Chapter 5

Examples of the consumer choice model

The theory of consumer choice is the foundation of a large range of economic models. Although we have introduced the theory using the language of consumption (where the consumer chooses between different ways of spending her money), some of the most interesting applications involve quite different interpretations of the goods the consumer chooses among.

In this chapter we will apply the model of consumer choice to develop theories of labor supply, insurance, and finance. For example, the theory of a worker's decision to supply labor by finding a job and earning a wage income is treated as a special case of the consumer choice model. Here the consumer is thought of as choosing between different ways to spend her time, in wage work, or not in wage work. Time spent outside of wage work may also be work (for example, childcare, house and car repair), but economists often refer to this time as "leisure". The consumer's choice is then between income (earned by working for a wage) and non-wage time.

Economic explanations of risk-taking and risk-avoidance, gambling and insurance rest on the consumer choice model as well. The consumer is thought of as choosing between consumption in different states of the world, corresponding to the outcome of various risky prospects. The relative price of consumption in different *states of the world* depends on the schedule of insurance premiums or gambling odds the consumer faces.

The modern theory of finance grows out of the general economic theory of risk. A wealth-holder choosing a portfolio of assets can also be viewed as choosing a pattern of income across different states of the world (say, recession or boom, high interest rates or low, oil embargo or no oil embargo). The price of an asset in this theory depends on the pattern of income it provides in different states of the world and on prices portfolio holders are willing to pay for income in those states.

5.1 The theory of labor supply

5.1.1 The labor supply of a single worker

In using the consumer choice model to explain the labor supply of an individual worker, we assume that the worker starts with an endowment of time, m_1, the number of hours in a week or year, and money income that does not depend on her working (interest or dividends, for example), m_2. For the moment, let us assume that $m_2 = 0$, so that the worker's only source of income is wages. This assumption is represented in Figure 5.1 by the location of the endowment point on the non-wage time axis.

If the worker can choose how many hours to work (which, or course, is not true of many lines of work, where workers must put in a set number of hours per week in order to hold the job), her income rises for each hour worked by the amount of the wage, w, she can earn. Remembering that $m_2 = 0$, her budget constraint is

$$wx_1 + x_2 = wm_1$$

The expression for the budget constraint is exactly like the general consumer choice model, and the wage, w, takes the place of the price of good 1, p. (Since the worker earns money income, $p_2 = 1$, just as in the two-good model where good 2 is other income.) You could think of the worker as selling all her time for the wage, and then buying back some part of it as non-wage time. The price of non-wage time (or "leisure") is the wage she can earn. The budget constraint is also shown in Figure 5.1. The consumer's preferences between working for a wage and non-wage time are represented by her utility function $u(x_1, x_2)$. If, for example, she has perfect complements preferences between income and non-wage time, she would insist on a rigid proportion between the time she had to herself and her money income. In Figure 5.1 some general indifference curves are drawn.

The marginal rate of substitution in this model can be thought of as the consumer's *reservation wage*, the lowest wage offer she will accept to work another hour. If she has convex preferences, her reservation wage will become higher and higher as she works more, as illustrated in Figure 5.1. In order to maximize her utility, given the wage she can earn and her time endowment, the worker continues to supply labor to the market until her reservation wage just equals the market wage, as shown in Figure 5.1.

As the wage changes, the worker's labor supply moves along her wage-offer curve, which depends on the exact utility function that represents her preferences between income and non-wage time. Figure 5.2 shows a situation in which a rise in the wage has no effect on the number of hours per week the worker works for a wage.

5.1.2 Income and substitution effects in labor supply

The theory of labor supply shows the importance of income and substitution effects in understanding how individual choice depends on market prices. A

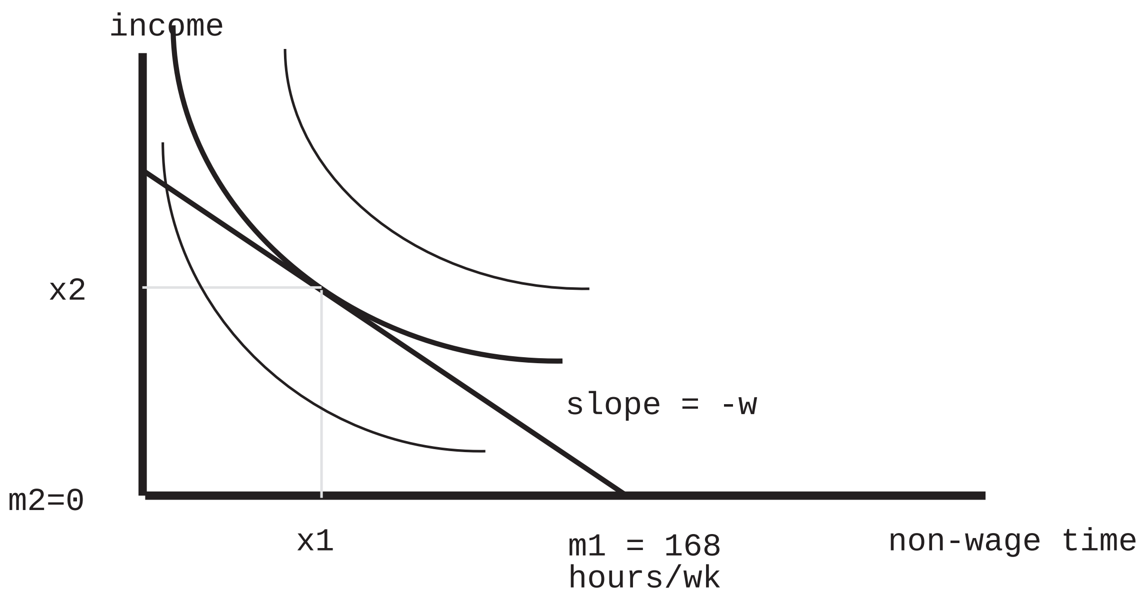

Figure 5.1: The labor supply decision can be thought of as a consumer choice between income and non-wage time, where the price of non-wage time is the wage, w. The worker equates her mrs to the wage at the point (x_1, x_2). She has x_1 non-wage time, and works $m_1 - x_1$ hours a week.

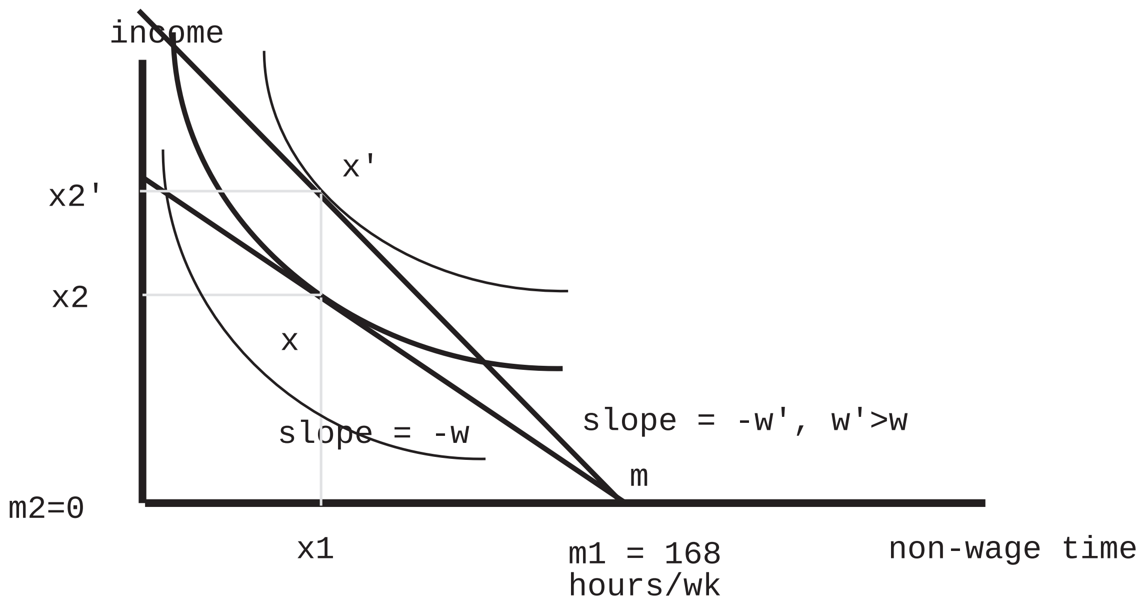

Figure 5.2: A rise in the wage rotates the worker's budget line through the endowment point. In this case the worker moves to higher indifference curve where she has higher income, and the same amount of non-wage time, because she works the same number of hours.

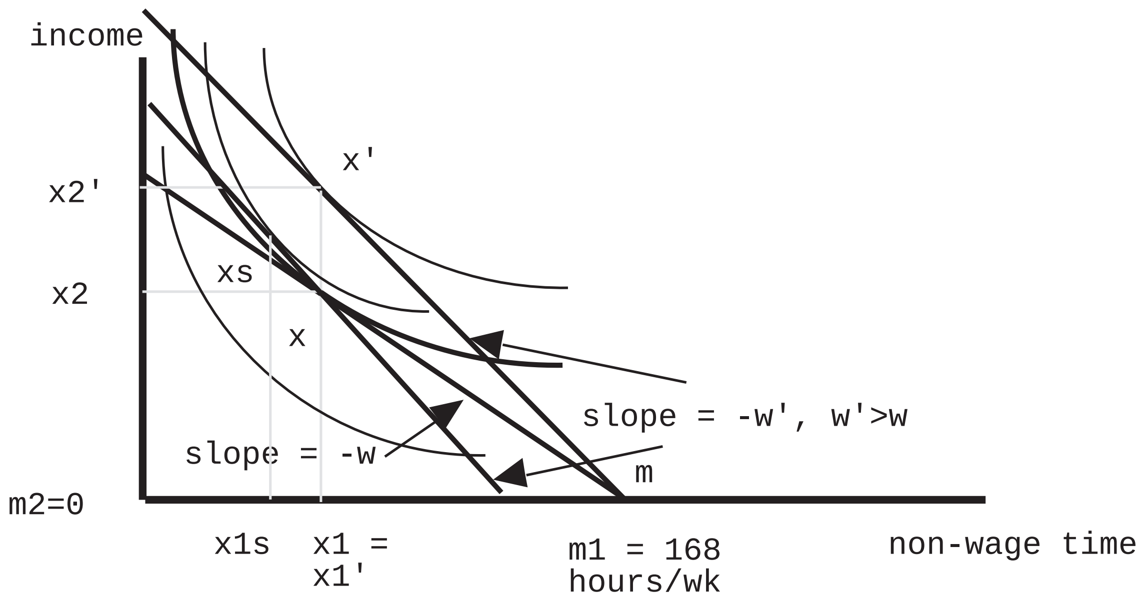

Figure 5.3: We find the substitution effect of a rise in the wage by rotating the budget line through the worker's original labor supply point. Notice that this requires us to lower the worker's income to make up for the rise in the wage. The substitution effect leads the worker to work more. The income effect leads the worker to consume more non-wage time, and with these preferences offsets the substitution effect exactly.

rise in the wage is a rise in the price of non-wage time. Therefore, we might suppose that the worker would choose to work more and have less of the higher-priced non-wage time. Indeed, that is exactly what would happen if the change depended only on the substitution effect. In Figure 5.3 the substitution effect is shown by rotating the budget line through the worker's original work-non-wage time choice. The effect is to reduce non-wage time and increase working hours.

The actual reaction of the worker to the rise in the wage also involves the income effect as well. The worker's endowment consists entirely of potential working time. A rise in the wage therefore increases her income potential income. At this higher income she chooses more non-wage time. As the Figure is drawn, the income effect exactly offsets the substitution effect, so that her labor supply is unaffected by the rise in the wage.

5.1.3 General equilibrium with identical workers

The model of labor supply can also be interpreted as a simple general equilibrium model of the entire economy. In order to make this interpretation, we have to assume that there are many workers in the economy, but that they all have the same preferences between income and non-wage time, the same endowment (in this case the same amount of time and zero other income) and all have the same waqe. In order to see this model as a general equilibrium, we must assume that an hour of labor by itself, without capital or land, can produce w units of money income (or the bundle of goods for which money income stands). In this case the wage will just be equal to the amount an hour of labor can produce. It would be impossible to pay a wage higher than labor's productivity, and profit making entrepreneurs would demand an unlimited amount of labor if the wage were smaller than labor productivity. In this interpretation w is both the average and the marginal product of labor.

With this interpretation the model shows how worker preferences in this simple economy determine the total output. If technological change were to raise average labor productivity in this model, the economy would move along the typical worker's price-offer curve. A rise in labor productivity might increase or decrease the typical worker's supply of labor, depending on the relative strength of the income and substitution effects in the typical worker's preferences.

5.1.4 Effects of a proportional income tax

The interpretation of the labor supply theory in terms of a whole economy of essentially identical workers can also be used to analyze the effects of a proportional income tax. If the government taxes wage income at rate t, with no exemptions or deductions, the budget constraint of the typical worker becomes

$$x_2 = w(1 - t)(m_1 - x_1)$$

As Figure 5.4 shows, the effect of the tax alone is definitely to make the typical worker worse off. The worker's labor supply and the total output change according to the wage-offer curve (which is the analogue of the price-offer curve) corresponding to her preferences between non-wage time and income.

In reality governments impose taxes in order to finance activities that have some benefit to the public. A complete analysis of the effect of a tax on the welfare of the typical worker depends on the use to which the tax revenue is put. In order to focus on the pure effects of taxes, economists often analyze the unrealistic case where the government simply gives the money back to the typical worker as a lump-sum subsidy. The difference between the income tax and the lump-sum subsidy is that the income tax a worker pays depends on her labor supply, while her lump-sum subsidy does not: she gets the same lump-sum no matter how much she works. Of course, because we have assumed that all the workers are identical, they will all wind up working the same amount and paying the same income tax as well, but in principle the amount of income tax

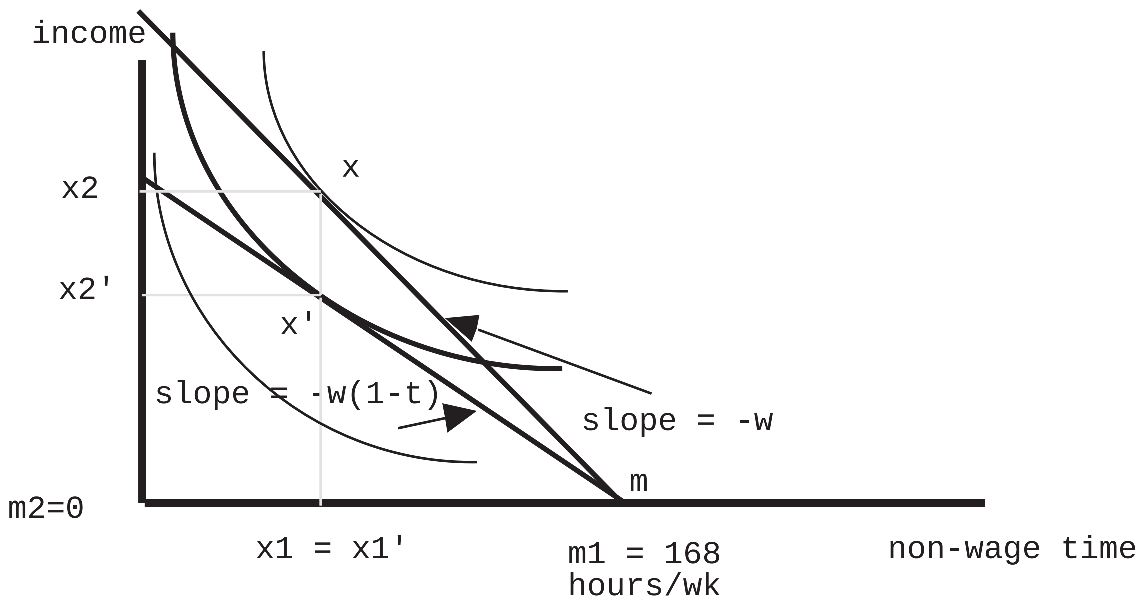

Figure 5.4: A proportional income tax in an economy with a large number of identical workers lowers the after-tax wage. The typical worker's reaction is to move along her wage-offer curve to a lower indifference curve. In the example shown here, the income effect of the tax offsets the substitution effect, and the worker's supply of labor is unchanged, while her after-tax income declines.

each worker pays depends on her labor supply. The lump-sum subsidy is just an increase in each worker's endowment of other income.

With the lump-sum subsidy, we need to consider three budget lines. The actual budget line the worker faces taking account of the tax and the subsidy m_s is

$$x_{2t}^s = (1 - t)w(m_1 - x_1) + m_s$$

The budget line she would face if there were no subsidy is

$$x_{2t} = (1 - t)w(m_1 - x_1)$$

The budget line she would face if there were neither tax nor subsidy is

$$x_2 = w(m_1 - x_1)$$

Notice that the tax the government collects, $tw(m_1 - x_1)$, is just the difference between the worker's income before taxes, x_2 and her income after tax but before receiving the subsidy, x_{2t}.

Figure 5.5 shows the analysis of an income tax rebated as a lump-sum subsidy.

In Figure 5.5 the typical worker's budget line with the tax and subsidy has a slope equal to $-w(1-t)$ and is shifted upward by the amount of the subsidy. The taxed equilibrium $x_t = (x_{t1}, x_{t2})$ is the point on this budget line that is tangent to the indifference curve passing through it. The tax paid when the worker supplies $(m_1 - x_1)$ hours of labor is $tw(m_1-x_1) = w(m_1-x_1)-(1-t)w(m_1-x_1)$, which is just the difference between the no-tax budget line and the budget line with the tax but without the lump-sum subsidy. Since we are assuming that the lump-sum subsidy is equal to the tax collected, the no-tax budget line must also go through the taxed equilibrium point xt, to make sure that the subsidy $m_s = tw(m_1 - x_{t1})$.

Once we see that the no-tax budget line passes through the taxed equilibrium point, we can also see that the worker is worse off with the tax and subsidy than she would be without either the tax or the subsidy. The no-tax budget line has a different slope from the budget line with the tax and the subsidy, so it must cut through the indifference curve at the taxed equilibrium. This implies that there are points on the no-tax budget line that the typical worker prefers to the taxed equilibrium.

This analysis is a way of demonstrating the *deadweight burden* associated with any tax that influences economic decisions. The loss of welfare to the typical worker under the proportional income tax exceeds the tax revenue collected by the government, as the analysis shows clearly by assuming that the tax revenue is just given back to the worker as a lump-sum. From an economic viewpoint, the government is justified in taxing the worker only if it can provide a service that gives the worker more utility than she would enjoy as a result of getting back the taxes as a gift

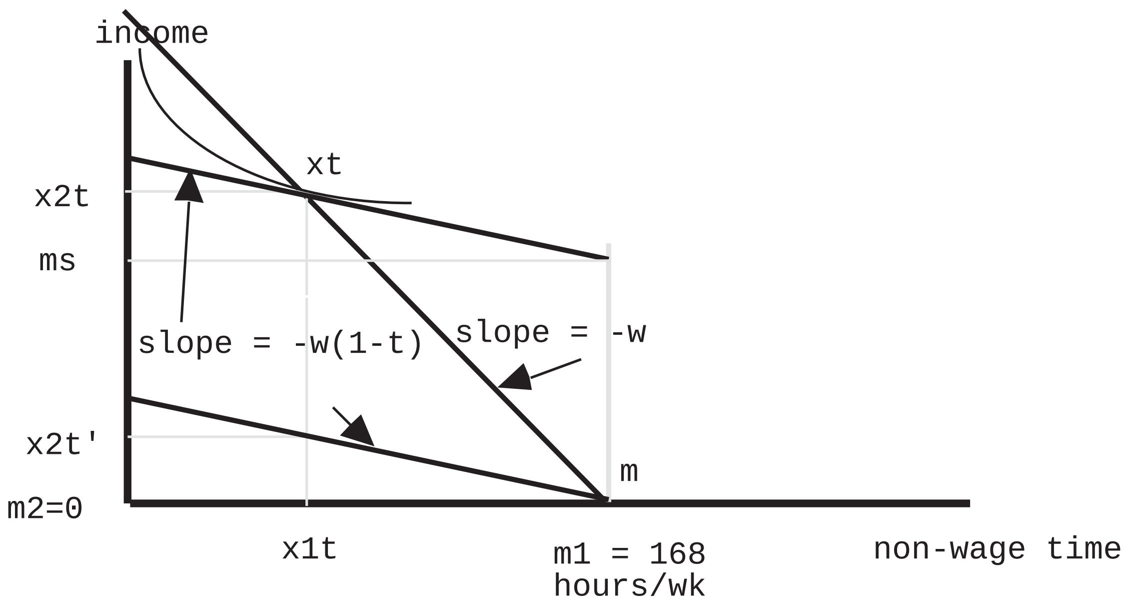

Figure 5.5: When the government gives the revenues from a proportional income tax back to the workers as a lump-sum subsidy, the typical worker's budget line has the slope $w(1-t)$, because she is paying income tax on her wages. The tax collected is $tw(m_1 - x_{1t})$. The equilibrium is at a point x_t where the worker's indifference curve is tangent to the after-tax budget line, and the after-tax budget line intersects the no-tax budget line. Since the no-tax budget line has a different slope, it cuts the indifference curve, showing that the worker is worse off with the tax and subsidy than she would be with no tax.

5.2 The theory of choice with uncertainty

In order to use the consumer choice model to study choice under uncertainty, we have to interpret the goods as contingent goods. Let us consider a consumer facing a choice involving an uncertain event, such as whether or not her house will burn down, or whether or not a certain horse will win a race. We can interpret x_2 to be the consumer's wealth if the event does not occur, and x_1 to be her additional wealth if the event occurs. When $x_1 = 0$ the consumer is bearing no risk at all. A movement of x_1 away from zero in either direction represents an increase in the consumer's risk.

With this interpretation, the consumer's endowment expresses her position before she has made any economic transaction, such as buying insurance, or betting on the horse race.

Consider a consumer facing the choice of gambling on a horse race. Suppose that she can buy tickets that will be worth \$1 if a certain horse wins the race, and that the price of a ticket is p. In this case x_2 is her wealth if the horse loses and x_1 is the number of tickets she purchases. It is natural to suppose that $m_1 = 0$, since if she does not purchase any tickets, the outcome of the horse race will make no difference to her wealth.

If the consumer buys x_1 tickets, she will have $x_2 = m_2 - px_1$ if the horse loses and x_1 extra wealth if the horse wins. Thus her budget constraint is just

$$px_1 + x_2 = m_2$$

The price p in this case represents the amount the consumer must bet in order to collect a dollar if the horse wins, or the *price* at which she can gamble.

Gamblers frequently express the price p of a bet as the ratio $(1-p)/p$, which they call the *odds against* the horse. For example, if $p = \$.50$, the odds are 1-to-1. If the odds are 3-to-1 against the horse, $p = \$.25$.

This situation is illustrated in Figure 5.6.

In the case of a consumer facing the risk of her house burning down, we would expect m_1 to be large and negative, since the consumer will have a major loss if the fire occurs, as in Figure 5.7.

The consumer who faces the risk of a fire can change her position by buying insurance. If the premium is p for \$1 of insurance, the situation is exacty analogous to betting on a horse race. In this case the consumer pays p for tickets (the insurance policy) paying \$1 in case her house burns down. (In fact, the insurance policy will specify a limit on its total liability.) When she buys $\$y$ of insurance coverage at a premium, p, the insurance company agrees to pay her $\$y$ if the fire does occur, thereby increasing her certain income. If she buys $\$y$ of coverage at a premium p her wealth in case the house does not burn down is $x_2 = m_2 - py$, and her extra wealth if the house does burn down is $x_1 = m_1 + y$. Her budget constraint (solving for $y = x_1 - m_1$) is

$$px_1 + x_2 = pm_1 + m_2$$

Thus the price of risk (wealth in the case the the house does burn down) is p.

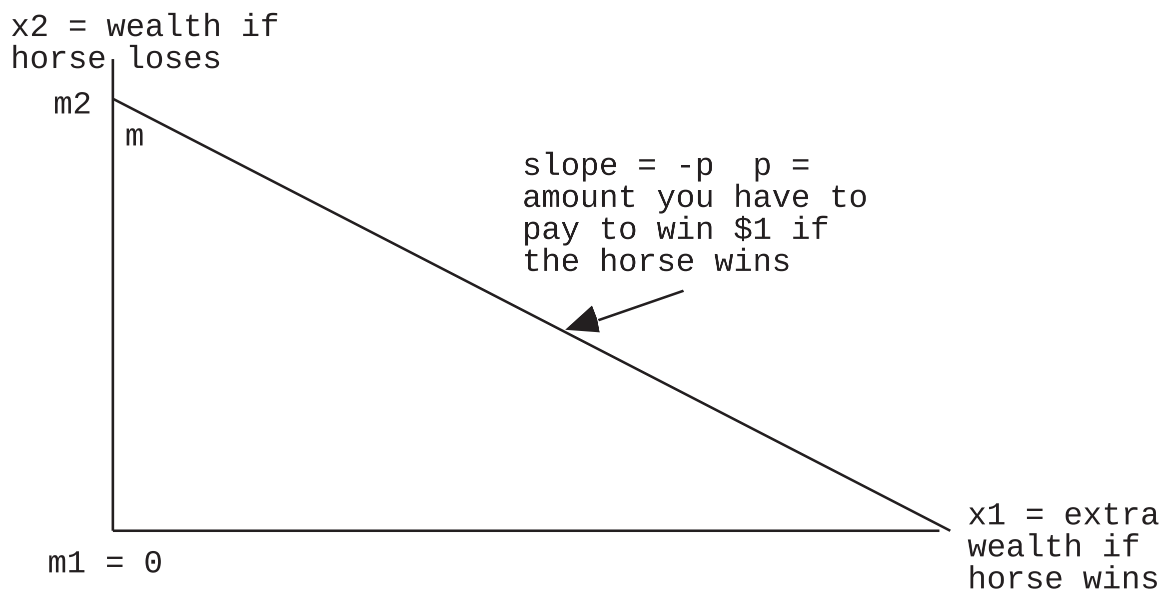

Figure 5.6: A consumer facing the choice of taking a gamble initially bears no risk, since $m_1 = 0$.

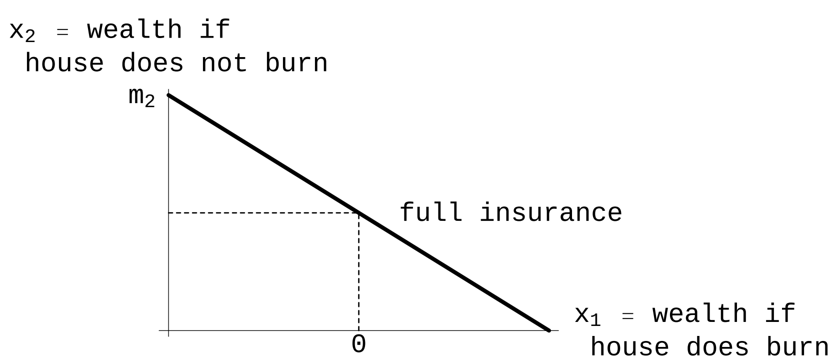

Figure 5.7: A consumer facing the risk that her house will burn down has an endowment of negative extra wealth in the contingency where the house burns down.

From an economic point of view the consumer is selling some of her risk to the insurance company.

5.3 The principle of expected utility maximization

When a consumer faces an uncertain choice, only one of the contingencies will actually occur. Therefore the consumer's utility in one contingency, some economists reason, should not depend on her prospects in other contingencies. If this is true, in evaluating the utility of a risky position the consumer should consider only how well off she will be in each contingency, and how likely she believes each contingency to be. A consumer who behaves in this way, and who estimates the probability of an event (the horse winning or the house burning down) as π, where $0 \leq \pi \leq 1$, will have a utility function of the form

$$u(x_1, x_2) = \pi v(x_1 + x_2) + (1 - \pi) v(x_2)$$

Here π is the consumer's estimate of the probability of the event (a fire, or the horse winning). She uses the same function $v(.)$ to evaluate her income in each contingency. As a result her utility before she knows which outcome happens is a weighted average of her welfare if the event happens, $v(x_1 + x_2)$, and her welfare if the event does not happen, $v(x_2)$. She weights each welfare by her estimate of the probability that she will actually be in that situation.

A consumer who behaves in this way obeys the principle of *expected utility maximization.*

An expected utility maximizer's marginal rate of substitution, which is her bid and asked price for wealth if the favorable event happens, is

$$\text{mrs}(x_1, x_2) = \frac{\partial u / \partial x_1}{\partial u / \partial x_2} = \frac{\pi v'(x_1 + x_2)}{\pi v'(x_1 + x_2) + (1 - \pi) v'(x_2)}$$

To see what this expression means, consider first the consumer who is thinking about betting on the horse race. Her initial endowment has $m_1 = 0$ and as a result $v'(m_1 + m_2) = v'(m_2)$. Thus before she has bet, her mrs, the amount she would be willing to bet in order to win \$1 if the horse wins, is just equal to

$$\text{mrs}(0, m_2) = \frac{\pi v'(m_2)}{(\pi v'(m_2) + (1 - \pi) v'(m_2))} = \pi$$

Remember that π is her estimate of the probability that the horse will win. If, for example, she is willing to bet \$.50 to get \$1 if the horse wins, she is saying that she believes $\pi = 1/2$ and $1 - \pi = 1/2$, that is, that there is an even chance the horse will win.

Thus we arrive at an important idea. The price at which an expected utility maximizer is willing to bet on an event when she is bearing no risk is equal to her private estimate of the probability that the event will occur.

Now consider the case of the consumer who faces the risk of a fire. In her case we have $m_1 << 0$, which means (if $v'(.)$, the marginal utility of wealth, declines as wealth increases), that $v'(m_1 + m_2) > v'(m_2)$, so that

$$\text{mrs}(m_1, m_2) = \frac{\pi v'(m_1 + m_2)}{\pi v'(m_1 + m_2) + (1 - \pi)v'(m_2)} > \pi$$

This means that the consumer will be willing to sell risk to the insurance company at a price higher than her estimate of the probability of a fire. She is willing to do this because wealth if the fire occurs is more valuable to her than wealth in the case the fire doesn't occur.

Thus we find a second important principle. The mrs of an expected utility maximizer expresses both her estimate of the probability the event will occur, and her wealth in that contingency relative to her wealth if the contingency does not occur.

5.4 Gambling and insurance

We can use the theory of consumer choice to understand the attitude of a consumer towards risk. Consider a consumer facing the choice of whether or not to bet on a horse, as in Figures 5.8 and 5.9. The consumer's initial endowment $(0, m_2)$ means that her wealth will be the same whether or not the horse wins the race, so that the outcome of the race won't make any difference to her if she doesn't bet. Suppose that the price at which she can bet is the same as her private estimate of the probability that the horse will win, so that her $\text{mrs}(0, m2) = p$. Economists sometimes call this situation a *fair gamble*, meaning that on the basis of the consumer's information the odds offered is the same as the price she herself would offer to take the gamble. In this case the budget line will be tangent to her indifference curve at the endowment point.

If the consumer has convex preferences, as in Figure 5.8, she will not gamble. Any movement along the budget line takes her to a lower indifference curve. With the principle of expected utility convexity of preferences means that $v'(w)$ declines as w increases (so that the consumer's bid price for the gamble will fall as x_1 increases), or that the consumer has decreasing marginal utility of wealth. Economists call this type of consumer, who will refuse gambles at a price equal to her private estimate of the probability of the event, *risk-averse*. A risk-averse consumer prefers to bear no risk, that is, to keep her consumption close to the line $x_1 = 0$.

If the consumer has non-convex preferences, as in Figure 5.9, she will gamble, because a movement along the budget line takes her to higher indifference curves. This consumer must have $v'(w)$ increasing with w, or increasing marginal utility of wealth. Economists call this type of consumer *risk-loving* because she will accept a gamble at a price equal to her private estimate of the probability of the event. A risk-loving consumer prefers, in fact, to bear risk, that is, to have $|x_1| > 0$.

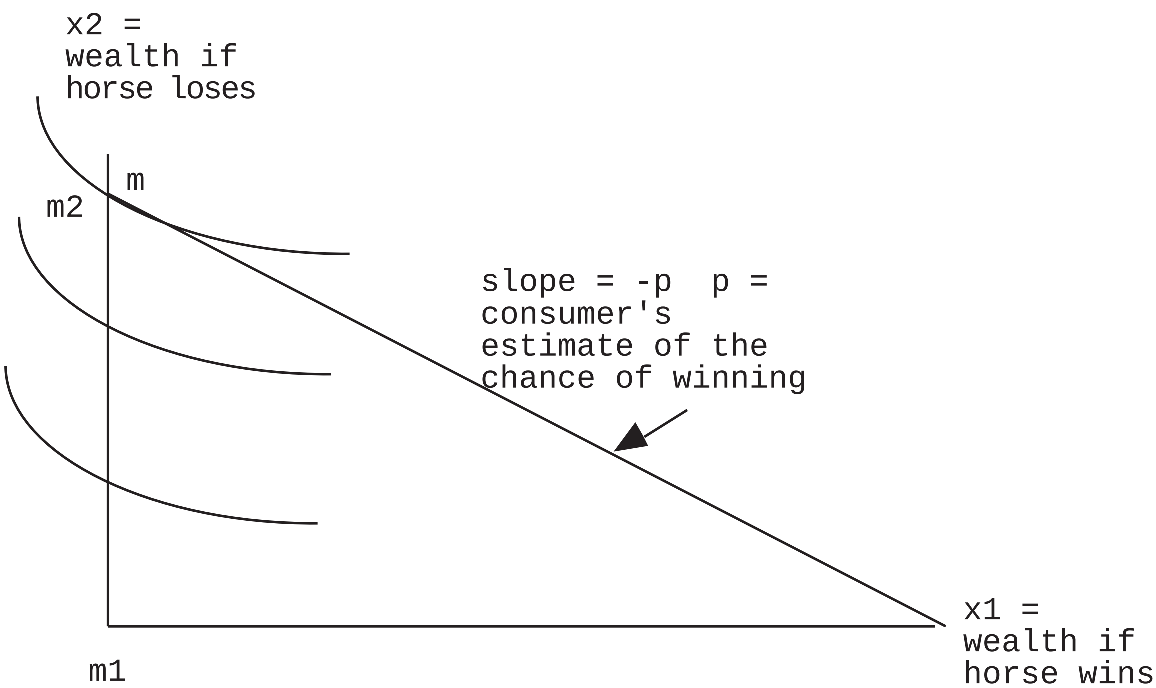

Figure 5.8: A consumer with convex preferences is risk-averse and will not accept a gamble at odds she considers correct. According to the principle of expected utility maximization, such a consumer has $v'' < 0$, that is, decreasing marginal utility of wealth.

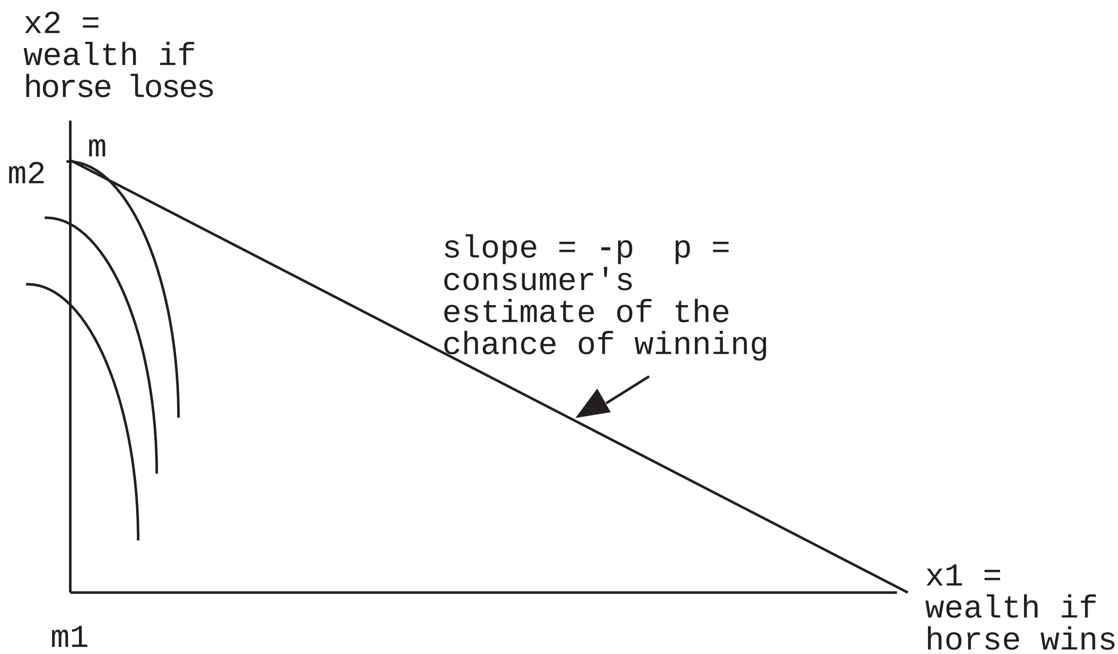

Figure 5.9: A consumer with non-convex preferences is risk-loving and will accept a gamble at odds she considers correct. According to the principle of expected utility maximization, a risk-loving consumer has $v'' > 0$, that is, increasing marginal utility of wealth.

A similar analysis shows that a risk-averse consumer will purchase fire insurance, even if the insurance premium is higher than her estimate of the probability of the fire. A risk-loving consumer, conversely, will sell insurance even at a premium below her estimate of the probability of the fire.

5.5 The theory of finance

The theory of choice under uncertainty is the foundation of the modern theory of finance, which tries to explain how markets value financial assets like corporate stocks and bonds. The key to understanding this approach is to consider an asset, like a corporate stock, as being worth different amounts in different contingencies. Thus a share of stock is like a collection of tickets that pay $1 in various contingencies. The market is viewed as establishing a price for wealth in different contingencies, and the price of a stock is equal to the sum of the values of its constituent risks.

5.5.1 Contingent commodities and assets

For example, suppose that there are just two important contingencies facing asset holders. Suppose the contingency the market is worried about is the imposition of an oil embargo on the country. (It might just as well be the occurrence of a serious recession, or the outbreak of sustained inflation.) The oil embargo will have different effects on the value of the assets of different companies. Suppose Company A's assets, for example, will not be affected by the embargo at all, so that a share of its stock will represent assets worth $50 in both contingencies. We can write the pattern of returns on the stock as (w_1, w_2), where w_2 is the amount of wealth represented by the stock, if the oil embargo occurs, and w_1 is the extra wealth in the case the embargo does not occur. Company A's stock can be described by as $(0, \$50)$. But suppose the value of Company B's assets will be very much affected by the oil embargo. Its stock might then might be described as $(\$50, \$25)$. This means B's assets will be worth $25 if the oil embargo occurs, but will be worth $75 = $25 + $50 in case the embargo does not occur.

5.5.2 General equilibrium with identical portfolio holders

Suppose now that Company A and Company B (and others just like them) are of equal size, and own the whole wealth of the economy. Suppose also that there are many wealth-owners, but that they are all identical in their endowments, utility functions for wealth, and estimates of the probability of an oil embargo. Suppose the representative wealth holder starts off owning one share of A and one share of B. Her endowment, in terms of certain wealth and contingent wealth, is $(\$50, \$75)$. In other words, whether or not there is an embargo the assets in the two companies will be worth $75 ($50 for Company A and $25 for

Company B), and if there is no embargo, the assets will be worth an extra \$50 (all due to the gain for Company B.)

Now, let us consider what happens if these identical wealth-holders trade their stocks on a stock market. Since all the wealth-holders are identical, their bid and asked prices for the stocks must be the same. (In fact, each wealth-holder will have to wind up holding one share of A and one share of B). What determines these bid and asked prices? The representative wealth-holder's mrs for extra income in the case that there is no embargo is

$$\text{mrs}(\$50, \$75) = \frac{\pi v'(\$125)}{\pi v'(\$125) + (1-\pi)v'(\$75)} = p$$

Thus if we knew the utility function of the representative wealth holder and her estimate of the probability that the embargo would not occur, we could predict the market price, p, of wealth in the case of no embargo. The price of Company A's stock would be \$50, since Company A represents no risk, and the price of Company B's stock would be $\$50p + \$25$. This situation is illustrated in Figure 5.10.

5.5.3 Example: logarithmic wealth function

If the representative wealth holder has $v(w) = \log(w)$, we can easily calculate the equilibrium market price of risk in the economy. We have

$$\text{mrs}(\$50, \$75) = \frac{\pi v'(\$125)}{\pi v'(\$125) + (1-\pi)v'(\$75)}$$

$$= \frac{\pi/\$125}{\pi/\$125 + (1-\pi)/\$75} = \frac{\pi}{\pi + (1-\pi)(\$125/\$75)} = p$$

If it is generally believed, for example, that there is a $\pi = 50\%$ chance of an oil embargo, the price of risk would be

$$p = \frac{.5}{(.5) + (.5)(1.6666)} = .375$$

The market will pay \$.375 for the promise of \$1 of wealth in case there is no embargo. Company B's stock will sell at $\$25 + \$50p = \$25 + \$18.75 = \$43.75$. If political developments reduced the chance of an oil embargo to zero, Company A's stock would still trade at \$50, but Company B's assets would now be worth \$75 with certainty, so its stock would trade at \$75.

This example illustrates some important ideas from the modern theory of finance. The price of any stock depends on the pattern of risk attached to that stock in relation to the whole pattern of risk in the market. A good exercise is to work out the market price of stock if Company A's pattern of risk is $(-\$50, \$25)$ to convince yourself of this. But the price of stock also depends on the typical wealthholders attitude toward risk. If the typical wealthholder is very risk-averse, the market price of risky wealth will be low, and the price of risky stocks will also be low.

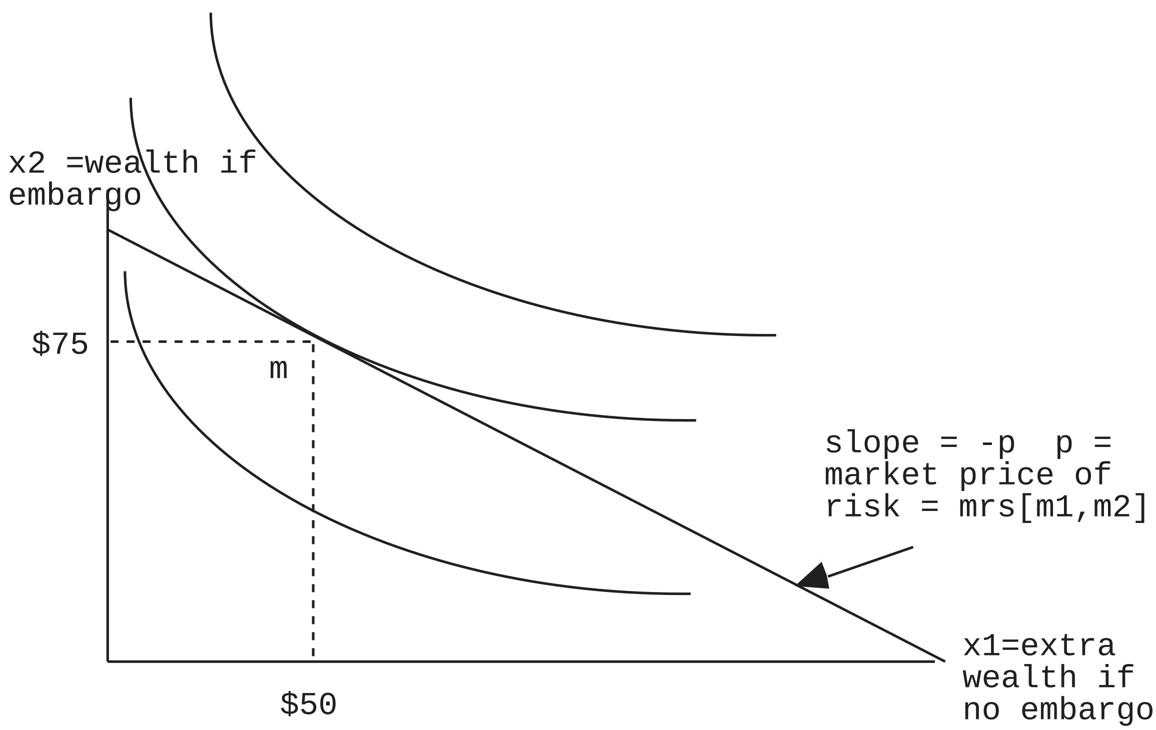

Figure 5.10: If all the wealth holders are alike, the representative wealth holder has to wind up holding whatever risk is present, even if markets open. The market price of risk reflects the representative wealth holder's bid and asked price for risk.

Chapter 8

General equilibrium: markets and efficiency with a single type of consumer

8.1 A simple model of the whole economy

The theories of consumer choice and production can be put together as a theory of *general equilibrium* to analyze the allocation of resources in the economy as a whole.

The simplest case to analyze is an economy where all the consumers are identical. In order to use the graphical tools and insights we have developed we will assume that there are two goods in the economy. The first good is a specific output, like food, housing, or automobiles, and its price is p. The second good is general output, which is like money, and will be measured in money units so that its price is 1. We might think of this output as consisting of all the other goods in the gross national product, and we will refer to it as money income. There is a large number of consumers and firms, so that we can assume strong competition among them. This means that each consumer and firm assumes that its transactions in the market will have no effect on the market price.

Each consumer has the same utility function, $u(x_1, x_2)$, where x_1 and x_2 are the consumption of good 1 and money income, and each consumer has the same endowment of good 1, m_1, and money income, m_2. Since all the consumers are the same, we can take one of them as the *typical consumer* and represent the whole economy in terms of this consumer, as in Figure 8.1. To find the money income endowment of the whole economy, for example, we would simply multiply m_2, the money income endowment of the typical consumer, by the number of consumers. If there are 100 million consumers, and each starts with an endowment of \$50,000 of money income and no good 1, the endowment of the economy would be \$5 trillion dollars (and zero good 1).

Figure 8.1: If all the consumers in an economy can be regarded as alike in tastes and endowments, we can represent the whole economy in terms of a typical consumer. The economy's endowment is just equal to the number of consumers multiplied by the typical consumer's endowment of the two goods, and the economy's consumption is just the typical consumer's consumption multiplied by the number of consumers.

In this model, then, as in the consumer choice model, utility functions and endowments are exogenous parameters, and the consumption pattern of the typical consumer (and hence of the whole economy) is an endogenous variable. But in the consumer choice model we took the price of good 1 as exogenous, while in the general equilibrium model, the price of good 1 will be endogenous, that is, explained by the model. To see how this works, we need to represent production and cost on the same graph as the typical consumer's endowment and preferences.

8.2 Supply curves and the production possibility frontier

The supply of the industry producing good 1 can also be viewed in general equilibrium terms.

The production possibility frontier for this economy shows the combinations of x_1 and x_2 that can be produced given the available resources (say, the money income endowment of the economy), as shown in Figure 8.2. In plotting this production possibility frontier we scale it to the size of the typical consumer. If the economy as a whole, consisting of 100 million consumers, each endowed with \$50,000, could produce \$4 trillion dollars of money income and 10 million cars per year, we plot this point as \$40,000 of money income and 1/10 of a car for the typical consumer.

Figure 8.2: The production possibility frontier for a two-good economy shows the combinations of money income, good 2, and good 1 that can be produced, given the technology and endowment of the economy.

Figure 8.3: The production possibility frontier for a two-good economy where good 1 can be produced at constant long run average cost is a straight line with slope = - marginal cost = - average cost.

8.2.1 Constant long-run average costs

If the long run supply curve for x_1 is horizontal, this means that it is possible to transform money income into x_1 at a constant rate equal to the long-run average cost. The production possibility frontier for the economy, showing the combinations of money income and x_1 that can be produced (per typical consumer) will be a straight line, as in Figure 8.3.

In this case the price of good 1, p, must be equal to the marginal cost of producing good 1, which will also equal the long-run average cost, since marginal cost equals minimum average cost. The slope of the production possibility frontier will be equal to $-\text{mc} = -p$. If the endowment of the typical consumer consists of money income then the production possibility frontier will also be her budget constraint. The typical consumer rents resources to good 1 manufacturers, and receives the whole value of the output as wages and dividends. The general equilibrium for the whole economy will be at the point where the typical consumer's indifference curve is tangent to the production possibility

Figure 8.4: The general equilibrium of an identical agent economy with constant long run average cost of production of good 1 is the point x where the production possibility frontier is tangent to the indifference curve of the typical consumer.

frontier, as in Figure 8.4.

8.2.2 Rising long-run average costs and rents

To understand the relation between an upward-sloping long-run supply curve, the production possibility frontier, and rent, consider a model where x_2 is money income and x_1 is food. Suppose that there is a limited amount of good land, where food can be grown cheaply, and an unlimited amount of ordinary land, where it costs more to grow food. Suppose the typical consumer has an endowment of m_2 units of money income and 1 unit of good land. The ordinary land is open to anyone who wants to produce on it. If we think of the typical consumer as a firm deciding on the profit- maximizing level of food production, her cost curve will consist of two linear parts, as in Figure 8.5.

The corresponding average cost, average variable cost, and marginal cost curves are shown in Figure 8.6.

Since there is a limited amount of good land in the economy, the number of firms is fixed, and the long-run supply curve is identical with the short-run supply curve, which in turn is identical to the marginal cost curve. The production possibility frontier is kinked as in Figure 8.7. If the demand for food is relatively low, as in this Figure, the equilibrium will require only production on the good land. The price will be equal to marginal cost on the good land, and there will be no economic profit, so the typical consumer's budget line passes through her original endowment point.

If the demand for food is relatively high, the equilibrium will be as in Figure 8.8.

The equilibrium of Figure 8.8 can be visualized in terms of the profit maximizing position of the typical firm as in Figure 8.9. The competitive firm's cost curve is tangent to the production possiblity frontier at the equilibrium because the individual firm sees rent as part of its marginal cost. Rent makes the private

Figure 8.5: A firm that has access to a limited amount of good land will have a total cost curve with two segments. The first segment has a slope equal to the lower marginal cost of production on the good land; the second segment has a slope equal to the higher marginal cost of production on ordinary land.

Figure 8.6: The marginal and average variable cost curves for a firm with one unit of good land coincide up to the maximum output on one unit of good land. Then the marginal cost curve jumps to a higher level, and the average cost rises gradually to become asymptotic to it.

Figure 8.7: The production possibility frontier has two segments, the first with a slope equal to the low marginal cost on good land, the second with a slope equal to the high marginal cost on ordinary land. In this economy the demand for food is so low that only good land is used. The price is equal to the marginal cost on good land, and there is no rent.

Figure 8.8: If the demand for food is relatively high, production will take place on ordinary land. Then the price will have to be equal to the higher marginal cost of production on ordinary land, and the firm will have positive economic profit. Since the typical consumer owns the firm, this profit increases her endowment of money income.

Figure 8.9: If the equilibrium price is equal to marginal cost on ordinary land, the typical firm has a positive economic profit on its production on good land. Since this profit cannot be competed away by entry, it is a *rent*.

marginal cost of production as seen by the firm equal the social marginal cost of production. As output increases, the rise in rent shifts all the competitive firm cost curves upwards. The economy moves along the production possibilities frontier, but the firm sees this as a shift of its own cost curve.

One can think of a rent equilibrium equally well as representing a situation where the typical consumer owns a unit of good land that she rents out to the firm. If the equilibrium price is equal to marginal cost on the good land, there is no rent on that land. If the equilibrium price is equal to the higher marginal cost on ordinary land, the consumer will bargain for a rent on her unit of good land equal to the economic profit the firm could earn at the higher price. Competition among firms will force each firm to pay this rent. Including the rent, the marginal cost on the good land will be just equal to the marginal cost on the ordinary land, so each firm will see itself as having constant marginal cost. In this case the typical consumer's income increases by the amount of the rent.

8.2.3 Rent and consumer surplus in competitive equilibrium

Whenever the production possibility frontier is curved because of the existence of a scarce resource that is owned by the typical consumer and cannot be reproduced, the scarce resource will correspond to a positive rent, which can be visualized by the distance from the x_2-intercept of the production possibility frontier to the x_2-intercept of the typical consumer's budget line in competitive equilibrium. If there is no fixed cost, this rent is equal to the producer's surplus. We can also visualize the consumer surplus in this diagram, since the x_2-intercept of the indifference curve through the equilibrium point measures the amount of money income the typical consumer views as equivalent to the competitive equilibrium consumption bundle. Figure 8.10 illustrates the way surpluses appear in this picture of competitive equilibrium.

Figure 8.10: Since the x_2-intercept of the indifference curve through the competitive equilibrium point shows the amount of money income equivalent to the competitive equilibrium consumption bundle, its distance from the x_2-intercept of the typical consumer's budget line measures the consumer surplus at the competitive equilibrium. The producer surplus in the absence of fixed costs is equal to rent, which is measured by the distance from the x_2-intercept of the typical consumer's budget line to the x_2-intercept of the production possibility frontier, which is the initial endowment point.

8.3 The Market and Individual Welfare

One of the most important issues in economic theory is the analysis of the conditions under which *laissez-faire* policies, under which governments do not interfere with the operation of markets, lead to desirable social outcomes. Economists study this problem as welfare economics. Our aim is to use the models we have studied to understand when and in what sense pure market forces lead to a socially desirable outcome.

8.3.1 Allocations

Given existing resources, labor-power, natural resources including land, and already produced means of production, we can think of any economy as choosing how to assign its resources to the production of different outputs, and how to distribute the outputs to the individuals who consume it. A plan of this kind is called an *allocation of resources.*

We can represent an allocation of resources in a very general economic model. An economy is a collection of n consumers, $i = 1, \ldots, n$, and k producers, $j = 1, \ldots, k$. If there are h goods in the economy $g = 1, \ldots, h$, we can describe the ith consumer's endowment by a list $m^i = (m^i_1, \ldots, m^i_h)$ and the consumer's consumption by a list $x^i = (x^i_1, \ldots, x^i_h)$ telling how much of each good she consumes. Similarly, we can describe the jth producer's production plan by a list $z^j = (z^j_1, \ldots, z^j_h)$ of the inputs the producer uses, and another list $y^j =$

$(y_1^j, \ldots, y_h^j)$ of the outputs the producer produces. Of course some elements of all of these lists may be zero. Notice that when we write x^i without a subscript it means the whole list of goods the consumer consumes, while x_g^i is just the amount of good g the consumer consumes.

Each consumer has a utility function $u^i(x^i)$ that describes her preferences over different patterns of consumption. Each producer has a production set $P^j = \{(z^j, y^j)$ such that the jth producer can produce the outputs y^j from the inputs $z^j\}$.

Then an allocation is a plan that tells what each producer will produce and what each consumer will consume. We write an allocation as

$$\{x^1, \ldots, x^n; (z^1, y^1), \ldots, (z^k, y^k)\}$$

Not all allocations are *feasible* given the resources and social-technological production possibilities of the economy. A feasible allocation is an allocation $\{x^1, \ldots, x^n; (z^1, y^1), \ldots, (z^k, y^k)\}$ that satisfies two conditions:

$$(z^j, y^j) \text{ is in } P^j \text{ for every producer (PF)}$$

$$\Sigma_{i=0}^n x_g^i + \Sigma_{j=0}^k z_g^j = \Sigma_{i=0}^n m_g^i + \Sigma_{j=0}^k y_g^j \text{ for } g = 1, ...h \text{ (RF)}$$

The first condition, *production feasibility*, says that the production plan is possible for the producers to achieve. The second condition, *resource feasibility*, makes sure that there are enough resources in the endowment to allow the plan to be carried out. The lefthand side of (RF) adds up all the uses of good g for consumption and as inputs to production, while the righthand side shows the total available amounts of good g from production outputs and the endowments of consumers.

8.3.2 Allocation with identical consumers

The feasible allocations in the identical consumer economy are just the (x_1, x_2) points on and under the production possibility frontier, as in Figure 8.11.

A feasible allocation that gives the same amount to each consumer can be described by a single point (x_1, x_2). For this to be a feasible allocation, it must be on or inside the production possibility frontier.

8.3.3 Evaluating allocations

Welfare economics assumes that the end goal of economic activity is individual consumption, and assumes that consumers' own preferences are the best way to judge their welfare. What ultimately interests us about an allocation is the levels of satisfaction reached by consumers. We can think of this in terms of the level of indifference curve each individual consumer reaches, or in terms of the utility indicator each consumer achieves. Different allocations of given resources will achieve different levels of utility for each of the consumers. Some allocations may be very good for some consumers, and very bad for others. In evaluating allocations, economists use two types of criteria.

Figure 8.11: In an identical agent economy we can visualize all the feasible allocations that give the agents the same consumption bundle as points on social production possibility frontier.

8.3.4 Distribution

First, we can judge an allocation by how equal or unequal the welfare levels achieved by different consumers are. In the identical consumer model, all the consumers receive exactly the same consumption bundle and have the same utility, so there is no issue of distribution.

8.3.5 Efficiency: the Pareto criterion

Second, we can judge an allocation by how *efficient* it is, that is, whether or not it wastes any resources. An allocation wastes resources if it is possible to produce more satisfaction from the given resources with a different allocation. A *Pareto-efficient allocation* (sometimes called a *Pareto-optimal allocation* or a *Pareto-optimum*) is one which does not waste resources. To test for Pareto-efficiency economists use the concept of a *Pareto-superior* allocation.

If one allocation gives the same or higher utility levels to all the consumers than a second allocation, and higher utility level to at least one, we say that the first allocation is Pareto-superior to the second. Notice that it may be impossible to compare two allocations by this method, if the first allocation gives higher utility levels to some consumers and lower utility levels to others. In this case we say the two allocations are *not Pareto-comparable*. The comparison of allocations by the Pareto method is not like the comparison of numbers by size. It is always possible to decide which of two numbers is bigger, but it is not always possible to determine which of two allocations is Pareto-superior.

An allocation is Pareto-efficient if no other allocation is Pareto-superior to it. This is a subtle definition that often confuses people. A Pareto-efficient allocation is not necessarily Pareto-superior to any allocation that isn't Pareto-efficient. All we know about a Pareto-efficient allocation is that there is no other allocation that beats it, not which allocations it beats according to the criterion of Pareto-superiority.

Figure 8.12: In the typical consumer model where all the consumers receive the same consumption bundle, Pareto-superior allocations are those that are on a higher indifference curve for the typical consumer. Allocation x^* is thus Pareto-superior to allocation x?. In fact, allocation x^* is the only Pareto-efficient allocation in this case.

You might mistakenly think that a Pareto-efficient allocation is Pareto-superior to every other feasible allocation. Many people think of Pareto comparisons as if they were the same as comparisons of numbers. If no number in a set of numbers is larger than a particular number, that number must be the largest one. The difference is that we can always compare two numbers by size, but we can not always compare two allocations by Pareto-superiority. So in the case of allocations when there are different consumers it is not necessarily true that a Pareto-efficient allocation is Pareto-superior to every other feasible allocation despite the fact that there is no feasible allocation that is Pareto-superior to the Pareto-efficient allocation, because one allocation may make some consumers better off and others worse off than the other.

A central question for welfare economic theory is under what conditions a market equilibrium will be efficient.

8.3.6 Example: identical agents with production

In the identical consumers model with production, suppose that we consider only allocations in which the consumers all receive the same consumption bundle. The only way to make any consumer better off (among these allocations) is to make all the consumers better off. We can easily judge which allocations make the typical consumer better off by drawing in her indifference curves, as in Figure 8.12.

In this model it is possible to compare all the allocations that give all the consumers the same consumption bundle by using the indifference curves of the typical consumer. As a result there is only one Pareto-efficient allocation and it is, in this model, Pareto-superior to all the other allocations. In Figure 8.12 the Pareto-efficient allocation is x^*.

Figure 8.13: Competitive equilibrium in the identical consumer economy with production occurs at x^*, where both the production possibility frontier and the indifference curve of the typical consumer are tangent to the price line. x^* is also the unique Pareto-efficient allocation. We can see that x^* maximizes the total of producer's surplus (rent) and consumer's surplus.

8.3.7 Competitive equilibrium allocations

A natural question to study is the relation between the competitive equilibrium allocations, which would be reached by perfectly competitive consumers and firms, and Pareto-efficiency. A central concern of welfare economics has been to specify the conditions under which a competitive equilibrium will be Pareto-efficient. The claim that competition leads to Pareto-efficiency, that is, avoids the waste of resources, has been a major argument for laissez-faire policy, under which government refuses to interfere in the workings of the market.

A competitive equilibrium is an allocation that is reached when consumers maximize utility subject to their budget constraints, and firms maximize profits, taking market prices as given.

8.3.8 Example: production with identical agents

In the identical agent economy with production shown in Figure 8.13, the profit maximizing point is the point on the production possibility frontier tangent to the price line. The utility maximizing point is the point on the price line tangent to the typical agent's indifference curve. The competitive equilibrium will be unique with decreasing or constant returns to scale in production and well-behaved indifference curves. Since the production possibility frontier and the indifference curve are both tangent to the same price line, they must be tangent to each other, so the competitive equilibrium is also the unique Pareto-efficient allocation in this case.

8.4 The Fundamental Welfare Theorems

Economists have summarized the analysis of the relation between Pareto-efficiency and competitive equilibrium in two *fundamental welfare theorems.* The first states the conditions under which a competitive equilibrium will be Pareto-efficient. The second states the conditions under which it is possible to achieve an arbitrary Pareto-efficient allocation as a competitive equilibrium.

8.4.1 A full information competitive equilibrium without externalities is a Pareto-optimum

In the identical consumer economy we have analyzed, it is clear that a competitive equilibrium is a Pareto-efficient allocation. This is the First Theorem of Welfare Economics. It holds whenever there is full information, so that problems like the market for lemons do not arise, and there are no *external* effects between agents, so that the aspects of production and exchange valued by the market represent the true social costs and benefits of economic activity.

8.4.2 When can a Pareto-optimum be supported as a competitive equilibrium?

In the identical consumer economy with production, there is only one Pareto-efficient allocation, which is also the competitive equilibrium. This is the Second Theorem of Welfare Economics in the identical consumer model.

If preferences are well-behaved and production has constant or decreasing returns to scale, there is full information, and there are no externalities, it is possible to achieve any Pareto-efficient allocation as a competitive equilibrium. Thus even a centrally planned socialist economy, when the consumers have well-behaved indifference curves and production does not have increasing returns to scale, if it achieved a Pareto-efficient allocation, would act as if it were following market determined prices, as long as we ignore externalities and problems of imperfect information.

8.4.3 Externalities

The First Welfare Theorem tells us that a competitive equilibrium is Pareto-efficient only if all the interactions between the consumers and producers are correctly priced by the market. If there are interactions that consumers and producers do not take account of because they are not priced, a competitive equilibrium allocation may not be Pareto-efficient.

One way to think about this problem is to distinguish the *social marginal cost* of an action from its *private marginal cost* as it appears to the firms and households. For example, a firm takes account of the costs of labor-power and other inputs when it makes its production decision, because it has to pay directly for these resources. But it does not take account of the costs of smoke pollution, if there is no tax on smoke pollution and no legislation limiting it. From a social point of view the smoke pollution is a cost of production, just like the use of labor-power and other inputs, but the firm will ignore this cost in making its decisions as long as it doesn't have to pay for the pollution. In this case the social marginal cost of production is higher than the private marginal cost because the social marginal cost includes the smoke costs, while the private marginal cost does not.

In competitive equilibrium firms equate private marginal cost to price, and consumers equate private marginal rates of substitution to price. If private marginal cost differs from social marginal cost, the market prices will give the wrong signals to consumers, and they will choose to consume too much or too little of the good in question. As a result the competitive equilibrium will not be Pareto-efficient.

These observations are the basis for the conventional economic recommendation that scarce resources should be made private property, that is, *appropriated*, so that their owners will exact a rent and price the externality created by their scarcity. This policy prescription has had wide influence in the pricing of such supposedly scarce resources as the radio-frequency spectrum, and proposals to create private-property rights in common-pool resources such as fisheries.

8.4.4 Example: a negative production externality

Consider a 2 good, identical consumer economy, with production. The second good is fish taken from an ocean fishery. The more boats go out, the higher the costs to each boat of catching a full load of fish, because each boat depletes the stocks and forces others to go farther to make their catch. To model this, suppose that the cost to each boat of catching a full load of fish when Y_1 boats go out is Y_1^{α}. Then the cost to society of Y_1 boatloads of fish is $Y_1^{\alpha} Y_1 = Y_1^{1+\alpha}$. The effect of more boats going out is to increase the marginal cost of each boat if $\alpha > 0$.

At the competitive equilibrium output Y_1^*, the private marginal cost perceived by each boat is $(Y_1^*)^{\alpha}$, but the social marginal cost, which can be calculated by adjusting the number of boats, is $(1+\alpha)(Y_1^*)^{\alpha}$ and is larger than the private marginal cost by $\alpha(Y_1^*)^{\alpha}$. Competition among the boats will result in a price of a boatload of fish equal to the private marginal cost, which is also the minimum average cost.

The competitive equilibrium with a negative externality will fail to be Pareto-efficient, as shown in Figure 8.14.

8.4.5 Example: a positive production externality

Now suppose that the second good is honey, and that the more total honey produced, the lower the costs to each beekeeper of producing honey, because each firm's bees can feed on the clover of the neighboring firms. We can model this in exactly the same way as the fishery example, by taking the coefficient $\alpha < 0$, so that an increase in the total number of hives in an area lowers the cost per hive.

Figure 8.14: With a negative externality the price of the good is below the social marginal cost, and the typical consumer is induced to buy more of it at the competitive equilibrium, x, than at the Pareto-efficient allocation x^*.

Figure 8.15: With a positive externality, the competitive equilibrium price, which is equal to private marginal cost, is too high, as at x.

In this case the private marginal cost of honey is $(Y_1^*)^\alpha$, which will be the competitive equilibrium price. But the social marginal cost of honey is $(1 + \alpha)(Y_1^*)^\alpha < (Y_1^*)^\alpha$ when $\alpha < 0$. Thus the competitive equilibrium price of honey will be too high, and the typical consumer will buy too little of it. This situation is illustrated in Figure 8.15.

8.5 Policies to correct externalities

The diagrams we have looked at suggest a way to correct allocation with externalities through *government intervention* by taxes or subsidies.

Since the problem is that the market price deviates from the social marginal cost, an excise tax or subsidy can raise or lower the price so as to make it equal to social marginal cost, and thus achieve the Pareto-efficient allocation.

Figure 8.16: A subsidy on honey equal to the size of the external effect will lower the cost of each producer and move the equilibrium allocation to the Pareto-optimum. The typical consumer must pay a lump-sum tax to finance the subsidy.

Figure 8.17: A negative externality can be corrected by an excise tax equal to the size of the externality. The tax revenues are returned to the typical consumer as a lump-sum subsidy in this example.

8.5.1 Example: a subsidy for a positive externality

Figure 8.16 illustrates how a subsidy could correct the positive externality in the honey example.

8.5.2 Example: a tax to correct a negative externality

In a similar fashion government intervention can correct a negative externality by imposing a tax equal to the size of the externality, the difference between social and private marginal cost. This possibility is illustrated in Figure 8.17.

8.5.3 The theory of the second best

The correction of externalities by the imposition of taxes and subsidies illustrates an extremely important principle called by Kelvin Lancaster and Richard

Lipsey the *Theory of the Second Best*. The First Welfare Theorem suggests that introducing imperfections into competitive equilibrium, like monopoly, externalities, or taxes and subsidies, will move the resulting equilibrium allocation away from a Pareto-efficiency. But if such an imperfection, like an externality, already exists, then the addition of other imperfections, like excise taxes can very possibly move the allocation *toward* Pareto-efficiency. If the competitive equilibrium is Pareto-efficient, that is, First Best, there is no case for changing it by government intervention on grounds of inefficiency (though there might be a case, as we shall see, on grounds of unequal distribution). If, on the other hand, the competitive equilibrium is not Pareto-efficient because of existing imperfections (as must often be the case in real economies), or Second Best, then there may be good arguments for disturbing it by additional intervention in order to move to a Pareto-superior allocation.