Chapter 2

Time, Uncertainty and Liquidity

Financial economics deals with the allocation of resources over time and in the face of uncertainty. Although we use terms like "present values," "states of nature," and "contingent commodities" to analyze resource allocation in these settings, the basic ideas are identical to those used in the analysis of consumer and producer behavior in ordinary microeconomic theory. In this chapter we review familiar concepts such as preferences, budget constraints, and production technologies in a new setting, where we use them to study the intertemporal allocation of resources and the allocation of risk. We use simple examples to explain these ideas and later show how the ideas can be extended and generalized.

2.1 Efficient allocation over time

We begin with the allocation of resources over time. Although we introduce some new terminology, the key concepts are the same as concepts familiar from the study of efficient allocation in a "timeless" environment. We assume that time is divided into two periods, which we can think of as representing the "present" and the "future." We call these periods **dates** and index them by $t = 0, 1$, where date 0 is the present and date 1 is the future.

2.1.1 Consumption and saving

Suppose a consumer has an **income stream** consisting of Y_0 units of a homogeneous consumption good at date 0 and Y_1 units of the consumption good at date 1. The consumer's utility $U(C_0, C_1)$ is a function of his **consumption stream** (C_0, C_1), where C_0 is consumption at date 0 and C_1 is consumption at date 1. The consumer wants to maximize his utility but first has to decide which consumption streams (C_0, C_1) belong to his **budget set**, that is, which streams are feasible for him. There are several ways of looking at this question. They all lead to the same answer, but it is worth considering each one in turn.

Borrowing and lending

One way of posing the question (of which consumption streams the consumer can afford) is to ask whether the income stream (Y_0, Y_1) can be transformed into a consumption stream (C_0, C_1) by borrowing and lending. For simplicity, we suppose there is a bank that is willing to lend any amount at the fixed interest rate $i > 0$ per period, that is, the bank will lend one unit of present consumption today in exchange for repayment of $(1 + i)$ units in the future. Suppose the consumer decided to spend $C_0 > Y_0$ today. Then he would have to borrow $B = C_0 - Y_0$ in order to balance his budget today, and this borrowing would have to be repaid with interest iB in the future. The consumer could afford to do this if and only if his future income exceeds his future consumption by the amount of the principal and interest, that is,

$$(1 + i)B \leq Y_1 - C_1.$$

We can rewrite this inequality in terms of the consumption and income streams as follows:

$$C_0 - Y_0 \leq \frac{1}{1 + i}(Y_1 - C_1).$$

Conversely, if the consumer decided to consume $C_0 \leq Y_0$ in the present, he could save the difference $S = Y_0 - C_0$ and deposit it with the bank. We suppose that the bank is willing to pay the same interest rate $i > 0$ on deposits that it earns on loans, that is, one unit of present consumption deposited with the bank today will be worth $(1 + i)$ units in the future. The consumer will receive his savings with interest in the future, so his future consumption could exceed his income by $(1 + i)S$, that is,

$$C_1 - Y_1 \leq (1 + i)S.$$

We can rewrite this inequality in terms of the consumption and income streams as follows:

$$C_0 - Y_0 \le \frac{1}{1+i}(Y_1 - C_1).$$

Notice that this is the same inequality as we derived before. Thus, any feasible consumption stream, whether it involves saving or borrowing, must satisfy the same constraint. We call this constraint the **intertemporal budget constraint** and write it for future reference in a slightly different form:

$$C_0 + \frac{1}{1+i}C_1 \le Y_0 + \frac{1}{1+i}Y_1. \tag{2.1}$$

— Figure 1 here —

Figure 1 illustrates the set of consumption streams (C_0, C_1) that satisfy the intertemporal budget constraint. It is easy to see that the income stream (Y_0, Y_1) must satisfy the intertemporal budget constraint. If there is neither borrowing nor lending in the first period then $C_0 = Y_0$ and $C_1 = Y_1$. The endpoints of the line represent the levels of consumption that would be possible if the individual were to consume as much as possible in the present and future, respectively. For example, if he wants to consume as much as possible in the present, he has Y_0 units of income today and he can borrow $B = Y_1/(1+i)$ units of the good against his future income. This is the maximum he can borrow because in the future he will have to repay the principal B plus the interest iB, for a total of $(1+i)B = Y_1$. So the maximum amount he can spend today is given by

$$C_0 = Y_0 + B = Y_0 + \frac{Y_1}{1+i}.$$

Conversely, if he wants to consume as much as possible in the future, he will save his entire income in the present. In the future, he will get his savings with interest $(1+i)Y_0$ plus his future income Y_1. So the maximum amount he can spend in the future is

$$C_1 = (1+i)Y_0 + Y_1.$$

Suppose now that consumption in the first period is increased by ΔC_0. By how much must future consumption be reduced? Every unit borrowed in the first period will cost $(1+i)$ in the second because interest must be paid.

So the decrease in second period consumption is $\Delta C_1 = -(1+i)\Delta C_0$. This shows that he can afford any consumption stream on the line between the two endpoints with constant slope $= -(1+i)$. (See Fig. 1).

We have shown that any consumption stream that can be achieved by borrowing and lending must satisfy the intertemporal budget constraint. Conversely, we can show that any consumption stream (C_0, C_1) that satisfies the intertemporal budget constraint can be achieved by some feasible pattern of borrowing or lending (saving). To see this, suppose that the intertemporal budget constraint is satisfied by the consumption stream (C_0, C_1). If $C_0 > Y_0$ we assume the consumer borrows $B = C_0 - Y_0$. In the future he has to repay his loan with interest, so he only has $Y_1 - (1+i)B$ left to spend on consumption. However, the intertemporal budget constraint ensures that his planned future consumption C_1 satisfies

$$\begin{aligned} C_1 &\leq (1+i)(Y_0 - C_0) + Y_1 \\ &= Y_1 - (1+i)(C_0 - Y_0). \end{aligned}$$

So the consumer can borrow B units today and repay it with interest tomorrow and still afford his planned future consumption. The other case where $C_0 \leq Y_0$ is handled similarly. Thus, we have seen that the income stream (Y_0, Y_1) can be transformed into the consumption stream (C_0, C_1) through borrowing or lending at the interest rate i if and only if it satisfies the intertemporal budget constraint (2.1).

Wealth and present values

Another way of thinking about the set of affordable consumption streams makes use of the concept of present values. The present value of any good is the amount of present consumption that someone would give for it. The present value of one unit of present consumption is 1. The present value of one unit of future consumption is $1/(1+i)$, because one unit of present consumption can be converted into $1+i$ units of future consumption, and vice versa, through borrowing and lending. Thus, the present value of the income stream (Y_0, Y_1), that is, the value of (Y_0, Y_1) in terms of present consumption is

$$PV(Y_0, Y_1) \equiv Y_0 + \frac{1}{1+i} Y_1$$

and the present value of the consumption stream (C_0, C_1) is

$$PV(C_0, C_1) \equiv C_0 + \frac{1}{1+i} C_1.$$

The intertemporal budget constraint says that the present value of the consumption stream (C_0, C_1) must be less than or equal to the present value of the consumer's income stream.

The present value of the income stream (Y_0, Y_1) is also called the consumer's **wealth**, denoted by $W \equiv Y_0 + \frac{1}{1+i} Y_1$. The intertemporal budget constraint allows the consumer to choose any consumption stream (C_1, C_2) whose present value does not exceed his wealth, that is,

$$C_0 + \frac{C_1}{1+i} \leq W. \tag{2.2}$$

Dated commodities and forward markets

There is still a third way to interpret the intertemporal budget constraint (2.1). We are familiar with the budget constraint of a consumer who has to divide his income between two goods, beer and pizza, for example. There is a price at which each good can be purchased and the value of consumption is calculated by multiplying the quantity of each good by the price and adding the two expenditures. The consumer's budget constraint says that the value of his consumption must be less than or equal to his income. The intertemporal budget constraint (2.1) can be interpreted in this way too. Suppose we treat present consumption and future consumption as two different commodities and assume that there are markets on which the two commodities can be traded. We assume these markets are perfectly competitive, so the consumer can buy and sell as much as he likes of both commodities at the prevailing prices. The usual budget constraint requires the consumer to balance the value of his purchases and expenditures on the two commodities. If p_0 is the price of present consumption and p_1 is the price of future consumption, then the ordinary budget constraint can be written as

$$p_0 C_0 + p_1 C_1 \leq p_0 Y_0 + p_1 Y_1.$$

Suppose that we want to use the first-period consumption good as our **numeraire**, that is, measure the value of every commodity in terms of this good. Then the price of present consumption is $p_0 = 1$, since one unit of the good at date 0 is worth exactly one (unit of the good at date 0). How much is the good at date 1 worth in terms of the good at date 1? If it is possible to borrow and lend at the interest rate i, the price of the good at date 1 will be determined by **arbitrage**. If $p_1 > \frac{1}{1+i}$, then anyone can make a riskless arbitrage profit by selling one unit of future consumption for p_1, using the proceeds to buy $\frac{1}{1+i}$ units of present consumption, and investing the $\frac{1}{1+i}$ units at the interest rate i to yield $(1+i)\frac{1}{1+i} = 1$ unit of future consumption. This strategy yields a profit of $p_1 - \frac{1}{1+i}$ at date 0 and has no cost since the unit of future consumption that is sold is provided by the investment at date 0. Such a risk free profit is incompatible with equilibrium, since anyone can use this arbitrage to generate unlimited wealth. Thus, in equilibrium we must have $p_1 \leq \frac{1}{1+i}$.

A similar argument can be used to show that if $p_1 < \frac{1}{1+i}$, it is possible to make a risk free profit by borrowing $\frac{1}{1+i}$ units of present consumption, buying one unit of future consumption at the price p_1, and using it to repay the loan at date 1. Thus, in equilibrium, we must have $p_1 \geq \frac{1}{1+i}$.

Putting these two arbitrage arguments together, we can see that, if it is possible to borrow and lend at the interest rate i and present consumption is the numeraire, the only prices consistent with equilibrium are $p_0 = 1$ and $p_1 = \frac{1}{1+i}$.

Substituting the prices $p_0 = 1$ and $p_1 = \frac{1}{1+i}$ into the budget constraint above, we see that it is exactly equivalent to the intertemporal budget constraint (2.1). Borrowing and lending at a constant interest rate is equivalent to having a market in which present and future consumption can be exchanged at the constant prices (p_0, p_1). The same good delivered at two different dates is two different commodities and present and future consumption are, in fact, simply two different commodities with two distinct prices. From this point of view, the intertemporal budget constraint is just a new interpretation of the familiar consumer's budget constraint.

Consumption and saving

Since the consumer's choices among consumption streams (C_0, C_1) are completely characterized by the intertemporal budget constraint, the consumer's decision problem is to maximize his utility $U(C_0, C_1)$ by choosing a consumption stream (C_0, C_1) that satisfies the budget constraint. We represent this decision problem schematically by

$$\begin{array}{ll} \max & U(C_0, C_1) \\ \text{s.t.} & C_0 + \frac{C_1}{1+i} = W. \end{array}$$

Note that here we assume the budget constraint is satisfied as an equality rather than an inequality. Since more consumption is preferred to less, there is no loss of generality in assuming that the consumer will always spend as much as he can on consumption. The solution to this maximization problem is illustrated in Figure 2, where the optimum occurs at the point on the budget constraint where the indifference curve is tangent to the budget constraint.

— Figure 2 —

The slope of the budget constraint is $-(1+i)$ and the slope of the indifference curve at the optimum point is

$$-\frac{\frac{\partial U}{\partial C_0}(C_0^*, C_1^*)}{\frac{\partial U}{\partial C_1}(C_0^*, C_1^*)},$$

so the tangency condition can be written as

$$\frac{\frac{\partial U}{\partial C_0}(C_0^*, C_1^*)}{\frac{\partial U}{\partial C_1}(C_0^*, C_1^*)} = (1+i).$$

It is easy to interpret the first-order condition by re-writing it as

$$\frac{\partial U}{\partial C_0}(C_0^*, C_1^*) = (1+i)\frac{\partial U}{\partial C_1}(C_0^*, C_1^*).$$

The left hand side is the marginal utility of consumption at date 0. The right hand side is the marginal utility of $(1+i)$ units of consumption at date 1. One unit of the good at date 0 can be saved to provide $1+i$ units of the good at date 1. So the first-order condition says that the consumer is indifferent between consuming one unit at date 0 and saving it until date 1 when it will be worth $(1+i)$ units and then consuming the $(1+i)$ units at date 1.

An alternative to the graphical method of finding the optimum is to use the method of Lagrange, which requires us to form the Lagrangean function

$$\mathcal{L}(C_0, C_1, \lambda) = U(C_0, C_1) - \lambda\left(C_0 + \frac{1}{1+i}C_1 - W\right)$$

and maximize the value of this function with respect to C_0, C_1, and the Lagrange multiplier λ. A necessary condition for a maximum at $(C_0^*, C_1^*, \lambda^*)$ is that the derivatives of $\mathcal{L}(C_0, C_1, \lambda)$ with respect to these variables should all be zero. Then

$$\frac{\partial \mathcal{L}}{\partial C_0}(C_0^*, C_1^*, \lambda^*) = \frac{\partial U}{\partial C_0}(C_0^*, C_1^*) - \lambda^* = 0,$$

$$\frac{\partial \mathcal{L}}{\partial C_1}(C_0^*, C_1^*, \lambda^*) = \frac{\partial U}{\partial C_0}(C_0^*, C_1^*) - \frac{\lambda^*}{1+i} = 0,$$

$$\frac{\partial \mathcal{L}}{\partial \lambda}(C_0^*, C_1^*, \lambda^*) = C_0^* + \frac{1}{1+i}C_1^* - W = 0.$$

The first two conditions are equivalent to the tangency condition derived earlier. To see this, eliminate λ^* from these equations to get

$$\frac{\partial U}{\partial C_0}(C_0^*, C_1^*) = \lambda^* = (1+i)\frac{\partial U}{\partial C_0}(C_0^*, C_1^*).$$

The last of the three conditions simply asserts that the budget constraint must be satisfied.

As before, the optimum (C_0^*, C_1^*) is determined by the tangency condition and the budget constraint.

Clearly, the optimal consumption stream (C_0^*, C_1^*) will be a function of the consumer's wealth W and the rate of interest i. If the pattern of income were $(W, 0)$ instead of (Y_1, Y_2) the value of wealth would be the same and hence the budget line would be the same. So the same point (C_1^*, C_2^*) would be chosen. In fact (Y_1, Y_2) could move to any other point on the budget line without affecting consumption. Only savings or borrowing would change.

On the other hand, an increase in W to W', say, will shift the budget line out and increase consumption. The case illustrated in Fig. 3 has the special property that the marginal rate of substitution is constant along a straight line OA through the origin. The slope of the budget line does not change so in this case the point of tangency moves along the line OA as W changes. In this special case, C_1 is proportional to W.

— Figure 3 —

Problems

1 An individual consumer has an income stream (Y_0, Y_1) and can borrow and lend at the interest rate i. For each of the following data points, determine whether the consumption stream (C_0, C_1) lies within the consumer's budget set (i.e., whether it satisfies the intertemporal budget constraint).

(C_0, C_1)	(Y_0, Y_1)	$(1 + i)$
$(10, 25)$	$(15, 15)$	2
$(18, 11)$	$(15, 15)$	1.1
$(18, 11)$	$(15, 15)$	1.5
$(10, 25)$	$(15, 15)$	1.8

Draw a graph to illustrate your answer in each case.

2 An individual consumer has an income stream $(Y_0, Y_1) = (100, 50)$ and can borrow and lend at the interest rate $i = 0.1\dot{1}$. His preferences are represented by the additively separable utility function

$$U(C_0, C_1) = \log C_0 + 0.9 \log C_1.$$

The marginal utility of consumption in period t is

$$\frac{d \log C_t}{dC_t} = \frac{1}{C_t}.$$

Write down the consumer's intertemporal budget constraint and the first-order condition that must be satisfied by the optimal consumption stream. Use the first-order condition and the consumer's intertemporal budget constraint to find the consumption stream (C_0^*, C_1^*) that maximizes utility.

How much will the consumer save in the first period? How much will his savings be worth in the second period? Check that he can afford the optimal consumption C_1^* in the second period.

2.1.2 Production

Just as we can cast the consumer's intertemporal decision into the familiar framework of maximizing utility subject to a budget constraint, we can cast the firm's intertemporal decision into the form of a profit- or value-maximization problem.

Imagine a firm that can produce outputs of a homogeneous good in either period subject to a production technology with decreasing returns. Output at date 0 is denoted by Y_0 and output at date 1 is denoted by Y_1. The possible combinations of Y_0 and Y_1 are described by the production possibility curve illustrated in Figure 4.

— Figure 4 here —

Note the following properties of the production possibility curve:

- the curve is downward sloping to the right because the firm must reduce output tomorrow in order to increase the output today;

- the curve is convex upward because of the diminishing returns — as the firm decreases output today, each additional unit of present output foregone adds less to output tomorrow.

The production technology can be represented by a transformation function $F(Y_0, Y_1)$. A pair of outputs (Y_0, Y_1) is feasible if and only if it satisfies the inequality

$$F(Y_0, Y_1) \leq 0.$$

The function F is said to be **increasing** if an increase in Y_0 or Y_1 increases the value $F(Y_0, Y_1)$. The function F is said to be **convex** if, for any output streams (Y_0, Y_1) and (Y_0', Y_1') and any number $0 < t < 1$,

$$F\left(t\left(Y_0, Y_1\right) + (1-t)\left(Y_0', Y_1'\right)\right) \leq tF\left(Y_0, Y_1\right) + (1-t)\, F\left(Y_0', Y_1'\right).$$

If F is increasing, then the production possibility curve is downward sloping. If the function F is convex, the production possibility curve is convex upward. In other words, if (Y_0, Y_1) and (Y_0', Y_1') are feasible, then any point on the line segment between them is feasible.

To illustrate the meaning of the transformation curve, suppose that the firm's past investments produce an output of $\bar{Y}_0$ at the beginning of period 0. The firm can re-invest K_0 units and sell the remaining $Y_0 = \bar{Y}_0 - K_0$. The investment of K_0 units today produces $Y_1 = G(K_0)$ in the future (assume there is no investment in the future because the firm is being wound up). Then the firm can produce any combination of present and future goods

(Y_0, Y_1) for sale that satisfies $Y_0 \leq \bar{Y}_0$ and $Y_1 = G(\bar{Y}_0 - Y_0)$. Then the transformation curve $F(Y_0, Y_1)$ can be defined by

$$F(Y_0, Y_1) = Y_1 - G(\bar{Y}_0 - Y_0).$$

Which combination of outputs Y_0 and Y_1 should the firm choose? In general, there may be many factors that will guide the firm's decision, but under certain circumstances the firm can ignore all these factors and consider only the market value of the firm. To see this, we need only recall our discussion of the consumer's decision. Suppose that the firm is owned by a single shareholder who receives the firm's outputs as income. If the consumer can borrow and lend as much as he wants at the rate i, all he cares about is his wealth, the present value of the income stream (Y_0, Y_1). The exact time-profile of income (Y_0, Y_1) does not matter. So if the firm wants to maximize its shareholder's welfare, it should maximize the shareholder's wealth. To make these ideas more precise, suppose that the firm has a single owner-manager who chooses both the firm's production plan (Y_0, Y_1) and his consumption stream (C_0, C_1) to maximize his utility subject to his intertemporal budget constraint. Formally, we can write this decision problem as follows:

$$\begin{array}{ll} \max & U(C_0, C_1) \\ \text{s.t.} & F(Y_0, Y_1) \leq 0 \\ & C_0 + \frac{1}{1+i} C_1 \leq W \equiv Y_0 + \frac{1}{1+i} Y_1. \end{array}$$

Then it is clear that the choice of (Y_0, Y_1) affects utility only through the intertemporal budget constraint and that anything that increases the present value of the firm's output stream will allow the consumer to reach a more desirable consumption stream. Thus, the joint consumption and production decision above is equivalent to the following two-stage procedure. First, have the firm maximize the present value of outputs:

$$\begin{array}{ll} \max & W \equiv Y_0 + \frac{1}{1+i} Y_1 \\ \text{s.t.} & F(Y_0, Y_1) \leq 0. \end{array}$$

The present value of outputs is also known as the **market value of the firm**, so this operational rule can be re-phrased to say that firms should always maximize market value. Then have the consumer maximize his utility taking the firm's market value as given:

$$\begin{array}{ll} \max & U(C_0, C_1) \\ \text{s.t.} & C_0 + \frac{1}{1+i} C_1 \leq W, \end{array}$$

where $W = Y_0 + \frac{1}{1+i}Y_1$. Figure 5 illustrates this principle for the case of a single shareholder.

— Figure 5 here —

In fact, this argument extends to the case where there are many shareholders with different time preferences. Some shareholders may be impatient and want to consume more in the present while others are more patient and are willing to postpone consumption, but all will agree that a change in production that increases the present value of output must be a good thing, because it increases the consumers' wealth. Figure 6 illustrates the case of two shareholders with equal shares in the firm. Then they have identical budget constraints with slope $-(1+i)$, i.e., their budget constraint is parallel to the maximum value line that is tangent to the production possibility frontier. Each shareholder will choose to consume the bundle of goods (C_0, C_1) that maximizes his utility subject to the budget constraint. Because they have different time preferences, represented here by different indifference curves, each shareholder will choose a different bundle of goods, as indicated by the different points of tangency between indifference curves and budget constraint. Nonetheless, both agree that the firm should maximize its market value, because maximizing the value of the firm has the effect of putting both shareholders on the highest possible budget constraint. This is known as the **separation theorem** because the firm's decision to maximize its value is separate from shareholders' decisions to maximize their utility.

— Figure 6 here —

Problem

3 A firm has 100 units of output at the beginning of period 0. It has three projects that it can finance. Each project requires an input of I units of the good at date 0 and produces Y_1 units of the good at date 1. The projects are defined in the following table:

Project	Investment I	Output Y_1
1	20	30
2	30	48
3	50	70

Which projects should the firm undertake when the interest factor is

$$1 + i = 2, 1.5, 1.1?$$

Trace the firm's production possibility curve (the combinations of Y_0 and Y_1 that are technologically feasible) assuming that the firm can undertake a fraction of a project. Use this diagram to illustrate how changes in the discount factor change the firm's decision.

2.2 Uncertainty

In the same way that we can extend the traditional analysis of consumption and production to study the allocation of resources over time, we can use the same ideas to study the allocation of risk bearing under uncertainty. Once again we shall use a simple example to illustrate the general principles.

2.2.1 Contingent commodities and risk sharing

We assume that time is divided into two periods or dates indexed by $t = 0, 1$. At date 0 (the present) there is some uncertainty about the future. For example, an individual may be uncertain about his future income. We can represent this uncertainty by saying that there are several possible **states of nature**. A state of nature is a complete description of all the uncertain, exogenous factors that may be relevant for the outcome of an individual's decision. For example, a farmer who is planting a crop may be uncertain about the weather. The size of his crop will depend on choices he makes about the time to plant, the use of fertilizers, etc., as well as the weather. In this case, we identify the weather with the state of nature. Each state would be a complete description of the weather — the amount of rainfall, the temperature, etc., — during the growing season. The outcome of the farmer's choices will depend on the parameters he determines — time to plant, etc. — and the state of nature. In other words, once we know the farmer's choices and the state of nature, we should know the size of the crop; but even after the farmer has determined all the parameters that he controls, the fact that the state is still unknown means that the size of the crop is uncertain.

In what follows we assume that the only uncertainty relates to an individual's income, so income is a function of the state of nature. The true state of nature is unknown at date 0 and will be revealed at date 1. For simplicity, suppose that there are two possible states indexed by $s = H, L$. The letters H and L stand for "high" and "low".

Commodities are distinguished by their physical characteristics, by the date at which they are delivered, and by the state in which they are delivered. Thus, the consumption good in state H is a different commodity from the consumption good in state L. We call a good whose delivery is contingent on the occurrence of a particular state of nature a **contingent commodity**. By adopting this convention, we can represent uncertainty about income and consumption in terms of a bundle of contingent commodities. If Y_H denotes future income in state H and Y_L denotes income in state L, the ordered pair (Y_H, Y_L) completely describes the future uncertainty about the individual's income. By treating (Y_H, Y_L) as a bundle of different (contingent) commodities, we can analyze choices under uncertainty in the same way that we analyze choices among goods with different physical characteristics. We assume that $Y_H > Y_L$, that is, income is higher in the "high" state.

An individual's preferences over uncertain consumption outcomes can be represented by a utility function that is defined on bundles of contingent commodities. Let $U(C_H, C_L)$ denote the consumer's utility from the bundle of contingent commodities (C_H, C_L), where C_H denotes future consumption in state H and C_L denotes future consumption in state L. Later, we introduce the notion of a state's probability and distinguish between an individual's probability beliefs and his attitudes to risk. Here, an individual's beliefs about the probability of a state and his attitudes towards risk are subsumed in his preferences over bundles of contingent commodities.

Complete markets

There are two equivalent ways of achieving an efficient allocation of risk. One approach to the allocation of risk assumes that there are **complete markets** for contingent commodities. An economy with complete markets is often referred to as an Arrow-Debreu economy. In an Arrow-Debreu economy, there is a market for each contingent commodity and a prevailing price at which consumers can trade as much of the commodity as they like subject to their budget constraint. Let p_H and p_L denote the price of the contingent commodities corresponding to states H and L, respectively. The consumer's income consists of different amounts of the two contingent commodities, Y_H units of the consumption good in state H and Y_L units of the consumption good in state L. We can use the complete markets to value this uncertain income stream and the consumer's wealth is $p_H Y_H + p_L Y_L$. Then the consumer's budget constraint, which says that his consumption expenditure must be less than or equal to his wealth, can be written as

$$p_H C_H + p_L C_L \leq p_H Y_H + p_L Y_L.$$

The consumer can afford any bundle of contingent commodities (C_H, C_L) that satisfies this constraint. He chooses the consumption bundle that maximizes his utility $U(C_H, C_L)$ subject to this constraint. We illustrate the consumer's decision problem in Figure 7.

— Figure 7 here —

Arrow securities

The assumption of complete markets guarantees efficient allocation of risk, but it may not be realistic to assume that every contingent commodity, of which there will be a huge number in practice, unlike in our example, can literally be traded at the initial date. Fortunately, there exists an alternative formulation which is equivalent in terms of its efficiency properties, but requires far fewer markets. More precisely, it requires that securities and goods be traded on spot markets at each date, but the total number of spot markets is much less than the number of contingent commodities.

The alternative representation of the allocation of risk makes use of the idea of **Arrow securities**. We define an Arrow security to be a promise to deliver one unit of money (or an abstract unit of account) if a given state occurs and nothing otherwise. In terms of the present example, there are two types of Arrow securities, corresponding to the states H and L respectively. Let q_H denote the price of the Arrow security corresponding to state H and let q_L denote the price of the Arrow security in state L. In other words, q_H is the price of one unit of account (one "dollar") delivered in state H at date 1 and q_L is the price of one unit of account delivered in state L at date 1. A consumer can trade Arrow securities at date 0 in order to hedge against income risks at date 1. Let Z_H and Z_L denote the excess demand for the Arrow securities corresponding to states H and L respectively.[1] If

[1] Each agent begins with a zero net supply of Arrow securities. Then agents issue securities for one state in order to pay for their purchase of securities in the other. The

$Z_s > 0$ then the consumer is taking a long position (offering to buy) in the Arrow security and if $Z_s < 0$ the consumer is taking a short position (offering to sell). We assume the consumer has no income at date 0 — this period exists only to allow individuals to hedge risks that occur at date 1 — so the consumer has to balance his budget by selling one security in order to purchase the other. Suppose that the consumer chooses a portfolio $Z = (Z_H, Z_L)$ of Arrow securities. Then at date 1, once the true state has been revealed, his budget constraint will be

$$\hat{p}_H C_H \leq \hat{p}_H Y_H + Z_H, \tag{2.3}$$

if state H occurs, and

$$\hat{p}_L C_L \leq \hat{p}_L Y_L + Z_L, \tag{2.4}$$

if state L occurs. The consumer will choose the portfolio Z and the consumption bundle (C_H, C_L) to maximize $U(C_H, C_L)$ subject to the date-0 budget constraint

$$q_H Z_H + q_L Z_L \leq 0$$

and the budget constraints (2.3) and (2.4). Since $Z_H = \hat{p}_H (C_H - Y_H)$ and $Z_L = \hat{p}_L (C_L - Y_L)$, the budget constraint at date 0 is equivalent to $q_H \hat{p}_H (C_H - Y_H) + q_L \hat{p}_L (C_L - Y_L) \leq 0$ or

$$q_H \hat{p}_H C_H + q_L \hat{p}_L C_L \leq q_H \hat{p}_H Y_H + q_L \hat{p}_L Y_L.$$

This looks just like the standard budget constraint in which we interpret $p_H = q_H \hat{p}_H$ and $p_L = q_L \hat{p}_L$ as the prices of contingent commodities and C_H and C_L as the demands for contingent commodities.

Now that we have seen how to interpret the allocation of risk in terms of contingent commodities, we can use the standard framework to analyze efficient risk sharing. Figure 8 shows an Edgeworth Box in which the axes correspond to consumption in state H and consumption in state L. A competitive equilibrium in which consumers maximize utility subject to their budget constraint leads to an efficient allocation of contingent commodities, that is, an efficient allocation of risk.

— Figure 8 here —

---

vector Z represents the agent's net or excess demand of each security: a component Z_s is negative if his supply of the security (equals minus the excess demand) is positive and positive if his net demand is positive.

As usual, the conditions for efficiency include the equality of the two consumers' marginal rates of substitution, but here the interpretation is different. We are equating the marginal rates of substitution between consumption in the two states rather than two different (physical) goods. The marginal rate of substitution will reflect an individual's subjective belief about the probability of each state as well as his attitude toward risk.

2.2.2 Attitudes toward Risk

To characterize individuals' attitudes toward risk we introduce a special kind of utility function, which we call a **von Neumann-Morgenstern** (VNM) utility function. Whereas the standard utility function is defined on bundles of contingent commodities, the VNM utility is defined on quantities of consumption in a particular state. Von Neumann and Morgenstern showed that, under certain conditions, a rational individual would act so as to maximize the expected value of his VNM utility function. If individuals satisfy the assumptions of the VNM theory, they will always make choices so as to maximize the value of their expected utility. To see what this means in practice, let $U(C)$ denote the VNM utility of consuming C units of the good at date 1 and suppose that the probability of state s occurring is $\pi_s > 0$ for $s = H, L$. Then the expected utility of a consumption plan (C_H, C_L) is $\pi_H U(C_H) + \pi_L U(C_L)$. The decision problem of the consumer we encountered above can be re-written as

$$\begin{array}{rl} \max & \pi_H U(C_H) + \pi_L U(C_L) \\ \text{s.t.} & q_H C_H + q_L C_L \leq q_H Y_H + q_L Y_L. \end{array}$$

The first-order conditions for this problem are

$$\pi_s U'(C_s) = \mu q_s,$$

for $s = H, L$, where $U'(C_s)$ is the marginal utility of consumption in state s and μ (the Lagrange multiplier associated with the budget constraint) can be interpreted as the marginal utility of money. Notice that the marginal utility of consumption in state s is multiplied by the probability of state s and it is the product — the *expected* marginal utility — which is proportional to the price of consumption in that state. Then the first-order condition can be interpreted as saying that the expected marginal utility of one unit of consumption in state s is equal to the marginal utility of its cost.

We typically think of individuals as being **risk averse**, that is, they avoid risk unless there is some advantage to be gained from accepting it. The clearest evidence for this property is the tendency to buy insurance. We can characterize risk aversion and an individual's attitudes to risk generally in the shape of the VNM utility function. Figure 9 shows the graph of a VNM utility function. Utility increases with income (the marginal utility of consumption is positive) but the utility function becomes flatter as income increases (diminishing marginal utility of consumption).

— Figure 9 here —

A VNM utility function has diminishing marginal utility of consumption if and only if it is strictly concave. Formally, a VNM utility function U is **strictly concave** if, for any consumption levels C and C' ($C \neq C'$) and any number $0 < t < 1$, it satisfies the inequality

$$U\left(tC + (1-t)\,C'\right) > tU(C) + (1-t)U(C'). \tag{2.5}$$

Concavity of the VNM utility function can be interpreted as an attitude towards risk. To see this, suppose that an individual is offered a gamble in which he receives C with probability t and C' with probability $1-t$. If his VNM utility is strictly concave, he will prefer to receive the expected value $tC + (1-t)C'$ for sure, rather than take the gamble. This is because the expected utility of the gamble (the right hand side of the inequality) is less than the utility of the sure thing. This relationship is illustrated in Figure 9. The utility function will always be strictly concave if the individual exhibits diminishing marginal utility of income. An individual who satisfies the assumption of diminishing marginal utility of income is said to be **risk averse**. (Draw the graph of a utility function with increasing marginal utility of income and see how the comparison of the two options changes. An individual with these preferences is called a **risk lover.**)

The conclusion then is that, faced with a choice between a risky income distribution and a degenerate distribution with the same expected value, a risk averse individual will always choose the one without risk. In what follows, we assume that the VNM utility function is concave and, hence, individuals are risk averse.

We have seen that risk aversion is associated with the curvature of the utility function, in particular, with the fact that the marginal utility of income is decreasing. Mathematically, this means that the second derivative of the utility function $U''(C)$ is less than or equal to zero. It would be tempting to take the second derivative $U''(C)$ to be the measure of risk aversion. Unfortunately, the VNM utility function is only determined up to an affine transformation, that is, for any constants α and $\beta > 0$, the VNM utility function $\alpha + \beta U$ is equivalent to U in terms of the attitudes to risk that it implies. Thus, we must look for a measure that is independent of α and $\beta > 0$. Two such measures are available. One is known as the degree of *absolute risk aversion*

$$A(C) = -\frac{U''(C)}{U'(C)}$$

and the other is the degree of *relative risk aversion*

$$R(C) = -\frac{U''(C)C}{U'(C)}.$$

There is a simple relationship between the degree of risk aversion and the risk premium that an individual will demand to compensate for taking risk. Suppose that an individual has wealth W and is offered the following gamble. With probability 0.5 he wins a small amount h and with probability 0.5 he loses h. Since the expected value of the gamble is zero, the individual's expected income is not changed by the gamble. Since a risk averse individual would rather have W for sure than have an uncertain income with the same expected value, he will reject the gamble. The risk premium a is the amount he would have to be given in order to accept the gamble. That is, a satisfies the equation

$$U(W) = \frac{1}{2}U(W + a - h) + \frac{1}{2}U(W + a + h).$$

A Taylor's expansion of the right hand side shows that, when h is small,

$$U(W) \approx U(W) + U'(W)a + \frac{1}{2}U''(W)h^2.$$

Thus,

$$a \approx -\frac{U''(W)}{U'(W)}\frac{h^2}{2} = A(W)\frac{h^2}{2}.$$

So the risk premium a is equal to the degree of absolute risk aversion times one-half the variance of the gamble (a measure of the risk). A similar interpretation can be given for the degree of relative risk aversion when the gamble consists of winnings of hW and $-hW$ with equal probability.

If the degree of absolute risk aversion is constant, the VNM utility function must have the form

$$U(C) = -e^{-AC},$$

and $A > 0$ is the degree of absolute risk aversion. If the degree of relative risk aversion is constant and different from 1, then

$$U(C) = \frac{1}{1-\sigma} C^{1-\sigma},$$

where $\sigma > 0$ is the degree of relative risk aversion. This formula is not defined when the degree of relative risk aversion is $\sigma = 1$; however, the limiting value of the utility function as $\sigma \to 1$ is well defined and given by

$$U(C) = \ln C,$$

where $\ln C$ denotes the natural logarithm of C.

The higher the degree of (relative or absolute) risk aversion, the more risk averse the individual with the VNM utility function $U(C)$ is.

2.2.3 Insurance and risk pooling

Returning to the example of efficient risk sharing studied earlier, we can use the assumption that consumers have VNM utility functions to characterize the efficient risk sharing allocation more precisely. Suppose that there are two individuals A and B with VNM utility functions U_A and U_B and income distributions (Y_{AH}, Y_{AL}) and (Y_{BH}, Y_{BL}) respectively. If the efficient allocation of consumption in the two states is given by $\{(C_{AH}, C_{AL}), (C_{BH}, C_{BL})\}$, then the condition equality between the marginal rates of substitution can be written as

$$\frac{U_A'(C_{AH})}{U_A'(C_{AL})} = \frac{U_B'(C_{BH})}{U_B'(C_{BL})}.$$

The probabilities do not appear in this equation because, assuming A and B have the same probability beliefs, they appear as multipliers on both sides and so cancel out.

It is interesting to consider what happens in the case where one of the consumers is **risk neutral**. We say that a consumer is risk neutral if his VNM utility function has the form

$$U(C) \equiv C.$$

A risk neutral consumer cares only about the expected value of his income or consumption. In other words, his expected utility is just the expected value of consumption. Suppose that consumer B is risk neutral. Then his marginal utility is identically equal to 1 in each state. Substituting this value into the efficiency condition, we see that

$$\frac{U'_A(C_{AH})}{U'_A(C_{AL})} = 1,$$

which implies that $C_{AH} = C_{AL}$. All the risk is absorbed by the risk neutral consumer B, leaving consumer A with a certain level of consumption.

Risk neutrality is a very special property, but there are circumstances in which risk averse consumers can achieve the same effects. First, let us consider more carefully the way in which the optimal consumption allocation depends on income. First, note that the efficiency equation implies that $U'_A(C_{AH}) < U'_A(C_{AL})$ if and only if $U'_B(C_{BH}) < U'_B(C_{BL})$. Since marginal utility is decreasing in consumption, this implies that $C_{AH} > C_{AL}$ if and only if $C_{BH} > C_{BL}$. This immediately tells us that the optimal consumption allocation depends only on the aggregate income of the pair. Let $Y_s = Y_{As} + Y_{Bs}$ for $s = H, L$. Then feasibility requires the total consumption equals total income in each state:

$$C_{As} + C_{Bs} = Y_s,$$

for $s = H, L$. Thus, the consumption of each consumer rises if and only if aggregate income rises. This property is known as coinsurance between the two consumers: they provide insurance to each other in the sense that their consumption levels go up and down together. In particular, if the aggregate income is the same in the two states, $Y_H = Y_L$, then $C_{AH} = C_{AL}$ and $C_{BH} = C_{BL}$. So if aggregate income is constant, the consumption allocation will be constant too, no matter how individual incomes fluctuate.

When there are only two consumers, constant aggregate income depends on a rather remarkable coincidence: when A's income goes up, B's income goes down by the same amount. When there is a large number of consumers, however, the same outcome occurs quite naturally, thanks to the Law of Large Numbers, as long as the incomes of the different consumers are assumed to be **independent**. This is, in fact, what insurance companies do: they pool large numbers of independent risks, so that the aggregate outcome becomes almost constant, and then they can ensure that each individual gets a constant level of consumption. Suppose there is a large number of consumers

$i = 1, 2, ...$ with random incomes that are independently and identically distributed according to the probability distribution

$$Y_i = \begin{cases} Y_H & \text{with probability } \pi_H \\ Y_L & \text{with probability } \pi_L. \end{cases}$$

Then the law of large numbers ensures that the average income is equal to the expected value of an individual's income $\bar{Y} = \pi_H Y_H + \pi_L Y_L$ with probability one, that is, with certainty. Then an insurance company could ensure every one a constant consumption level because the average aggregate income is almost constant.

2.2.4 Portfolio choice

The use of Arrow securities to allocate income risk efficiently is a special case of the portfolio choice problem that individuals have to solve in order to decide how to invest wealth in an uncertain environment. We can gain a lot of insight into the portfolio choice problem by considering the special case of two securities, one a safe asset and the other a risky asset.

As before, we assume there are two dates $t = 0, 1$ and two states $s = H, L$ and a single consumption good at each date. Suppose that an investor has an initial income $W_0 > 0$ at date 0 and that he can invest it in two assets. One is a **safe asset** that yields one unit of the good at date 1 for each unit invested at date 0. The other is a **risky asset**: one unit invested in the risky asset at date 0 yields $R_s > 0$ units of the good in state $s = H, L$. We assume that the investor's risk preferences are represented by a VNM utility function $U(C)$ and that the probability of state s is $\pi_s > 0$ for $s = H, L$.

The investor's portfolio can be represented by the fraction θ of his wealth that he invests in the risky asset. That is, his portfolio will contain θW_0 units of the risky asset and $(1 - \theta) W_0$ of the safe asset. His future consumption will depend on his portfolio choice and the realized return of the risky asset. Let C_H and C_L denote consumption in the high and low states, respectively. Then

$$C_s = R_s \theta W_0 + (1 - \theta) W_0,$$

for $s = H, L$. The investor chooses the portfolio that maximizes the expected utility of his future consumption. That is, his decision problem is

$$\begin{array}{rl} \max_\theta & \pi_H U(C_H) + \pi_L U(C_L) \\ \text{s.t.} & C_s = R_s \theta W_0 + (1 - \theta) W_0, \;\; s = H, L. \end{array}$$

Substituting the expressions for C_H and C_L into the objective function we see that the expected utility is a function of θ, say $V(\theta)$. The optimal portfolio $0 < \theta^* < 1$ satisfies the first-order condition $V'(\theta^*) = 0$, or

$$\pi_H U'(C_H)(R_H - 1) + \pi_L U'(C_L)(R_L - 1) = 0.$$

The optimum is illustrated in Figure 10.

— Figure 10 here —

The set of attainable consumption allocations (C_H, C_L) are represented by the line segment with endpoints (W_0, W_0) and $(R_H W_0, R_L W_0)$. If the investor puts all of his wealth in the safe asset $\theta = 0$, then his future consumption in each state will be $C_H = C_L = W_0$. If he puts all his wealth in the risky asset, then his future consumption with be $R_H W_0$ in the high state and $R_L W_0$ in the low state. If he puts a fraction θ of his wealth in the risky asset, his consumption bundle (C_H, C_L) is just a convex combination of these two endpoints with weights $1-\theta$ and θ respectively. In other words, we can trace out the line joining these two endpoints just by varying the proportion of the risky asset between 0 and 1.

The optimal portfolio choice occurs where the investor's indifference curve is tangent to the consumption curve. The tangency condition is just a geometric version of the first-order condition above.

Depending on the investor's risk preferences and the rates of return, the optimal portfolio may consist entirely of the safe asset, entirely of the risky asset, or a mixture of the two. It is interesting to see under which conditions each of these possibilities arises. To investigate this question, we need to find out more about the slopes of the indifference curves and the feasible set.

The slope of the feasible set is easily calculated. Compare the portfolio in which all income is invested in money with the portfolio is all invested in bonds. The change in C_H is $\Delta C_H = W_0 - W_0 R_H$ and the change in C_L is $\Delta C_L = W_0 - W_0 R_L$. So the slope is

$$\frac{\Delta C_L}{\Delta C_H} = \frac{W_0 - W_0 R_L}{W_0 - W_0 R_H} = \frac{1 - R_L}{1 - R_H}.$$

The slope is negative if $R_L < 1 < R_H$.

An indifference curve is a set of points like (C_H, C_L) that satisfy an equation like

$$\pi_H U(C_H) + \pi_L U(C_L) = \text{constant}.$$

Consider a "small" movement (dC_H, dC_L) along the indifference curve: it must satisfy the equation

$$\pi_H U'(C_H) dC_H + \pi_L U'(C_L) dC_L = 0.$$

We can "solve" this equation to obtain the slope of the indifference curve:

$$\frac{dC_L}{dC_H} = -\frac{\pi_H U'(C_H)}{\pi_L U'(C_L)}.$$

Note that at the point A, the slope of the indifference curve is $-\pi_H/\pi_L$.

Now we are ready to characterize the different possibilities. For an interior solution, $0 < \theta^* < 1$, the slope of the indifference curve must equal the slope of the feasible set, or

$$\frac{\pi_H U'(C_H)}{\pi_L U'(C_L)} = \frac{1 - R_L}{R_H - 1}.$$

A necessary and sufficient condition for this is that the indifference curve is flatter than the feasible set at $(R_H W_0, R_L W_0)$ and steeper than the feasible set at (W_0, W_0). That is,

$$\frac{\pi_H U'(W_0)}{\pi_L U'(W_0)} > \frac{1 - R_L}{R_H - 1} > \frac{\pi_H U'(R_H W_0)}{\pi_L U'(R_L W_0)}.$$

The left hand inequality can be simplified to

$$\frac{\pi_H}{\pi_L} > \frac{1 - R_L}{R_H - 1}$$

or $\pi_H R_H + \pi_L R_L > 1$. In other words, the investor will hold a positive amount of the risky asset if and only if the expected return of the risky asset is greater than the return to the safe asset. This makes sense because there is no reward for bearing risk otherwise. The right hand inequality implies that

$$\frac{1 - R_L}{R_H - 1} > 0,$$

or $R_L < 1 < R_H$. In other words, the investor will hold the safe asset only if the risky asset produces a capital loss in the low state. Otherwise, the risky asset dominates (always pays a higher return than) the safe asset. Notice that even if the risky asset sometimes yields a loss, the investor may choose to invest all his wealth in the risky asset. It just depends on his attitude toward risk and the risk-return trade-off.

2.3 Liquidity

The word liquidity is used in two senses here. First, we describe assets as liquid if they can be easily converted into consumption without loss of value. Secondly, we describe individuals as having a preference for liquidity if they are uncertain about the timing of their consumption and hence desire to hold liquid assets.

Liquid assets

Once again, we illustrate the essential ideas using a simple example. Time is divided into three periods or **dates** indexed by $t = 0, 1, 2$. At each date, there is a single all-purpose good which can be used for consumption or investment.

There are two assets that consumers can use to provide for future consumption, a short-term, liquid asset and a long-term, illiquid asset. In what follows, we refer to these as the **short** and **long** assets, respectively. Each asset is represented by a constant-returns-to-scale investment technology. The short asset is represented by a storage technology that allows one unit of the good at date t to be converted into one unit of the good at date $t + 1$, for $t = 0, 1$. The long asset is represented by an investment technology that allows one unit of the good at date 0 to be converted into $R > 1$ units of the good at date 2. We assume the return of the long asset is known with certainty. This assumption simplifies the analysis and allows us to focus attention on the other source of uncertainty, that is, uncertainty about individual time preferences.

There is a trade-off between an asset's time to maturity and its return. The long asset takes two periods to mature, but pays a high return. The short asset matures after one period but yields a lower return. This trade-off is characteristic of the yield curve for bonds of different maturities, where we see that bonds with short maturities typically have lower returns than bonds with long maturities. The higher returns on the longer-dated assets can be interpreted both as a reward for the inconvenience of holding illiquid assets and as a reflection of the greater productivity of roundabout methods of production.

Liquidity preference

We model preference for liquidity as the result of uncertainty about time preference. Imagine a consumer who has an endowment of one unit of the good at date 0 and nothing at the future dates. All consumption takes place in the future, at dates 1 and 2, but the consumer is uncertain about the precise date at which he wants to consume. More precisely, we assume there are two types of consumers, **early consumers** who only want to consume at date 1 and **late consumers** who only want to consume at date 2. Initially, the consumer does not know his type. He only knows the probability of being an early or a late consumer. Let λ denote the probability of being an early consumer and $1-\lambda$ the probability of being a late consumer. The consumer learns whether he is an early or late consumer at the beginning of date 1.

Uncertainty about time preferences is a simple way of modeling what economists call a "liquidity shock," that is, an unanticipated need for liquidity resulting from an event that changes one's preferences. This could be an accident that requires an immediate expenditure, the arrival of an unexpected investment opportunity, or an unexpected increase in the cost of an expenditure that was previously planned. We can think of λ as measuring the degree of a consumer's liquidity preference. Other things being equal, he will want to earn the highest return possible on his investments. But if he is uncertain about the timing of his consumption, we will also care about liquidity, the possibility of realizing the value of this assets at short notice. If λ is one, the consumer's liquidity preference will be high, since he cannot wait until date 2 to earn the higher return on the long asset. If λ is zero, he will have no preference for liquidity, since he can hold the long asset without inconvenience. For λ between zero and one, the consumer's uncertainty about the timing of his consumption poses a problem. If the consumer knew that he was a late consumer, he would invest in the long asset because it gives a higher return. If he knew that he was an early consumer, he would hold only the short asset in spite of its lower return. Since the consumer is uncertain about his type, he will regret holding the short asset if he turns out to be a late consumer and he will regret holding the long asset if he turns out to be an early consumer. The optimal portfolio for the consumer to hold will depend on both his risk aversion and his liquidity preference and on the return to the long asset (the slope of the yield curve).

Investment under autarky

Suppose the consumer has a period utility function $U(C)$ and let C_1 and C_2 denote his consumption at date 1 (if he is an early consumer) and date 2 (if he is a late consumer). Then his expected utility from the consumption stream (C_1, C_2) is

$$\lambda U(C_1) + (1-\lambda)U(C_2).$$

His consumption at each date will be determined by his portfolio choice at date 0. Let θ denote the proportion of his wealth invested in the short asset. Recall that he has an initial endowment of one unit of the good at date 0 so he invests θ in the short asset and $1-\theta$ in the long asset. Then his consumption at date 1 is given by

$$C_1 = \theta$$

since he cannot consume the returns to the long asset, whereas his consumption at date 2 is given by

$$C_2 = \theta + (1-\theta)R,$$

since the returns to the short asset can be re-invested at date 1 and consumed at date 2. Note that $C_1 < C_2$ except in the case where $\theta = 1$, so the consumer faces some risk and if he is risk averse this will impose some loss of expected utility compared to a situation in which he can consume the expected value $\bar{C} = \lambda C_1 + (1-\lambda)C_2$ for sure. Of course, he can choose $\theta = 1$ if he wants to avoid uncertainty altogether, but there is a cost to doing so: his average consumption will be lower.

The consumer's decision problem is to choose θ to maximize

$$\lambda U(\theta) + (1-\lambda)U(\theta + (1-\theta)R).$$

At an interior solution, the optimal value of θ satisfies

$$\lambda U'(\theta) + (1-\lambda)U'(\theta + (1-\theta)R)\,(1-R) = 0.$$

For example, if $U(C) = \ln C$, then the first-order condition becomes

$$\frac{\lambda}{\theta} + \frac{(1-\lambda)}{\theta + (1-\theta)R}(1-R) = 0,$$

or

$$\theta = \frac{\lambda R}{R-1}.$$

So the fraction of wealth held in the short asset is greater if λ is greater and lower if R is greater. Note that $\theta = 0$ for any value of λ greater than $1 - 1/R$.

Risk pooling

As we have seen already, the consumer's attempt to provide for his future consumption needs is bound to lead to regret as long as he cannot perfectly foresee his type. If he could ensure against his liquidity shock, he could do better. Suppose that there is a large number of consumers, all of whom are ex ante identical and subject to the same shock, that is, they all have a probability λ of being early consumers. If we assume further that their liquidity shocks are independent, then the Law of Large Numbers assures us that there will be no aggregate uncertainty. Whatever happens to the individual consumer, the fraction of the total population who become early consumers will be λ for certain. This suggests the potential for pooling risks and providing a better combination of returns and liquidity.

To see how this would work, suppose a financial institution were to take charge of the problem of investing the endowments of a large number of consumers and providing for their consumption. The financial institution would take the endowments at date 0 and invest a fraction θ in the short asset and a fraction $1-\theta$ in the long asset. At date 1 it would provide consumption equal to C_1 units of the good to early consumers and at date 2 it would provide C_2 units of the good to late consumers. The important difference between the financial institution and the individual consumers is that the company faces no uncertainty: it knows for sure that a fraction λ of its clients will be early consumers. Consequently, it knows for sure what the demand for the consumption good will be at date 1 and date 2. At date 1 it needs to provide λC_1 per capita and at date 2 it needs to provide $(1-\lambda)C_2$ per capita. Because the return to the short asset is lower than the return to the long asset, the financial institution will hold the minimum amount of the short asset it needs to provide for the early consumers' consumption at date 1, that is, $\theta = \lambda C_1$ and it will hold the rest of the portfolio in the long asset. Then the company's plans are feasible if

$$\lambda C_1 = \theta$$

and

$$(1-\lambda)C_2 = (1-\theta)R.$$

The company's decision problem is

$$\begin{array}{ll} \max & \lambda U(C_1) + (1-\lambda)U(C_2) \\ \text{s.t.} & \lambda C_1 = \theta \\ & (1-\lambda)C_2 = (1-\theta)R. \end{array}$$

If we substitute for the C_1 and C_2 in the objective function, we get

$$\lambda U\left(\frac{\theta}{\lambda}\right)+(1-\lambda)U\left(\frac{(1-\theta)\,R}{1-\lambda}\right)$$

and the first-order condition for maximizing this expression with respect to θ is

$$U'\left(C_1\right)-U'\left(C_2\right)R=0.$$

Notice that the terms involving λ have cancelled out. In terms of the earlier example, where the period utility function is $U(C)=\ln C$, the first-order condition implies that $C_2=C_1R$, or

$$\theta=\lambda.$$

2.4 Concluding remarks

This chapter has provided a foundation in terms of the basic finance and economics that will be needed to understand this book. Those wishing to pursue these topics further can consult a textbook such as Mas-Collel, Whinston, and Green (1995).

References

Mas-Collel, A., M. Whinston, and J. Green (1995). *Microeconomic Theory*, Oxford: Oxford University Press.

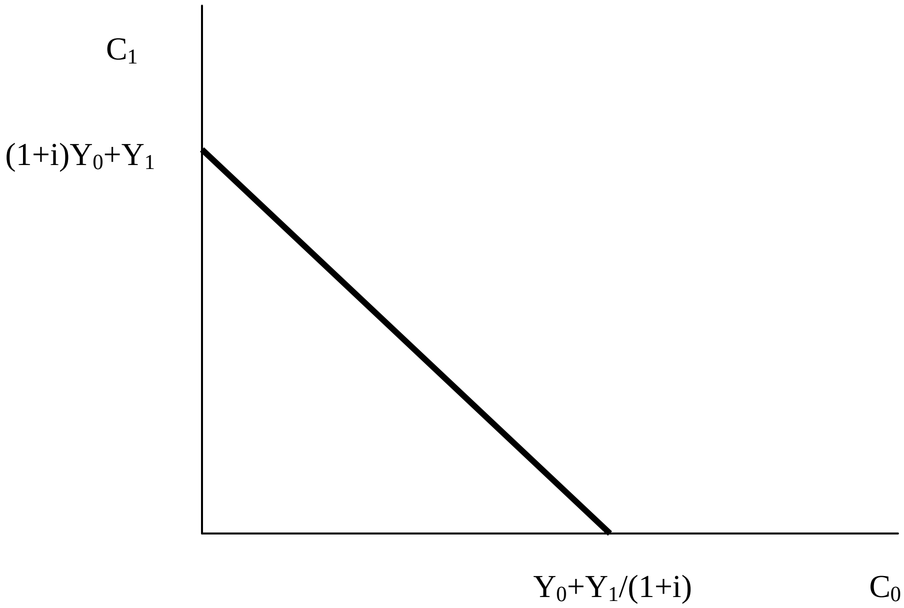

Figure 1
Intertemporal budget constraint

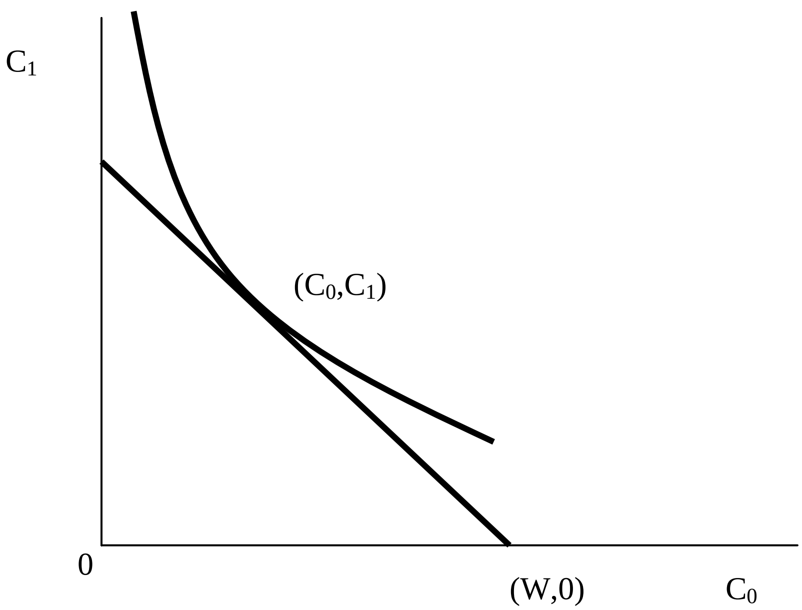

Figure 2
Consumption and saving

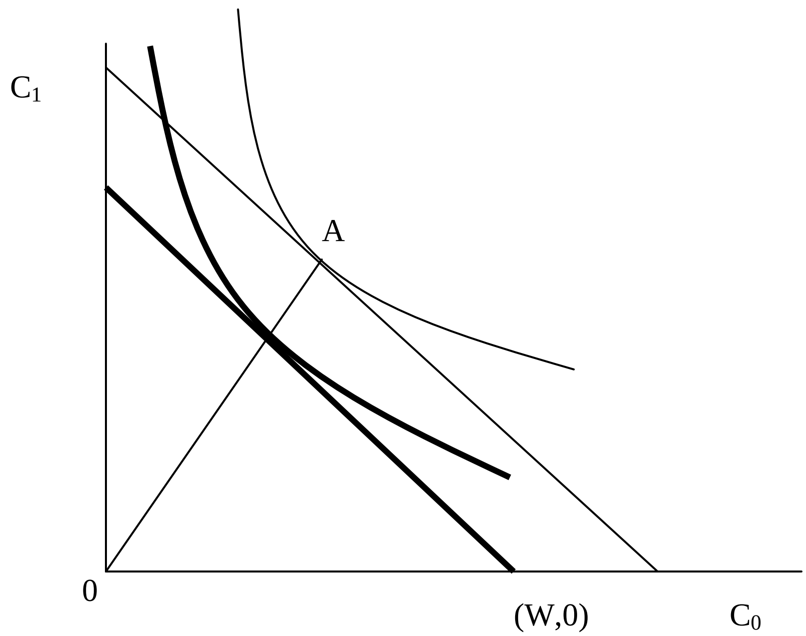

Figure 3
The effect of an increase in wealth

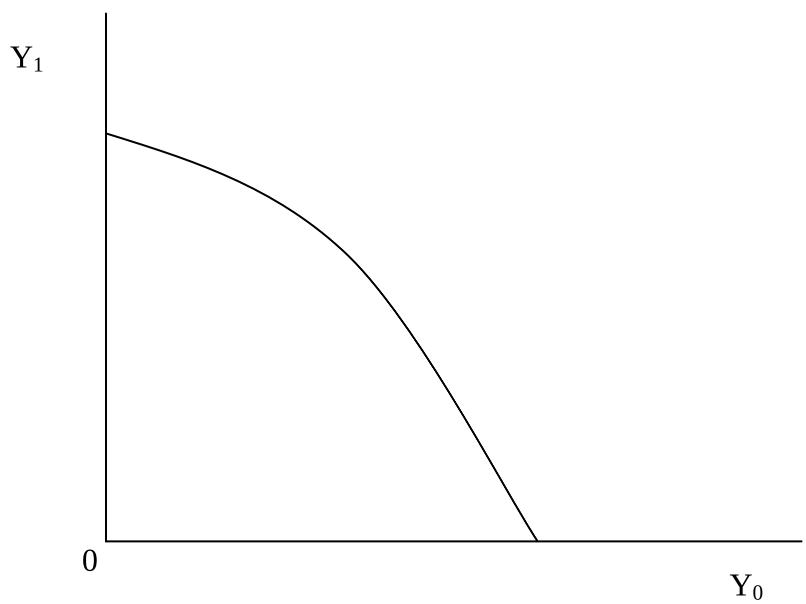

Figure 4
The production possibility curve

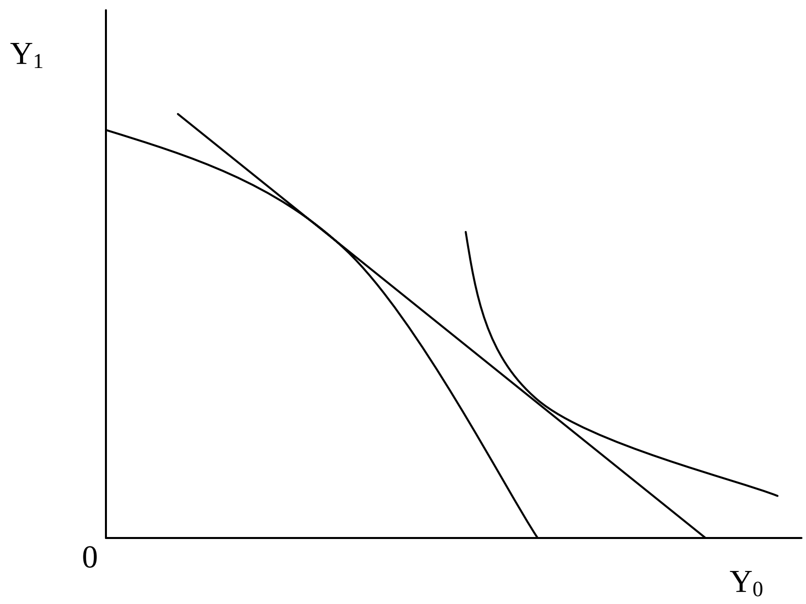

Figure 5
Value maximization and utility maximization

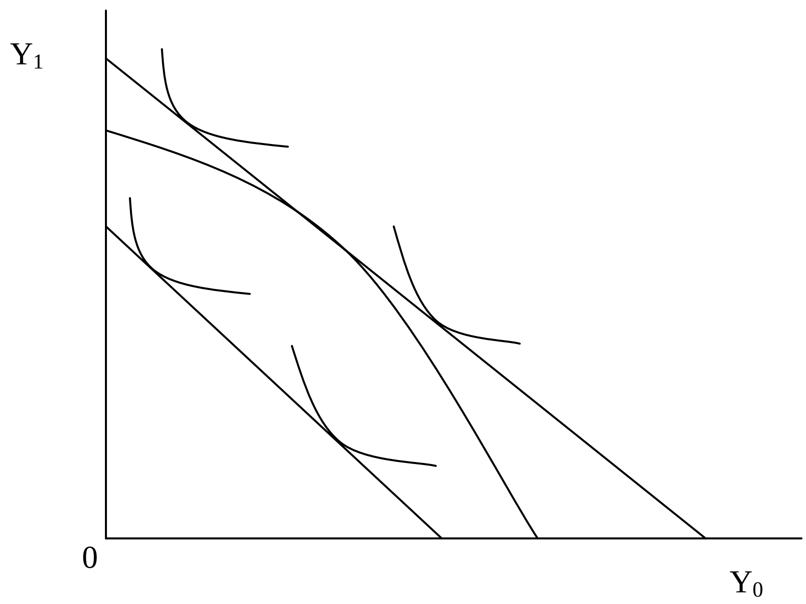

Figure 6
The separation theorem

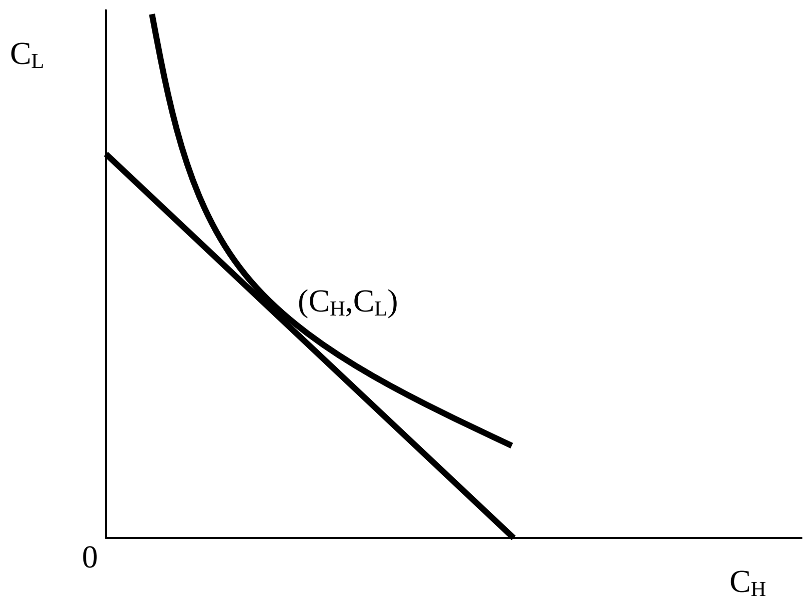

Figure 7
Maximizing utility under uncertainty

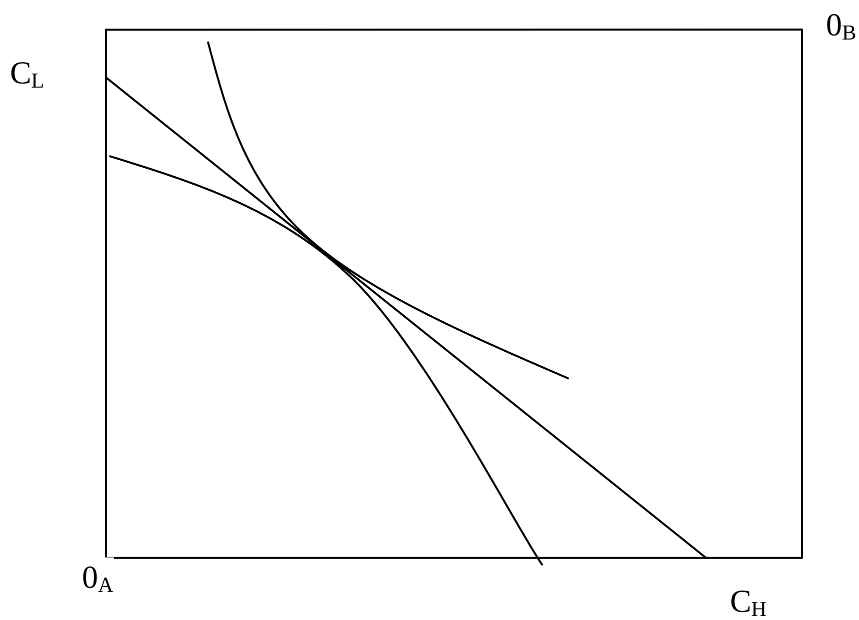

Figure 8
Efficient allocation of risk between two agents

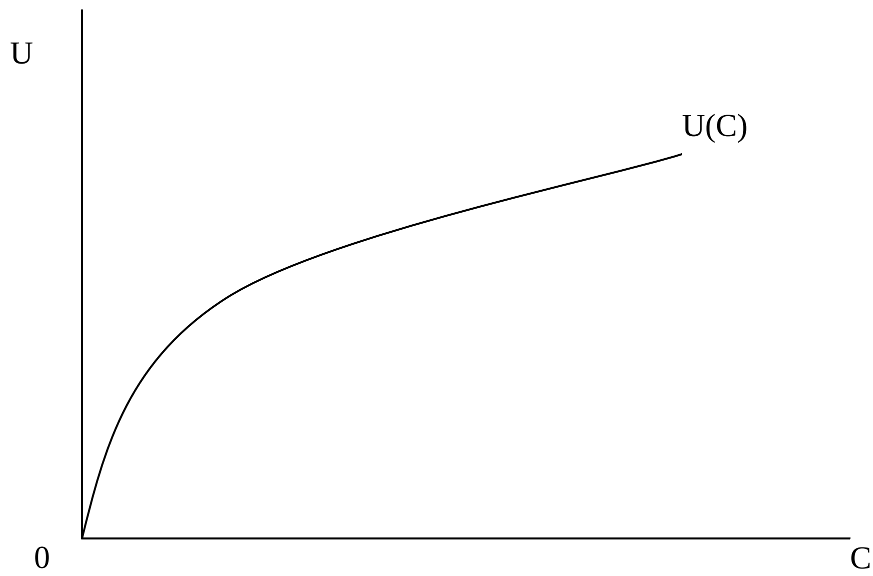

Figure 9
The von Neumann-Morgenstern utility function

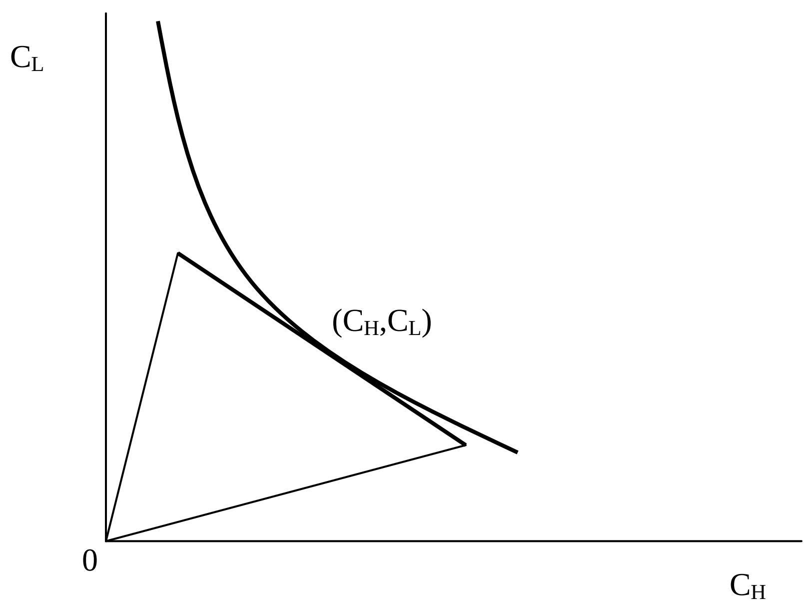

Figure 10
Optimal choice of portfolio with two assets