \text{ for at least one } i.$$

If there is no such state and alternative allocation, x is *ex-post efficient.*

The question that ex-post efficiency answers is whether one could perform an improvement on the entire economy after uncertainty has disappeared, by operating on the markets for goods. One important fact is the relationship between these two concepts of efficiency, sometimes referred to as the *no-trade theorem.* This says that, if the economy is at an ex-ante efficient allocation and information arrives so that uncertainty vanishes, no trade will be observed in the ex-post markets, because the allocation is also ex-post efficient. Thus, additional information does not result in new trade. We prove this result next:

No-Trade Theorem. If an allocation x is ex-ante efficient, it is also ex-post efficient.

Proof: We argue by contradiction. Suppose that x is ex-ante efficient, but not ex-post efficient. Since it is not ex-post efficient, there exists a state s' and an alternative allocation y that makes everyone in that state s' at least as good as does x, and at least one individual strictly better.

Now define the following allocation z:

$$z = \begin{cases} y_s \text{ if } s = s', \\ x_s \text{ otherwise} \end{cases}$$

Since z changes with respect to x only in state s', it follows that

$$\sum_{s=1}^{k} q_{is}u_{is}(z_{is}) \geq \sum_{s=1}^{k} q_{is}u_{is}(x_{is}) \text{ for all } i = 1, 2, \ldots, n$$
$$\sum_{s=1}^{k} q_{is}u_{is}(z_{is}) > \sum_{s=1}^{k} q_{is}u_{is}(x_{is}) \text{ for at least one } i.$$

To see the strict inequality, recall that for all i and for all s, $q_{is} > 0$. But this means that x is not ex-ante efficient either, which is a contradiction. Q.E.D.

Similarly, one can also provide definitions for the ex-ante core and the ex-post core of the economy, as a function of whether coalitional improvements are envisioned before or after the resolution of uncertainty. We define the ex-ante core next, which will be relevant for our result:

An allocation y of state contingent commodities is *ex-ante feasible* for coalition S if for every state s, $\sum_{i \in S} y_{is} = \sum_{i \in S} \omega_{is}$. That is, if the group S of individuals get together, they may think of trading state contingent commodities that add up to their endowments of goods in each state.

Then, if x is a proposed allocation of state contingent commodities and S is a coalition, we will say that S can *ex-ante block* x if there is a feasible allocation y for S, such that:

$$\sum_{s=1}^{k} q_{is}u_{is}(y_{is}) \geq \sum_{s=1}^{k} q_{is}u_{is}(x_{is}) \text{ for all } i \in S$$
$$\sum_{s=1}^{k} q_{is}u_{is}(y_{is}) > \sum_{s=1}^{k} q_{is}u_{is}(x_{is}) \text{ for at least one } i \in S$$

The *ex-ante core* is the set of allocations that cannot be ex-ante blocked by any coalition.

To define the notion of ex-post core, one should think that the coalition formation process may start before uncertainty, but coalitions need not act until it is resolved. Ex-post feasibility for a coalition in a state is equivalent to feasibility in an economy with no uncertainty. Thus, if x is an allocation of state contingent commodities and S is a coalition, we will say that S can *ex-post block* x if there is a state s and an ex-post feasible allocation y for S, such that:

$$u_{is}(y_{is}) \geq u_{is}(x_{is}) \text{ for all } i \in S$$
$$u_{is}(y_{is}) > u_{is}(x_{is}) \text{ for at least one } i \in S$$

The *ex-post core* is the set of allocations that cannot be ex-post blocked by any coalition.

Using similar steps as in the proof of the no-trade theorem, one can establish that the ex-ante core is a subset of the ex-post core: if coalition S finds that they can ex-post block x with y in state s', at the ex-ante stage they can commit to stay with the grand coalition at x if the realized state is other than s', and move to y otherwise. In addition, the ex-ante core is contained in the set of ex-ante efficient allocations, and the ex-post core is contained in the set of ex-post efficient allocations. This is so because efficiency concepts rely on the coalition of all agents finding an improvement, while core notions use every possible coalition.

By making the assumptions on preferences of Chapter 3, one can then establish the two fundamental theorems of welfare economics. We next provide their statements. We shall not prove them, although we note that the proofs are identical to the corresponding theorems seen in that chapter, once one takes account of the different notation (essentially, the distinction between goods versus state contingent commodities, and utility versus expected utility).

First Fundamental Theorem of Welfare Economics for Economies with Uncertainty. Suppose uncertainty is symmetric and there are complete markets (i.e., there is a market for each state contingent commodity). Suppose also that all traders have monotonic self-interested utility functions in each state, and are expected utility maximizers. If (x, p) is an Arrow-Debreu equilibrium, then x is in the ex-ante core (and is, therefore, ex-ante efficient as well).

For the second theorem, we shall need the following property of utility functions: a utility function in state s u_{is} is *concave* whenever for every pair of bundles x_{is}, y_{is} and for every $\alpha \in [0,1]$,

$$u_{is}([\alpha x_{is} + (1-\alpha)y_{is}]) \geq \alpha u_{is}(x_{is}) + (1-\alpha)u_{is}(y_{is}).$$

This just says that the utility derived from a bundle that is a convex combination of two bundles is no less than the convex combination of the utilities of those bundles. Concavity implies the condition of convex indifference curves, assumed earlier in the book. We are now ready to state the second theorem.

Second Fundamental Theorem of Welfare Economics for Economies with Uncertainty. Suppose uncertainty is symmetric and there are complete markets (i.e., there is a market for each state contingent commodity). Suppose also that all traders have self-interested, monotonic and concave utility functions in each state, and are expected utility maximizers. Let y be any ex-ante efficient allocation. Then there exists a vector of bank balance transfers $(T_1, T_2 \ldots, T_n)$ and a price vector $p = (p_{11}, \ldots, p_{km})$ such that (y, p) is an Arrow-Debreu equilibrium given the transfers. That is, for all i, y_i maximizes person i's expected utility $\sum_{s=1}^{k} q_{is} u_{is}(z_{is})$ subject to $p \cdot z_i \leq p \cdot \omega_i + T_i$.

Note that most of the assumptions made here are the same as in Chapter 3. The only difference is the requirement of concave ex-post utility functions, instead of the weaker assumption of convex indifference curves. This is related to the fact that expected utility is preserved only under positive affine transformations of utility functions. Thus, we are led to assume concavity, a cardinal property of a function, instead of convex indifference curves, a property that would be preserved through any monotone transformation. The good news is that concave ex-post utility functions admit a nice economic interpretation, because they amount to risk aversion or risk neutrality (i.e., they rule out risk lovers from the analysis, those "nutty" people for whom facing risky situations is exciting; for most people, risk is something to be avoided).

We now know from the first welfare theorem that if the economy has enough markets - complete markets for all state contingent commodities - the market equilibrium, driven by the forces of demand and supply, will yield an ex-ante core allocation. That means that, even after the arrival of information, after the resolution of uncertainty, no coalition will be able to benefit from reallocating its endowments to improve upon the equilibrium allocation. (Recall that the ex-ante core is contained in the ex-post core). The markets for state contingent commodities perform an efficient transfer of income across states for each agent, as a function of how much consumption each of them wishes to have in each state. This is the role of the ex-ante budget constraint. From the second welfare theorem, we know that any ex-ante efficient allocation can be supported with Arrow-Debreu equilibrium prices, provided that one makes the appropriate cash transfers (which amount to reallocations of state contingent commodities).

4. Examples

To understand the concept of Arrow-Debreu equilibrium and its welfare properties, it is useful to look at some examples. First, the reader can easily check that in the example of Figure 6.1 the Arrow-Debreu equilibrium allocation is x^*, which allocates the state contingent bundle $(1/2, 1/2)$ to each trader, associated with prices $p = (1, 1)$. That is, given the concavity of the ex-post utility functions (which is equivalent to risk aversion) and given that both traders expect each state to be as likely as the other, they provide complete insurance to one another by promising to split equally the endowment of the good in each state.

Consider next the example shown in Figure 6.2. In it, we still have two agents, two states and one good per state. The ex-post utility function of each trader continues to be $\sqrt{x_{is}}$. Endowments are also as before, so that trader i will only have one unit of the good in state i. However, suppose that now traders have different probability assessments about the states. Let's say that trader 1 believes $q_1 = (1/3, 2/3)$, while trader 2 believes $q_2 = (2/3, 1/3)$. Now the Arrow-Debreu equilibrium allocation, labeled as x^* in the figure, assigns the bundle $(1/5, 4/5)$ to trader 1 and $(4/5, 1/5)$ to trader 2. The equilibrium relative price is still 1 (exactly half way between each agent's relative probability assessment – the market takes those assessments into account to calculate the value of state contingent commodities). Note how now the allocation is strongly driven by the different beliefs: each trader plans a higher consumption in that state that he believes more likely.

To finish this section, consider Figure 6.3. Here, we have an economy with the same endowments as before. Agents hold the same beliefs as in the last example, so that $q_1 = (1/3, 2/3)$ and $q_2 = (2/3, 1/3)$. the ex-post utility function of agent 1 continues to be the same as before, but the one for agent 2 is now linear, so that he is risk neutral (he cares about the expected consumption, even if it varies greatly from state to state). That is, the expected utility functions for each of them are $(1/3)\sqrt{x_{11}} + (2/3)\sqrt{x_{12}}$ for agent 1, and $(2/3)x_{21} + (1/3)x_{22}$ for agent 2.

How does this change in agent 2's preferences affect the equilibrium? (To find the equilibrium, you should draw the *offer curves* of both traders and find their intersection. We leave this as an exercise). Note that at the equilibrium allocation of Figure 6.2, now the marginal rate of substitution of agent 2, which is 2, exceeds the old equilibrium relative price of 1. If the relative price were 1, agent 2 would now demand more units of the good in state 1 and less in state 2. This causes the price of

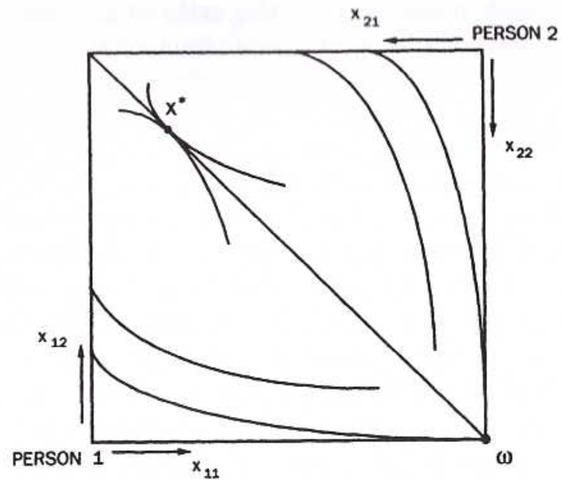

Figure 6.2.

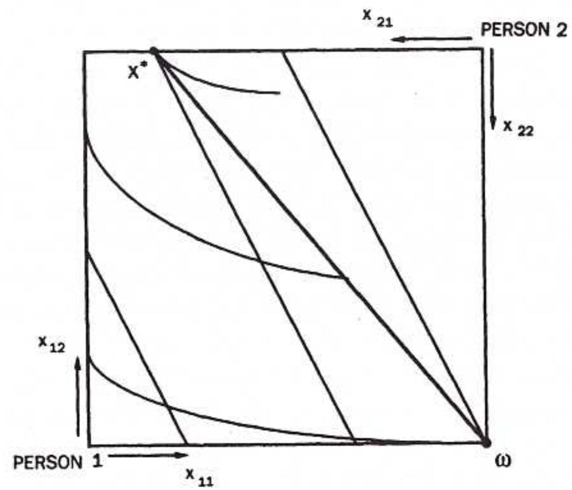

Figure 6.3.

good 1 in terms of good 2 to go up. The new equilibrium has a relative price of $p = 1/(\sqrt{8} - 2) > 1$, and the equilibrium allocation assigns the bundle $(3 - \sqrt{8}, 1)$ to agent 1 and $(\sqrt{8} - 2, 0)$ to agent 2. That is, agent 2 is happy to give up all his consumption in state 2 to fully insure agent 1: note that only for agent 1 is it true that his marginal rate of substitution in equilibrium equals the relative price of state contingent commodities (thus, he is fully insuring himself in the sense that he is planning his consumption in both states so that the ratio of his marginal utilities exactly coincides with the market equilibrium relative price). On the other hand, agent 2, at the equilibrium price, would be content with demanding more consumption in state 1 and less in state 2, but he can't because he is already planning a zero consumption in state 2.

In all of these examples, the equilibrium allocation is in the ex-ante core. First, it is ex-ante efficient, thanks to the coordination in trade via the signals provided by prices. And second, no individual can block, since for each of them the equilibrium bundle maximizes expected utility over the budget set, which contains the endowment point.

5. Asymmetric Information

Now we move on to the study of economies in which some people have more information than others. For example, suppose the two states of the world concern tomorrow's weather (either sunny or rainy), one of the agents is a weather man who already knows today which of the two will occur, while the rest of agents are still uncertain. Or the government of a country knows whether a foreign dictator holds "nukular" weapons, while other countries' governments do not know.

At this point, we shall focus on a model with three dates. At date -1, no one has private information. Date -1 is the *ex-ante* stage (similar to date 0 in the previous sections). At date 0, some people have private information, and so this is the stage in which information is asymmetric. This is called the *interim* stage. At date 1 the true state of the world is resolved, it is verifiable by all agents and any information superiority disappears. This is the *ex-post* stage. We shall concentrate on models in which decisions are made at the interim stage, that is, in the presence of asymmetric information.

Formally, the way we model the arrival of information at date 0 is by assuming that each agent i receives a signal σ_i. Of course, in different states of the world, agent i can receive different signals. Recall that the set of possible states of the world is $\{1, \ldots, k\}$. The signal σ_i will allow agent i to rule out some states for sure, and hence his beliefs will be updated over the set of remaining states. We shall model the signal as simply a subset of the original set of states (the set of states that have not been ruled out): $\sigma_i \subseteq \{1, \ldots, k\}$. The interpretation is that agent i some how finds out that the true state is in σ_i, thereby being able to rule out states not in σ_i. We shall assume that there is no false information in the signals (a state s that has been zeroed out by the signal σ_i does not become ultimately the actual state at date 1).

An example will serve to clarify these ideas. Suppose there are five possible states: $\{1, 2, 3, 4, 5\}$, corresponding to the number of guests that will show up to a party that you are hosting. (Suppose you care only about the number of people that show up, and not who those are). Let your ex-ante probability assessment about the five states be $q_i = (0.1, 0.2, 0.2, 0.1, 0.4)$. That is, at time -1 you believe that the most likely state is that all five guests show up, and this has probability 0.4. This is four times more likely that only one of them does not show up or that only one shows up. The intermediate cases (two or three guests) are believed to be equally likely, with probability 0.2 each.

To help you figure out how many will attend, they are supposed to RSVP to your invitation. Let's say that only three possible events might happen (and hence, there are only three possible signals that you might receive), which exhaust all possibilities: either you get exactly three affirmative RSVP's, or more than three, or less than three. Indeed, suppose you get a signal that informs you that three of them are for sure not coming. (E.g., only three phone calls, all of them declining your invitation). Let's call this signal σ_i. This signal allows to update your beliefs, which now are $q_i|\sigma_i = (1/3, 2/3, 0, 0, 0)$. This expression is read "belief q_i given signal σ_i," and updating is done using Bayes' rule. That is, given the signal σ_i, we can safely rule out the last three states; and for the rest (one or two guests are still possible), you continue to believe that a two-guest party is twice as likely as a party with only one; since the revised probabilities must add up to 1, these are $q_i(1|\sigma_i) = 1/3$ and $q_i(2|\sigma_i) = 2/3$. If the party ends up having three or more guests, you would not have received signal σ_i, since we are assuming that there are no false signals. As we said above, in that case you would have received either signal σ_i' or signal σ_i'': signal σ_i' is fully informative (you get exactly three affirmative and two negative RSVP's to the party), while signal σ_i'' still leaves you with some uncertainty (you receive four affirmative RSVP's and the fifth person forgets to call). Of course, the use of Bayes' rule to update your ex-ante beliefs gives $q_i|\sigma_i' = (0, 0, 1, 0, 0)$ and $q_i|\sigma_i'' = (0, 0, 0, 1/5, 4/5)$, respectively.

As a function of whether you receive signal σ_i, σ_i' or σ_i'', you may want to make different decisions regarding the number of potato chip bags, cans of soda, and pounds of other junk food to get.

We now return to our exchange economy with three dates, but where no trade takes place at date -1. At date 0, the interim stage, each trader may have received an informative signal (not receiving any signal is also a possibility, in which case we shall say that σ_i coincides with the entire set of states). Trade may take place and contracts may be written at this stage. The execution of these contracts and the consequent consumption of goods takes place at date 1, after uncertainty vanishes. (Alternatively, as we will point out, one can also think of consumption taking place at the interim stage, at date 0).

The notion of market equilibrium most used in these contexts is the *rational expectations equilibrium* (REE), which we will define in the next paragraph. The idea of an REE is the following: we'll have an equilibrium price function p^*, which will assign a vector of prices of goods for each state s, i.e., $p^*(s) = (p_1^*(s), \ldots, p_m^*(s))$ for $s = 1, \ldots, k$. That is, in different states, given the different information held by traders, prices could incorporate parts of this information. We shall say that a price function is *fully revealing* if for every two distinct states s and s', $p^*(s) \neq p^*(s')$. In this case, all traders, who rationally expect the equilibrium price function, will be able to update their beliefs and learn the state completely just by looking at the market prices. Otherwise, such complete learning will not be possible and we will speak of a *nonrevealing* price function. (Note the slight change of notation with respect to earlier sections: to emphasize the concept of price function, we write the state as an argument, and not in the subscript as before; we shall do this also for allocations.) Given an REE price function, traders formulate their demands by maximizing their conditional expected utility (conditional on the information that they have received, via the private signal and the market prices). Furthermore, the prices must be such that in each state markets must clear. Many economists contributed to this notion, but among the first we should name Jerry Green, Sanford Grossman, Robert Lucas and Roy Radner. We give the definition next.

A price function p^* that consists of prices $(p_1^*(s), \ldots, p_m^*(s))$ in each state s, and an allocation $x = (x_1, \ldots, x_n)$, where each x_i describes for each state s the bundle $x_i(s) = (x_{i1}(s), \ldots, x_{im}(s))$ assigned to trader i is a *rational expectations equilibrium* (REE) if two conditions are satisfied:

Conditional expected utility maximization: given signal σ_i and p^*, let T be the set of states that agent i still assigns positive probability. Then, for each agent i and each signal σ_i, the bundles x_i maximize his conditional expected utility $\sum_{s=1}^k q_i(s|\sigma_i, p^*)u_{is}(z_i(s))$ subject to the budget constraint $p^*(s) \cdot z_i(s) \leq p^*(s) \cdot \omega_i(s)$ for all $s \in T$, and $z_i(s) = z_i(s')$ for every $s, s' \in T$.

Market clearing: in the market for each good j in each state s, $\sum_{i=1}^n x_{ij}(s) = \sum_{i=1}^n \omega_{ij}(s)$.

Note how the first requirement incorporates the updating of beliefs given the signal and the market prices. Also, we require that, if an agent cannot rule out either of two states, his consumption bundles in those two states must be the same. This is why the concept can also be understood as one in which consumption may take place at date 0, before the resolution of uncertainty.

In terms of evaluating the welfare properties of a concept like REE, we have to introduce another definition of efficiency, in which improvements can be performed at the interim stage, given a fixed structure of signals σ_i's. The new notion is called *interim efficiency* and we define it now.

Given a fixed structure of signals, an allocation x is not interim efficient if there is another allocation y such that

$$\sum_{s=1}^{k} q_i(s|\sigma_i)u_{is}(y_i(s)) \geq \sum_{s=1}^{k} q_i(s|\sigma_i)u_{is}(x_i(s))$$

for all $i = 1, 2, \ldots, n$ and for all σ_i,

$$\sum_{s=1}^{k} q_i(s|\sigma_i)u_{is}(y_i(s)) > \sum_{s=1}^{k} q_i(s|\sigma_i)u_{is}(x_i(s))$$

for at least one pair i, σ_i.

If there is no such alternative allocation, x is an *interim efficient* allocation.

That is, at an interim efficient allocation, one cannot make an improvement on the entire economy at the time each agent has received some private information. If at an interim efficient allocation one wishes to perform any redistribution of goods, at least one of the agents, given his private information, will suffer from the redistribution and obtain a lower interim or conditional expected utility.

It is important to note that the set of allocations that is interim efficient falls between the ex-ante and the ex-post efficiency sets. That is, a no-trade theorem holds for the asymmetric information model. We state it here and leave the proof as an exercise.

No-Trade Theorem; Asymmetric Information: If an allocation x is ex-ante efficient, it is also interim efficient. If an allocation x is interim efficient, it is also ex-post efficient.

Our next question is whether one can extend the welfare theorems to the asymmetric information model. The results here will be largely negative. We shall concentrate on the first welfare theorem, and show the following results:

First Fundamental Nontheorem of Welfare Economics; Asymmetric Information – Part I. Under the usual assumptions on preferences, the conclusions of the first fundamental theorem of welfare economics may be violated under asymmetric information in terms of ex-ante efficiency. Specifically, a rational expectations equilibrium allocation may fail to be ex-ante efficient.

Proof: We shall provide an example to make this point. Consider the following economy. There are two agents, two goods per state, and two equally likely states $s = 1, 2$. Agent 1 is fully informed, which means he can receive two possible signals at time 0, σ_1 (i.e., state $s = 1$ will hold) and σ_1' (i.e., $s = 2$ will hold): therefore, $q_1|\sigma_1 = (1, 0)$ and $q_1|\sigma_1' = (0, 1)$. Agent 2 is uninformed, his signal does not provide any new information so that $q_2|\sigma_2 = (1/2, 1/2)$. Preferences in state $s = 1, 2$ are represented by the utility functions

$$u_{1s}(x_1) = x_{11} + s \cdot \ln(x_{12}), \quad u_{2s}(x_2) = \ln(x_{21}) + s \cdot \ln(x_{22}).$$

As for endowments, they are state independent:

$$\omega_1 = (1, 1), \quad \omega_2 = (2, 1).$$

Figure 6.4 illustrates this economy in two Edgeworth box diagrams. Note that Figure 6.4a is for state 1, and Figure 6.4b is for state 2.

It turns out that there is only one REE in this economy (this uniqueness is not essential for our current argument, and we leave it as an exercise). Moreover, it is fully revealing. Thus, it corresponds to the unique ex-post competitive equilibrium in each of the states. We call the unique REE allocation x. Normalize the price of good 2 to 1 in each state. Then, the equilibrium price function is $p^*(1) = (3/4, 1)$ and $p^*(2) = (2/5, 1)$. We provide the description of the economy, the REE allocation x and the utility levels it yields in each state in the table below.

State 1	State 2
$u_{11}(\cdot) = x_{11}(1) + \ln(x_{12}(1))$	$u_{12}(\cdot) = x_{11}(2) + 2\ln(x_{12}(2))$
$u_{21}(\cdot) = \ln(x_{21}(1)) + \ln(x_{22}(1))$	$u_{22}(\cdot) = \ln(x_{21}(2)) + 2\ln(x_{22}(2))$
$\omega_1(1) = (1,1), \omega_2(1) = (2,1)$	$\omega_1(2) = (1,1), \omega_2(2) = (2,1)$
$x_1(1) = (x_{11}(1), x_{12}(1)) = (4/3, 3/4)$	$x_1(2) = (x_{11}(2), x_{12}(2)) = (3/2, 4/5)$
$x_2(1) = (x_{21}(1), x_{22}(1)) = (5/3, 5/4)$	$x_2(2) = (x_{21}(2), x_{22}(2)) = (3/2, 6/5)$
$u_{11} = (4/3) + \ln(3/4) = 1.0457$	$u_{12} = (3/2) + 2\ln(4/5) = 1.0537$
$u_{21} = \ln(5/3) + \ln(5/4) = 0.7340$	$u_{22} = \ln(3/2) + 2\ln(6/5) = 0.7701$

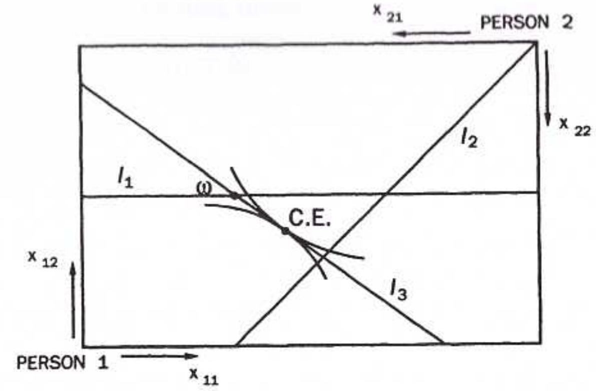

Figure 6.4a.

State 1. Along line l_1, $MRS_1 = 1$; along line l_2, $MRS_2 = 1$; along line l_3, absolute slope = the price ratio = 3/4.

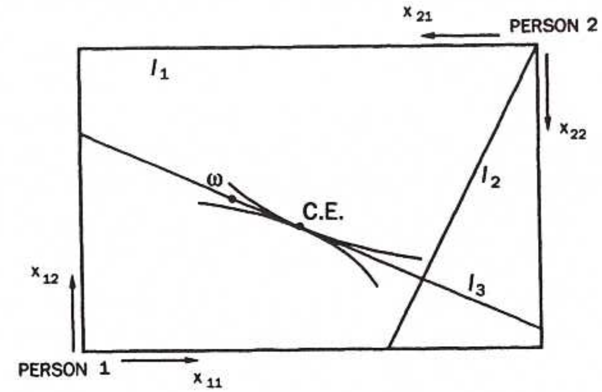

Figure 6.4b.

State 2. Along line l_1, $MRS_1 = 1$; along line l_2, $MRS_2 = 1$; along line l_3, absolute slope = the price ratio = 2/5.

The ex-ante expected utility for person 1 from the REE allocation x is

$$(1/2) \cdot (1.0457) + (1/2) \cdot (1.0537) = 1.0497.$$

And the ex-ante expected utility for person 2 from the REE allocation x is

$$(1/2)\cdot(0.7340)+(1/2)\cdot(0.7701)=0.7525.$$

We need to show that this REE allocation is not ex-ante efficient. We propose the improvement x'. This new allocation and the utility levels it yields in each state is also shown in table form:

State 1	State 2
$x'_1(1)=(x'_{11}(1),x'_{12}(1))=(17/12,3/4)$	$x'_1(2)=(x'_{11}(2),x'_{12}(2))=(17/12,4/5)$
$x'_2(1)=(x'_{21}(1),x'_{22}(1))=(19/12,5/4)$	$x'_2(2)=(x'_{21}(2),x'_{22}(2))=(19/12,6/5)$
$u'_{11}=(17/12)+\ln(3/4)=1.1290$	$u'_{12}=(17/12)+2\ln(4/5)=0.9704$
$u'_{21}=\ln(19/12)+\ln(5/4)=0.6827$	$u'_{22}=\ln(19/12)+2\ln(6/5)=0.8242$

Now the ex-ante expected utility for person 1 from the allocation x' is

$$(1/2)\cdot(1.1290)+(1/2)\cdot(0.9704)=1.0497.$$

And the ex-ante expected utility for person 2 from x' is

$$(1/2)\cdot(0.6827)+(1/2)\cdot(0.8242)=0.7534.$$

Since the improvement does not change the allocation of good 2, only changes in utility through good 1 must be evaluated. Since that part of 1's utility function is linear, note how he is exactly indifferent in ex-ante terms (i.e., before he receives his signal, before he knows the state) because 17/12 is the average of the amounts of good 1 that he was receiving in each state at the REE allocation (4/3 and 3/2). However, since the utility of agent 2 for good 1 is a logarithmic function, which is strictly concave, agent 2 is strictly better off in ex-ante terms with this improvement.

Therefore, the allocation x' is an ex-ante improvement over the REE allocation x. Q.E.D.

First Fundamental Nontheorem of Welfare Economics; Asymmetric Information – Part II. Under the usual assumptions on preferences, the conclusions of the first fundamental theorem of welfare economics may be violated under asymmetric information in terms of interim efficiency. Specifically, a rational expectations equilibrium allocation may fail to be interim efficient.

Proof: We shall modify the example in the earlier proof slightly to make this point.

Allocation x' is not yet an interim improvement over the REE allocation x. This is because agent 1 in state $s=2$ is made actually worse off in interim terms: $17/12<3/2$. Note, however, how agent 1 in $s=1$ is made better off in interim terms at x'. Agent 2 is also made better off in interim terms (for him, ex-ante and interim terms are the same).

Here's how we modify the example: suppose we look at the economy with $k+1$ copies of agent 1 and $k+1$ copies of agent 2 in terms of preferences and endowments in the two states. However, only the original agent 1 is informed about the state. All other copies of agent 1 and all copies of agent 2 (including the original one) are uninformed.

For this economy, the original REE continues to be an REE, and it is fully revealing (that is, the ex-post competitive equilibrium allocation of this replica economy is the replica of the original equilibrium allocation). To see this, notice that the fully revealing nature of the equilibrium makes the apparent asymmetry in the different copies of agent 1 disappear.

But now we can take away an amount ϵ of good 1 from the improvement made on each copy of agent 2 to make him receive $(19/12)-\epsilon$ in each state. We choose ϵ small enough so that this still constitutes an interim improvement for each of the copies of agent 2 (and note that this ϵ does not depend on k). Now give this extra amount of good 1 to the original agent 1. The one we were concerned with, that is, the original agent 1 in state $s=2$ is now receiving an amount $(17/12)+(k+1)\epsilon$. Thus, it suffices to choose k large enough to guarantee that $(17/12)+(k+1)\epsilon\geq 3/2$. Now the resulting allocation is indeed in interim improvement, so that the REE allocation of this replica economy is not interim efficient. Q.E.D.

First Fundamental Theorem of Welfare Economics; Asymmetric Information – Part III. Under the usual assumptions on preferences, the conclusions of the first fundamental theorem of welfare economics almost always hold under asymmetric information in terms of ex-post efficiency. Specifically, a rational expectations equilibrium allocation in almost every economy is ex-post efficient.

Proof: To prove the "almost always" bit rigorously is beyond the scope of this book. However, here's an intuition. Constructing a non-revealing REE is difficult. In it, one must require that an uninformed agent who cannot distinguish among several states (recall that he faces the same prices over those states) maximizes expected utility subject to the budget constraint in each state. The solution to this problem will typically yield different bundles in each such state, violating the REE requirement that consumption must be the same.

Thus, we learn that almost every REE must be fully revealing. But then, in each state, the projection of the REE to that state is an ex-post Walrasian equilibrium. By the standard version of the first welfare theorem, the corresponding ex-post allocation is then ex-post efficient.

Q.E.D.

A few remarks are in order:

First, what we have shown is a strong violation of the first welfare theorem in these settings. Given that REE is a concept that concerns interim trade, the most relevant result of the ones we have just shown is the second nontheorem. Since REE allocations are not necessarily interim efficient, they will also fail to be elements of any interim core. The interim core will be contained in the set of interim efficient allocations. Defining interim cores, though, is not a straightforward task, as one needs to take account of the amount of information transmission that may take place within coalitions that consist of asymmetrically informed individuals.

Second, the first part of our nontheorem similarly implies that REE allocations may fail to be in the ex-ante core. It is important to point out that the reason for the lack of interim and ex-ante efficiency of REE's is that "too much" information is revealed by the equilibrium price function, which kills insurance possibilities.

Third, as argued above, an REE is "almost always" fully revealing. That is, in almost every economy all REE's yield a fully revealing price function. Whenever this is the case, the corresponding REE allocations will be in the ex-post core, simply because they correspond to ex-post competitive equilibria. The fact that an REE typically gives away all the information and allows people to learn the state perfectly is a very intriguing result. To try to understand how all the private information ends up being incorporated in the price function, different authors have investigated alternative trading procedures, in which traders choose how much information to transmit to the market (see for example Wolinsky (1990) and Gottardi and Serrano (2005)).

6. Exercises

1 Show that the set of ex-ante efficient allocations is contained in the set of interim efficient allocations. Show that the set of interim efficient allocations is contained in the set of ex-post efficient allocations.

2 Consider the economy depicted in Figure 6.1. Find the entire set of ex-ante efficient allocations. Find the entire set of ex-post efficient allocations.

3 Consider again the economy of Figure 6.1, but now allow any beliefs. That is, let trader i's belief be $q_i = (q_{i1}, q_{i2})$. Calculate the Arrow-Debreu equilibrium in each of these economies.

4 In the economy with asymmetric information depicted in Figure 6.4, show that there is a unique REE (which is the one given in the text).

7. Selected References

(Items marked with an asterisk (*) are mathematically difficult.)

*1. K. Arrow, "The Role of Securities in the Optimal Allocation of Risk Bearing," *Review of Economic Studies* V. 31, 91-96, 1964.

In this article, Arrow introduces the so-called Arrow securities, a simpler way to get the results obtained under complete markets. Arrow argues that, instead of having km state contingent commodities, it suffices to have as many assets as states (i.e., only k assets) to be able to transfer wealth across states.

*2. G. Debreu, *Theory of Value*, Wiley, New York, 1959.

Debreu's book also covers the case of symmetric uncertainty as part of the model of a general equilibrium economy.

*3. F. Forges, E. Minelli and R. Vohra, "Incentives and the Core of an Exchange Economy: A Survey." *Journal of Mathematical Economics*, V. 38, pp. 1-41, 2002.

This article is an excellent survey of the recent literature on the core with asymmetric information. Ex-ante and interim cores are described, and the emphasis is on economies in which states are not verifiable, even ex-post. For these cases, incentive constraints must be imposed, as we shall do in the last chapter of this book.

*4. P. Gottardi and R. Serrano "Market Power and Information Revelation in Dynamic Trading," *Journal of the European Economic Association* V. 3, December 2005.

B can do so as well. It wouldn't be practical to screen it and charge to let people inside the screen. On the other hand, it is practically possible to exclude people from the inside of the Washington Monument. The Statue of Liberty, viewed from outside, is a nonexclusive good. If it's available for one to see, it's available to all.

National defense is an important nonexclusive good. If the person and property of U.S. resident A are being protected from foreign armies, then the person and property of U.S. resident B are also necessarily being protected from those armies.

Another important example is scientific and technological knowledge. Some technological knowledge is patentable and its use can be restricted, but a larger part is not. The technology of the internal combustion engine is nonexclusive. If it is available in person A's library, it is probably available in B's. Medical knowledge is partly nonexclusive. If a cure for infection, like penicillin, is known to A's doctor, then it is known to B's. When someone discovers a cure for some form of cancer, that cure will ultimately be public knowledge, and nonexclusive in its use. We say ultimately because it will likely be patented and restricted in use initially.

In each of these cases, when the good is there for one, it is necessarily there for all. Goods with this property are called public goods.

What are the efficiency implications of public goods? Like externalities, public goods undermine optimality in a standard competitive equilibrium. What then should be done? What are the optimal quantities of public goods? How should public goods be financed? We shall explore these and similar questions in this chapter.

2. The Public Goods Model

In this section we shall develop a rather special, and rather different model to analyze the problems of public goods. A good is public if it is by nature available to all: if one man uses it, everyone can use it. Public goods can be viewed as goods with extreme external effects: if person i's consumption of the good is X, then X appears in each and every person's utility function. However, we won't continue the externality and Pigouvian tax/subsidy analysis of the last chapter: it is more convenient to start anew.

The model we use here has both production and consumption, because one principal question we want to answer is this: How much of the public good should be produced? And the answer to the question depends both on people's demand for it, and on the nature of the productive sector of the economy. But in order to avoid notational and analytical complexity, our model will be exceedingly simple. We assume that there are only two goods, one private, and one public. Also we assume that the productive sector of the economy can transform units of the private good into units of the public good, in the ratio of one to one. And, therefore, we assume the equilibrium prices of the two goods are 1 and 1.

Our model will also make a crucial simplifying assumption about the nature of utility functions. We assume that person i's utility is the sum of the quantity of the private good he consumes, plus a well-behaved function of the quantity of the public good produced and available to all, including i. Such a utility function is said to be *separable* between private and public consumption. Some of the analysis below hinges on this special assumption; some does not. We make the assumption for two reasons: (1) it greatly simplifies all the mathematics, and (2) the discussion of demand revealing taxes breaks down without it.

Now let's develop some of the notation. It should be observed that this notation differs slightly from what is used in the exchange and production models treated above. First, we let

$$x = \text{ the quantity of the public good.}$$

Note that x is a scalar, not a vector. Also, note that x can be viewed as the quantity (or size) of the public good in physical units, or in dollars, since we assume that the prices of both the public and the private good are one. Second, we let

$$y_i = \text{ person } i\text{'s quantity of the private good.}$$

Note that y_i is a scalar. We assume that person i's utility function u_i can be written

$$u_i = v_i(x) + y_i.$$

That is, i's utility is the sum of the function v_i, which depends only on x, plus i's quantity of the private good. We also assume that v_i is continuous, smooth, monotonic, and concave; that is, it looks like the one in Figure 8.1. (Actually, monotonicity can largely be relaxed. We could assume instead that only *some* individuals' utility functions are monotonic.)

In the figure, the function v_i is smooth and concave, that is, it bends downward. The intercept a might be positive or negative, or v_i might even be asymptotic to the (negative half of the) vertical axis. At the point P, the ratio $\Delta v_i/\Delta x$ is person i's *marginal utility* from the public good, or approximately the amount by which his utility rises if the quantity of the public good is increased by 1, while his private consumption remains fixed. Note that for sufficiently small Δ's, $\Delta v_i/\Delta x$ and the slope of the v_i function at P are equal. Instead of writing MU of the

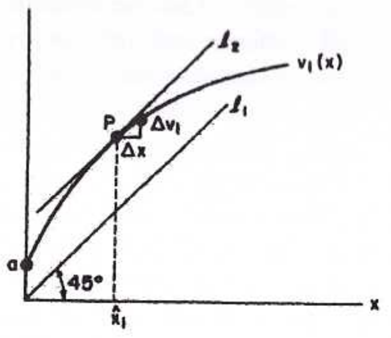

Figure 8.1.

public good for i, we now use this simpler notation:

$$v_i'(x) = \text{ MU of the public good for } i, \text{ when the quantity of the public good is } x.$$

We also assume that each person starts with an initial endowment of the private good. We let

$$\omega_i = i\text{'s initial quantity of the private good.}$$

Note that ω_i is also a scalar.

In order to be *feasible*, a vector of public and private good consumption levels $(x, y_1, y_2, \ldots, y_n)$ must satisfy this condition:

$$x + \sum_{i=1}^{n} y_i = \sum_{i=1}^{n} \omega_i.$$

(Later in this chapter, we will relax the feasibility condition to $x + \sum_{i=1}^{n} y_i \leq \sum_{i=1}^{n} \omega_i$.)

3. The Samuelson Public Good Optimality Condition

Let's now see how a Pareto optimal output for the public good can be found. In order to do this, we start by considering the inequality

$$v_1'(x) + v_2'(x) + \ldots + v_n'(x) > 1.$$

That is, we consider an output x of the public good, such that the sum of the marginal utilities of all individuals at x exceeds 1. Let's assume for mathematical simplicity that each person's private consumption y_i exceeds his marginal utility from public consumption $v_i'(x)$. Now suppose we reduce each person's consumption of the private good by an amount $v_i'(x)$; that is, we define a new amount of private good consumption for each i:

$$\bar{y}_i = y_i - v_i'(x).$$

If we stop here, then each person's utility must drop by the amount of the private good he loses, $v_i'(x)$. But we have extracted

$$v_1'(x) + v_2'(x) + \ldots + v_n'(x) = 1 + \Delta$$

units of private good from the economy, where Δ is some number greater than zero. Now suppose we take 1 unit of private good from this total (leaving Δ) and send it to the productive sector of the economy (or the firm) to be transformed into 1 unit of public good. Then we get a new public good output of

$$\bar{x} = x + 1.$$

But increasing the public good available by a unit increases each person's utility by an amount approximately equal to the marginal utility of the public good, or, for i, $v_i'(x)$. Therefore, at $(\bar{x}, \bar{y}_1, \bar{y}_2, \ldots, \bar{y}_n)$ each person is as well off as he was at $(x, y_1, y_2, \ldots, y_n)$. But Δ units of the private good are left over. This amount can be redistributed among the individuals to make some (or all) better off than they used to be at $(x, y_1, y_2, \ldots, y_n)$. Consequently, if

$$v_1'(x) + v_2'(x) + \ldots + v_n'(x) > 1,$$

it is possible to make some people better off and no one worse off through an appropriately financed increase in the public good x. Therefore, x is not a Pareto optimal output for the public good.

This result shows us when the output of the public good ought to increase. But as x increases, $v_i'(x)$ decreases for every i, because of the shapes of the v_i functions assumed in Figure 8.1. Unless all the private good is exhausted first, if x is continually increased, we eventually reach a point where

$$v_1'(x) + v_2'(x) + \ldots + v_n'(x) < 1.$$

An argument similar to the one made above establishes that, when this inequality holds, it is possible to make some people better off and no one worse off through a decrease in public good x, with the savings appropriately distributed among the individuals. Therefore, x is not the Pareto optimal level of output for the public good.

We have seen that when

$$v_1'(x) + v_2'(x) + \ldots + v_n'(x) > 1$$

x is too small, and is not Pareto optimal, and that when

$$v_1'(x) + v_2'(x) + \ldots + v_n'(x) < 1$$

x is too large, and is not Pareto optimal either. A necessary condition for Pareto optimality, therefore, is

$$v_1'(x) + v_2'(x) + \ldots + v_n'(x) = \sum_{i=1}^{n} v_i'(x) = 1.$$

That is, the sum of the marginal utilities for the public good should equal the marginal cost of producing the public good, in terms of foregone private good. Named after Paul Samuelson, this is called the *Samuelson optimality condition for public goods.*

There is another simple way to view the Samuelson optimality condition. Consider the expression

$$v_1(x) + v_2(x) + \ldots + v_n(x) - x = \sum_{i=1}^{n} v_i(x) - x.$$

This can be interpreted as the *aggregate net benefit* of the public good output level x. Now it's rather clear that if x does not maximize aggregate net benefit it cannot be Pareto optimal: If there is an alternative level of output $\bar{x}$ that gives a higher aggregate net benefit than does x, there must be some way to shift from x to $\bar{x}$ and make everyone better off in the shift. In fact, a simple application of calculus indicates that maximizing

$$v_1(x) + v_2(x) + \ldots + v_n(x) - x$$

leads directly to the Samuelson condition.

In short, maximizing aggregate net benefit

$$\sum_{i=1}^{n} v_i(x) - x$$

is necessary for achieving the Pareto optimal output of the public good, just as the Samuelson condition

$$\sum_{i=1}^{n} v_i'(x) = 1$$

is necessary for achieving the Pareto optimal x. Moreover, the two conditions are also sufficient, under our assumptions of concave utility for the public good and the equality in the definition of feasible allocations. Consequently, the Pareto optimal output of the public good is determined by the Samuelson condition, or, equivalently, the condition that aggregate net benefit be maximized.

4. Private Financing of the Public Good and the Free Rider Problem:

Now let's consider how the public good is financed. In this section we assume that it is privately purchased. This means individual i might pay for a certain amount of the public good, which would then be available to all. (Some examples of privately purchased public goods are privately owned parks open to the public, privately owned pieces of art on display in a public museum, and private contributions to charitable organizations.) We are assuming, then, that the unmodified (private) market mechanism is being used to supply the public good. In subsequent sections we'll analyze public (that is, government) financing of the public good.

To start the analysis, we suppose that person i takes the lead; he makes the first purchase of the public good. How much does he buy? He wants to choose an x to maximize his utility

$$u_i = v_i(x) + y_i$$

subject to his budget constraint

$$1 \cdot x + 1 \cdot y_i = 1 \cdot \omega_i.$$

Substituting for y_i, person i wants to maximize $u_i = v_i(x) - x + \omega_i$. The graphical solution to the problem can be seen in Figure 8.1. In that figure, the lines ℓ_1 and ℓ_2 have slope 1. Maximizing $v_i(x) - x + \omega_i$ is equivalent to maximizing $v_i(x) - x$, the vertical distance between the v_i function and the line ℓ_1. This vertical distance is greatest at the point P, where v_i is tangent to the line ℓ_2. At the tangency point, the slope of the v_i function equals the slope of the line ℓ_2, or

$$v_i'(x) = 1.$$

We shall let $\hat{x}_i$ be the quantity of the public good that i would choose to purchase privately. Note that the subscript i serves here only to remind us that i is making a private purchasing decision; the good is still public; and i cannot exclude others from enjoying the benefits of his purchase.

Now if person i has purchased $\hat{x}_i$ units of the public good what do the others do? Each of the others is enjoying the benefits from i's purchase without paying for those benefits. To be more precise, let's suppose first that $v_j'(\hat{x}_i) < 1$: That is, the marginal utility to person j from the public good is less than the marginal cost of the public good, given person i's (generous) decision to purchase $\hat{x}_i$ (and provide it to all). In this case, if j were to curtail his own private consumption by 1 unit, and purchase an additional unit of the public good for his (and everyone else's) extra enjoyment, he would be worse off. So he won't do it. That is, if $v_j'(\hat{x}_i) < 1$, person j will be content to let i buy the public good, he will not buy any himself. He will take a free ride on i's consumption of the public good. On the other hand, if $v_j'(\hat{x}_i) > 1$, then j would benefit if he curtailed his private consumption and purchased more units of the public good for his (and everyone else's) extra enjoyment. How much would he purchase? He would want the quantity of the public good increased until $v_j'(x) = 1$, since whenever $v_j'(x) > 1$, person j makes himself better off by reducing his consumption of the private good and increasing his (and everyone else's) consumption of the public good.

The final equilibrium in the unmodified private market for public and private goods will have these properties: For (at least) one person i, we will have $v_i'(x) = 1$. For all i, we will have $v_i'(x) \leq 1$. And the financing — the details of who pays how much — will largely be determined by who made the first purchase, who the second, and so on. But if $v_i'(x) = 1$ for some i, and if (as we are obviously assuming) $v_i'(x) > 0$ for the other i's, then

$$v_1'(x) + v_2'(x) + \ldots + v_n'(x) > 1.$$

That is, the sum of the marginal utilities will exceed the marginal cost of the public good and x will not be Pareto optimal. It will generally be possible to make some people better off and no one worse off through an appropriately financed increase in the public good. There will not be enough of the public good produced.

To get a clear intuitive idea of what's causing this insufficiency of public good production, think of the case where v_i is the same function for all i. In other words, everyone's tastes are the same. Now if a particular person, say person 1, takes the first step, and purchases $\hat{x}_1$ of the public good, what do the others do? They all take a free ride, since $\hat{x}_1 = \hat{x}_2 = \ldots = \hat{x}_n$. Persons 2 through n are free riders on the purchase of the public good by person 1: they enjoy the benefits and pay none of the costs. It's no surprise that 1 doesn't choose the optimal quantity of the public good; he figures his own benefit against the total cost, and pays no attention to the benefits enjoyed by the free riders. And in this particular instance, there is a very large discrepancy between

$$\sum_{i=1}^{n} v_i'(x) \text{ and the marginal cost of the public good, since}$$

$$\sum_{i=1}^{n} v_i'(x) = 1 + 1 + \ldots + 1 = n \gg 1.$$

Consequently, there is probably a large difference between $\hat{x}_1$ and the optimal quantity of the public good.

To sum up this section, the problem with the private provision of public goods is the problem of free riders. Those who enjoy the public good without paying for it never signal their desire for it. Consequently, not enough of the public good is provided. This is why public goods ought to be, in some way, publicly financed.

5. The Wicksell-Lindahl Tax Scheme

We now focus on a system in which the public good is publicly financed. A central government authority called the Public Good Board decides on the output of the public good x, and collects taxes to pay for x. What would we like this system to accomplish? First, it should somehow provide for a Pareto optimal output of the public good. Any x that is not Pareto optimal is unambiguously unsatisfactory. And, second, it ought to link a person's taxes to the benefits he receives. This is partly a matter of common sense and partly a matter of justice. A person's benefits from his private consumption are clearly linked to his payments for private consumption. And that linkage works well: it produces a Pareto optimal outcome in a private-good economy. So common sense suggests the linkage might be useful in an economy with a public good. Moreover, ever since Knut Wicksell wrote about "just" taxation in 1896, economists have occasionally suggested linking taxes and benefits because it's "just" to do so. Why should a person who gets little or no utility from the public good pay the same taxes as a person who gets lots of utility from it? As Wicksell wrote, "it would seem to be a blatant injustice if someone should be forced to contribute toward the costs of some activity which does not further his interests or may even be diametrically opposed to them."

Optimality and linkage are, then, the guiding principles in the taxation system named after Wicksell and Erik Lindahl. To explain the scheme we need a little more notation. Again, this notation is specific to our discussion of public goods, and should not be confused with similar notation we have used before.

We shall let T_i stand for person i's total tax payment to the Public Good Board. With the total tax T_i, i's budget becomes

$$y_i + T_i = \omega_i.$$

We are continuing to assume that the prices of the public and private goods are 1 and 1, respectively, so T_i can be viewed as a payment in dollars (or currency), or as a payment in units of the private good. We shall assume in this section that each person i has to bear a fractional share of the expenditure on the public good, and we shall let t_i stand for person i's share. Since the sum of the fractional shares of the individuals must be 1, we have

$$\sum_{i=1}^{n} t_i = 1.$$

Also, if i's share is t_i and the total quantity of (or expenditure on) the public good is x, then necessarily

$$T_i = t_i x.$$

Now suppose we confront individual i with this question: "If your share $t_i = 1/4$, what quantity of the public good do you want produced?" Individual i thinks to himself: "I shall pay $T_i = 1/4 \cdot x$, so the x I want is the one that maximizes $u_i = v_i(x) + y_i$ subject to $y_i + T_i = \omega_i$, or $y_i + \frac{1}{4}x = \omega_i$. In short, I want the one that maximizes $v_i(x) - \frac{1}{4}x$." The graphical solution to this problem can be found by redrawing Figure 8.1 in such a way that the lines ℓ_1 and ℓ_2 have slope $= 1/4$, rather than slope $= 1$. If we confront individual i with this question: "If your share $t_i = 1/2$, what quantity of the public good do you want produced?" He will go through the same calculations, except with 1/2 where 1/4 used to be. And he will obviously come up with a different answer. In Figure 8.2 we have drawn a v_i function and lines ℓ_1 and ℓ_2 for the general case: ℓ_1 and ℓ_2 have slope $= t_i$. The symbol $\hat{x}_i(t_i)$ represents the quantity of the public good i wants produced, given that his share of the cost is t_i.

For the general case, i wants to maximize $u_i = v_i(x) + y_i$ subject to $y_i + T_i = \omega_i$ or $y_i + t_i x = \omega_i$. In short, he wants to maximize $v_i(x) - t_i x + \omega_i$, or, equivalently, $v_i(x) - t_i x$. This quantity is greatest when the vertical distance between the v_i function and the line ℓ_1 is greatest, which occurs at the point P where ℓ_2, a line parallel to ℓ_1, is tangent to v_i. At that tangency point, the slope of ℓ_2 equals the slope of v_i. But the slope of ℓ_2 is t_i, and the slope of v_i is i's marginal utility from the public good, or $v_i'(x)$. Therefore, the point P, and i's desired quantity of the public good $\hat{x}_i(t_i)$, are determined by the equation

$$v_i'(x) = t_i.$$

A careful examination of Figure 8.2 should convince the reader of this general result: The higher is person i's share t_i, the lower is the quantity of the public good $\hat{x}_i(t_i)$ he wants produced. This makes intuitive sense; it's analogous to the Law of Demand for private goods: the higher the price of a private good, the less the individual wants to purchase, all else equal.

It is crucial to note that if the actual output of the public good happens to coincide with i's desired output of the public good, or if

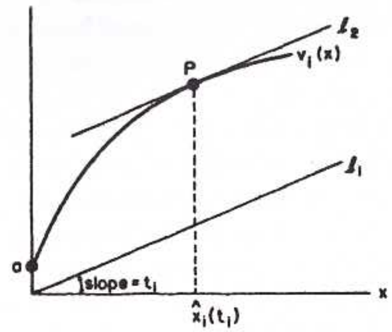

Figure 8.2.

$x = \hat{x}_i(t_i)$, then i's marginal utility from the public good $v_i'(x)$ equals his tax share t_i. That is, his tax share is linked to his marginal benefit.

But how can the actual output be made to agree with i's desired output? After all, each of the n people has his own $\hat{x}_i(t_i)$ function, his own schedule of desired outputs contingent on t_i, and given any list of fractional shares $(t_i, t_2, \ldots, t_n)$, each person will have his own particular desired output of the public good. How can the Public Good Board find a way to insure that each and every person's tax share is linked to his marginal benefit? The Wicksell-Lindahl tax scheme answers this question.

The trick of this tax scheme is to adjust the tax shares until every person agrees on the desired output of the public good. For instance, suppose there are just two people, and, when $t_1 = t_2 = 1/2$, person 1's desired public good output is $\hat{x}_1(1/2) = 10$, while person 2's desired public good output is $\hat{x}_2(1/2) = 20$. Given the shares $(1/2, 1/2)$, they disagree about the best level of output. No matter what level of output is actually chosen, for at least one of the two, his tax share will diverge from his marginal benefit. The solution? Gradually decrease t_1, and increase t_2. As t_1 decreases, 1 wants more and more of the public good produced. As t_2 increases, 2 wants less and less of the public good produced. Eventually, a point is reached where each thinks the same quantity of the public good should be produced. Say that point is reached when $t_1 = 1/3$ and $t_2 = 2/3$ and say $\hat{x}_1(1/3) = 14 = \hat{x}_2(2/3)$. Then, the Public Good Board assigns person 1 a tax share $t_1 = 1/3$, and person 2 a tax share $t_2 = 2/3$, and it has 14 units of the public good produced. Under these circumstances, each person's tax share is linked to the marginal utility he gets from the public good.

Formally, a *Lindahl equilibrium* is defined to be a vector of tax shares $(t_1, t_2, \ldots, t_n)$ and a level of output $\hat{x}$ for the public good, such that, for all i, when i's tax share is t_i his desired level of public good output equals $\hat{x}$. That is, for all i, $\hat{x}$ maximizes $u_i = v_i(x) + y_i$ subject to the budget constraint $y_i + t_i x = \omega_i$.

It is clear from its definition that a Lindahl equilibrium, or the Wicksell-Lindahl tax scheme, equates tax shares and marginal benefits for every person. That is, what people pay is connected to what they get. It should be emphasized that the connection is with marginal utility rather than total utility, but at least the linkage is there. What about optimality? Does a Lindahl equilibrium level of public good output $\hat{x}$ have to be optimal? The answer is Yes.

Recall that i's desired quantity of the public good $\hat{x}_i(t_i)$ is determined by the equation

$$v_i'(x) = t_i.$$

Therefore, since a Lindahl equilibrium $\hat{x}$ equals $\hat{x}_i(t_i)$ for every i, we must have

$$v_i'(\hat{x}) = t_i \text{ for all } i.$$

Summing over all the i's gives

$$v_1'(\hat{x}) + v_2'(\hat{x}) + \ldots + v_n'(\hat{x}) = t_1 + t_2 + \ldots + t_n = 1.$$

Consequently, the Samuelson condition is satisfied by $\hat{x}$, and $\hat{x}$ is the optimal quantity of the public good.

The Wicksell-Lindahl scheme then accomplishes the two things we set out for it: optimality and linkage. But it does have several drawbacks.

First, a Lindahl equilibrium might not exist. The adjustment process described above may not lead to a successful end. For example, if the utility functions of the two individuals are $v_1(x) = (3/4)x + \ln x$ and $v_2(x) = (1/4)x + \ln x$, there is no equilibrium. This is because the sum of the marginal utilities from the public good of the two individuals never equals 1. One can add an extra condition on preferences to guarantee existence: for example, that the sum of the marginal utilities of the individuals when x becomes large enough be less than 1. (This amounts to saying that as the public good x becomes very abundant, the social valuation for it falls below its marginal cost, something that sounds plausible).

Second, if Lindahl tax shares are viewed as "personalized prices" attached to "personalized commodities" (person i would be purchasing his amount of the public good), the equilibrium concept only makes good sense if the consumption of the public good is excludable, something that goes against the nature of many public goods. See the discussion in Mas-Colell, Whinston and Green (1995, Chapter 11) on this point.

Third, the Lindahl equilibrium concept is one in which individuals are "price-takers," i.e., they take their tax shares, or personalized prices, as given. As we will see below, price-taking behavior is a heroic assumption in this setting: in a sense, there is a problem of bilateral monopoly between the Public Good Board and each individual, to determine $\hat{x}_i$.

Fourth and finally, there is a problem of incentives. When the Public Good Board calculates the Lindahl equilibrium, when it calculates the appropriate vector of tax shares $(t_1, t_2, \ldots, t_n)$ and the output for the public good $\hat{x}$, it relies on information it receives from the individuals. It needs each individual's $\hat{x}_i(t_i)$ schedule, which we might call i's demand schedule for the public good. Or, equivalently, it needs each individual's marginal utility function v_i', or his total utility function v_i. The three functions are really interchangeable, but the Public Good Board must ask for one of them. Now put yourself in the place of a bright individual who has just been asked to provide one of these schedules, one of these functions. Suppose you know how the Public Good Board operates; you know that you will end up with a tax share t_i equal to your marginal utility $v_i'(\hat{x})$. What will you report to the Board?

If you are at all devious, you will lie about your utility, or marginal utility, or demand function. You'll say the public good is useless to you; or you'll say that your marginal utility from the public good is zero for all relevant levels of output. You won't reveal your true demand. You'll misrepresent your preferences and take a free ride. The incentives here are not compatible with truthful answers.

And, of course, if everyone is lying like mad about his demand or marginal utility function, the Public Good Board is not likely to reach an equilibrium that links tax shares to real marginal utilities, or that is Pareto optimal given the real utility functions.

Let us note that there is also an incentive compatibility or a demand-revealing problem in a model with only private goods. But it's much less severe. It might well be to a person's advantage to understate his demand for a private good in order to put downward pressure on the price of that good. But if he does so he'll probably end up with significantly less of the good. And if n is large, if there are many people, his misrepresentation will have very little effect on the price. So if he says "I don't want the good, it's of no use to me" he will in fact end up with a lot less of it than he would have otherwise, and he'll pay almost as much per unit as he would have otherwise. Contrast this situation to the public good model we've just analyzed. Suppose there are lots of people, and a particular person says to the Public Good Board "I don't want the good, it's of no use to me." His falsehood will induce the Board to produce a slightly smaller $\hat{x}$; but only slightly smaller, when n is large. And under the Wicksell-Lindahl tax scheme, his tax share will drop dramatically. So the person who is willing to conceal his demand for the public good will end up with almost as much of it as he would have otherwise, and he'll pay a lot less per unit than the would have otherwise!

The critical problem with the Wicksell-Lindahl tax scheme is this problem of demand-revelation. People will not want to reveal their true feelings to the Public Good Board. This is why we must look elsewhere for an ideal public finance scheme. We do so in the rest of this chapter, and more generally, in the last chapters of the book, on the theory of implementation.

6. Fixed Tax Shares and Majority Voting

Before continuing the search for a theoretically ideal tax scheme, let's digress slightly, and carefully examine a realistic method for financing the public good and determining the quantity that ought to be produced. In this scheme the tax shares $t_1, t_2, \ldots, t_n$ are fixed. And the amount to produce, the quantity x, is determined by majority voting.

As an example of a fixed tax shares scheme, we might have $t_i = 1/n$, for all i. That is, each person might pay an equal share of the cost of producing x. This is obviously a commonly used scheme. As another example, the t_i's might be proportional to the ω_i's. That is, we might have

$$t_i = \frac{\omega_i}{\sum_{i=1}^{n} \omega_i} \text{ for all } i.$$

The person who starts out the richest pays the highest tax; the poorest pays the lowest. Note that in both of these examples the t_i's sum to unity.

Now let's analyze what individual i wants under this system of finance. When informed that his share is t_i, i thinks of how to maximize $u_i = v_i(x) + y_i$ subject to $y_i + t_i x = \omega_i$. That is, he wants to maximize $v_i(x) - t_i x$. We've already seen in the analysis of the Wicksell-Lindahl tax scheme that person i would most prefer the level of public expenditure $\hat{x}_i(t_i)$ shown in Figure 8.2. However, if given a choice between any two public good levels x_1 and x_2, and asked to vote between those two, he would very probably vote for the one for which $v_i(x) - t_i x$ is greater.

How is an equilibrium found? Each person i has his favorite level of output $\hat{x}_i(t_i)$. For some i, this is small, for some i, it is large. For instance, if persons 1 and 2 have the same v_i function, but t_1 is big, while t_2 is small, then 1 will have a small $\hat{x}_1(t_1)$, and 2 will have a large $\hat{x}_2(t_2)$. The one who pays a bigger share will want a smaller project. Now for the sake of mathematical simplicity, let's assume here that all the $\hat{x}_i(t_i)$'s are distinct, that the people are numbered in such a way that

$$\hat{x}_1(t_1) < \hat{x}_2(t_2) < \ldots < \hat{x}_n(t_n)$$

and that the number of people n is odd. Let person M be the *median person*, the person whose $\hat{x}_M(t_M)$ is in the middle. That is, there are as many $\hat{x}_i(t_i)$'s less than $\hat{x}_M(t_M)$ as there are $\hat{x}_i(t_i)$'s greater than $x_M(t_M)$.

In the fixed tax shares, majority voting system of public finance, the Public Good Board conducts a sequence of elections, elections in which the candidates are levels of output for the public good. For simplicity, we shall assume that the list of candidates is just the set $\{\hat{x}_1(t_1), \ldots, \hat{x}_n(t_n)\}$. The Public Good Board conducts these elections until it finds a level of output $\hat{x}_i(t_i)$ which wins a majority over any other level of output $\hat{x}_j(t_j)$. The winning $\hat{x}_i(t_i)$ is the *fixed tax shares, majority voting equilibrium.*

And it turns out that the equilibrium must be $\hat{x}_M(t_M)$, that is, the desired level of output of the median voter. Let's just briefly indicate why this ought to be the case. (For a fuller treatment see Chapter 12). In what follows we will write $\hat{x}$ instead of $\hat{x}_M(t_m)$. Now consider a vote between $\hat{x}$ and some $\hat{x}_i(t_i) < \hat{x}$. Some reflection should convince you that any person $j \geq M$ will want to vote for $\hat{x}$. But since M is the median, the people numbered M or above make up a majority. So $\hat{x}$ wins a majority vote over $\hat{x}_i(t_i)$. Similarly, in a vote between $\hat{x}$ and any $\hat{x}_i(t_i) > \hat{x}$, $\hat{x}$ wins another majority, since it gets the votes of everyone numbered M or below.

What are the advantages of this system of public finance? First of all, it is relatively simple and comprehensible. It can be easily understood by the people reporting their desired quantities of the public good, and voting on those quantities. Second, unlike the Wicksell-Lindahl scheme, the incentives for misrepresentation and duplicity don't stand out like a sore thumb. But if subtle, these incentives might still be there. We have deliberately been vague about the exact nature of the voting process, or what agenda the Board uses. It is possible that, given certain agendas, people might vote against a preferred expenditure level at one stage, in order to end up with a better outcome at a later stage. This possibility becomes a virtual certainty if there are two or more public good expenditure levels being chosen simultaneously. Nonetheless, the incentives to lie are not as glaring in the fixed tax share majority voting scheme as they are in the Wicksell-Lindahl scheme.

information that each of them holds. This behavior serves to define the fine core, which is shown to be empty in many economies.

*17. A. Wolinsky "Information Revelation in a Market with Pairwise Meetings," *Econometrica* V. 58, 1990, pp. 1-23.

In this article, Wolinsky studies a trading procedure in which transactions take place in pairwise meetings of agents. Information is heavily restricted: traders have no access to observing any signal of market performance, such as a price function. Each trader tries to learn the state from his past history of observed prices when he visits different stores. The result is that, even when frictions - search costs - are removed from the model, its equilibria fail to converge to an REE.

Chapter 7

EXTERNALITIES

1. Introduction

The fundamental results of the previous chapters, the results that link optimality and competition, depend on assumptions that (1) people have self-interested utility functions, and (2) firms' production sets are unaffected by other firms' production decisions. In many actual cases these assumptions break down. When person A's utility depends on what person B consumes, or when A's technological production possibilities depend on what firm B does, there is an external effect. The decisions of one person or firm have a tangible, nonmarket impact on a different person or firm.

For example, suppose person A knows person B, and feels that person B eats too little or too much. Then person B's eating has a direct effect on A's utility level. This we call an external effect. On the other hand, if A does not know B and does not care how much B eats, B's consumption might still have some impact on A. For whenever B buys food he affects the equilibrium price of food (perhaps by a minute amount), and therefore, B's appetite indirectly influences A's utility level. However, we do not call this an external effect. In the case of consumers, an external effect is a direct effect of one person's consumption on another person's utility level, not an indirect one that operates via the price mechanism.

What is the difference? When one person's consumption affects another's welfare through the price mechanism, when B bids up the price of food for A, the system is working in the way assumed by the two fundamental theorems. If B gets more food at the expense of A because he is willing to pay more, then the price mechanism is directing the food to the person who wants it most. And the distribution of goods that results is efficient. However, if B's consumption of goods affects A's utility directly, irrespective of prices, then the price mechanism gives inappropriate signals. When B consumes food, he thinks only about his utility; he looks at the prices, and then makes a decision. But this decision has a direct impact on A's utility, and the price that B pays for food does not reflect this impact on A. Consequently, the price mechanism does not tell B of the total social benefits and costs of his actions, and the resulting distribution of goods is not, in fact, efficient.

Once we know what to look for, we can discover external effects all around us. Many of us are directly affected when we learn that other people don't have enough food to eat. When we hear that a child a thousand miles away is starving, we are worse off. When we learn that people at the other end of the country are living in tar paper shacks, we are worse off. There are externalities, then, in the consumption of food, of housing, perhaps even of medical care. Many of us are worse off when we discover that others are in severe distress because of inadequate consumption of some vitally important good.

There are also myriad mundane consumption externalities. Nonsmokers are bothered by smokers. When A smokes, B's utility level drops. Nondrinkers are occasionally bothered by drinkers. When A drinks to excess, B, a member of the Women's Christian Temperance Union, feels worse off. Those who prefer classical music are bothered by those who play rock music. Many of us are bothered by loud exhaust noises of cars, trucks, and motorcycles. When A drives his car with modified (amplified) exhaust pipes through town, hundreds or thousands of B's might be made briefly worse off.

Externalities among firms are common. One standard story has two firms located on the same river. Firm A, the upstream firm, dumps its wastes in the river, while firm B, the downstream firm, uses river water for washing and otherwise processing its outputs. If firm A increases its output (and its wastes), firm B's production suffers. To produce the same output with dirtier water, firm B must use more chemical agents, more labor, and more electricity. That is, firm B's production set shrinks. If firm A pays nothing for dumping wastes in the river, it receives no information from the price system about the external costs it is imposing on firm B. The consequence is that the price mechanism no longer ensures efficiency. (Note that if firm A were downstream, this externality problem might not arise.)

A second standard story has two firms sharing the same air. Firm A is an old-fashioned electric generating facility that burns coal and uses no scrubbing or antipollution devices. Consequently, every kilowatt hour produced results in a belch of black smoke. Firm B is a laundry located nearby. When firm A produces more electricity (and more smoke), firm B has to cope with more dirt and grime settling down on its plant, in its machines, on its tables and presses, and on the clothing being cleaned. So firm B must use more soap, more labor, more wrapping paper, and so on, to produce the same output of clean garments. In other words, when A's output rise, firm B's production set shrinks. But A does not take these costs for B into account in its decisions, so the price mechanism provides it with misleading information. It acts as if the air is free.

A very important type of externality occurs when a firm's production decisions have direct nonmarket effects on a person's utility level. For example, a firm that stripmines for coal without reclaiming the land affects the utility levels of people who see the results. A firm that produces smoke affects the utility levels of people who breathe the smoke. The residents of Chicago, Illinois are directly affected by the output decisions of steel mills in Gary, Indiana, whenever the wind is from the east. People who live near the Three Mile Island Nuclear Reactor in Pennsylvania might be affected by the production decisions of the firm that operates that plant. People who live near coal-burning electric generators are often affected by the output decisions of those firms.

Not all externalities of this firm-person type are harmful or negative. People who live downwind from a bakery might be happier when bread production is high. Firms that build attractive plants or office buildings make people who look at those buildings better off. Many of the impressive and exciting sights of a large city are the skyscrapers built by private firms, such as the Empire State Building and the Chrysler Building in New York, the Sears and John Hancock towers in Chicago. Much of New England is dotted with handsome nineteenth century mill buildings, which still provide viewing pleasure long after the firms that built them went bankrupt or moved away.

In all of these cases, whether the externalities are in consumption or in production, whether they are positive (beneficial) or negative (hurtful), the price mechanism does not provide complete enough information to the decision maker. In the case of negative externalities, the price mechanism does not tell the decision maker how much his decision really costs. In the case of positive externalities, the price mechanism does not tell the decision maker how much his decision really helps. And it follows that the link between competition and optimality is broken.

In this chapter we shall carefully analyze two examples of external effects, one in an exchange model and one in a production model. The examples will illustrate how the external effects destroy the optimality of a competitive equilibrium.

However, the existence of externalities does not mean that markets must be disbanded. Abolishing the price mechanism because air pollution is bothersome and because flower gardens are pleasurable would be throwing the baby out with the bath water.

Since the early twentieth century economists have advocated taxes and subsidies to correct important externality-induced inefficiencies. The idea is that those who harm others through their production or consumption decisions should pay a tax to reflect that harm. The size of the tax should depend on the extent of the harm: these are not lump sum taxes like the bank balance transfers of the Second Fundamental Theorem. With the tax in their figuring, the decision makers would be led, via the tax-modified price system, to the right decisions. They would take into account the real social costs of their decisions. Similarly, consumers and firms that create external benefits should be subsidized to reflect those benefits. Again, the extent of the subsidies should depend on the extent of the benefits. With the subsidies in their figuring, the decision makers would be led, again, to the right decisions.

We will show how the appropriate taxes or subsidies ought to be figured in each of our examples, and we will show how the tax- or subsidy-modified price mechanism once again produces an optimal distribution of goods, or an optimal production plan.

2. Externalities in an Exchange Economy: An Example

We now look at what happens when the self-interestedness assumption is relaxed in an exchange economy. To illustrate the problem, we construct a simple two-person two-goods example.

Let

$$u_1(x) = x_{11}x_{12} + x_{21} \qquad \omega_1 = (10, 0)$$
$$u_2(x) = u_2(x_2) = x_{21}x_{22} \; \omega_2 = (0, 10).$$

Person 1 is altruistic; he gets some pleasure out of 2's consumption of good 1. Person 2, on the other hand, is self-interested.

To start the analysis, we solve for the set of Pareto optimal allocations. First, we rewrite u_1 as

$$u_1(x) = x_{11}x_{12} + x_{21} = x_{11}x_{12} + (10 - x_{11}).$$

Next we calculate the marginal rate of substitution of good 2 for good 1 for person 1:

$$\text{MRS person 1} = \frac{\text{MU of good 1}}{\text{MU of good 2}} = \frac{x_{12} - 1}{x_{11}}$$

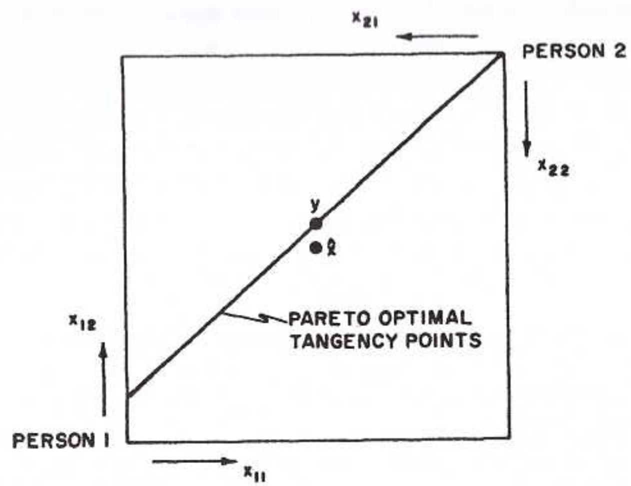

Figure 7.1.

and for person 2:

$$\text{MRS person 2} = \frac{\text{MU of good 1}}{\text{MU of good 2}} = \frac{x_{22}}{x_{21}}.$$

Then, to find the locus of tangency points of the two individuals' indifference curves, we set MRS for person 1 equal to MRS for person 2:

$$\frac{x_{12} - 1}{x_{11}} = \frac{x_{22}}{x_{21}} = \frac{10 - x_{12}}{10 - x_{11}}.$$

Solving this equation for x_{12} in terms of x_{11} gives

$$x_{12} = 1 + \frac{9}{10}x_{11}.$$

Figure 7.1 shows this locus of Pareto optimal tangency points and two other allocations to which we shall soon refer.

The next step is to calculate a competitive equilibrium. We assume that 1 and 2 act as price takers, and at this point we also suppose that 1 does not know that $x_{11} + x_{21} = 10$, and does not know that

$$u_1(x) = x_{11}x_{12} + (10 - x_{11}).$$

We must make this somewhat artificial assumption in order to discuss external effects in the simple two-person case; if we do not make it, all apparently altruistic (or malevolent) utility functions could be rewritten and solved as self-interested ones. If there were three or more people this artificiality would disappear.

Person 1 wants to maximize $x_{11}x_{12} + x_{21}$ subject to his budget constraint. Since he cannot choose x_{21} himself (2 has something to say about it too), we suppose he contemplates buying a quantity g of the second good, and giving it to 2. Person 2's consumption of the first good will be the quantity he purchases, which 1 cannot control, plus the gift g. Therefore, 1 wants to maximize $x_{11}x_{12}$ plus g, subject to the budget constraint

$$p_1x_{11} + p_2x_{12} + p_1g \leq p_1 \cdot 10 + p_2 \cdot 0 = 10p_1.$$

He will clearly want to spend all his income 10 p_1, and we can therefore ignore the inequality.

Person 2, on the other hand, simply wants to maximize $u_2(x_2) = x_{21}x_{22}$, subject to $p_1x_{21} + p_2x_{22} = p_1 \cdot 0 + p_2 \cdot 10 = 10p_2$.

A competitive equilibrium in this economy is a price vector $p = (p_1, p_2)$ and consumption (+ gift) vectors $(\hat{x}_{11}, \hat{x}_{12}, g)$, $(\hat{x}_{21}, \hat{x}_{22})$, such that $(\hat{x}_{11}, \hat{x}_{12}, g)$ maximizes u_1 subject to 1's budget equation, and $(\hat{x}_{21}, \hat{x}_{22})$ maximizes u_2 subject to 2's budget equation.

To solve for the equilibrium, note that if x_{12} is grater than 1, person 1 will choose $g = 0$, that is, he will give nothing. This is so because the marginal utility of 1's private consumption of good 1 is x_{12}, whereas the marginal utility of 1's charity is 1. But it is obvious from inspection that x_{12} will be chosen greater than 1. Therefore, 1 will attempt to maximize the private part of his utility function $x_{11}x_{12}$ subject to the constraint

$$p_1x_{11} + p_2x_{12} = 10p_1.$$

The perfect symmetry of the two individuals' maximization problems, and the symmetry of ω, lead to the conclusion that the competitive equilibrium is

$$p = (1,1)$$
$$\hat{x}_{11} = \hat{x}_{12} = \hat{x}_{21} = \hat{x}_{22} = 5$$
$$g = 0.$$

The allocation $\hat{x}$ is shown in Figure 7.1. Note that $\hat{x}$ is not on the locus of tangencies in that figure; it does not satisfy the tangency condition

$$x_{12} = 1 + \frac{9}{10}x_{11}.$$

Consequently, it is not Pareto optimal. When there are externalities present, a competitive equilibrium allocation need not be Pareto optimal, and the First Fundamental Theorem breaks down.

Recall that the exchange version of the Second Fundamental Theorem of Welfare Economics says that any Pareto optimal point can be viewed as a competitive equilibrium allocation given appropriate cash transfers. Now consider the allocation y defined by

$$y_1 = (5, 5.5), y_2 = (5, 4.5).$$

Note that

$$5.5 = y_{12} = 1 + \frac{9}{10}y_{11} = 1 + \frac{9}{10} \cdot 5.$$

That is, y satisfies the tangency condition; it is on the locus of tangencies in Figure 7.1. Therefore, y is Pareto optimal. Do there exist cash transfers T_1 and T_2, such that individual 1 will finish at y_1, and individual 2 will finish at y_2, when they maximize their utilities subject to their (adjusted) budget constraints

$$p_1x_{11} + p_2x_{12} + p_1g \leq 10p_1 + T_1$$

and

$$p_1x_{21} + p_2x_{22} \leq 10p_2 + T_2?$$

It should again be clear that the gift 1 chooses will be zero, so g may be ignored. To see whether we can induce 1 and 2 to go to y via the competitive mechanism with cash transfers, we first see what is required to induce person 2 to go to y_2:

Person 2 wants to maximize $u_2(x_2) = x_{21}x_{22}$
subject to $p_1x_{21} + p_2x_{22} \leq 10p_2 + T_2$.

He sets his marginal rate of substitution x_{22}/x_{21} equal to the price ratio. But his MRS at y_2 is 4.5/5, so we have

$$\frac{4.5}{5} = \frac{9}{10} = \frac{p_1}{p_2}.$$

In order to get person 2 to the point y_2, the price ratio must be 9/10. Let's normalize prices by setting $p_2 = 10$. Then p must equal (9, 10). Substituting these prices and quantities in 2's budget equation gives

$$9 \cdot 5 + 10 \cdot 4.5 = 10 \cdot 10 + T_2$$

so $T_2 = -10$. In short, person 2 will choose y_2 if $p = (9, 10)$ and $T_2 = -10$.

However, if the prices are (9, 10) person 1 will *not* choose the point $y_1 = (5, 5.5)$. Since he will not give a gift, he will attempt to maximize the private part of his utility function, that is, $x_{11}x_{12}$. His private marginal rate of substitution at y is $5.5/5 = 11/10$, which differs from the price ratio $9/10$. He will never choose y_1 when $p = (9, 10)$. Since person 2 will choose y_2 only if $p = (9, 10)$ and person 1 will never choose y_1 when $p = (9, 10)$, the two people cannot be induced to move to y through a cash transfer arrangement. When there are externalities present, a Pareto optimum need not be achievable via the cash transfer modified competitive mechanism, and the Second Fundamental Theorem breaks down.

3. Pigouvian Taxes and Subsidies: The Exchange Example Continued

Externalities weaken the link between competition and optimality. But all is not lost. There is remedy that is consistent with a decentralized price mechanism. It is not necessary to have a Central Bureau (a super Environmental Protection Agency) to direct the consumption decisions of all individuals. The decentralized remedy is the introduction of per-unit (or marginal) *taxes* or *subsidies* on the consumption of the goods that induce the external effects. These taxes or subsidies cannot be of the lump sum cash transfer variety; we saw above that simple cash transfers won't get us to an optimal allocation. They must depend on the quantities actually consumed, for they must affect the relative prices paid by different individuals. They must be designed to encourage a person's consumption of a good if that consumption has positive external effects, and to discourage a person's consumption of a good if that consumption has negative external effects.

With this general motivation in mind, let's return to the example of the previous section. In that example, person 2's consumption of good 1 confers an external benefit on person 1. But when person 2 makes his consumption decisions in the standard competitive equilibrium model, he does not take the external benefit into account. This suggests that this consumption of good 1 ought to be subsidized.

Let s be a subsidy paid to person 2 for each unit of good 1 that he consumes. The link between competition and optimality will be rebuilt if there exist an s, a price vector p, and cash transfers T_1 and T_2, so that, when 1's budget constraint is

$$p_1 x_{11} + p_2 x_{12} \leq 10p_1 + T_1$$

and 2's budget constraint is

$$(p_1 - s)x_{21} + p_2 x_{22} \leq 10p_2 + T_2$$

the two will move to a Pareto optimal allocation through the competitive mechanism. In fact, there do exist such s, p, T_1, T_2, and, with the appropriate choice of T_1 and T_2, one can move to whatever interior Pareto optimal allocation one desires.

For example, suppose

$$y_1 = (5, 5.5) \text{ and } y_2 = (5, 4.5)$$

is again the goal. It can be achieved this way. First consider person 1. He will again maximize the self-interested part of his utility function. The marginal rate of substitution condition is

$$\frac{x_{12}}{x_{11}} = \frac{p_1}{p_2}$$

and his budget equation is

$$p_1 x_{11} + p_2 x_{12} = 10p_1 + T_1.$$

If we let $T_1 = 0$, the two equations imply $x_{11} = 5$, and $x_{12} = 5p_1/p_2$. We want him to choose $y_{11} = 5$, $y_{12} = 5.5$. Therefore, if we normalize prices by setting $p_2 = 10$, we must have $p_1 = 11$.

Now turn to person 2. The marginal rate of substitution condition and the budget constraint for 2 are

$$\frac{x_{22}}{x_{21}} = \frac{p_1 - s}{p_2}$$

and

$$(p_1 - s)x_{21} + p_2 x_{22} = 10p_2 + T_2.$$

But $p_1 = 11$ and $p_2 = 10$. Moreover, we want person 2 to choose the point $y_{21} = 5$, $y_{22} = 4.5$. Substituting these values in the above equations gives

$$\frac{4.5}{5} = \frac{11 - s}{10}$$

and $(11 - s) \cdot 5 + 10 \cdot 4.5 = 10 \cdot 10 + T_2$. Consequently, we can let $s = 2$ and $T_2 = -10$.

In short, if $T_1 = 0$, $T_2 = -10$, and if $s = 2$, the competitive mechanism, modified by T_1, T_2 and s, will take the economy to the Pareto optimal allocation y. Thus the introduction of the subsidy s re-establishes the link between competition and optimality.

These particular calculations are not especially intuitive; for policy applications we ought to have some simpler concepts to guide the choice of taxes and subsidies. To derive those concepts, we shall carry the example through a few steps further.

Recall that person 1's utility function is $u_1(x) = x_{11}x_{12} + x_{21}$. The marginal utility to person 1 of person 2's consumption of good 1 is therefore 1. The marginal utility to person 1 of his own consumption of good

1 is, in general, equal to x_{12}. At the allocation $y = ((5, 5.5), (5, 4.5))$, the marginal utility to person 1 of his own consumption of good 1 is then 5.5. Now if person 2 reduced his consumption of good 1 by a unit, person 1 would have to increase his consumption of good 1 by 1/5.5 units in order to remain as well off as before. That is, there is a well-defined marginal rate of substitution of 1's own consumption of good 1 for 2's consumption of good 1, equal to person 1's marginal utility from his own consumption of good 1 divided by person 1's marginal utility from person 2's consumption of good 1. We call this particular marginal rate of substitution the *marginal external benefit* or MEB of person 2's consumption of good 1:

$$\text{MEB} = \frac{\text{MU to person 1 of person 2's consumption of good 1}}{\text{MU to person 1 of his own consumption of good 1}}.$$

This gives us a measure of the benefit provided to 1 by 2's consumption, in terms of units of good 1. At y, the MEB $= 1/5.5$.

Now suppose we ask this question: How many dollars (or other units of currency) would person 1 have to be given to just compensate him for person 2's reducing his consumption of good 1 by 1 unit? The answer, of course, is $p_1 \cdot$ MEB, which we shall call MEB in dollars. At the point y, if $p_1 = 11$, then

$$\text{MEB in dollars } = p_1 \cdot \text{ MEB } = 11 \cdot \frac{1}{5.5} = 2.$$

But we found above that the required subsidy s was 2 (that is, \$2/unit). In fact, the intuitive rule for finding the right subsidy is given by the formula

$$\text{MEB in dollars } = s.$$

The subsidy should just equal the value (in dollars) of the (marginal) external benefit. This makes sense: if person 2 is causing \$2 worth of external good for every extra unit of good 1 he consumes, then the appropriate way to achieve optimality is through a subsidy of \$2 per unit on each extra unit he consumes. The externality is in this way internalized; it is plugged back into the calculation of the decision maker.

What would happen if the externality were negative? Suppose, for instance, that person 1's utility function were

$$u_1(x) = x_{11}x_{12} - x_{21}.$$

Then every extra unit of good 1 that person 2 consumes would make person 1 worse off, and person 1 would need to be given more of good 1 to compensate him for an increase in good 1 consumption by person 2. We would then have a *marginal external cost* associated with 2's consumption, defined as follows:

$$\text{MEC } = \frac{\text{Marginal disutility to person 1 of person 2's consumption of good 1}}{\text{MU to person 1 of his own consumption of good 1}}.$$

(Or, equivalently, we could have a negative MU in the numerator and therefore a negative MEB.) We would find MEC in dollars the same way as before:

$$\text{MEC in dollars } = p_1 \cdot \text{ MEC}.$$

And the tax required to correct the externality problem would be

$$\text{MEC in dollars } = t.$$

For each extra unit of good 1 that he consumes, person 2 would be required to pay a tax of t, and this tax, like the subsidy above, would internalize the externality. That is, it would be plugged into the calculations of the decision maker who is responsible for the externality.

To summarize these results: In order to re-establish the connection between competition and optimality when there are externalities present, per-unit taxes or subsidies can be imposed on consumption, and they should be chosen so that

$$\text{MEC in dollars } = t$$

or

$$\text{MEB in dollars } = s.$$

This idea of introducing taxes to take care of the market failure caused by the presence of externalities was developed by Arthur Pigou in the 1920's.

4. Pigouvian Taxes and Subsidies: A Production Example

Now let's work through a simple production externality example. We assume there are two firms and two goods. Both firms use good 1 to produce good 2. Firm 1 can be viewed as the "upstream" firm. Its production set Y_1 is determined by

$$y_{12} \leq \sqrt{-y_{11}}.$$

For instance, with nine units of the input good 1 ($y_{11} = -9$), it can produce up to three units of the output good 2. Firm 2 can be viewed