Chapter 5: Taste

The tragedy of the commons develops in this way. Picture a pasture open to all. It is to be expected that each man will try to keep as many cattle as possible on the commons. Such an arrangement may work reasonably satisfactorily for centuries because tribal wars, poaching and disease keep numbers of both man and beast well below the carrying capacity of the land. Finally, however comes the day of reckoning ... the only sensible course for him is to add another animal to his herd. And another; and another. This is the conclusion reached by each and every rational herdsman sharing a commons. Therein is the tragedy. Each man is locked into a system that compels him to increase his herd without limit - in a world that is limited. Ruin is the destination toward which all men rush in a society that believes in the freedom of the commons. Freedom in a commons brings ruin to all.

Garrett Hardin, 1968

The main story in chapter 3 can be summarized in two propositions: that there exists a competitive equilibrium and that it is efficient. Existence of a competitive equilibrium means that there is some set of prices at which the total amount of each good that all consumers want to buy is just equal to the total amount of that good that all producers want to sell, as long as everybody looks upon prices as determined by the market and as independent of his own actions. Efficiency of a competitive equilibrium means that no planner, however knowledgeable, however powerful, however benevolent, could rearrange the economy to make everybody better off simultaneously. These propositions are far from obvious and are generally disbelieved by non-economists. Together, they constitute a significant part of the contribution of economics to the understanding of the world. The propositions have strong political implications. They suggest that the government might best stay out of the running of the economy, restricting itself to the protection of property rights and the broad redistribution of income from rich to poor. The main preoccupation in this chapter and is whether and to what extent these implications of the simple model are sustained when account is taken of important aspects of the world that have so far been assumed away.

Concealed in the demonstration of the virtues of the competitive economy were four critical assumptions: that all goods are *private*, that utility is timeless, that information is complete, and that everybody is unreservedly selfish. The main task of this chapter is to relax the first of these assumptions. Ordinary private goods are contrasted with public goods, externalities and personal goods to be defined presently. There is some discussion of risk. Preferences among goods available at different times are considered by replacing the bread and cheese in the utility function with consumption today and consumption tomorrow. Information and advertising are discussed briefly at the end of the chapter. Altruism is put off until chapter 8.

A private good is a good such that its total output is allocated among people, where each person benefits from his own portion exclusively with no impact, favourable or unfavourable, of one person's consumption upon the rest of the population. With only two goods, bread and cheese, total outputs, B and C, are divided up among the N people in the economy. Person j consumes b^j loaves of bread and c^j pounds of cheese. His utility is $u^j(b^j, c^j)$. The sum of all b^j is equal to B and the sum of all c^j is equal to C. This characteristic of private goods - that total output is divided up among people, each of whom consumes his portion exclusively - extends easily from a world with only two good to a world with many goods, as long as we adhere to the assumption that all goods are private.

Cheese is a private good in that the pound of cheese I consume is necessarily denied to you. We can share a piece of cheese, but we cannot eat the same mouthful. That is true of bread as well, but not of all goods. We can both watch the same television program, you on your set and I on mine, without interfering with one another. We are both protected simultaneously by the army and by the police. There is a class of *public* goods that people can consume together, all at once. The distinction between private and public goods is important because the virtues of a competitive economy do not extend to the provision of public goods. The market cannot be relied upon to supply the right amount of public goods. Typically though not invariably, public goods have to be supplied collectively by the government if they are to be supplied at all.

Half way between private goods and public goods are *externalities.* One and the same good may convey some benefits to its owner, while at the same time conveying other benefits or harm to the rest of the community. My car is a private good in the sense that you and I cannot both drive the same car at the same time, but my use of the car harms everybody at once when exhaust fumes are released into the atmosphere. The term externalities implies that the harm to society is external to my purpose in using the car, an undesirable but unintended consequence of driving.

There was no provision for risk in the economy as described in chapters 3 and 4. There was no uncertainty about how much bread and cheese each person would consume or about how consumption would be affected by one's choices in production and trade. The world is not like that at all. The future is intrinsically uncertain. Looking ahead, one must recognize that actions today can influence consumption tomorrow but cannot determine consumption exactly. Representation of taste by a utility function can be reconstituted to allow for uncertainty. Utility can be made to depend not just on amounts of different goods, but on amounts of income in equally likely *states of the world.* The accounting for risk gives rise to a utility of income function incorporating a person's degree of risk aversion.

The assumption about time in chapters 3 and 4 was that everything happens simultaneously. Utility was defined in an instantaneous world with no past and no future, or, equivalently, in a world without change, where past, present and future are all identical. That is implied when we assume that $u = u(b, c)$. Later on in the chapter, the notion of utility will be enlarged to cover goods consumed at different times. An intertemporal utility function will treat bread consumed at different dates as distinct commodities. This reinterpretation of the utility function does not overturn the demonstration of the virtues of the competitive economy.

There follows brief discussions of personal goods and advertising. Personal goods - notably leisure and life expectancy - are private goods like bread and cheese except that their valuations differ from one person to the next. There was no place for advertising in the world of chapters 3 and 4 where taste was not malleable, people were fully rational and all relevant information was available automatically to everybody in the economy. Advertising becomes profitable to the advertiser when that is not completely so. Advertising may provide knowledge about products, certify quality, influence taste, and supply an incentive for the private provision of some public goods.

The chapter concludes with an explanation of the concept of *real income*. Real income was exemplified in chapter 1 by the time series of gross national product per head in Canada, interpreted as indicating the improvement in the Canadian standard of living over a long period of time. In chapters 3 and 4, utility and income were treated as distinct concepts. In statistics of real income, these concepts are fused. The discussion of real income explains how this is done.

Types of Goods

The utility function is considerably more flexible than has been supposed so far. It can be made to encompass many kinds of goods: private goods which have been the exclusive focus of our analysis so far, public goods, such as the army, that convey benefits to everybody simultaneously, externalities, exemplified by exhaust fumes from automobiles, where one person's consumption harms another person or the rest of society, goods consumed at different times, and risk. Types of goods will be discussed in turn, with emphasis on how each type can be fitted into the utility function, on the extent to which the strong propositions about the existence and efficiency of the competitive equilibrium carry over from a world where all goods are private to a world where other types of goods are important, and on the implications for the role of government.

<u>Many Private Goods</u>: [$u = u(b, c, d, e, f, g, h \ldots)$]

Though the utility function has been looked upon so far as representing a person's tastes in an economy with only two goods, the function is easily expanded to account for the virtually infinite array of goods and services that people consume: bread, cheese, dates, electric bulbs, fruit, games, houses, and so on. Though not demonstrated here, it can be shown that the nice properties of the competitive economy extend from two to many goods. A demand curve for any good whatsoever shows how the amount of that good consumed is affected by changes in its price when prices of all other goods remain the same. All of the types of demand curves discussed at the end of the last chapter can be reformulated in a multi-good context. The extension of the constant income demand curve requires only that money income remain unchanged as the price of some good varies. The production demand curve and the compensated demand curve can be extended to a multi-good context by combining the change in the price of some good with a change in money income just sufficient to keep the consumer on the production possibility frontier or on the same indifference curve as the case may be.

<u>Public Goods and Private Goods</u>: [u = u(b, G) where b is butter and G is guns]

As discussed in chapters 3 and 4, bread and cheese were *private* goods, defined by the property that society's total supply of these goods must somehow be apportioned among all consumers and that each person's utility is affected by his consumption of these goods, regardless of the amounts consumed by anybody else. Not all goods are like that. Some goods, such as television signals and protection by the army and the police, carry benefits to everybody at once. Such goods are called *public*.

The distinction between public goods and private goods is exemplified by guns and butter. Guns (shorthand for national defence as a whole) are public because each person's benefit depends on the total supply of the nation as a whole. Butter is a private good because each person's benefit depends not on the total supply, but on the portion assigned to him exclusively. Public goods are non-rivalrous; one person's benefit from a public good does not detract from the benefit of any other person. Private goods are rivalrous; the more that accrues to me the less left over for you when total production is invariant. When the army buys a tank, that tank protects us all. When one of us consumes a pound of butter, the rest of us are automatically deprived of the pound of butter he consumes. To account for public goods, the arguments in the utility function are transformed from "b and c" into "b and G" where b, now mnemonic for butter rather than bread, is a person's consumption of a private good, and G, mnemonic for guns, is the total output of a public good.

Consider tanks rather than guns. A country with a population of 30 million buys tanks at a cost of $1.5 million each. When a tank is purchased by the government on behalf of the entire population, the cost of a tank per person is 5 cents [1.5 million/30 million]. A government acting on behalf of its citizens buys the tank if and only if the value to the typical citizen of the extra protection that the tank provides against the enemies of one's country is at least 5 cents. Value in this context is the amount of money the citizen would be prepared to pay for the extra protection that the tank supplies. If the citizen is prepared to pay in 5 cents, the tank should be purchased. Otherwise not.

The value of the tank to the typical citizen would presumably depend on the number of tanks the army has already. If the army is well supplied with tanks, the extra tank may not add much to the safety of the citizen. If the supply of tanks on hand is small, the value to the citizen of an extra tank may be quite large. The right procedure for a government buying tanks in the interest of its citizens is to acquire tanks up to the point where the value to the typical citizen of the last tank bought is just a nickel.

It would, on the other hand, be ludicrous for a citizen to supply tanks for the army all by himself. He might perhaps do so as an act of charity, but never in his own self-interest, for the disparity between the cost of a tank ($1.5 million) and the benefit he, personally, can expect from the tank (5 cents) is just too large. When he can spend his income as he pleases on guns (that is, on tanks) or butter, the citizen spends his entire income on butter even though he can be made better off when the government takes some of his money in taxation for the purchase of guns. The rational citizen votes for the compulsory purchase of items he would never buy on his own.

In sharp contrast to the bread-and-cheese economy in chapter 3 where the government can never rearrange production to make everybody better off, the guns-and-butter economy requires the active intervention of the government in determining what to produce. In one economy, it is sufficient for the government to enforce property rights. In the other, the government must intervene in production as well.

To formalize this example, imagine a society where everybody is just like everybody else and where supply prices are invariant so that the production possibility curve is a downward-sloping straight line and all supply curves are flat. In a bread-and-cheese economy with N identical people, the economic problem is to place the typical person on the highest possible indifference curve, maximizing his utility u(b,c) where b and c are consumption of bread and cheese per person. When, in addition, the production possibility curve is a downward-sloping straight line, each person can be thought of as producing y loaves of bread which he can convert to cheese at a rate of trade-off in production of p. His choice of b and c is then confined by the equation

$$b + pc = y \tag{1}$$

As the nation is the individual writ large, the national production possibility curve becomes

$$B + pC = I \tag{2}$$

where B = Nb, C = Nc and I = Ny which can be thought of as the national income in units of bread. Equivalently,

$$BP^{\$B} + CP^{\$C} = Y \tag{3}$$

where $P^{\$B}$ and $P^{\$C}$ are money prices of bread and cheese, and Y is the national income in units of money where, by definition, $Y = IP^{\$B}$.

By contrast, in a guns and butter economy, the economic problem is to maximize the typical person's utility u(b,G) when b and G are consumption per person of butter and the total supply of guns, and when the economy may be represented by the national production possibility curve

$$BP^{\$B} + GP^{\$G} = Y \tag{4}$$

where B is now the total production of butter, $P^{\$B}$ and $P^{\$G}$ are invariant supply prices of butter and guns, and Y is, once again, the national income. Equation (4) shows the options in production for the nation as a whole. Suppose, for convenience, that $P^{\$B}$ equals \$1 per pound so that the term $P^{\$B}$ drops out of equation (4) and $P^{\$G}$ can be interpreted equally well as the money price of guns and as the relative price of guns. The equivalent set of options for the individual, obtained by dividing all terms by N, becomes

$$b + G(P^{\$G}/N) = y \qquad (5)$$

where y is a person's income in units of butter, and $(P^{\$G}/N)$ is the price of guns as seen by the individual who can share the cost of guns with the rest of the population. Without collective provision of guns, each person would have to pay the full price of guns if he wished to add guns to the nation's stock. His options in productions become

$$b + GP^{\$G} = y \qquad (6)$$

which is the same as equation (5) except that *his* price of guns rises from $P^{\$B}/N$ to $P^{\$G}$.

Replacing tanks with guns in our example, the price of butter, $P^{\$B}$, is \$1 per pound, the price of guns, $P^{\$G}$, is \$1.5 million per gun and the population, N, is 30 million. Assume each person's income to be \$50,000. He produces 50,000 pounds of butter which can be converted to guns at a rate of $P^{\$G}$ pounds of butter per gun. On these assumptions, a person's options in production are represented by the two production possibility curves in figure 1. The lower and steeper production possibility curve shows a person's options for the purchase of guns and butter all by himself. The higher and flatter production possibility curve shows that person's options when he buys butter privately and buys guns collectively. All by himself, this person could acquire no more than a than one-thirtieth of a gun, even if he devoted his entire income to the purchase of guns. Collectively, he can acquire as many as a million guns [50,000 x 30 million / 1.5 million] if he, along with everybody else, devoted all of his income to the purchase of guns. Presumably, society would choose many fewer guns than that.

Each person's preference is represented in figure 1 by a set of indifference curves, comparable to the indifference curves for bread and cheese in chapter 3. A society of identical people unanimously chooses a combination of guns and butter, shown in the figure as b* and G*, to place each person on the highest possible indifference curve. Public choice when people have different preferences and different incomes is discussed under the general heading of voting in

Figure 1: Guns and Butter

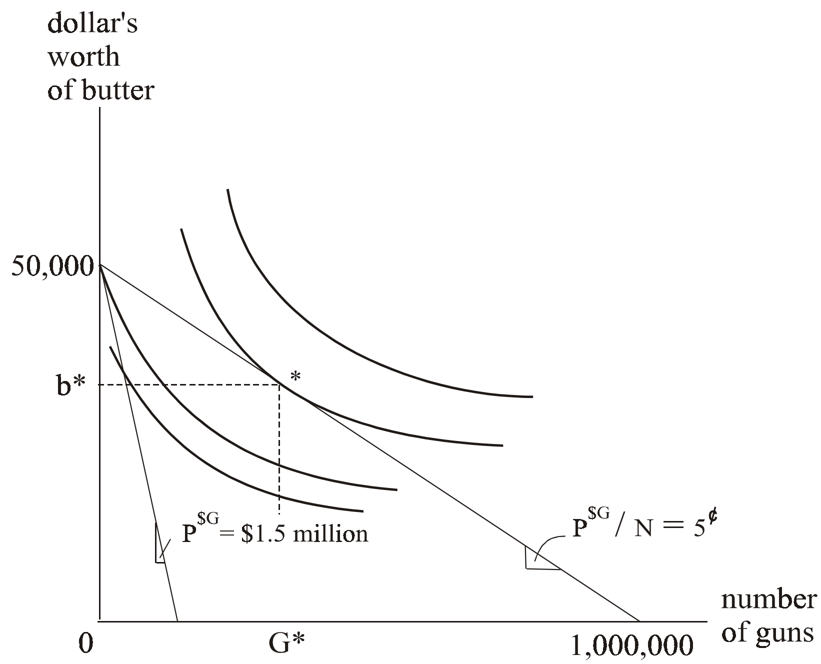

chapter 9. As the indifference curves are drawn, a person with a stock of 50,000 pounds of butter - representative, of course, of all private goods - cannot make himself better off by buying guns individually, but he can do so by buying guns collectively.

<u>Externalities</u>: u = u(b, c, C) where b and c are a person's consumption of bread and cheese and C is society's total consumption of cheese.

The general form of an externality is that one person's activity is harmful or beneficial to somebody else where the harm or benefit is not reflected by the price mechanism. A person's consumption of cheese might convey some harm to the rest of society if manure from cattle farms polluted the drinking water. A person's consumption of cheese might convey some benefit to the rest of society if cheese production gave out a delicious smell that everybody enjoyed. In either case, each person's utility would be affected by the total consumption of cheese. A somewhat different example illustrates what may be at stake.

Fifty students in a lecture hall. Every student likes to smoke but dislikes sitting in a smoky room. All students are alike. Each student who smokes emits 1,000 particles of smoke into the room. In deciding to smoke, a student acquires the pleasure of smoking together with the displeasure of occupying a slightly more smoky room, and his decision to smoke or not to smoke is necessarily a balancing of these considerations.

Every student values the pleasure of smoking at one dollar. Every student values the displeasure of smoky rooms at (1/100)¢ per particle of smoke. To say that the student values the pleasure of smoking at $1 is to say that - taking account of the price of cigarettes and the risk to health - he would pay up to $1 for the right to smoke if, miraculously, his own smoking had no effect on the amount of smoke in the room, that is, if the amount of smoke in the room were the same regardless of whether he smokes or not.

Two social arrangements are to be compared. In the first, each student smokes or desists from smoking as he pleases. In the second, students vote to permit smoking or to forbid it altogether. Voting is by majority rule, but the details of the voting mechanism do not matter here because identical students would vote identically as long as voting is restricted to rules treating everybody alike. Students are assumed to be entirely selfish, full-fledged "economic men" who maximize their own welfare exclusively, in choosing to smoke or not to smoke or in voting to allow or disallow smoking.

Suppose, first, that each student smokes or desists from smoking as he pleases. In deciding whether to smoke, a student must weigh the pleasure of smoking against the extra unpleasantness in the lecture hall from the extra smokiness he would create. The monetary equivalent of the pleasure of smoking is $1.00. The monetary equivalent of his displeasure from the extra smokiness he generates by smoking is 10¢, [1,000 particles @ 1/100¢ per particle]. His net gain from smoking is therefore 90¢, and he chooses to smoke.

Suppose instead that smoking is permitted or forbidden altogether in accordance with the outcome of a vote. There is a proposal to ban smoking, and each of the fifty students must vote yes or no. How does a student vote? Once again, he compares the benefit and cost of smoking, but now it is on the understanding that, by giving up his own right to smoke, he is stopping his classmates from smoking as well. In voting to ban smoking, a student forgoes the pleasure of smoking in return for an entirely smoke-free room. He values the loss of the pleasure of smoking at $1. He values a smoke-free room as compared with a room where all 50 students smoke at $5.00. [When each of 50 students emits 1,000 particles of smoke, there are 50,000 particles of smoke in the room. A student who values the harm from each particle at (1/100)¢, must value the harm to himself from 50,000 particles at $5.00.] Since every student's gain from a smoke-free room ($5.00) exceeds his benefit from smoking ($1.00), the students vote unanimously to have smoking banned. In his own self-interest, a rational person votes to ban smoking though he personally would smoke if the ban were not enacted! The moral of the story is that selfish people may vote to prohibit behaviour that would be advantageous were it not prohibited.

This is no great paradox. It is a commonplace example of an "externality." The bread and cheese story in the last two chapters was about goods that convey benefits to consumers without affecting anybody else. Other goods or actions convey costs (such as smokiness in our example) or benefits (such as the security from the presence of an army) to people without compensation. Such benefits or costs to other people are called externalities because they are external to the price mechanism. The market does not require the purveyor of a negative externality to compensate his victims, or provide the purveyor of a beneficial externality with compensation from his beneficiaries. Smoke in the air exemplifies a great many types of pollution where some common, nationwide or worldwide aspect of life is adversely affected by actions of people who are not held responsible for the consequences of their behaviour. Depletion of fish stocks by fishermen, the effect of hydrofluorocarbons on the ozone layer of the atmosphere, smog in cities, the greenhouse effect of carbon dioxide in the atmosphere, the extinction of species of plants and animals and the destruction of rainforests are all, with individual variations, negative externalities of private actions. Positive externalities are conveyed by a beautiful garden that gives pleasure to passers-by, by television programs that may be watched by anybody who chooses to turn on his set, or by education that increases the taxable earnings of the student.

Returning to smoking and smokiness, note that, when smoking is allowed, it is individually advantageous for each person to smoke no matter how many or how few others do so. If n out of N students choose to smoke, the net benefit or cost, of smoking and smokiness together, to a student who chooses to smoke is the monetary values of his pleasure in smoking less his displeasure when n students smoke. As each smoker imposes a cost of 10¢ on every student in the room, the net benefit or cost to a student who smokes is $(1 - n/10). A student who chooses not to smoke bears the same cost of smokiness without the pleasure of smoking himself. He bears a cost of $(n/10) and is always worse off than a student who smokes regardless of the number of students who smoke. This is illustrated in figure 2 with the number

Figure 2: How Universal Selfishness May be Worse than Cooperation

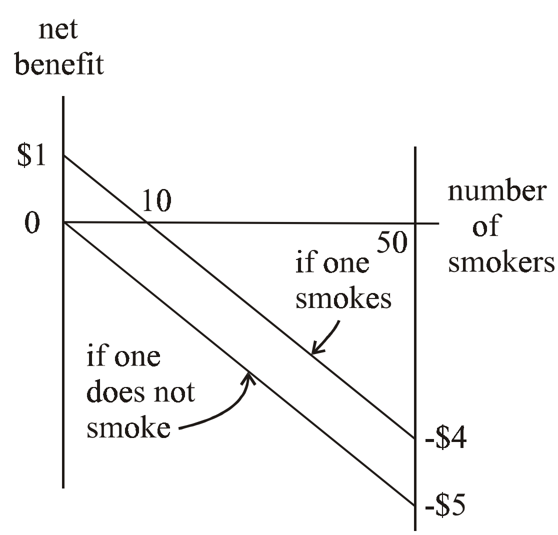

of smokers on the horizontal axis and the net benefit per person on the vertical axis. The two lines in the figure show the net benefit per student, if he smokes and if he does not smoke, depending on the number of smokers in the class. Net benefit if one smokes falls steadily with the number of smokers from $1 if it were somehow possible to smoke without smoking up the room at all, to -\$4.00 if all fifty students smoked. Net benefit if one does not smoke falls steadily from 0 to -\$5.00. The diagram shows that one is always better off smoking than not smoking, though everybody would be better off if nobody smoked at all. The "tragedy of the commons" in the quotation at the outset of this chapter is that uncoordinated self-interested behaviour makes everyone worse off than he might be if the students could coordinate their behaviour in the common interest. The tragedy is that the Adam Smith's invisible hand - referred to in the quotation at the beginning of the introduction of this book - may sometimes be unreliable.

The conclusion is dependent on the parameters of the example. It was assumed that the harm per particle of smoke in the room was 1/100¢ per particle. Had it been assumed instead that the harm is only (1/1,000)¢ per particle, the harm per student from smoke in the air when all 50 students emit 1,000 particles each would have been only 50¢ rather than \$5.00, and every student would be better off when everybody smokes than when nobody smokes. Not all externalities should be banned or even curtailed. It is sometimes in the common interest to ignore externalities because the private gain from an externality-bearing activity exceeds the social cost. My neighbour may resent my neglect of my front lawn, but, within limits, we think it best for each person to look after his lawn as he pleases.

Alternatively, the common interest might best be served by a reduction in smoking rather than by a total ban. Suppose a student's benefit from smoking depends on how many cigarettes he smokes. Let p represent the student's valuation of his pleasure from smoking one extra cigarette; p is a student's demand price per cigarette, excluding the monetary equivalent of his displeasure from the extra smokiness brought about by his own smoking. Suppose p diminishes with q, the number of cigarettes smoked, in accordance with the demand curve

$$q = 5 - 5p \tag{7}$$

For convenience, the quantity of cigarettes is assumed to be continuous rather than discrete, as though one could buy fractions of a cigarette. The student smokes five cigarettes if they are free, no cigarettes if they cost a dollar each, one cigarette if they cost 80¢ each, and 3.95 cigarettes if they cost 21¢ each.

The demand and supply sides of the market for cigarettes are illustrated in figure 3 on the assumptions that the price of a cigarette at the smoke shop is 20¢, that each cigarette emits 100 particles of smoke and that the monetary value of each person's displeasure from smokiness is (1/100)¢ per particle. The demand curve is a direct representation of equation (7), but there are three distinct supply curves, corresponding to different interpretations of the price of cigarettes. All three curves are flat because prices per cigarette are independent of the number of cigarettes smoked. The height of the lowest supply curve is 20¢, the price of cigarettes at the smoke shop. The height of the next supply curve up is 21¢, the full cost of a cigarette to the smoker: the sum of the price at the smoke shop and the cost to the smoker himself [100 x (1/100)¢] of the extra smokiness in the lecture hall. The height of the top curve is 70¢, the entire cost of a cigarette to society: the sum of the cost of the cigarette at the smoke shop and the combined cost to all 50 students [50 x 100 x (1/100)¢] of the extra smokiness when one student smokes one extra cigarette.

Figure 3: Private and Social Cost of Smoking

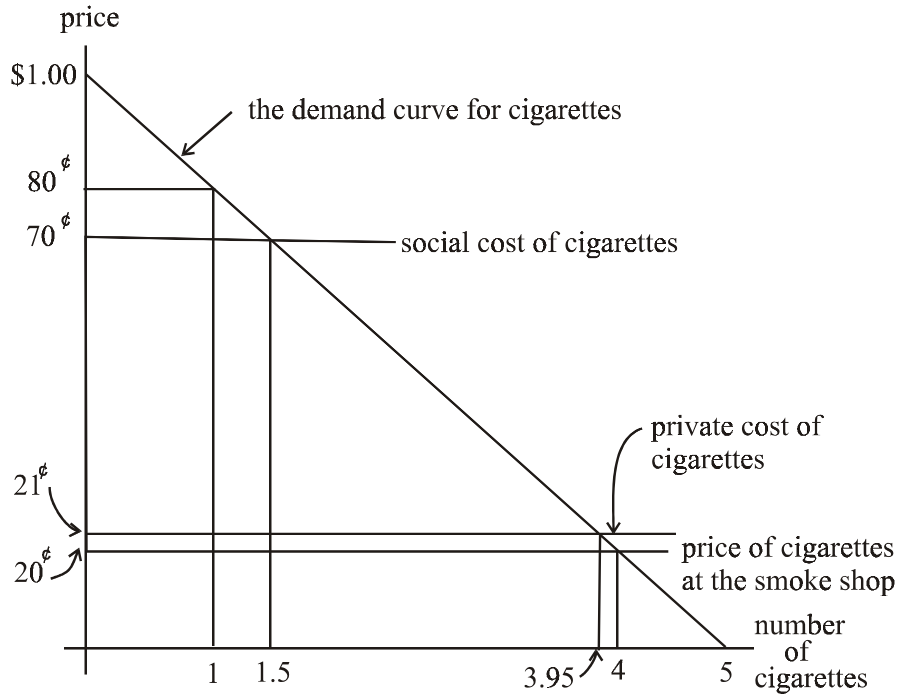

A student whose smoking is unconstrained, chooses to smoke 3.95 cigarettes at which private benefit from one's very last puff is just equal to private cost. By contrast with the preceding variant of this example, the class would not now vote to ban smoking altogether. Consider a vote to allow each student to smoke one cigarette but no more. The benefit to the smoker is 90¢ [half way between the value of the first puff which, as may be seen from the demand curve, is worth $1.00 per cigarette, and the last puff which is worth only 80¢]. The cost to each person of the smoke generated when everybody smokes one cigarette is 50 ¢ and the cost of the cigarette itself is 20¢. The benefit per person when everybody smokes one cigarette, and no more, is 95¢. The corresponding cost is only 70¢. Thus everybody is better off smoking one cigarette than no cigarettes, and student vote unanimously to allow the first cigarette.

Self-interested students vote to allow one another to smoke one and a half cigarettes. Imagine a vote between two bills in the student parliament. The first bill allows all students to smoke 1.5 cigarettes. The second bill allows students to smoke a number, q* of cigarettes, where q* is different from 1.5. Students vote unanimously for the first bill. When q = 1.5, the corresponding p in accordance with equation 1 is 70¢ which is precisely the full social cost of a cigarette. Below 1.5, the benefit of an extra puff exceeds the social cost. Above that, the benefit of an extra puff falls short of the social cost.

In the original version of this example, all students, left to themselves, would smoke, but the appropriate public policy was to ban smoking because everybody was better off when nobody smoked than when everybody smoked. Now all students, left to themselves, would smoke 3.95 cigarettes, but the appropriate public policy is to reduce smoking to 1.5 cigarettes. Students would vote unanimously to reduce smoking to 1.5 cigarettes. The question now arises of how such a rule might be policed. When smoking was banned altogether, it would presumably be obvious to everybody when the ban was violated. It may be less obvious whether a particular student is exceeding his quota.

The standard remedy is taxation. Suppose the class party costs $150 in total or $3 per student. The class might vote to impose a tax of $3 per student to finance the party. Alternatively, it could levy a tax of 49¢ per cigarette, raising the tax-inclusive cost of a cigarette 70¢ [20¢ at the smoke shop, plus 1 cent for the cost of smokiness to the smoker, plus 49¢ tax to compensate other students for the harm from the smoke emitted by a cigarette]. With this tax in place, each student reduces his smoking to 1.5 cigarettes, the income from the cigarette tax is $36.75 [49¢ x 50 x 1.5] or 73.5¢ per student. The extra $2.26 per student to pay for the party could be acquired by direct taxation. The revenue to finance the party remains at $3 per student, but the change in the form of the tax alters smoking patterns making everybody better off than he would be if smoking were unconstrained. [While it is true that taxing students may be cumbersome in the context of this example, it should be evident to the reader that what we are really talking about is an excise tax on the producer of an externality-bearing activity.]

The moral of the new version of the story is that it is better to tax than to ban when the common good is best served by a reduction, rather than a complete termination, of the externality-bearing activity.

Now consider once again our bread-and-cheese economy with the additional assumption that there is an externality in the production of cheese. Cheese factories belch smoke which pollute the atmosphere, and each person's well-being depends upon his own consumption of bread and cheese and upon the total amount of smoke in the atmosphere. There are N people in the economy, production of a pound of cheese emits s particles of smoke, and the cost to each person of a particle of smoke (the amount of bread he would give up to procure a reduction of one particle of smoke in the atmosphere) is k. Then the "private" supply price of cheese to the buyer of cheese is

$$p^p = ()b + sk)/)c \qquad (8)$$

where the superscript p is mnemonic for private, where)b is the number of loaves of bread that must be sacrificed when resources are diverted to make)c extra pounds of cheese, and where sk is the cost (to the buyer of cheese alone) of the extra smokiness in the atmosphere that he generates in arranging for the production of one extra pound of cheese. The "social" supply price of cheese is different. It incorporates the cost to every person when any person consumes one extra pound of cheese. It is

$$p^s = (\,)\, b + skN)\,)\, c \tag{9}$$

where the the superscript s is mnemonic for social and where N is introduced because everybody has to put up with extra smoke when anybody consumes an extra pound of cheese. The story is illustrated in figure 4 which, like figure 3 above, shows demand and supply curves with price on the vertical axis and quantity on the horizontal axis. The important difference between the two figures is that the price in figure 3 was a money price (cents per cigarette) while the price in figure 4 is a relative price (loaves of bread per pound of cheese as in the demand and supply curves in chapter 4). The demand curve for cheese in figure 4 is just like the demand curves in the last two chapters. Now, however, there are two supply curves: the private supply curve shows the cost of cheese in terms of bread as the sum of the ordinary cost of production and the cost to the consumer himself of the smoke from his extra consumption of cheese, and the social supply curve shows the cost of cheese in terms of bread as the sum of the ordinary cost of production and the cost to everybody of the smoke from some person's extra consumption of cheese. The heights of the supply curves above any arbitrarily chosen point c are shown. If the externality is uncorrected, then the equilibrium, market-clearing consumption of cheese is c^e pounds (where the superscripts e is mnemonic for equilibrium). This is too large because everybody would become better off if each person reduced his consumption of cheese to c^o pounds (where the superscript o is mnemonic for optimal) which is what a society of identical people would vote for unanimously. As in the classroom example, the optimal outcome could be obtained by taxing the externality-bearing activity.

Figure 4: How the Equilibrium Quantity of Cheese Differs from the Optimal Quantity when there is an Uncorrected Externality

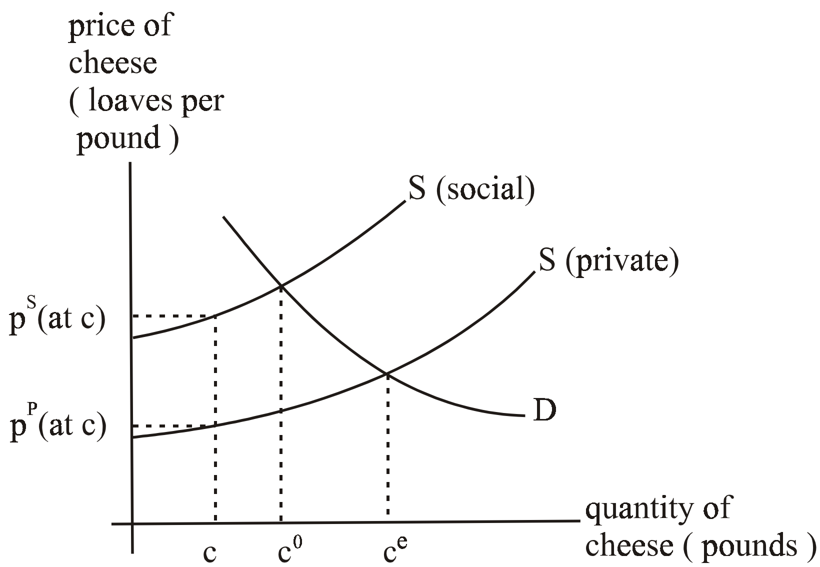

On the face of it, the moral of the smoke and smoking story told here appears to be the exact opposite of the taxation story in chapter 4. There, it was shown that, if an amount of money is to be extracted from Robinson Crusoe by taxation, he is better off when that money is taken as a lump sum regardless of how he spends his remaining income, or when both goods are taxed at equal rates, when cheese is taxed but not bread. The discrepancy is only apparent. In the context of chapter 4, taxation of cheese but not bread created a distortion in an otherwise distortion-free economy. Here, taxation of cheese corrects for an externality-borne distortion. The overriding moral is to use taxation to correct distortions in the market, but, otherwise, to tax all goods alike, perhaps by means of an income tax. In either case, the welfare of the typical citizen is the appropriate guide to public policy.

Finally, an important aspect of many externalities is abstracted from, and even hidden by, the smoke-and-smoking story. The story is strictly *atemporal*. Actual externalities are often *intertemporal*. The distinction is best introduced by an extension of the story. In telling the story, it was implicitly supposed that the private benefits from smoking and the common harm occur all at once. The quality of the air today depended on how many people smoke today regardless of whether or not people smoked yesterday or at any time in the past. Nor does smoking today affect the quality of the air tomorrow. It is as though air quality restores itself automatically over night. The assumption was atemporal (without time) in the sense that cause and effect are simultaneous. In this example as in many others, the atemporal assumption is very convenient and effective in conveying the essence of a problem clearly. It may nevertheless conceal important considerations.

The smoke-and-smoking story could have been told differently. It might have been supposed that smoke hangs in the air for days or even years. The harm from smoking today would then linger on and on, perhaps forever. This would be an intertemporal (among different times) variant of the story, more difficult to tell but capturing aspects of the world that might turn out to be important. Many actual environmental problems have a crucial intertemporal dimension. Grazing of my cattle on the commons today diminished the nutrition available for your cattle not just today but for years to come as the soil is, perhaps permanently, depleted. Overfishing today depletes the stock of fish tomorrow. Greenhouse gasses emitted today do little harm immediately but hang in the atmosphere for centuries gradually warming the planet. Extraction of minerals today has no impact on human welfare until such time as the world's mineral resources become more expensive to extract or run out altogether. All of these phenomena are like the smoke-and-smoking example in that private gain yields public loss, but their timing is not the same.

The next section passes from atemporal to intertemporal analysis. Externalities from population growth are discussed in some detail in the next chapter.

<u>Consumption Today and Consumption Tomorrow</u>: [$u = u(b_1, b_2)$ where b_1 and b_2 are loaves of bread consumed at different periods of time]

The models in chapters 2, 3 and 4 were strictly atemporal. Production and consumption took place in an isolated moment of time with no links to the future or to the past. One can think of those models as either abstracting from connections between moments of time or pertaining to a world where every yesterday and every tomorrow is just like today, and where people's actions remain the same forever. The simplest way to introduce time into the model is to replace "bread and cheese today" with "bread today and bread tomorrow", and to replace the relative price of cheese with the rate of interest as the mediator between the quantities of the two goods consumed. There is a straightforward analogy between (1) the relative price of cheese as a mediator between the demand for and supply of cheese at a moment of time and (2) the rate of interest as a mediator between consumption of bread today and consumption of bread tomorrow.

To see this, it may be helpful to begin with the ordinary rate of interest on money. Consider the deposit of money in a bank. I deposit $1,000 at 5% interest. In doing so, I am trading $1,000 today for $1,050 next year, or, equivalently, I am buying dollars next year with dollars this year at a price of approximately 95¢. Recall that the price of cheese in terms of bread is $\Delta b/\Delta c$ where Δb is the amount of bread I must give up to acquire an amount of Δc of cheese. By analogy, the price of dollars next year in terms of dollars this year has to be (Δ dollars this year)/(Δ dollars next year) which in our example is $(1000)/(1050) \approx 0.95$. That is how much money I must give up this year to acquire a dollar next year. More generally, when the rate of interest is that r (r = 0.05 means a rate of interest of 5%), a dollar this year exchanges for (1+r) dollars next year, and the price of a dollar next year in terms of dollars this year is 1/(1+r). The rate of interest is not some mysterious entity that is altogether different from prices. It is just an ordinary price transformed. Similarly, if I leave money in the bank indefinitely and collect the interest each year, then 1/r becomes the price in terms of current dollars of a stream of $1 per year forever.

Our main concern here is not with the relative prices of money available at different periods of time, but with the comparable prices of goods. We are concerned here with the relative price of bread next year and bread today. To buy a pound of cheese with money is to exchange dollars for cheese. To buy a pound of cheese with bread is to exchange bread for cheese. To buy a loaf of bread next year with bread today is to give up as many loaves of bread today as the market requires in exchange for a loaf of bread next year. Designate this year by 1 and next year by 2. Suppose the price of bread this year is $P^{\$B}(1)$, the price of bread next year will be $P^{\$B}(2)$ and the ordinary rate of interest on money is r, or, equivalently, 100r%. To exchange a loaf of bread today for bread available next year, here is what I must do. I sell the loaf for $P^{\$B}(1)$, the going price of bread today, and then I lend the $P^{\$B}(1)$ dollars at the going rate of interest. At the end of the year, I receive back $P^{\$B}(1)(1 + r)$ dollars with which I purchase bread at the going price at that time, acquiring $P^{\$B}(1)(1 + r)/ P^{\$B}(2)$ loaves of bread next year. If, by this process, one loaf of bread today can be exchanged for $P^{\$B}(1)(1 + r)/ P^{\$B}(2)$ loaves delivered next year, then, by definition, the relative price of bread next year, with the bread this year as the numeraire, must be the inverse of that fraction. With bread this year as the numeraire, the price today of bread available next year - the number of loaves one must give up this year in order to acquire extra one loaf next year - becomes $P^{\$B}(2)/ [P^{\$B}(1)(1 + r)]$.

Now define the *own rate of interest on bread*, r_B, as the amount of extra bread I can acquire next year by postponing the consumption of one loaf by one year. From the definition r_B, it follows at once that a loaf today exchanges for $1 + r_B$ loaves next year, and that the relative price of bread next year in terms of bread this year is $1/(1 + r_B)$. It then follows immediately that

$$1/(1 + r_B) = P^{\$B}(2)/ [P^{\$B}(1)(1 + r)] \tag{10}$$

for the two sides of the equation are different expressions for one and the same thing. Equation (10) has two straightforward but important consequences. First, if the price of bread is the same in both years - if $P^{\$B}(2) = P^{\$B}(1)$ - then the money rate of interest and the own rate of interest on bread must be the same, i.e.

$$r_B = r. \tag{11}$$

Second, when the rate of inflation of the price of bread is i, that is, when $P^{\$B}(2)/ [P^{\$B}(1) = 1 + i$, then

$$(1 + r) = (1 + r_B)(1 + i) \tag{12}$$

or, equivalently, when the cross-product, r_R i, is considered small enough to ignore,

$$r = r_B + i \tag{13}$$

If prices of goods are changing over time at different rates, then each good has its own forward price and its own rate of inflation, but equation (13) remains valid for each good individually as long as the term r_B is reinterpreted as the own rate of interest for the good in question.

If prices of all goods increase at the same rate, then the term i in equation (12) could be interpreted as the common rate of inflation. However, the equation remains approximately valid when the term i is reinterpreted as the average rate of inflation as measured, for example, by the consumer price index. In this context the term r_B becomes the average of the own rates of interest on all goods together and is called the "real rate of interest". Ignoring cross-products, equation (13) can be interpreted as the proposition that

"the money rate of interest" equals "the real rate of interest" plus "the rate of inflation." (14)

If the money rate of interest is 5% and if prices of goods and services are rising at 3%, then a dollar invested today yields me 2% more goods next year than I must give up this year to acquire them. A person saving for his old age is of course concerned with the real rate of interest on his money, rather than with the money rate of interest.

The introduction of time requires us to differentiate between *flow* prices and *stock* prices. The relative price of cheese is a flow price, a rate of exchange between two goods produced and consumed at a moment of time. Money prices of bread and cheese are also flow prices because bread and cheese are short-lived even though money is not. The price of land is a stock price because land persists through time and because the purchase of a plot of land today is really the purchase of the stream of goods that the stream of goods produced by that land every year until the end of time. Stock prices are connected to flow prices by interest rates. If a plot of land yields $1,500 per year forever and if the money rate of interest is 5%, then the price of land must be $30,000 [1,500 x 1/(0.05)]. If a plot of land yields 500 loaves of bread each year forever, if the price of bread today is $3 a loaf, and if the own rate of interest on bread is 2%, then the price of that land - the stock price - has to be $75,000[500 x 3/(0.02)]. Note that the own rate of interest on bread can only differ from the money rate of interest when the price of bread is expected to increase over time, so that a constant flow of bread becomes the equivalent of a steadily increasing flow of money. At these interest rates, one would need to deposit $30,000 in the bank, to provide oneself with an annual income of $1,500, but one would need $75,000 to provide oneself with an annual flow of 500 loaves of bread.

Though interest rates are usually positive, they are not always so. Money rates of interest cannot be negative as long as gold or paper money can be stored costlessly. Real rates of interest are usually positive because land and machinery are productive. Expenditure on machinery today can be expected to yield a positive return, after provision for depreciation, for as long ahead as one can see. But real rates of interest can be negative in some circumstances. The five year own rate of interest on grain was negative in ancient Egypt at the end of the seven fat years, when the seven lean years were due to begin, and when pharaoh, on the advice of Joseph with his flair for economic planning, had accumulated stocks of grain some of which would surely be eaten by mice or accidentally burned from time to time. If, of every bushel of grain stored the, only half a bushel remains available to be consumed in five years time, then the own rate of interest on grain would be -13%, the solution to the equation $(1+r_B)^5 = ½$. Similarly and for the same reason, the own rate of interest on grain is typically negative over the six months after the harvest.

Nothing has been said so far about how the rate of interest is determined. As the rate of interest is a price, one would expect it to be determined by the same interactions between taste and technology that determine the relative price of cheese. So it is, but the mechanism is complex because the rate of interest today is conditioned by anticipations of future for as far ahead as one can see.

A simplified version of the mechanism determining the rate of interest is illustrated in a reinterpretation of the Robinson Crusoe story as set out in chapters 3. Now, Robinson Crusoe consumes only bread, not bread and cheese as had been assumed. Instead, he lives for two years, youth and age, produces 10 loaves of bread when he is young, and invests of his produce so that he has something to consume when he is old. In chapter 3, Robinson Crusoe was confronted with a trade off in production between bread and cheese. He could produce more cheese and less bread or more bread and less cheese. Now he is confronted with a similar trade off between quantities of bread at two periods of time. Suppose - no matter how - he can transform 1 loaf of bread in the first year for 1.1 loaves of bread in the second, so that his technologically determined rate of interest on bread is 10%. There is no explanation within the model of why the technologically given rate of interest is 10% or of why it is positive at all. One might think of this as representative of more elaborate economies where investment - the giving up of consumption this year to acquire consumption in the future - is productive.

Robinson Crusoe's options for consumption in the two years can be represented by the intertemporal production possibility curve

$$10 = b_1 + (1/1.1)b_2 \tag{15}$$

where b_1 his production of bread in the first year and b_2 is his production of bread in the second. Among his options are to consume 10 loaves per day in first year and nothing in the second, to consume 5 loaves per day in the first year and 5.5 loaves in the second, or to consume nothing in the first year and 11 loaves in the second. It follows immediately that the supply price of bread next year in terms of bread this year is ratio $\Delta b_1/\Delta b_2$ where Δb_1 is a small decrease in consumption of bread this year and Δb_2 is the resulting increase in consumption of bread next year in accordance with equation (15), and that $\Delta b_1/\Delta b_2 = 1/1.1$. Note also that $\Delta b_1/\Delta b_2$ is the same for all consistent values of b_1 and b_2 because equation (15) is a downward-sloping straight line.

Robinson Crusoe's tastes, or preferences, can be represented by a set of indifference curves with a corresponding utility function

$$u = u(b_1,b_2) \tag{16}$$

analogous to his earlier choice of bread and cheese. He chooses consumption each year, b_1 and b_2 , to place himself on the highest possible indifference curve, as specified in equation (16), attainable with the available technology, as specified by the production possibility curve in equation (15). His chosen combination of bread this year and bread next is illustrated in figure 5, analogous to figure 5 in chapter 3, with technology and taste in part A and with demand and supply in part B.

Indifference curves and the production possibility curve are shown in part A of figure 5 with b_1 on the vertical axis and b_2 on the horizontal axis. Clearly, Robinson Crusoe has attained the highest possible indifference curve at a combination of b_1 and b_2 for which that indifference curve is just tangent to the production possibility curve. In the special case where indifference curves can be represented by the simple function $u = b_1b_2$, the slope of the indifference curve is equal (by analogy with the earlier demonstration for bread and cheese) to b_1/b_2 . Equating the slope of the indifference curve to the slope of the production possibility curve, we see that $b_1/b_2 = 1/1.1$. Since any combination of b_1 and b_2 must lie on the production possibility curve, the two equations $b_1/b_2 = 1/1.1$ and $10 = b_1 + 1/1.1 b_2$ imply that $b_1 = 5$ and $b_2 = 5.5$.

Figure 5: Intertemporal Choice

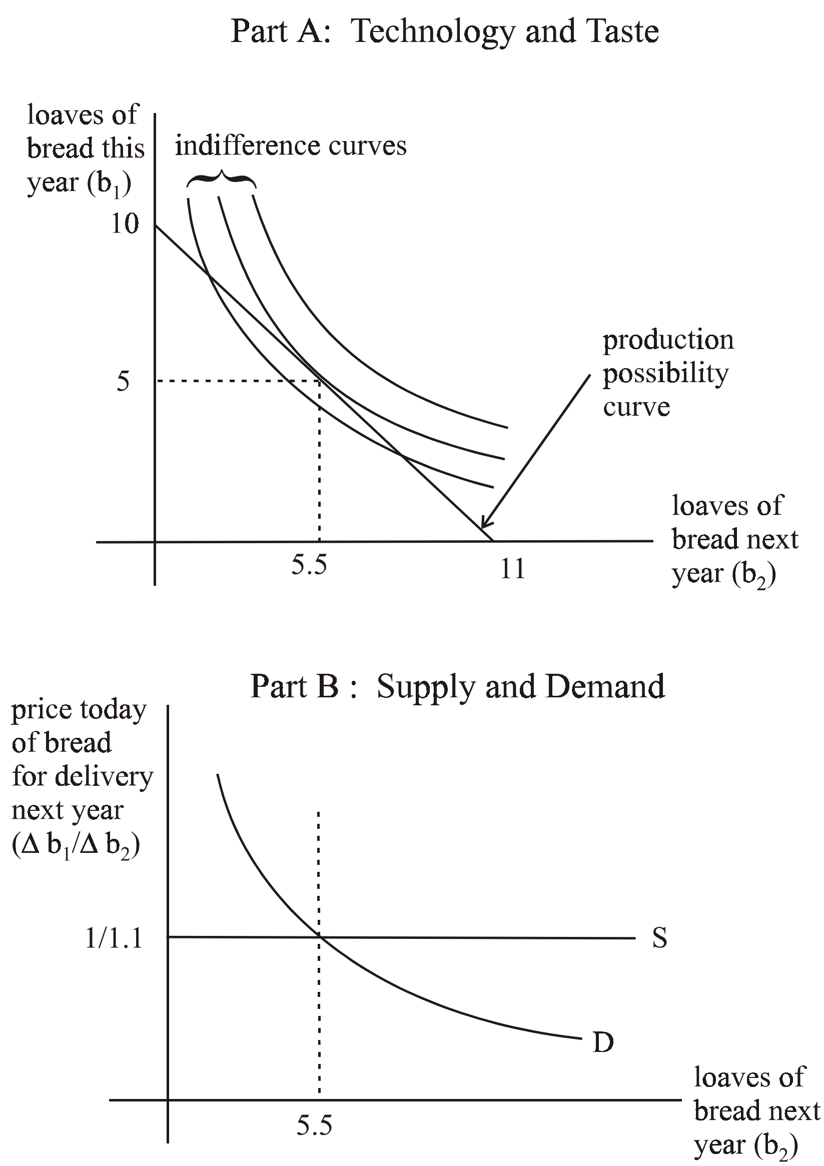

Exactly the same story is told in part B with b_2 (consumption in the second year) and the price of b_2 in terms of b_1 on the vertical axis. The price of b_2 in terms of b_1 is the ratio $)b_1/)b_2$, the amount of bread this year that must be given up to acquire a loaf of bread next year. The supply curve, S, shows the supply price of bread in the year 2 as a function of the quantity supplied, where the supply price is, as in earlier chapters, the rate of substitution in production as indicated by the slope of the production possibility curve in part A. Here the supply curve is flat because the production possibility curve is a downward sloping straight line. The demand curve shows the slope of the indifference curve along the production possibility curve for various values of b_2. The demand curve is downward sloping reflecting the assumed curvature of the indifference curves. The intersection of the demand and supply curves shows the amount of bread produced and consumed in the second year together with the relative price of bread in the second year from which the real rate of interest (i.e. the own rate of interest on bread) can be determined.

The rate of interest and the allocation of consumption over time are determined by the same market forces that determine the prices of goods and the choice of a mix of goods at a moment of time. Just as the atemporal story was told about an economy with only two goods, so that intertemporal story is told about an economy with only two periods of time. Both stories can be generalized to many goods and many periods of time.

Risk: [$u = u(Y_V, Y_M, Y_U)$ where Y_V, Y_M and Y_U are incomes in different *states of the world*.]

A person is choosing a career. He has narrowed his options to law and medicine, and his only concern in this choice is his annual income once his career is under way. If he knew for certain what his annual incomes would be in the two careers, he would automatically choose the career with the higher income, but both careers are somewhat risky. Specifically, in each career, he is equally likely to be very successful, moderately successful or unsuccessful. His incomes in each career and in each eventuality are shown in table 1. The main difference between these careers is that the expected income is higher in law than in medicine ($110,000 as compared with $100,000), but law is more risky. One may earn as little as $30,000 or as much as $190,000 in law, as compared with $90,000 and $110,000 in medicine. What does this person do?

Table 1: Annual Incomes in Law and Medicine Depending on whether one is Very Successful, Moderately Successful or Unsuccessful.

	unsuccessful	moderately successful	very successful	expected income	dispersion of income
medicine	$90,000	$100,000	$110,000	$100,000	$10,000
law	$30,000	$110,000	$190,000	$110,000	$80,000

The person's choice between law and medicine depends on his attitude toward risk. If he is indifferent to risk or only slightly risk averse, he chooses law which provides him with the higher expected income. If he is quite risk averse, he chooses medicine with a lower expected income but a smaller gap between the best and the worst outcome. Our object here is to describe this choice precisely.

An uncertain world can be modeled precisely as consisting of a number of *states of the world*, each of which will occur with a certain probability. In the toss of a coin, there are two states of the world, heads and tails, each of which will occur with a probability of one-half. In the toss of a dice there are six states of the world, each of which will occur with a probability of one-sixth. In each profession as described above, there are three states of the world, unsuccessful, moderately successful, and very successful, each of which will occur with a probability of one-third.

To compare professions, each yielding a different set of incomes depending on the state of the world, we would like to construct a utility function - over incomes in different states of the world rather than over different amounts of bread and cheese - such that the preferred profession yields the greater utility. We would like to construct a utility function

$$u = u(Y_V, Y_M, Y_U) \qquad (17)$$

where Y_V, Y_M and Y_U are a person's incomes in equally likely "states of the world", that the person is very successful (V), moderately successful (M) and unsuccessful (U). We would hope that the function which reflects choice among risky prospects with three states of the world can be generalized to any number of equally likely states of the world.

Suppose initially that the person is risk neutral. To say that a person is risk neutral that, in any choice between gambles or between a gamble and a sure thing, he always chooses the option with the largest *expected income*, E, defined with reference to our example as

$$E = (1/3)Y_V + (1/3)Y_M + (1/3)Y_U \qquad (18)$$

The expected income of doctors is what every doctor would obtain under an agreement among all doctors to share their incomes, whatever they turn out to be. The expected income of lawyers is defined accordingly. If incomes in each profession were to be shared among all practitioners, then everybody would choose law over medicine. A person is who is risk neutral chooses law regardless, even when income is not shared. A person who is strongly risk averse might choose medicine over law because the larger expected income in law does not compensate such people for the one-third chance of an income of only $30,000 which is only a third of the lowest possible income in medicine.

Corresponding to every risky prospect is a *certainty equivalent* that differs from one person to the next according to their degrees of risk aversion. Consider the choice between a risky prospect with expected income E and a sure thing with an income of Y. A risk-neutral person chooses whichever is the larger. A risk-averse person might choose the sure thing even when Y is less than E. A person's certainty equivalent of a risky prospect is an income Y^c (where c is mnemonic for certainty) such that the person is indifferent between the risky prospect and the sure thing. If one is risk neutral then $Y^c = E$ where Y^c is the certainty equivalent and E is the expected income of the risky prospect. If one is risk averse, then $Y^c < E$. The size of the gap depends on the person's degree of his risk aversion. The greater his risk aversion, the greater the gap. When the risky prospect is law or medicine and when the person choosing a profession is risk averse,

$$Y^c < E = (1/3)Y_V + (1/3)Y_M + (1/3)Y_U \tag{19}$$

Consistency of choice requires that one choose the profession with the higher certainty equivalent as assessed in accordance with one's taste for risk

Since all choice can be interpreted as the maximizing something and since a risk-averse person does not maximize expected income in his choice among risky prospects, there must be some function of income that a risk-averse person does maximize. We call that function the *utility of income*, u(Y), which may be thought of as a measure of utility to account for risk. Recall that, in the world without uncertainty in chapters 3 and 4 where people consume only bread and cheese, money income is a completely satisfactory measure of utility as long as prices are invariant, for, no matter what the shape of one's indifference curves, one is better off at any given set of prices with more money rather than less. Recall also the discussion of the numbering of indifference curves at the beginning of this chapter where it was shown that *any* numbering of indifference curves is satisfactory as an indicator of utility as long as higher curves are assigned higher numbers, and that any monotonicly increasing function of a satisfactory utility indicator is itself a satisfactory utility indicator. For example, if indifference curves conform to the function u = bc, so that the product bc is a satisfactory utility indicator, then so too are $(bc)^2$ and $10(bc)^{12}$. The utility of income function we seek is a renumbering of utility to account for a person's behaviour toward risk.

Utility is initially represented by income. What we are seeking in a utility of income function, u(Y), is a monotonic transformation of income such that, if Y^c is the certainty equivalent of equal chances of incomes of Y_U, Y_M and Y_V, then

$$u(Y^c) = (1/3)u(Y_V) + (1/3)u(Y_M) + (1/3)u(Y_U) \tag{20}$$

Having to choose between any two risky prospects, a person always chooses the prospect with the larger expected utility of income. For a risk neutral person, this means no more than that he chooses the prospect with the higher expected income, for u(Y) = Y, the inequality in equation (19) becomes an equality, and equations (18) and (20) are essentially the same. Otherwise the shape of u(Y) has to be discovered.

Since the utility of income function is designed to represent a person's willingness to bear risk, and since one's willingness to bear risk is an aspect of taste, the only way to discover a person's utility of income function is to ask him, directly or indirectly, to tell you what it is. A person's utility function is discovered in essentially the same process that was used to discover his ordinary indifference curves over different combinations of bread and cheese. He must be asked a long series of questions of the general form, "Do you prefer this to that?" until the entire shape of his utility of income function is revealed. The trick is to ask a person about his preferences in a simple, well-specified risky situation, and then to rely upon the assumed consistency of "economic man" to infer his behaviour in more complex risky situations.

The demonstration is a bit complex but can be broken down into stages: (1) A person's attitude toward risk is elicited in a long series of questions not about preferences among combinations of bread and cheese, but about preferences between a standard gamble and a sure thing. The sure thing is a fixed income attained with certainty. The gamble yields a big prize or a little prize, with a given probability of each. A sort of indifference curve will be derived connecting the probability of winning the big prize and the size of the income as a sure thing. This curve encapsulates a person's attitude toward risk. (2) Along this curve, the probability of the big prize must be an increasing function of income as a sure thing. Since income is a valid indicator of utility in a world of invariant prices and since any increasing function of a utility indicator is a utility indicator too, the probabilities along this curve can be interpreted as utilities of the corresponding incomes. This is the utility of income function we seek if it can be shown that the person to whom the curve refers always seeks to maximize expected utility in accordance with this curve, not just in the simple context where the curve is derived, but in more complex risky situations. That turns out to be so. (3) Once identified through the linking of utility and probability, the utility of income function can be employed to rationalize choice in complex situations. The prospects of law and medicine can be equated to simple gambles that can at once be ranked in accordance with their expected utilities. The higher the expected utility, the more desirable the gamble in the assessment of the person to whom the utility of income function refers.

By these steps, there is established a utility of income function encapsulating a person's degree of risk aversion, while at the same time ranking all risky prospects. Knowing a person's utility of income function, one knows how he will behave in any risky situation such as the choice between law and medicine in table 1.

The questions are framed as follows: A person is confronted with a choice between a sure thing and a simple risky prospect. The sure thing is an annual income of Y for the whole of one's working life. The risky prospect has two possible outcomes: one relatively good, the other relatively bad. The good outcome is an annual income of $250,000 for the whole of one's working life. The bad outcome is an annual income of $20,000 for the whole of one's working life. The probability of the good outcome is B.

The person is asked to choose between Y and B. He may, for instance, be asked to choose between a sure income of $100,000 and a risky prospect where B = ½, that is, with a 50 percent chance of an annual income of $250,000 and a 50 percent chance of an annual income of $20,000. Whatever he chooses, the value of B can then be adjusted, up or down as need be, until the person is indifferent between the risky prospect and the sure thing. If, for example, he is indifferent between a sure income of $100,000 and a 73% chance of the good outcome, we say that B(100,000) = 0.73. Then we change the value of the sure income (for example, from $100,000 to $110,000) and repeat the process over and over again until we know B(Y) for every value of Y from $20,000 to $250,000.

Necessarily, B(250,000) = 1 because, otherwise, one could only lose by choosing the risky prospect over the sure thing when Y is as high as \$250,000. Similarly, B(20,000) = 0 because, otherwise, one could only lose by choosing the sure thing over the risky prospect when Y is as low as \$20,000. The function B(Y) must increase with Y because a risky prospect with a higher probability of the good outcome is preferred to a risky prospect with a lower probability of the good outcome as long as the outcomes themselves remain the same.

The function B(Y) connecting each income Y to a probability of winning the big prize in the standard lottery is precisely the utility of income function we seek, an increasing function of income capturing a person's deportment toward risk. The function B(Y) is defined so that, if a person is indifferent between a sure income and a gamble, the utility of the sure income and the expected utility of the gamble are necessarily the same. It follows that equation (20) is automatically valid for this utility function, that is, when the general utility function u(Y) is replaced by the specific utility function B(Y) as defined here. Each of the three utilities on the right hand side of equation (20) - $u(Y_V)$, $u(Y_M)$ and $u(Y_U)$ - becomes a probability of winning the big prize rather than the small one, and, since the events V, M, and U are mutually exclusive, the entire right hand side of equation (20) is the probability of winning the big prize that is implicit in the gamble with equal probabilities of the three events. To say that a rational self-interested person seeks to maximize the expected probability of the big prize in any choice among alternative gambles, is to say that he maximizes expected utility as defined in equation (20) with B(Y) as the utility of income function.

Figure 6: The Utility of Income Function when a Person is Risk Neutral And when a Person is Risk Averse `

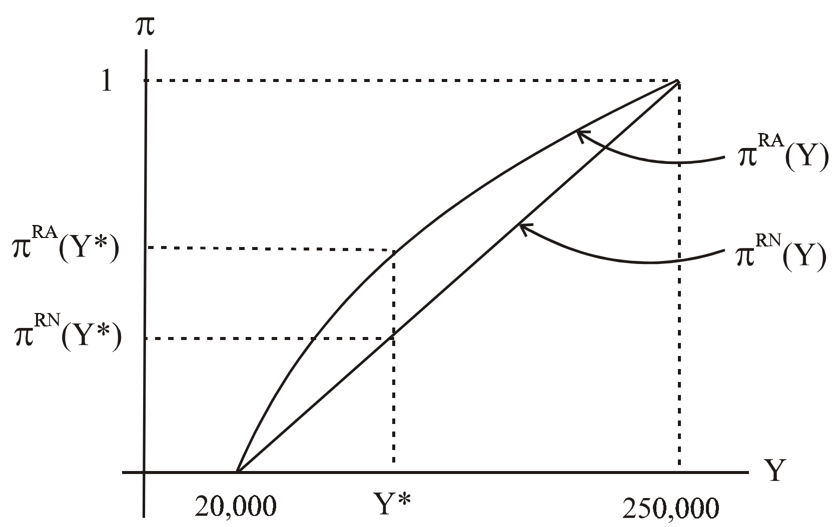

Shapes of alternative utility of income functions B(Y) are illustrated in figure 6 with Y on the horizontal axis and B on the vertical axis. Two utility of income functions are shown, one for a risk neutral person and the other for a risk averse person. A risk neutral person is indifferent between “a sure income of y” and “a probability p of winning the big prize of $250,000 where the alternative is winning a small prize of $20,000" if and only if

$$\begin{aligned} Y &= p(250,000) + (1-p)(20,000) \\ &= 20,000 + p(230,000) \end{aligned} \tag{21}$$

or, replacing p with the function $B^{RN}(Y)$ where the superscript RN stands for risk neutral,

$$B^{RN}(Y) = Y/230,000 - 20,000/230,000 \tag{22}$$

as shown in the upward-sloping straight line in figure 6. On the other hand, if a person is risk averse, it must be the case that

$$Y < 20,000 + p(230,000) \tag{23}$$

or, replacing p with $B^{RA}(Y)$ where the superscript RA stands for risk averse,

$$B^{RA}(Y) > Y/230,000 - 20,000/230,000 \tag{24}$$

A utility of income function of a risk averse person - a utility function consistent with equations (23) and (24) - is illustrated by the curved line in figure 6. For any point Y* that is greater than

\$20,000 but less than \$250,000, $B^{RA}(Y^*)$ must be greater than $B^{RN}(Y^*)$. Thus, since $B^{RA}(Y)$ and $B^{RN}(Y)$ have to be the same at the upper and lower limits, the function $B^{RA}(Y)$ must be concave as shown in the figure.

As constructed, the function B(Y) is *a* satisfactory utility of income function but not the *only* satisfactory function. Obviously, our choice of limits, \$250,000 and \$20,000, was entirely arbitrary. Any other limits would have done as well. It can be shown that any increasing linear function of B(Y) would be equally acceptable as a utility of income function reflecting a person's behaviour toward risk.

Since the utility of income function, B(Y), is discovered in a conceptual experiment, one cannot say *a priori* that it will conform to any arbitrarily chosen functional form. Nevertheless, just as indifference curves in a bread and cheese economy were illustrated by the function u = bc, so too can taste for risk be illustrated by supposing that a person's choice under uncertainty just happens to conform to the function

$$B(Y) = AY^{"} + B \tag{25}$$

where " is an indicator of a person's risk aversion and the parameters A and B are chosen for each value of " so that B(20,000) = 0 and B(250,000) = 1. This function has the right limits, the right general shape for risk aversion, and is easy to manipulate. If " = 1, the person to whom the function B(Y) belongs is risk neutral because the expected value of B(Y) would automatically be maximized whenever the expected value of Y is maximized. Otherwise, if " < 1, a person is risk averse, the more so the smaller the value of ". Note, however, that there is nothing in the conceptual experiment requiring a person's risk aversion to be representable by an exponential function. A person's aversion to risk might be representable by that function, but need not be. The exponential function is just useful for exposition.

Table 2: Utility as Probability

[The utility of income is the minimal probability of a good outcome required to induce a person to accept a risky prospect in preference to a sure thing]

Income, Y ($ thousand)	Utility of Income: Risk Neutral, " = 1	Utility of Income: Moderate Risk Aversion, " = 3/4	Utility of Income: Strong Risk Aversion, " = 1/4
20	0	0	0
30	.0435	.0629	.1212
50	.1304	.1744	.2924
90	.3034	.3610	.5185
100	.3478	.4150	.5627
110	.3913	.4588	.6037
150	.5652	.6254	.7439
190	.7391	.7810	.8584
250	1	1	1

Note: Utility and income are related by the general formula, $B(Y) = AY^{"} + B$ where A and B are selected for each value of " to ensure that $B(250{,}000) = 1$ and $B(20{,}000) = 0$. One can easily check that $A = 1/ [(250{,}000)^{"} - (20{,}000)^{"}]$ and $B = (20{,}000)^{"}/[(250{,}000)^{"} - (20{,}000)^{"}]$. With risk neutrality, the values of A and B must be $.434783\text{x}10^{-5}$ and -.0869565. With moderate risk aversion such that " = 3/4, the values of A and B must be $.1052793\text{x}10^{-3}$ and -.1770579. With strong risk aversion such that " = 1/4, the values of A and B must be .09552368 and -1.1359734.

For Y varying from a low of $20,000 to a high of $250,000, the B(Y) is shown in table 2 for three possible values of " signifying risk neutrality (" = 1), moderate risk aversion (" = 3/4) and strong risk aversion (" = 1/4). The first column lists alternative incomes. The second, third and fourth columns show the corresponding minimal probabilities of success in the gamble required to induce a person with the indicated degree of risk aversion to accept the gamble rather than the sure thing.

A person's choice between law and medicine can now be modeled as the maximization of expected utility where the utility of income function is deemed to represent that person's behaviour toward risk in a simple experiment. Each occupation yields equal probabilities of incomes Y_V, Y_M, and Y_U as shown in table 1 above. Expected utility of these prospect becomes

$$\text{Exp}(B) = (1/3)B(Y_V) + (1/3)B(Y_M) + (1/3)B(Y_U) \tag{26}$$

and one chooses between law and medicine according to the values of Exp(B) in equation (26). Consistency requires that one choose the prospect with the larger value of Exp(B). One chooses the prospect with the larger value of Exp(B) because Exp(B) is is equivalent in every respect to a probability of the good outcome in the simple conceptual experiment as described above. As explained in the discussion of equation (20), the crux of the matter is that a risky prospect with equal chances of incomes Y_V, Y_M and Y_U - or any other risky prospect, no matter how complex - is necessarily equivalent to a simple risky prospect where the good outcome is \$250,000, the bad outcome is \$20,000. The logic of this assertion can be seen with the aid of figure 7. Each of the three possible values of Y is equivalent in a person's assessment to a probability of the good outcome in our simple conceptual experiment. That being so, equal chances of each of the three outcomes must be equivalent in that person's assessment to a compound probability of the good outcome of Exp(B) as shown in in equation (26) above.

Figure 7: The Expected Utility of a Risky Prospect

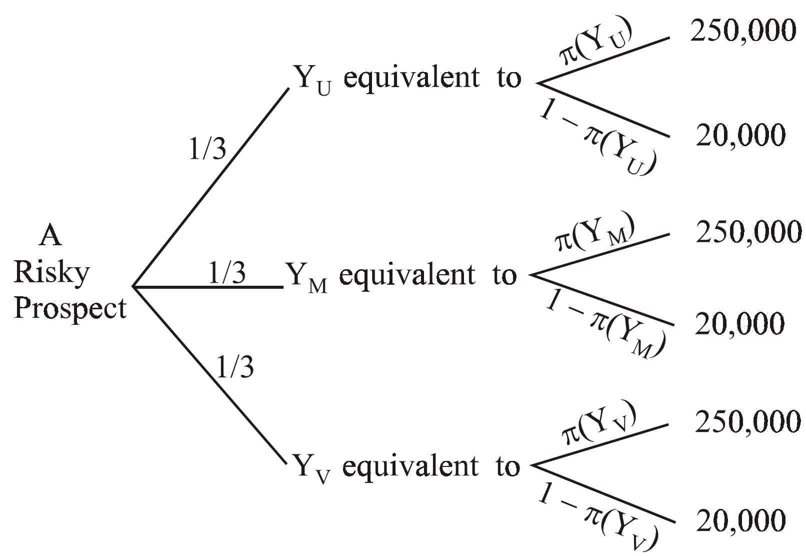

The choice between law and medicine is a choice between risky prospects with different outcomes in each of three equally likely states of the world. A person's choice between these prospects can be predicted from his utility of income function, summarizing his supposedly observed behaviour in the conceptual experiment from which the function was derived. The key to the reduction is consistency. Each risky prospect can be reduced to a simple choice between a standard gamble and a sure thing. Within this framework, there is imposed the assumption that a person's utility of income function conforms to equation (25) for some value of ". For that value of ", each of the three "prizes" in medicine - \$90,000, \$100,000 and \$110,000 - is itself equivalent in the person's assessment to a lottery with a certain probability of the big prize of $250,000, and where to lose is to acquire the small prize of $20,000. For instance, as may be read off table 2, an income of $90,000 is equivalent in the assessment of a person for whom " = 3/4 to a lottery with a probability of 0.361 of winning the big prize. Thus, equal chances of incomes of $90,000, $100,000 and $110,000 in the medical profession are like equal chances of entering one of three lotteries with the same prizes - $250,000 and $20,000 - but different probabilities of attaining the big prize, probabilities dependent, as shown in table 2, on one's degree of risk aversion. That in turn is equivalent in every respect to entering one grand lottery with the same prizes and a probability of winning, B(medicine), which is the average of the probabilities attached to the grand lottery with the same prizes and a probability of winning, B(medicine), which is the average of the probabilities attached to the three possible incomes.

The choice boils down to a comparison between B(medicine) and B(law), each defined as expected utility in accordance with equation (25). For three values of " - 1, 3/4 and 1/4 - the expected utility of each profession can be determined from the information in table 2. If one is risk neutral so that " = 1, the appropriate values of B(Y) can be read off the second column of the table.

$$B(\text{medicine}) = (1/3)(0.3034) + (1/3)(0.3478) + (1/3)(0.3913) = 0.3475 \quad (27)$$

and

$$B(\text{law}) = (1/3)(0.0435) + (1/3)(0.3913) + (1/3)(0.7391) = 0.3913 \quad (28)$$

With risk neutrality, B(law) > B(medicine), and one chooses the profession of law. If one is moderately risk averse (" = 3/4), the appropriate values of B(Y) can be read off the third column.

$$B(\text{medicine}) = (1/3)(0.3610) + (1/3)(0.4150) + (1/3)(0.4588) = 0.4116 \quad (29)$$

and

$$B(\text{law}) = (1/3)(0.0629) + (1/3)(0.4588) + (1/3)(0.7810) = 0.4342 \quad (30)$$

and law is still the preferred profession. But if one is strongly ly risk averse (" = 1/4), the appropriate values of B(Y) can be read off the last column.

$$B(\text{medicine}) = (1/3)(0.5185) + (1/3)(0.5627) + (1/3)(0.6037) = 0.5614 \quad (31)$$

and

$$B(\text{law}) = (1/3)(0.1212) + (1/3)(0.6037) + (1/3)(0.8584) = 0.5278 \quad (32)$$

and one chooses medicine instead. The choice between the professions depends on one's degree of risk aversion as represented by the curvature of the utility of income function. By asking himself about his preferences between risky and riskless options, the reader may discover his utility of income function and may discover whether he prefers law or medicine in the circumstances of table 1.

Almost everybody is risk averse, though some people are more risk averse than others. Consider a family of four people assessing a risky prospect with 50% chance of $190,000 per year and a 50% chance of $30,000 per year. Though its expected income is $110,000 per year, it would almost always prefer a sure income of $110,000 to the risky prospect. One can live comfortably with an income $110,000. A 50% chance of earning an extra $80,000 would not normally be worth the risk of poverty at an income of $30,000. Suppose the family's degree of risk aversion could be represented by a value of " of 1/4 in equation (25). If so, the certainty equivalent of the risky prospect is $83,907, over $26,000 less than the expected income. [This may be deduced from equation (25) and the data in table 2. The family's expected utility is 0.4898, that is, ½(0.1212 + 0.8584). As explained above, the parameters A and B of the utility function in equation (25) are determinate for any given value of ", and, as indicated in the note to table 2, A = 0.09552368 and B = -1.1359744 when " = 1/4. Plugging these numbers into equation (24), we see that the certainty equivalent of the risky prospect with an expected utility of 0.4898 is $83,907.] Confronted by a risky prospect with a 50% chance of an income of $30,000 and a 50% chance of an income of $190,000, a family with a value of " of 1/4 would be prepared to pay about $106,000 in the fortunate state in return for a grant of only $54,000 in the unfortunate state, these amounts being just sufficient to convert the risky prospect into its certainty equivalent. If that family could take out fair insurance - costing one dollar in the good state for one dollar in the bad state - it would certainly do so. Conversion of the risky prospect with an expected value of $110,000 into a sure income of $100,000 would be clearly advantageous. If necessary, the family would be willing to pay a "risk premium" of up to $26,000 for the right to take out fair insurance. Insurance companies convert risk into certainty, or reduce the intensity of risk, by pooling where the misfortune of one family is set against the good fortune of another. Insurance companies cannot offer "fair" insurance because they must cover a cost of administration. Insurance remains advantageous as long as the cost of administration is less than the risk premium.

People cover risks with insurance when they can. We insure houses and personal property against fire and theft. We insure against illness. We insure against unemployment. Pensions and annuities insure us against poverty through survival beyond what our accumulated assets can finance. Programs of welfare for the very poor are insurance against destitution. Sometimes insurance is private. Sometimes it is undertaken through the intermediary of the government.

Recognition that insurance is sometimes private and sometimes public raises the general question of whether insurance is a commodity like bread and cheese that is supplied efficiently by the market, or a commodity like guns for the army that must be provided by the government if they are to be provided at all. Fire insurance is invariably private. Insurance against unemployment, destitution, and natural disasters is invariably supplied through the intermediary of the government. Insurance against illness and against poverty on surviving to a very old age are partly public and partly private, the mix varying at different times and places. Our concern here is not so much with what we do, but with our reasons for doing so. To what extent do the virtues of the competitive economy extend to the provision of insurance? Is there something about insurance that calls for collective action?

There are three major impediments to insurance in a competitive economy: (1) We learn too soon about impending misfortune; (2) Insurance destroys the incentive to avert misfortune; (3) The insured may have private information that is concealed from the insurer.

It is normally supposed that information is desirable. The more we learn, the better off we become. That is true of most kinds of information, but there is at least one important exception. As an example, consider the date of your death. Suppose an angel offered to announce the date of your death not just privately to you, but to everybody who might want to know. The information would be advantageous by helping you to plan the rest of your life. The information may be disadvantageous by condemning you to poverty in the event that you will live to a ripe old age. Whether public or private, pensions have an aspect of insurance. One's annual income from a pension is set in accordance with the average life expectancy of the participants in the pension fund, so that people who live long lives automatically gain at the expense of people who do not. Participants are content with that arrangement as long as nobody knows in advance who the long-lived people will turn out to be. But if the angel announced the date of your death, you would not participate in the pension fund. If you are to have a short life, you would not wish to participate. If you are to have a long life, you would not be allowed to do so.

Similarly, if every doctor's prospects are as indicated in table 1, if nobody knows in advance who the very successful doctors and who the unsuccessful doctors will be and if the difficulties discussed in the next two paragraphs can be ignored, all doctors would wish to pool their incomes, transforming the risky prospect into a sure income of $100,000 per year. But if it is known in advance who will fall into each category, then pooling of incomes becomes impossible. All medical students agree to pool their incomes at any time prior to the moment when their incomes as doctors are revealed. Afterwards, once incomes are revealed, it is too late for pooling because the wealthy doctors will no longer participate. To destroy insurance, knowledge need not be complete. It may be sufficient to know that some people have a better than average chance of being successful.

The problem is compounded by *adverse selection* when information is private. Suppose there were two types of medical students, those with a 60% chance of being very successful and those with a 20% chance of being very successful, and suppose everybody knew who belonged to each type. Incomes could still be pooled within each type separately. People with a 20% chance of success would want to enter a pool with people with a 60% chance of success, but they would not be allowed to do so. Pooling could break down completely if each person knew his type and if the information were private, so that potentially successful people could not identify themselves as a basis for excluding people of the less successful type.

Another impediment to pooling is that *moral hazard* may destroy the incentive to work and save. Pooling may be destroyed even though there is no telling at the time of the insurance contract who the prosperous or fortunate people will be. Consider the doctors once again, but with a slight change in the interpretation of table 1. Suppose the numbers in table 1 refer to the incomes of doctors who work hard. A doctor who does not work hard is guaranteed an income of $90,000 (the income of an unsuccessful doctor) regardless. A doctor who works hard has equal chances of each of the three incomes in the table. Doctors work hard because they consider the chance of higher incomes to be worth the sacrifice. Under these circumstances, it is likely that pooling of incomes would destroy the incentive to work hard unless doctors could compel one another to do so. Without pooling, hard work is worth a one-third chance of an extra income of $10,000 and a one-third chance of an extra income of $20,000. With pooling and with many doctors in the pool, hard work gets you nothing because the gains from your hard work are shared equally among all the doctors in the pool. Analogously to the smoke and smoking example, nobody works hard and everybody is worse off in expectation than he would be without pooling. Pooling tends to emerge in sectors of the economy where disincentives are small. Pooling may be thwarted when disincentives are large.

The moral of the story is that people may favour public insurance in circumstances where the market supplies no private insurance or where they themselves would not take out private insurance if it were available. Public insurance cannot circumvent disincentives. Public insurance can, in effect, be contracted "earlier" than private insurance. Public provision of health care, of assistance for the unemployed and of pensions for the old and destitute can be thought of as a contract established by our grandparents on our behalf, signed in our names well before any of us knows what his circumstances will be and long before we could arrange a contract for ourselves. Insurance is not the only motive for public provision of services. Other motives will be discussed in the chapters ahead.

<u>Personal Goods</u>: [u = u(b, c, L, s) where b is bread, c is cheese, L is hours of leisure and s is survival probability]

Leisure and life expectancy are like ordinary private goods in that they belong to one person exclusively, but they differ from other private goods in that they cannot be purchased at a uniform price that is the same for everybody. My leisure and my life expectancy are mine alone, but I cannot trade them as I can trade bread or cheese. Consider leisure first. A person's utility function may be written as U = u(b, c, L) signifying that his well-being depends upon his consumption of bread, b, per day, his consumption of cheese, c, per day, and his hours of leisure, L per day. Endowed with 24 hours in the day and confronted with prices $P^{\$b}$ and $P^{\$c}$ of bread and cheese and with a wage w, he chooses b, c and L to maximize U subject to his budget constraint

$$bP^{\$b} + cP^{\$c} = w(24 - L). \qquad (33)$$

The value at market prices of his consumption of bread and cheese equals the value at his wage of his supply of hours of work, 24 - L. What differentiates leisure from bread and cheese is that the competition churns up market-wide prices, $P^{\$b}$ and $P^{\$c}$, of bread and cheese but not for leisure. The wage would serve as a market-wide price of leisure if everybody were equally skilled, but that is not normally the case. Highly skilled people have high prices of leisure. Less-skilled people have low prices of leisure. You give up more bread per extra hour of leisure when your wage is high than when when your wage is low. That is the sense in which leisure is a personal good while bread and cheese are not. Each person's wage becomes his own personal price of leisure in terms of goods.

Life-expectancy is like bread and cheese in two important respects. We strive to have more of it, and we are prepared to trade life expectancy for other goods or for money as, for example, when we desist from buying a safer car because it is too expensive. Life expectancy could be incorporated into the utility function, but only if the atemporal (without reference to the passage of time) function that we have been employing were generalized to account for goods consumed in different years with due allowance for the probability of being alive in each future year. We can avoid these complications and still account for survival probability by postulating an atemporal utility function U = u(b, c, L, s) where, once again, b, c, and L are bread, cheese, and leisure, where s is one's probability of survival over the course of the current year, and where U is an increasing function of all four terms. What makes survival probability into a personal good like leisure rather than an ordinary private good like bread or cheese is the form of the budget constraint.

$$bP^{\$b} + cP^{\$c} + V(s)s = w(24 - L) \tag{34}$$

where the price of survival, V(s), unlike prices of bread and cheese, is an increasing function of s. To see why V(s) is an increasing function, note first that

$$s + m = 1 \tag{35}$$

where m is one's mortality rate. Survival would be an ordinary private good like bread and cheese if everybody could buy reductions in mortality rates at a uniform market-determined price. Were that so, the function we are calling V(s) would reduce to a constant like $P^{\$b}$ and $P^{\$c}$.

The function V(s) is increasing with s rather than constant because we buy reductions in mortality rates indirectly. Why this is so is best explained by an example. Consider a person whose base mortality rate - if he takes no steps to reduce it - is 2%; there would be a 1-in-50 chance of his dying over the course of the year. However, as shown in table 3, he can reduce his mortality rate in various ways.

The table must be taken with a grain of salt. First, all the numbers are made up, with only the slightest connection to actual market costs. Second, and more importantly, many safety expenditures are joint products. For instance, one may trade in one's jalopy for a Ford, not just because the Ford is safer, but because it more pleasant to drive. One visits a doctor not just to prolong one's life, but to make one's life more comfortable and disease free. Think of the numbers in the table as costs net of other benefits from the purchase in question. Third, not all expenditure to reduce mortality rates is in the form of money. For instance, the decision whether or not to wear a seat belt on any particular trip is a trade-off between time and mortality in the first instance, though the evaluation of time can convert trade-off into one between mortality and money. Since leisure is a personal good too, the exact terms of the trade-off between mortality and money would vary from one person to the next, but there is some determinate rate of trade-off for each and every person. Typically thought not invariably, the monetary value of time is higher for the rich than for the poor. Fourth, certain expenditures may be appropriate for some people but not for others. For instance, the option of putting a radio in one's yacht is of no significance to somebody who cannot afford or does not want to own a yacht. Fifth, one's base mortality rate is not constant as we have supposed but depends on how one conducts one's life. For instance, with reference to the last row of the table, one's mortality rate may be higher if one flies one's own plane than if one does not, but that mortality rate is reduced nonetheless by having a co-pilot.

Table 3: The Cost of Survival and the Value of Life
[The mortality rate with no expenditure on safety is 2%]

Ways to reduce one's mortality rate	(1) Cost ($)	(2) Percentage reduction in mortality rate (%)	(3) Cost per life saved ($) 100(1)/(2)
Visit doctor once a year	1,000	.4	250,000
Trade in one's jalopy for a new Ford	1,000	.2	500,000
Install a fire alarm in one's house	200	.02	1,000,000
Test for cancer	1,000	.05	2,000,000
Trade in one's Ford for a Cadillac	10,000	.2	5,000,000
Put a radio in one's yacht	800	.004	20,000,000
Hire a co-pilot for one's private plane	30,000	.06	50,000,000

What remains valid and instructive in the table is that everybody is presented with an array of mortality-reducing options, that each option has a well-specified market price, and that there is a limit on how much mortality-reduction one can buy with each option. For instance, one can buy a 0.4 % reduction in mortality by visiting the doctor once a year, but one cannot buy an extra 0.4 % reduction in one's mortality rate by visiting the doctor a second time. Indeed, if one could acquire the same reduction in mortality rate with each trip to the doctor, there would be no need for any of the other, more expensive means of mortality reduction shown in the other rows of the table. One would go to the doctor over and over again until one's mortality rate is reduced to 0 and one lives forever. It is this unhappy feature of mortality reduction that makes it a personal good rather than an ordinary good like bread and cheese.

The important feature of the table is the third column showing the cost per life saved of each type of mortality-reducing expenditure as shown in column (3). The cost per life saved is the amount of money a group of people would need to spend to save one life among them. For instance, if one person's trip to the doctor reduced his mortality rate by 0.4 %, then the expectation is that one life would be saved per 250 people [100/(0.4) = 250] who visit the doctor. Their combined cost would be $250,000 [$1000 x 250]. The different ways of reducing one's mortality rate are ordered in the table from the lowest to the highest cost per life saved. One may think of the cost of life column (3) as a numerical correspondence to V(s) in equation (34). Specifically, the cost of per life saved is the increase in V required for a given increase in s.

Taking the table at face value, supposing that the mortality-reducing expenditures in the table are the only such expenditures available, and ignoring the qualifications in the preceding paragraphs, we may think of each person as choosing his preferred place in the table. For each person, there is a cut-off expenditure such that he buys all cheaper reductions in his mortality rate and desists from buying all more expensive reductions. For instance, one may install a fire alarm but not test for cancer, in which case one would also visit a doctor once a year and buy a Ford rather than a jalopy, but not buy a Cadillac, put a radio in one's yacht or hire a co-pilot.

Generalizing the example, we may suppose each person to be confronted with a schedule of ways of buying reductions in mortality rates, and each person may be said to have a *value of life* defined as the most money he is prepared to pay *per life saved* saved to decrease his overall mortality rate. Typically but not invariably, the value of life of the rich would exceed the value of life of the poor, meaning not that the life of a rich person is worth more than the life of a poor person in some absolute sense or that the government, in choosing among public projects, should pay more to save the life of a rich person than it is prepared to pay to save the life of a poor person, but simply that the rich are prepared to spend more than the poor are prepared to spend to reduce their own mortality rates. We shall return to this matter in the chapter on public administration.

Quality, Information and Advertising

It has been implicitly assumed so far, and will be assumed again in subsequent chapters, that people have full knowledge of their options, of the nature of the commodities they buy or might buy, and of their own preferences. These assumptions are all partly true, but not completely so. When there is only one kind of bread, one kind of cheese, and no other commodities, as was assumed in previous chapters, it is reasonable to suppose that choice is in conformity with a well-specified utility function. With thousands of different goods, each with an unlimited variety of qualities and textures, that is no longer so. I buy pills at a drug store not because I know by my own experience and expertise that they are on balance beneficial, but because I have been advised to do so by a doctor who, in turn, relies on the local medical journal reporting on research by the Food and Drug Administration whose scientists have come to believe that the good effects of the pills probably, for there may be no absolute certainty, outweigh the possible harm. I choose a brand of soap not because I have tested all brands, but by habit or because I was influenced by advertising on television. Neither the advice nor the advertising could be influential if the economist's model of taste were an accurate description of the world. There would be no advertising because no producer would have a financial incentive to advertise. A vast industry occupies the territory which the model postulates away.

The virtues of the competitive economy were demonstrated in chapter 3 for a model with perfect and complete knowledge of all products. Recognition of this limitation of the model raises the question of whether and to what extent these virtues carry over from the world in the model to the world at large. Does advertising promote efficiency in the economy, or could people become better off if advertising were curtailed? Does advertising foster the production of useful goods and the allocation of goods to people who need them most, or is advertising like piracy in the fishermen and pirates model, benefiting advertisers at the expense of consumers who pay the cost of advertising in the price of the goods they buy. Advertising is a complex phenomenon. The most that can be undertaken here is to list some of its impacts on society, without attempting to draw a balance between benefits and costs.

Advertising provides knowledge: A new product comes onto the market. There may be many people who would buy the product - and become better off by doing so - if they knew of its existence. Advertisements in newspapers and on television inform such people that the product has arrived and describe the product in enough detail that potential users are inclined to try it.

Advertising certifies quality: Consider a new brand of tomato juice. I may like it or I may not. Ordinarily, I might not be inclined to try it. But if the new tomato juice is heavily advertised, I may think to myself that it is probably quite good because the producer would not otherwise go to the expense of advertising. Let x be the percentage of those who try the new tomato juice who will like it and continue to use it. Presumably, the better the tomato juice, the higher x will turn out to be. Initially, the producer can be expected to know enough about the quality of the new tomato juice to make a shrewd guess about the magnitude of x, but potential consumers are entirely uninformed. Advertising may be the producer's only way of conveying to potential consumers the advertiser's true belief that x is rather high, for if x were low so that most people who tried the new tomato disliked it and resolved never to buy it again, the producer's expenditure on advertising would be wasted.

Both of these explanations would seem to be at least partly correct, and, in so far as either is correct, advertising promotes the welfare of the consumer. Though there is no role for advertising in the world of perfect knowledge as described in chapter 3, a role emerges when knowledge is imperfect and incomplete. Advertising transmits knowledge credibly from producer to potential consumer, causing actual markets to be closer than otherwise to the perfect markets we have postulated. In that role, advertising is fishing rather than piracy, productive rather than predatory. Yet it is hard to believe that the transmission of information is the whole of the matter. Too much of the advertising we encounter is for well-known products and adds nothing to our information about the products themselves.

Advertising transforms taste, creating a false association of the advertised product with happy times or conveying an aura of pleasure with no basis in the quality of the advertised product. A brand of cigarettes identifies the smoker with a cowboy riding a horse on the open range. Kids are encouraged to eat seriously twisted fries. We enjoy the antics of the energizer bunny. The province of Ontario associates gambling with wholesome small town life around the country store. No watcher of television can suppose that the hour or so of commercials to which he is subjected each day is anything but a play on his emotions, an attempt to influence his behaviour and a waste of his time. There is surely more here than the mere provision of information or assurance of quality.

Advertising is a waste of resources in inter-brand rivalry. Producers of different brands of essentially identical soap, cola, cigarettes, or financial services each tout the virtues of its brand and the superiority of its brand over all the rest. My advertising persuades consumers to buy my brand rather than yours. Your advertising persuades consumers to buy your brand rather than mine. Total sales of whatever it is we produce may be very little affected by the commotion, but resources that could be devoted to making things or to providing useful services are devoted instead to persuasion in circumstances where the net effect of each advertiser may be to neutralize the impact of its rivals. Advertising in this context is largely piracy, where the pirates take not just from fishermen, but from one another. There may be a prisoners' dilemma among advertisers within an industry. It may be profitable for each firm to advertise as long as other firms are free to advertise or not as they please, but every firms's profit might be higher if no firm advertised at all. Tobacco companies may secretly welcome a government-imposed ban on cigarette advertising, a ban that would be illegal collusion if the cigarette companies arranged it themselves.

Advertising finances public goods: The discussion of public goods earlier in this chapter began by citing the army and television as examples, but then focussed upon the army alone in the explanation of why public goods have to be supplied by the government if they are to be supplied at all. Television and the army are public goods because the entire expenditure on each of these goods conveys actual or potential benefits to a great many people at once, because each person's benefit from public goods flows from the entire expenditure rather than from a part reserved specially for him, and because beneficiaries do not interfere with one another. If you are protected by the army, then so too am I. Your pleasure from watching a television program is not diminished if I watch it too. But the army is provided by the government, while television is not. The army is provided by the government because nobody has an incentive to contribute to the cost of the army without some guarantee that other people will pay their share as well, a guarantee that can only be provided by compulsion. We vote for a government that forces each of us to pay for the army though our taxes. That is not true of television because there is an alternative source of finance. Television programs are provided free by firms as a vehicle for advertising. The programs are an inducement to watch ads we would never consent to watch if they were provided alone. It is difficult to say whether this is fishing or piracy: television programs offered in exchange for an opportunity to condition our tastes for the goods that the sponsors produce.

It is easier to list impacts of advertising - provision of information, certification of quality, creating artificial preferences, financing communicative public goods, distorting the quality of the public goods they finance - than to assess its overall impact upon the economy. Though a distinction can be drawn between commercial free speech and political free speech, a right of free expression blends at the edges into a right to advertise in that the latter cannot be abrogated altogether without affecting the former significantly. Yet the two are by no means identical. A right to advertise cigarettes may be denied without, at the same time, denying a right to discuss cigarettes, and to persuade people that cigarettes are not harmful to one's health or that the pleasure of smoking is worth the cost. All advertising could be taxed (or advertising could be reclassified as investment rather than as a cost of production) without violating the right of free speech.

Real Income as an Indicator of Utility

Time series of income per person were a substantial part of the evidence in chapter 1 of how dreadful life used to be. Table 9 of chapter 1 showed the Canadian "real" national income per head to have increased from $2,554 in the year 1870 to $11,343 in the year 1950 and then to $34,492 in the year 2000. The term "real" implies comparability of incomes over time. Incomes each year had to be expressed in dollars of a common *base* year, so that, for instance, the income in the year 1950 would not appear small for no other reason than that prices of most goods have increased over time. The choice of a base year was arbitrary. The year 2000 was chosen because it is the most recent year in the series. People alive today can best relate to that year and can be expected to have a sense of what it means for a person to have an income of $2,554, $11,343 or $34,492 in the year 2000. These numbers signify that, over entire the 130-year time span between 1870 and 2000, the rate of economic growth - assessed as the solution, r, to the equation $2{,}554e^{r130} = 34{,}492$ - has been almost exactly 2%.

Most readers of this book will be familiar with statistics of real income and economic growth. Newspapers regularly quote such statistics as evidence of how much better off the nation is today than it was in some former year or of how much better off one country is than another. In the presenting the numbers, it was simply assumed that the reader would have a general sense of what they mean. As already mentioned in chapter 1, real income each year would be the quantity of bread consumed if people consumed only bread. It is when many goods are consumed and when their prices change at different rates from year to year that the concept of real income becomes problematic. There was no attempt in chapter 1 to define real income precisely or to explain how the numbers were produced. We consider this matter now.

Statistics of real income are based on an analogy between a country and a person. If your annual income rises from $11,343 to $34,492, you have no difficulty in saying that you have become just over three times (3.05) as prosperous as you were before. The core meaning of the statement that Canadian real income per head grew from $11,343 in the year 1950 to $34,492 in the year 2000, is that people in the year 1950 were on average as prosperous as you would be today with an income of $11,343 and that people in the year 2000 were on average as prosperous as you would be today with an income of $34,492. Without that analogy, the statistics in table 9 of chapter 1 would be devoid of implications about our lives. Once the analogy is recognized, real income each year assessed with reference to prices in the year 2000 becomes the amount of money one would require in the year 2000 to be as well off as the typical person in that year.

To pin down the meaning of "as well off as," imagine a comparison between yourself today (presumed to be the year 2000) and your grandfather at about the same age in the year 1950, where both you and grandfather are representatives of the people of your times, where you are presumed to be enough of a chip off the old block that both of your preferences can be represented by one and the same set of indifference curves. What has changed between 1950 and 2000 is the technology of production.

To sharpen the analysis, consider an economy with only two goods, bread and cheese, where everybody is identical to everybody else so that an unambiguous representative consumer can be identified, where taste (the set of indifference curves) remains invariant over time, but where technology is changing so that more of both goods can be produced in the year 2000 than could be produced in the year 1950. The technology in each of the two years is illustrated in figure 8 as a pair of production possibility curves, the lower curve showing all possible combinations of bread and cheese that could be produced in the year 1950, and the higher curve showing all possible combinations of bread and cheese that could be produced in the year 2000. A common set of indifference curves in the two years is assumed to conform to the utility function $u = bc$ where b and c are loaves of bread per person and pounds of cheese per person. Recall from the discussion surrounding equation (2) of chapter 3 that the demand price of cheese, $p^D(b, c)$, the slope of the indifference curve at any combination of b and c, is just equal to b/c. Two indifference curves are shown, one tangent to the production possibility curve in the year 1950 and the other tangent to the production possibility curve in the year 2000.

Figure 8: Different Technologies and Common Tastes in the Years 1950 and 2000

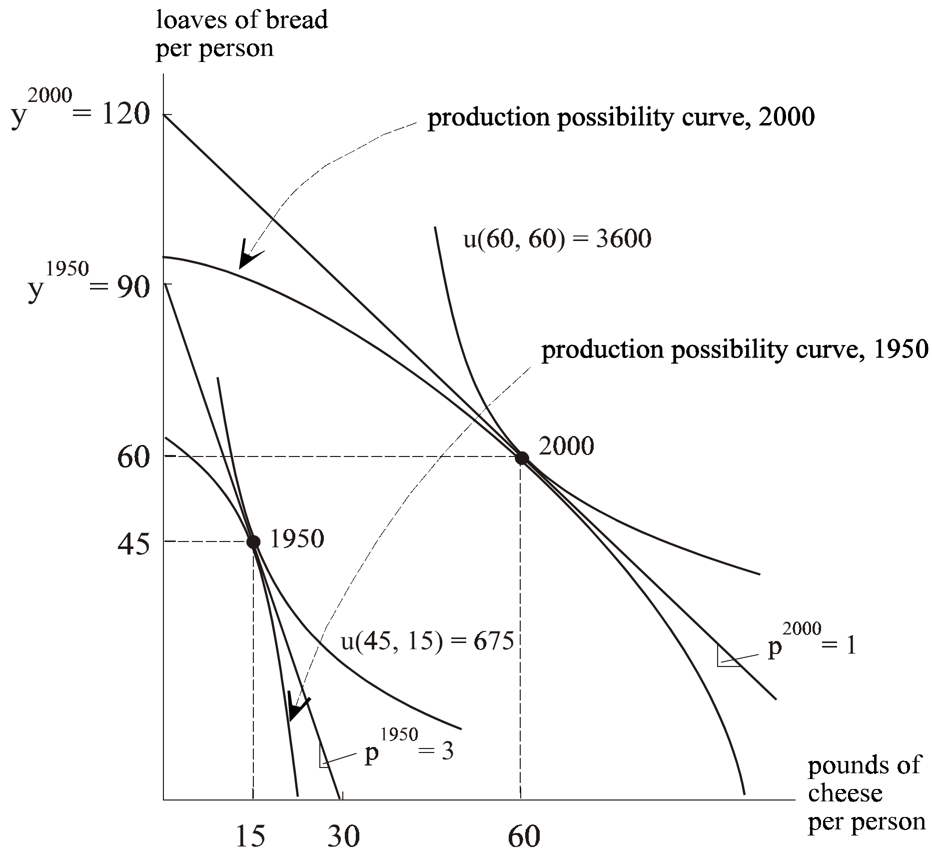

The chosen outputs of bread and cheese in the two years are labeled "1950" and "2000". As in chapter 3, these are characterized by the tangency between an indifference curve and the production possibility curve, signifying that people make themselves as well off as possible with the available technology. The common slopes of indifference curves and production possibility curves at these points are the relative prices of cheese in terms of bread in the two years. Had the production possibility curve in the year 2000 been a scaled-up version of the production possibility curve in the year 1950, the relative price of cheese would have been the same in both years, for the postulated indifference curves are all scaled up or scaled down versions of one another. Prices are not the same because the shapes of the production possibility curves are different. As the curves are drawn, technical change between the year 1950 and the year 2000 was biased toward cheese. The percentage increase in the output of cheese as it would be if only cheese were produced exceeds the percentage increase in the output of bread as it would be if only bread were produced. Consequently, the increase in cheese production exceeded the increase in bread production, and the relative price of cheese declined. As shown in the figure, production of cheese increased from 15 pounds to 60 pounds per person, production of bread increased from 45 loaves to 60 loaves per head, and the relative price of cheese fell from 3 loaves of bread per pound to 1 loaf of bread per pound. This information is also listed in the first three rows of table 4.

Table 4: Prices Quantities and Incomes
[The computation of money income and real income from prices and quantities as will be explained below.]

	1950	2000	growth rate (%)
Quantity of Bread (loaves per head)	45	60	0.6
Quantity of Cheese (pounds per head)	15	60	2.8
Price of Cheese (loaves per pound)	3	1	-
Price of Bread ($ per loaf)	40¢	$4.00	4.6
Price of Cheese ($ per pound)	$1.20	$4.00	2.4
Money Income ($ per head)	$36.00	$480.00	5.0
Quantities Revalued at Prices in the Year 2000	$240.00	$480.00	1.4
Real Income at Prices in the Year 2000 ($ per head)	$207.84	$480.00	1.7

In every year t, the income of the representative consumer is the amount of money required to buy the bread and cheese consumed per person in that year at the going market prices. Money income per person in the year t is

$$Y^{\$t} = P^{\$Bt}b^t + P^{\$Ct}c^t \tag{36}$$

where b^t and c^t are amounts of bread and cheese consumed per person, and where $P^{\$Bt}$ and $P^{\$Ct}$ are money prices of bread and cheese in the year t. For the years 1950 and 2000, money prices and money incomes are shown in table 4.

Money income cannot be represented on a bread and cheese diagram such as figure 8. What can be represented is income with bread rather than money as the numeraire, the amount of bread that would be acquired if one devoted one's entire money income to the purchase of bread. With an income of $Y^{\$t}$ dollars, one could purchase $Y^{\$t}/P^{\$Bt}$ loaves of bread. Dividing both sides of equation (36) by the price of bread, we obtain the income of the representative consumer with bread rather than money as the numeraire.

$$y^t = b^t + p^t c^t \tag{37}$$

where y^t is defined equal to $Y^{\$t}/P^{\$Bt}$ and p^t is defined equal to $P^{\$Ct}/P^{\$Bt}$, the price of cheese expressed as loaves (rather than dollars) per pound. For the years 1950 and 2000, income with bread as the numeraire is shown as a distance on the vertical axis of figure 8. It is the sum of actual consumption of bread and the amount of extra bread one could acquire by exchanging cheese for bread at the current relative price of cheese. It is the height of the projection onto the vertical axis of the point representing consumption of bread and cheese by a line with slope equal to the current relative price of cheese. Income in terms of bread is 90 loaves [45 + (15 x 3)] in the year 1950, and is 120 loaves [60 + (60 x 1)] in the year 2000.

With this machinery in place, we return to the original problem of measuring real income each year for the simple economy we have constructed. Set the year 2000 as the base year or standard of comparison, so that real income in the year 2000 and money income in the year 2000 are one and the same. Define real income per head in any other year t as the amount of money one would need in the year 2000 - and confronted with prices as they were in the year 2000 - to be as well off as one was in the year t with the average income and confronted with prices as they were in the year t.

Since real income is a measure of utility, one might suppose that real income could be measured as utility itself when, as assumed, u = bc. With the numbers we have chosen, the utility, u, of the representative consumer increases from 675 [15 x 45] in the year 1950 to 3,600 [60 x 60] in the year 2000. These numbers are clearly unsatisfactory s measures of real income in 1950 and 2000 because the percentage increase in real income between any two years must lie within the range of the percentage increases of the different commodities. Bread consumption increases by a third (from 45 loaves to 60 loaves), cheese consumption increases four-fold (from 15 pounds to 60 pounds), but utility increases by more than five fold.

The reason for the anomaly is that utility is *ordinal*. Recall the construction of Robinson Crusoe's indifference curves in chapter 3. Robinson Crusoe was asked a long series of questions of the general form "Do you prefer this to that?" He was never asked questions of the form, "By how much do you prefer this to that?" He could not be expected to answer such questions because there was no scale against which "by how much" could be determined. That is why utility is said to be *ordinal* (recognizing more or less) rather than *cardinal* (placing numbers on how much more or less). Only the shapes of indifference curves could be identified and their numbering was entirely artificial. A set of indifference curves conforming to the utility function u = bc would conform equally well to the function $u = (bc)^2$ or to the function $u = \sqrt{(bc)}$. In this respect, utility is like temperature before the invention of the thermometer. It is often said that you cannot attach numbers to indifference curves, but that is not quite right. A more accurate statement is that you can attach numbers too easily. Any numbering system will do as long as higher curves get higher numbers. However, just as the height of a column of mercury quantifies temperature, so too may one particular quantification of utility take precedence over the rest because it focuses on some special concern or supplies a unique answer to some precise question. What we are seeking is a thermometer of utility, a choice of one of the many possible cardinalizations to indicate in a natural and humanly meaningful way how much better off people are becoming over time.

One might suppose that real income in any year t could be measured by repricing goods consumed in the year t at prices as they became in the year 2000. For any year t and expressing income in terms of bread rather than money, this measure becomes

$$y(b^t, c^t, p^{2000}) = b^t + p^{2000}c^t \tag{38}$$

The income, $y(b^t, c^t, p^{2000})$, is what would be required to purchase the bread and the cheese actually consumed in the year t at prices as they became in the year 2000. For t = 1950, this measure of income is illustrated as $y(b^{1950}, c^{1950}, p^{2000})$ on the vertical axis of figure 9 which is a reproduction of figure 8 with the production possibility curves removed to avoid cluttering the diagram and with additional information. It is the projection of the point “1950" onto the vertical axis by means of a line with slope equal to 1 (rather than 3) because the relative price of cheese in terms of bread in the year 2000 was equal to 1. As shown in the figure, the value of $y(b^{1950}, c^{1950}, p^{2000})$ is 60 loaves of bread. Converted from loaves of bread into dollars, the value of 1950 quantities at 2000 prices becomes \$240 [60 x 4] because the price of bread in the year 2000 was \$4.00 per loaf.

Figure 9: Income at Current Prices, Income at Prices in the Year 2000, and Real Income

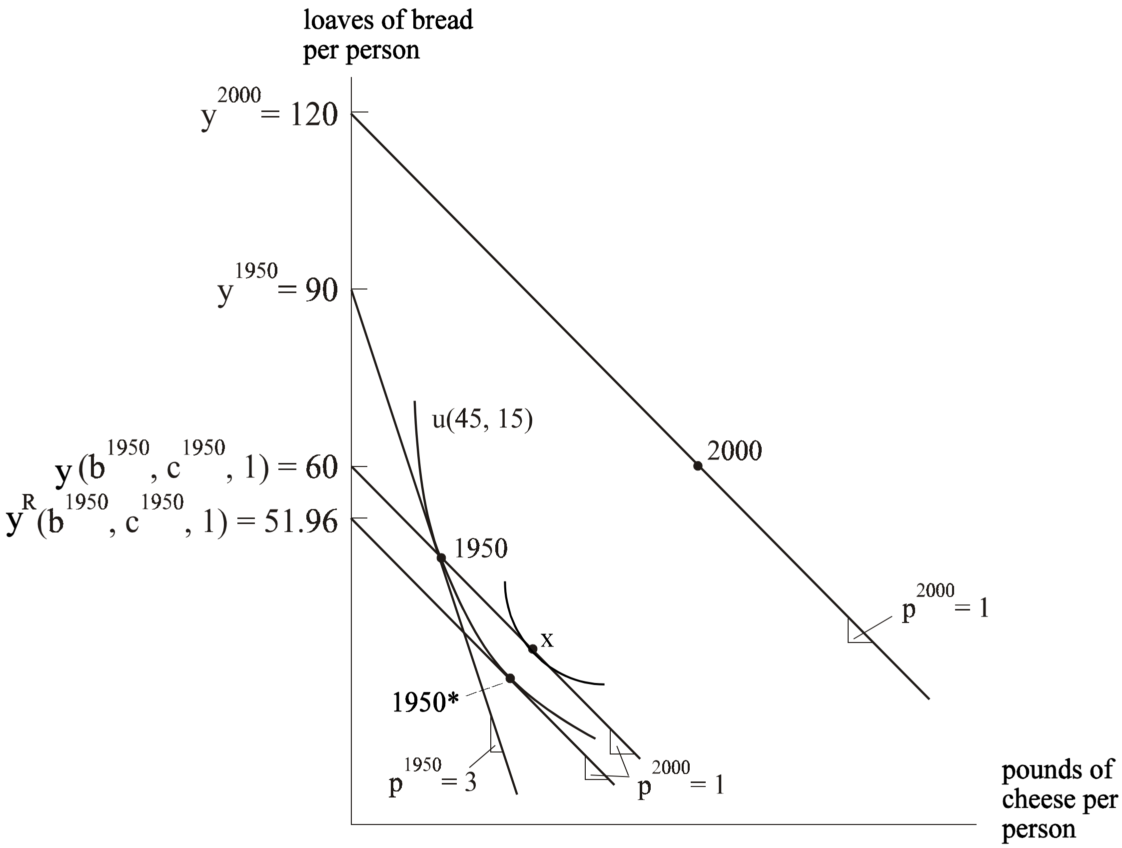

The value at prices in the year 2000 of quantities consumed in the year 1950 turns out to be a good approximation to the measure of real income we are seeking, and it is often employed in practice because the primary data are readily available. But it is not quite right. It is an overestimate of real income in the year 1950 with the year 2000 as the base year because a person provided with enough money at prices in the year 2000 to buy the bundle of goods actually purchased in the year 1950 could make himself somewhat better off than the representative consumer in the year 1950. He would buy a bit more cheese which has become relatively cheap and a bit less bread which has become relatively dear, moving to a point such as x on a somewhat higher indifference curve. With an income sufficient to buy 45 loaves of bread and 15 pounds of cheese when the price of cheese is 1 pound per loaf - an income of 60 loaves of bread - a person whose taste is represented by the utility function u = bc would devote equal amounts of income to each good. He would buy 30 loaves of bread and 30 pounds of cheese, yielding him a utility of 900 [30 x 30] as compared with a utility of only 675 [15 x 45] acquired by the representative consumer in the year 1950 . At prices in the year 2000, an income of 60 loaves is a bit too high.

With the year 2000 as the base year and with bread as the numeraire, the true measure of real income in 1950 is shown in figure 9 as the height on the vertical axis of the point $y^R(b^{1950}, c^{1950}, p^{2000})$. With that income and confronted with the relative price of cheese as it was in the year 2000, a person places himself on the highest attainable indifference curve by purchasing quantities of bread and cheese represented by the point 1950*. Real income $y^R(b^{1950}, c^{1950}, p^{2000})$ is a valid indicator of utility because, by construction, utilities at the points 1950 and 1950* are the same.

$$y^R(b^{1950}, c^{1950}, p^{2000}) = y(b^{1950*}, c^{1950*}, p^{2000}) \tag{39}$$

With our simple utility function, u = bc, we can easily compute the quantities b^{1950*} and c^{1950*} and the real income $y^R(b^{1950}, c^{1950}, p^{2000})$. These may be derived from two equations:

$$b^{1950*}c^{1950*} = b^{1950}c^{1950} \tag{40}$$

indicating that the two combinations of bread and cheese - b^{1950*} and c^{1950*}, and b^{1950} and c^{1950} - lie on the same indifference curve, and

$$b^{1950*}/c^{1950*} = p^{2000} = 1 \tag{41}$$

indicating a tangency between the notional budget constraint and an indifference curve - an equality between the demand price and slope of the budget constraint - when the consumer places himself on the highest attainable indifference curve. Together, equations (40 and (41) imply that real income in the year 1950 must be 51.96 loaves of bread[1]. Expressed in dollars rather than loaves of bread and with a price of bread of $4.00 per loaf in the year 2000, real income in the year 1950 becomes $207.84 [51.96 x 4] and shown in the bottom row of table 4.

For any year t and with the year 2000 as the base year, real income becomes

$$y^R(b^t, c^t, p^{2000}) = y(b^{t*}, c^{t*}, p^{2000}) \tag{42}$$

where b^{t*} and c^{t*} are determined by the procedure we have employed to determine b^{1950*} and c^{1950*}. This is a genuine utility indicator. Real income is the same for all points on the same indifference curve. It increases in passing from a lower to a higher indifference curve. As mentioned above, the choice of the year 2000 as the base year for a time series of real income is entirely arbitrary in that any other base year would have yielded equally valid measures of real income as an indicator of utility, but it is entirely appropriate in that the most recent year is "our" natural standard of comparison, telling us what we want to learn from the data.

---

[1]Together equations (40) and (41) imply that $b^{1950*} = c^{1950*} = /(b^{1950}c^{1950}) = /(15x45) =$ 25.98 so that $y^R(b^{1950}, c^{1950}, p^{2000}) = b^{1950*} + p^{2000}c^{1950*} = b^{1950*} + c^{1950*} = 51.96$ as shown in table 3.

Defined precisely for a person whose indifference curves are presumed to remain invariant, the concept of real income is put to work for comparisons between entire countries where people within each country have different preferences and where preferences differ from one time or one place to another. In constructing statistics of real income, such as the series for Canada in table 9 of chapter 1, the statistician has no option except to proceed as though the time series of prices and quantities from which statistics of real income are to be constructed reflect the preferences of a representative consumer whose circumstances change but whose tastes remain invariant over time. Without that presumption, the weighting of quantities by prices would be meaningless. Furthermore, even in circumstances where everybody's taste is the same (in the sense of having the same set of indifference curves) and even if tastes remained invariant over time, the statistician cannot observe what the representative consumer would buy at some arbitrarily chosen set of prices. Quantities per head of bread, cheese, and other goods consumed are averages over many people whose tastes are never quite the same and whose incomes differ substantially. The rich may consume relatively more bread and the poor may consume relatively more cheese, even though their tastes are the same in the sense that each would adopt the same consumption pattern at any given income. At best, the shapes of indifference curves can be estimated, never observed directly. In practice, the statistician may have to rely on a repricing of quantities as the best available approximation to real income, on a measure of $y(b^t, c^t, p^{2000})$ in equation (38) as the best available approximation to $y^R(b^t, c^t, p^{2000})$ in equation (39).

The most formidable difficulty in construction of a time series of real income per head is in accommodating the virtually infinite range of goods and services consumed. Any measure of real income must account for a greater diversity of goods than the statistician can ever hope to observe. Bread is not a uniform substance as we have so far assumed. It is a collective noun incorporating hundreds of varieties and qualities of rye bread, bagels, muffins, baguettes, sliced white bread, onion bread, pita bread, and so on. Cheese is a collective noun incorporating hundreds of varieties and qualities of cheddar, Swiss, camembert, stilton, feta, cream cheese, cottage cheese, and so on. And, believe it or not, consumption encompasses more than bread and cheese. Qualities as well as quantities are changing all the time. What is the poor statistician to do? His only recourse is to estimate quantities of broad classes of goods by value deflated prices of selected items.

At the statistician's disposal each year are current money values of the purchases of the different classes of goods - such as groceries, clothing, and housing - and prices of a list of goods specified in great detail. Broad categories of expenditure may be broken down into somewhat finer categories such as vegetables, bread and cheese, but without direct measures of quality change over time. Prices on the other hand may be very specific, but only for a selection of goods. For example, the price of a certain quality of cheddar cheese may be tracked over time, but prices of many varieties of cheese may not be tracked at all. Quantities may be inferred by "deflation" of categories of goods or of the national income as a whole. Suppose that, between 1950 and 2000, the dollar value of sales of cheese per head rose by a factor of 324% and that the price of a specific quality of cheddar cheese rose by a factor of 152%. If we knew that prices of all varieties of cheese rise and fall in step, we would infer that the quantity of cheese per head increased by a factor of 213%. Comparable information about quantities could be inferred for each and every category of goods. If all prices rose or fell proportionally over time, an accurate time series of real income could be obtained by deflating money income each year with the price of tooth picks.

But prices do not rise or fall proportionally. The price of tooth picks soared over the last fifty years by comparison with the price of personal computers, which is to say that computing power has become dramatically cheaper. Since prices of different goods change at different rates, statistics of real income are computed by deflating money income with a price index, a weighted average of prices. High weighting for prices of goods becoming relatively more expensive over time yields a relatively low rate of economic growth. High weighting for prices of goods becoming relatively less expensive over time yields a relatively high rate of economic growth.The problem of how to measure real income can be reformulated as a problem of choosing the appropriate price index.

Repricing quantities and deflating money income with a price index are two sides of the same coin, and all of the problems discussed above in the choice of price weights reappear in the choice of the appropriate price index. Conceptually, these procedures are identical. In practice, statistics of real national income, such as the Canadian time series in table 9 of chapter 1, are constructed by deflating money income with a price index because adequate data on income and prices are available but adequate data on quantities are not. The usual procedure for the construction of price indices is to weigh price changes by observed value shares of the different goods, and to change weights about once per decade to ensure that the estimated rate of economic growth each year does not depart too much from current valuations. This has the additional advantage of capturing some of the surplus from the introduction of new types of goods that tend to be expensive when first introduced and to become progressively less expensive over time. A unit of such goods is automatically given more weight at first and progressively less and less later on.

Statistics of real income must also take account of investment, depreciation, public expenditure, exports, and imports. We have so far been discussing real income as though the world were entirely static. To focus on the core meaning of statistics of economic growth, each year was looked upon as though it were entirely self-contained with no influence from the past and no preparation for tomorrow. By contrast, national statistical agencies construct income statistics as snapshots of economic activity each year. In table 9 of chapter 1, the concept of real income was referred to as "Gross Domestic Product at 2000 Prices". *Product* refers to all goods and services produced in the current year by the government as well as by the private sector, inclusive of goods for consumption, such as bread, health care, and cheese, and goods for investment, such as factories, roads, and machines. *Gross* means that there is no deduction for depreciation. A new machine counts as part of gross domestic product even if an identical old machine is taken out of service. The reason for the asymmetry is that real depreciation is difficult to measure accurately. *Domestic* refers to production in Canada regardless of the owners of the factors of production. The study of how data are collected and compiled for the construction of time series of real national income is beyond the scope of this book, but would be covered in a text on the national accounts.