Chapter 10 develops a theory of market equilibrium with risk on the hypothesis that producers hold rational expectations. That is, producers understand the way in which risk affects the outcome of their actions, and they use this knowledge in choosing their actions. Chapter 11 examines alternative theories of expectation formation, and distinguishes two different kinds of benefits which may flow from policies like price stabilization. Such policies may improve the quality of information available to decision-makers, who are led to make more rational decisions, and they may change the allocation of resources and hence the efficiency of equilibrium. If agents remain in rational expectations equilibrium before and after the policy change, then only the second kind of benefit is relevant, so the concept of rational expectations focuses attention on this kind of benefit.

This part of the book produces a logically self-contained theory of supply and demand under risk for a simple economy in which there is no choice between alternative crops, no futures or credit markets, and no storage. It lays the foundation for the subsequent parts in which these complications are introduced and analysed.

P. M. G. Newbery and J. E. Stiglitz
The Theory of Commodity Price Stabilization. Clarendon Press Oxford 1981

Chapter 5

Competitive Supply with Risk-Neutral Farmers

5.1 Introduction

In order to study the impact of commodity price stabilization we need a theory of why prices fluctuate and what effect risk has on farmers' decisions. The previous chapter discussed the various sources of price variability, while this and the following chapter analyses the farmer's choice of supply under risk and shows how his choice depends upon the source of the price variability. One popular method of describing the effect of risk on supply is to start from a theory which relates supply to the price of output, and then to argue that because risk makes production less attractive, the effect of risk can be captured by adjusting the return to production (the output price) downwards by a risk premium, this premium being the amount needed to compensate the farmer for undertaking the risk. Put another way, this approach would work in terms of a *certainty equivalent* price, that is, the perfectly certain price which would yield the same choices in the absence of risk as the farmer actually makes in the presence of risk. If such a price could be simply derived from the average price via the risk premium, then it would appear that the deterministic theory of supply could be simply translated into a theory of supply under risk.

We shall argue in the next two chapters that this is an unsatisfactory and potentially misleading approach to the study of risk, and needs to be replaced by a rigorous theory of behaviour under risk. One simple way to demonstrate this is to consider the behaviour of farmers who are risk neutral – that is, farmers who are indifferent between a risky return and a safe return yielding the same expected value. On the previous argument, for these farmers the risk premium should be zero, and hence we can ignore the effect of risk. We shall show in this chapter that even if farmers are risk neutral, risk nevertheless may have an important effect on their behaviour. In particular, the simple graphical methods of analysis which relate supply to price are quite misleading, so that the traditional Waugh–Oi–Massell analysis discussed in Chapter 2 rests on shaky foundations which need replacing.

The second objective of this chapter is to prepare for the fuller analysis of risk-taking when agents are not risk neutral, which we defer to the next chapter. We define precisely what is meant by the notion of a certainty equivalent price and assess the validity of this popular approach to the analysis of risk. We shall argue that it is important to distinguish between the *action certainty equivalent* price and the *utility certainty equivalent* price. The first is relevant to predicting the effects of risk on the level of supply, while the second is relevant for a welfare analysis of risk. We shall show that they may differ, and that they depend quite sensitively on the form and source of risk. In the next chapter, we shall further show that they also depend on attitudes to risk when farmers are not risk neutral, as assumed in the present chapter.

The first section of this chapter briefly reviews the theory of supply without risk, to provide both a bench-mark, and to introduce certain basic concepts which will be used repeatedly in what follows – concepts such as the production function, cost and profit functions, convexity and concavity. Readers already familiar with these may proceed directly to section 5.3 which presents the theory of supply for risk-neutral farmers.

5.2 Competitive supply without risk

Each farmer is assumed to have full information about his production possibilities, summarized by a production function relating inputs, **x**, to output, q, and about the prices of inputs, **w**, and output, p. It will be convenient to assume that the production function is differentiable and *concave*. Thus, if the production function is

$$q = f(\mathbf{x}), \qquad \mathbf{x} = (x_1, x_2, \ldots, x_n) \tag{5.1}$$

where **x** is a vector of inputs (land, labour, fertilizer, seed, tractor services, etc.), then f is concave if for any pair of input bundles $\mathbf{x}^1$ and $\mathbf{x}^2$

$$f(\lambda\mathbf{x}^1 + (1-\lambda)\mathbf{x}^2) \geqslant \lambda f(\mathbf{x}^1) + (1-\lambda) f(\mathbf{x}^2), \quad 0 < \lambda < 1. \tag{5.2}$$

Figure 5.1 illustrates this property for the case of a single input, for which the inequality can be written $OB \geqslant OA$. Equation (5.2) is equivalent to the statement

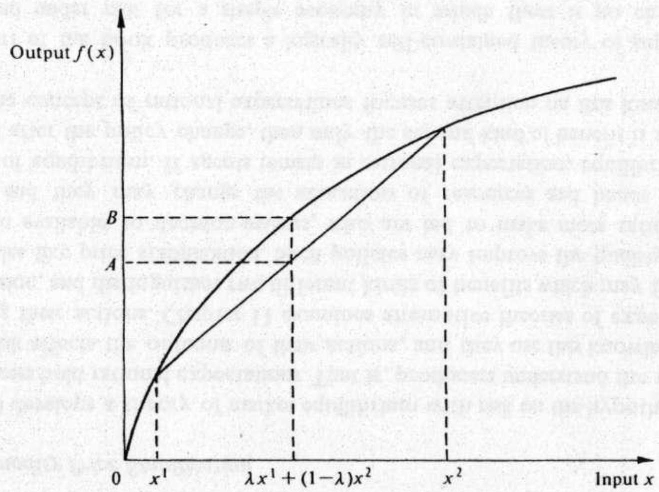

Fig. 5.1 Concave production function

that the chord joining points on the function lies below the function. Strict concavity requires a strict inequality. A *convex* function has the chord above the function, and the inequality reversed. The definition of concavity in equation (5.2) implies that

$$f(\mathbf{x}) \leqslant f(\bar{\mathbf{x}}) + \sum_i (x_i - \bar{x}_i) \frac{\partial f(\bar{\mathbf{x}})}{\partial x_i} \tag{5.3}$$[1]

and, if f is a twice continuously differentiable function of a single variable, it is concave if and only if

$$f'' \leqslant 0.$$

(For a vector function, the condition is that the matrix of second-order partial derivatives be negative semi-definite, while for convex functions the inequalities are reversed, or the matrix must be positive semi-definite.) These and other related properties of concave and convex functions will be used extensively in the rest of the book. Strict concavity holds if the inequality is strict and is the same as diminishing returns to scale, while concavity includes the case of constant returns to scale. Strict concavity is sufficient to ensure that the profit-maximizing choice of input levels is given by the first-order conditions, while for constant returns, input proportions are determined, but not their scale. Thus, if profits are Y

$$Y = pf(\mathbf{x}) - \sum_i w_i x_i, \tag{5.4}$$

the profit-maximizing choice of inputs x_i is the solution to

$$p\frac{\partial f}{\partial x_i} = w_i, \quad i = 1, 2, \ldots, n \tag{5.5}$$

(assuming an interior solution). Equation (5.5) is the familiar result that the value marginal product of the ith input, x_i, is set equal to its price, w_i. Although this approach is the more fundamental, most textbooks use it as a stepping-stone to the derivation of cost curves. To do this, define a *cost-function*, $C(q, \mathbf{w})$ as the minimum cost needed to produce output q at input prices **w**:

$$C(q, \mathbf{w}) = \min \sum w_i x_i \text{ subject to } f(\mathbf{x}) \geqslant q. \tag{5.6}$$

(It can be shown that the cost function is *concave* in **w**, using the same proof as proposition 1, Chapter 8.)

If the industry is competitive, the farmer will not be able to influence any prices (p, **w**) and will maximize profits:

$$Y = pq - C(q, \mathbf{w}) \tag{5.7}$$

when output q is such that marginal cost is equal to price:

[1] Take $\mathbf{x}^1 = x$, $\mathbf{x}^2 = \bar{x}$ in equation (5.2) and let λ tend to zero.

$$p = \frac{\partial C(q, w)}{\partial q} \qquad (5.8)$$

provided that price is above average variable costs, otherwise higher profits are obtained at zero output. If there are diminishing returns everywhere this is guaranteed.

In most textbooks the average cost curve is typically shown as U-shaped, which corresponds to initially increasing returns to scale, perhaps due to indivisibility, followed by diminishing returns. In such cases the marginal cost will intersect the average cost at its lowest point. The distinction between variable or avoidable costs, and fixed or inescapable costs is obvious, but important. If price is below average variable cost, then profits would be increased by avoiding the costs altogether, that is, by closing down this line of production (and shifting the resources to some alternative use, which, in agriculture, usually means producing an alternative crop). Thus to summarize: the supply curve of a competitive firm is the portion of the marginal cost curve above average variable costs.

We can immediately deduce two important additional properties of the supply curve which make it useful in economic analysis. First, the supply of the whole industry is simply the horizontal sum of the individual supply curves and, second, the total cost of industry supply is the area under the supply curve. This follows from the derivation of the supply curve as the marginal cost curve, whose integral (the area below the curve) is clearly the total cost. This implies that profit, or producers' surplus, is the area between the price line and the supply curve, as Fig. 5.2 illustrates.

As drawn, the marginal cost curve is *MEC*, the supply curve is *AEC*, that portion above the average variable cost curve, and total costs are *OAECD* (plus any inescapable, and hence irrelevant, costs).

Finally, we should draw attention to another feature of competitive, riskless markets. Even though the farmer may be interested in the utility of consumption,

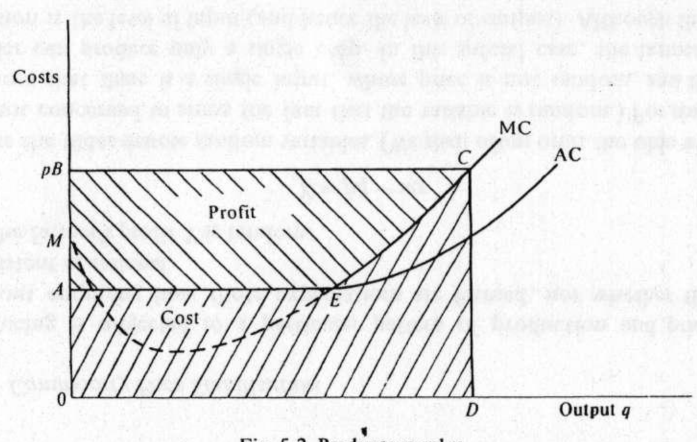

Fig. 5.2 Producer surplus

provided that all inputs and outputs can be traded on competitive markets, he can separate the two problems of making money (profits) and spending them. Even if he consumes the things he grows and some of the inputs he uses (especially his own time), he maximizes his utility by maximizing the cash value of his profits, since the more money he has, the more of every kind of consumption good he can buy. If, on the other hand, some inputs (such as entrepreneurial skill or effort) cannot be purchased on competitive markets, then farmers will not typically wish to maximize profits, but the utility of profits and effort. Thus, if effort is z, the farmer will choose x, z to maximize $U(Y, z)$, where U is his utility function. In this case traded input levels are determined as before:

$$\frac{\partial U}{\partial Y} \cdot \frac{\partial Y}{\partial x} = 0 \quad \text{or} \quad \frac{\partial Y}{\partial x} = 0$$

but effort, z, is found from

$$\frac{\partial U}{\partial Y} \cdot \frac{\partial Y}{\partial z} + \frac{\partial U}{\partial z} = 0.$$

The ratio $-U_z/U_Y$ (where subscripts denote partial derivatives) can be thought of as an implicit price for effort. For a further discussion of this point, see Scitovsky (1943).

To summarize, under riskless, perfect competition, both individual and aggregate supply can be expressed as a function of prices, and, in particular, a supply schedule can be drawn giving output as a function of output price, holding input prices constant. The area below this curve is total cost, and the area between the supply curve and the price is producer surplus, or profit. The farmer is interested in maximizing profits no matter what his consumption preferences are. None of these properties holds generally in the presence of supply risk.

5.3 Competitive supply with risk-neutral farmers

Agriculture is subject to all manner of risks. From the point of view of the individual farmer, these can be divided into two categories:

(i) Production risks: risks which affect his output and which arise because of variations in weather, the prevalence of pests and disease, and other natural causes, such as fire;

(ii) Price risks: risks which affect the prices he receives for the goods he produces or the inputs he plans to purchase (e.g. harvesting labour). Of these, output price risk appears to be more important for the farmer's decision-making and we shall henceforth ignore input price risk. As we noted earlier, price variability may be generated by supply variability or demand variability.

For the market as a whole, price and production variability are intimately connected: variations in output lead to variations in prices, and much of our subsequent analysis is concerned with the precise relationship between the two. For the moment, however, we consider a farmer who believes that the crop he is producing is subjected to a particular pattern of production and price risk, without enquiring how those expectations are formed, nor whether they are consistent or rational.

The farmer's profit Y is random:

$$\tilde{Y} = \tilde{p}\tilde{q} - wx \quad (5.9)$$

where the tildes denote random variables. (We shall often omit the tilde when we are not concerned to stress the fact that the variable is random.) For simplicity suppose that there is a single input, whose price is not random, and that the farmer can produce only a single crop. In this special case, the farmer's only decision is the level of input (and hence the level of output). Although there was no ambiguity about what the correct objective of the farmer was in the riskless situation (he maximizes profit, Y), there is considerable controversy about the farmer's objective in the presence of risk.

The simplest hypothesis from an analytical point of view is that instead of maximizing profits farmers maximize expected (or average) profits, which we write as $E\tilde{Y}$. In certain circumstances, for example if farmers are wealthy, have widely diversified crops, and access to capital markets, or if they can hedge most of their risks (which we shall describe in greater detail in Chapter 13) this is a plausible objective.

A farmer who maximizes $E\tilde{Y}$ is said to be *risk neutral*; he neither seeks to avoid risk (he is not risk averse) nor does he seek after risk (he is not a risk lover, a gambler). Even though the individual is risk neutral, risk may have important effects on his behaviour. The farmer does *not* maximize

$$Eq \cdot Ep - wx = \bar{p}\bar{q} - wx.$$

(Bars over the variable will be used systematically to denote the expected or average value of that variable.) Unless price and output are uncorrelated, expected revenues are not just the product of mean output and mean price. If, as is likely, price and output are correlated, then expected profits are

$$E\tilde{Y} = \bar{p}\bar{q} + \mathrm{Cov}(p, q) - wx \quad (5.10)$$

where the covariance of p with q is defined as

$$\mathrm{Cov}(p, q) = E(p - \bar{p})(q - \bar{q}) = Epq - \bar{p}\bar{q}.$$

If, for instance, there were a negative correlation between this particular farmer's output and the market price, were he to ignore the covariance term, he would overestimate the return to increasing his output.

It is sometimes convenient to refer to the *action certainty equivalent price* – the price which, if it prevailed on the market, and if there were no risk, would yield exactly the same supply response as does the random price. This is not the only useful certainty equivalent concept. It may be contrasted with the *utility certainty equivalent price*, the price which would generate the same level of expected utility in the absence of risk. The two are not in general equivalent, but

they may be when the individual is risk neutral, as we shall see.

To calculate the certainty equivalent price we need to be somewhat more precise in our specification of the production variability, and the relationship between inputs and outputs. It is convenient to write the production function as

$$q = f(\mathbf{x}, \tilde{\theta}, \xi). \quad (5.11)$$

Output is a function of the inputs, $\mathbf{x}$, the state of nature which is described by the random variable $\tilde{\theta}$, e.g. weather, the incidence of pests and diseases, etc., and the choice of technique of production ξ, e.g. the timing of planting, harvesting, etc., all of which may have an important effect on the variability of output.

This is, however, too general for analytical and practical purposes. Two specifications have been used extensively in theoretical work and both are natural counterparts of alternative econometric specifications. For any given technique we have:

(i) *Multiplicative risk*

$$q = \tilde{\theta}f(x), \quad E\tilde{\theta} = 1, \quad \mathrm{Var}\,\tilde{\theta} = \sigma^2. \quad (5.12)$$

Rain at harvest times leads to spoilage which is a constant fraction of the crop, regardless of its size, or disease affects a fraction of the crop. If all farms within a particular area face the same risks, then total supply will also experience multiplicative risk:

$$\tilde{Q} = \tilde{\theta}\bar{Q}, \quad \bar{Q} = \Sigma f(x).$$

(Here average total supply is the sum of average individual outputs.) Econometrically such functions are estimated logarithmically, in which case it is typically assumed that θ is log normal: that is for $\log\theta$ to be normal.

(ii) *Additive risk*

$$q = f(x) + \tilde{\theta}, \quad E\tilde{\theta} = 0, \quad \mathrm{Var}\,\tilde{\theta} = \sigma^2. \quad (5.13)$$

Rain destroys a constant amount regardless of the size of the total crop; disease wipes out a limited area of the crop independent of the total area. This specification is attractive to econometricians wishing to estimate linear supply functions, but it is difficult to justify on theoretical grounds. It is difficult to see how to aggregate to obtain total supply without making total risk proportional to total output – in which case we are back with multiplicative risk. In our view, multiplicative risk seems a better approximation than additive risk, especially for a microeconomic theory of an individual farmer's decisions. Additive risk is at best a simplification used at the aggregate level for econometric estimation. Unfortunately, this simplification is bought at a price, for, as we shall show later, a number of results about the effect of commodity price stabilization schemes depend critically on which assumption is made.

Since price variability is affected by supply variability, we need a simple theory to model this dependence. In Chapter 10 the relationship is examined more carefully, but for the moment it will be enough to suppose that the price depends on two sources of variability

$$\tilde{p} = \bar{p} + \beta(\tilde{\theta} - \bar{\theta}) + \tilde{\nu}, \quad E\tilde{\nu} = 0, E\tilde{\nu}\tilde{\theta} = 0. \quad (5.14)$$

The first source of risk is supply risk, and all remaining risk (that is, risk which is uncorrelated with supply risk) is placed in the residual term, $\tilde{\nu}$. For example, if the demand schedule is linear: $p = a - bQ$, supply risk is multiplicative, and if there is additional, independent, additive demand risk $\tilde{\nu}$, then the equation holds exactly, with $\beta = -b\bar{Q}$, where $-b$ is the slope of the inverse demand schedule. If supply risk is *multiplicative*, risk neutral producers maximize expected profit.

$$E\tilde{Y} = E\{\bar{p} + \beta(\tilde{\theta} - 1) + \tilde{\nu}\}\tilde{\theta}f(x) - wx, E\tilde{\theta} = 1,$$

$$= \hat{p}f(x) - wx,$$

where the certainty equivalent price is $\hat{p}$:

$$\hat{p} = \bar{p} + \beta\sigma^2 \quad (5.15)$$

or

$$\hat{p} = \bar{p} + \frac{\mathrm{Cov}(p, q)}{\bar{q}}.$$

(We have normalized θ to have mean unity, coefficient of variation σ.) It follows that

$$\hat{p} \gtrless \bar{p} \text{ as } \beta \gtrless 0, \quad \text{i.e. as } \mathrm{Cov}(p, q) \gtrless 0.$$

It is clear that in this case the action certainty equivalent price and the utility certainty equivalent are identical. In particular, the certainty equivalent price exceeds or is less than the mean price as price is positively or negatively correlated with output.

If supply risk is *additive*, producers maximize

$$E\tilde{Y} = E\{\bar{p} + \beta\tilde{\theta} + \tilde{\nu}\}\{f(x) + \tilde{\theta}\} - wx, E\tilde{\theta} = 0$$

$$= \bar{p}f(x) + \beta\sigma^2 - wx. \quad (5.16)$$

In this case the action certainty equivalent price *is* the mean price though average profits are affected by risk and hence the utility certainty equivalent is greater or less than the mean price as $\beta \gtrless 0$. The specification of risk evidently has an important bearing on the choice of supply.

We can obtain an expression for $\hat{p}$ in terms of observable parameters. In what follows σ_x is to be interpreted as the coefficient of variation of the variable x. With multiplicative supply risk we have, from equation (5.12)

$$\sigma_q = \sigma,$$

and from (5.14)

$$\mathrm{Var}(p) = \bar{p}^2\sigma_p^2 = \mathrm{Var}\,\nu + \beta^2\sigma_q^2.$$

If $\beta < 0$, then this can be substituted into (5.15) and solved to give

$$\hat{p} = \bar{p}\left\{1 - \frac{\sigma_q}{\bar{p}}\sqrt{(\bar{p}^2\sigma_p^2 - \mathrm{Var}\,\nu)}\right\}. \quad (5.17)$$

If, for instance, there is no pure demand risk, so $\mathrm{Var}\,\nu = 0$, then

$$\hat{p} = \bar{p}(1 - \sigma_p\sigma_q).$$

In Chapter 10 we derive explicit expressions for the certainty equivalent price for particular demand functions. For example, if demand is stable and of constant elasticity ϵ, then equation (10.19) gives the action certainty equivalent price as approximately

$$\hat{p} = p(\bar{Q})\{1 + \tfrac{1}{2}(1 - \epsilon)\sigma_p^2\}, \quad (5.18)$$

where $p(\bar{Q})$ is the price when quantity is $\bar{Q}$.

5.3.1 *Pitfalls in graphical analysis*

In the first section we showed that in the absence of risk supply could be graphed as a function of price, independently of demand, to find the equilibrium output and level of profits. Many authors have been tempted to use similar graphical methods in the presence of risk. Even if we assume risk-neutral producers (the most favourable case) average supply is a function of the certainty equivalent price, which cannot be found independently of the demand schedule. In Fig. 5.3 inputs are chosen at the start of the year, and there is no demand risk.

The representative farmer will choose an expected level of supply $\bar{q}$ (corresponding to total supply $\bar{Q}$), such that the marginal cost of producing $\bar{Q}$ is equal

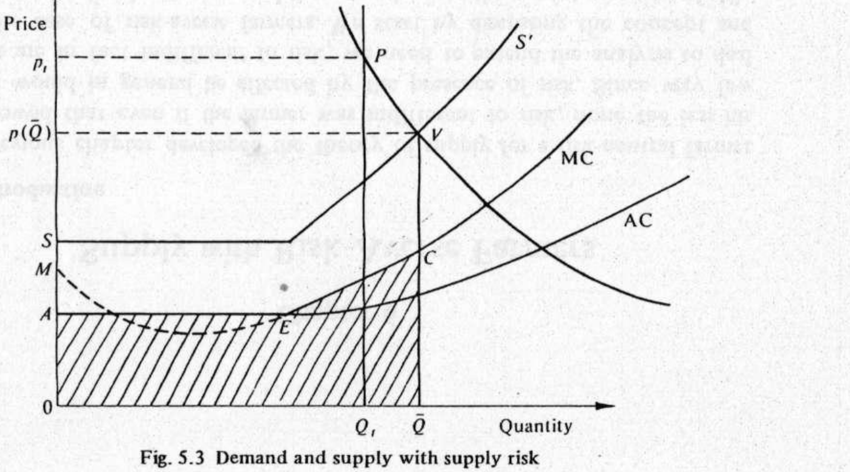

Fig. 5.3 Demand and supply with supply risk

to the action certainty equivalent price, or

$$MC = \hat{p}.$$

If the demand schedule has constant elasticity, this can be written, using equation (5.18), as

$$p(\bar{Q}) = \frac{MC}{1 - \frac{1}{2}(\epsilon - 1)\sigma_p^2} = (1 + m)MC,$$

where

$$m = \frac{\frac{1}{2}(\epsilon - 1)\sigma_p^2}{1 - \frac{1}{2}(\epsilon - 1)\sigma_p^2}$$

and is graphed as SVS'. (The curves MC and AC are riskless, and are derived as for the riskless case of section 5.2.) SVS' is a long-run pseudo supply curve, derived by a 'mark up' m, equal to the ratio $VC/C\bar{Q}$ in Fig. 5.3 on MC. The average total supply will be $\bar{Q}$, where SVS' intersects the demand curve. The short-run supply curve will be vertical and fluctuate around $\bar{Q}$, assuming positions such as Q_tP at date t, which leads to a market clearing price of p_t. Costs will be the area under the non-random MC curve, $OAEC\bar{Q}$, while revenue will be random, for example, Op_tPQ_t at date t. Profit will be the difference between random revenue and non-random cost, and not the area between the price line and the (long-run pseudo) supply schedule SVS'.

To summarize, the concept of the certainty equivalent price is a useful one for describing how farmers make their decisions. But it is one to be used with caution, particularly when we are concerned with policies (such as commodity price stabilization schemes) which will affect the distribution of prices. Changes in certainty equivalent prices are in general not equal to (nor even proportional to) changes in mean prices. Different policies may have differential effects on mean prices and on certainty equivalent prices. Moreover, the certainty equivalent price for one farmer may be quite different from that for another farmer. Farmers in a region which is the primary source of a particular commodity (for example West African cocoa farmers) may have a high negative correlation between output and quantity (and hence a low certainty equivalent price), while farmers in a small region (Brazilian cocoa farmers) far away from the main supply region may face a nearly zero correlation between output and quantity. Then, for risk-neutral farmers, a change in the price distribution for the latter is only important in so far as the mean price is changed, whereas for the former, effects on the covariance are crucial, as we shall see later.

Chapter 6

Supply with Risk-Averse Farmers

6.1 Introduction

The previous chapter developed the theory of supply for a risk-neutral farmer and showed that even if the farmer was indifferent to risk, none the less his actions would in general be affected by the presence of risk. Since very few farmers are in fact indifferent to risk, we need to extend the analysis to deal with the case of risk-averse farmers. We start by discussing the concept and measurement of risk aversion, and then enquire into the precise meaning of risk. The remainder of the chapter applies these concepts to a variety of problems which will recur throughout the book. The eventual objective is to develop a method of analysing the effects of a change in the distribution of prices resulting from a commodity price stabilization scheme. This turns out to be a fairly difficult question, and it is helpful to break the problem down into a number of more manageable questions, some of which we discuss in the later sections of this chapter.

In section 6.4 we analyse the effect of changes in risk on the behaviour of producers. In the following section we examine the important special case of the mean-variance model, which has been widely employed because of its analytical simplicity. In this model the analysis of the effects of risk on farmer's behaviour is remarkably simple; unfortunately, as we shall argue, the conditions under which the model may reasonably be used are very restrictive. However, we also show that for many problems the impact of risk can be approximated by a Taylor series expansion in which only the mean and variances of the distribution appear, and we shall discuss when this approximation is legitimate.

While these two sections discuss the effect of changes of risk on the *behaviour* of producers, the final section is devoted to analysing the impact on the welfare of producers. It derives formulae which will be refined and applied later in assessing the costs and benefits of price stabilization. Of course, as this is an introductory chapter it necessarily simplifies, but it lays the foundation for the more comprehensive analysis to be developed later.

6.2 The meaning of risk aversion

Economists have for a long time modelled consumer behaviour on the assumption that the consumer's preferences between goods can be represented by an ordinal utility function defined over these goods. This is possible if the consumer has stable preferences and if he is rational, that is, consistent, in making choices. Given enough observations, a skilful econometrician would be able to derive this utility function by observing the consumer's choices at different prices and levels of money income, and then use it to predict the consumer's behaviour once his income and prices have been specified. (In Chapter 8 we summarize this theory and derive some of its implications.) If we also assume that the consumer is well informed about the consequences of his choices (which in any case is almost required if his preferences are to remain stable), and if we assume that the consumer is concerned with his own satisfaction, then we can draw certain welfare conclusions of the form 'if the consumer chooses A rather than B when both choices are feasible, then his welfare or satisfaction is higher with A than B'. If, further, we accept an individualist welfare ethic, then certain normative consequences follow, and the utility function can be used not only to describe behaviour, or for prediction, but also for welfare analysis and to evaluate social choices.

So much is basic in elementary welfare economics, and it is natural to extend this reasoning to choices involving risk. Most agents would prefer an action which has a sure return, Y, to another action which yields a risky return with the same expected value, and it seems reasonable that a rational agent should be able to compare alternative risky choices, balancing gains against risk, just as the rational consumer compares alternative baskets of goods. Indeed, under certain assumptions discussed in Chapter 8 agents will act as though they had a utility function, defined over the consequences of their choices, and will choose the action which maximizes the expected value, not of the outcome, but of the utility of the outcome. We thus postulate that individual behaviour in the face of risk can be described as if the individual

$$\text{maximizes } EU(\tilde{Y})$$

the expected value of the utility, U, of the risky outcome, $\tilde{Y}$. Essentially, agents need to know the consequences of their choices, to have beliefs about the probabilities of these consequences, and to be concerned only with the consequences (and not with the process by which they are brought about). Just how reasonable these assumptions are will be left until Chapter 7 when we examine some of the empirical evidence, but for the moment we shall explore some of the consequences of this approach. Since we are interested in the choices of producers, we shall suppose that they are primarily concerned with their money *income*, which they are free to spend on commodities whose prices are fixed. To keep the story simple we assume that the world comes to an end after the farmers have spent their income and enjoyed the consumption allowed. We shall discuss the strength of this assumption in Chapter 7 and the consequences of relaxing it in Chapter 14.

We shall proceed in our discussion as follows. First, we shall define what we mean by risk aversion and relate it to the properties of the utility function $U(Y)$. We shall then discuss a more precise quantification of risk aversion. Finally, we shall present several specific utility functions which have played a major role in the recent literature and which we will find useful in the subsequent analysis.

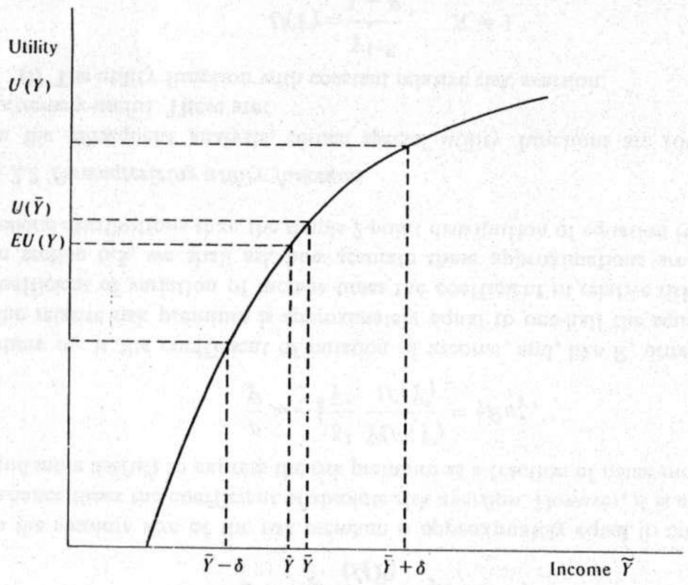

Fig. 6.1 The value of risky income

For a risk-averse individual, the utility function $U(Y)$ appears as in Fig. 6.1, where U is concave. To see that this does in fact correspond to our intuitive notions of risk aversion, calculate the expected utility associated with a random income

$$\tilde{Y} = \begin{cases} \bar{Y} + \delta & \text{with probability } \tfrac{1}{2} \\ \bar{Y} - \delta & \text{with probability } \tfrac{1}{2} \end{cases} \tag{6.1}$$

The expected utility is given by

$$EU(\tilde{Y}) = \tfrac{1}{2}\{U(\bar{Y} + \delta) + U(\bar{Y} - \delta)\},$$

i.e. is half-way between the two utility levels. But note from the diagram that with a concave utility function, this is less than $U(\bar{Y})$, the utility associated with the sure income of $\bar{Y}$.

The difference between the two is a measure of the cost of the risk in terms of the loss of expected utility. We can also measure this cost by asking how much of his sure income would he be willing to give up, and still prefer the sure income to the risky income. That is, what sure income is equivalent (in the utility that it yields) to the random income. In the diagram, $\hat{Y}$ gives the same utility, and is referred to as the *certainty equivalent* income. It can be defined formally by the equation

$$EU(\tilde{Y}) = U(\hat{Y}). \qquad (6.2)$$

The difference between the mean income $\bar{Y}$, and its certainty equivalent is sometimes referred to as the *risk premium* (or the cost of the risk):

$$\rho = \bar{Y} - \hat{Y}. \qquad (6.3)$$

The magnitude of the risk premium (the cost of the risk) can be related to the shape of the utility function and the probability distribution function of returns. We would expect that an increase in risk would increase the risk premium and so would an increase in risk aversion. There are simple representations for both of these. If, for instance, we increase δ, clearly this is an increase in riskiness. From Fig. 6.1 we immediately see that this does increase the size of the risk premium (it reduces the certainty equivalent income).

Similarly, greater risk aversion is associated with a more 'curved' utility function. In the limiting case of utility function which is a straight line ($U'' = 0$) there is no risk aversion (we call such an individual risk neutral and discussed his behaviour in the previous chapter). The risk premium is identically zero, regardless of the size of the risk.

These concepts have been made more precise in a series of papers (by Arrow, 1965; Pratt, 1964; Rothschild and Stiglitz, 1970, 1971; and Diamond and Stiglitz, 1974). Rather than repeat their detailed derivations, we shall simply summarize their results and attempt to provide an intuitive motivation for them.

6.2.1 *Measuring risk aversion*

Since greater risk aversion is associated with a more curved utility function, it is natural to relate risk aversion to the curvature of the utility function. One simple measure of this is the elasticity of marginal utility, or the *coefficient of relative risk aversion*, R, defined as

$$R(Y) = -\frac{YU''(Y)}{U'(Y)} \qquad (6.4)$$

and evaluated at some chosen level of income, Y. As an elasticity it is dimensionless, and hence a very convenient way in which to describe risk aversion. The other simple measure is the *coefficient of absolute risk aversion*, A, defined as

$$A(Y) = -\frac{U''(Y)}{U'(Y)}. \qquad (6.5)$$

It is not dimensionless and depends on the units in which income is measured. Therefore, if risk aversion is to be described by the numerical measure of the coefficient of absolute risk aversion, the income level must also be given to make sense of the value. Notice that the two measures are trivially related:

$$R(Y) = YA(Y). \qquad (6.6)$$

To show that these measures are appropriate and useful, observe that $U(Y)$ can be expanded in a Taylor series. If $Y = \bar{Y} + h$, then

$$U(Y) = U(\bar{Y}) + hU'(\bar{Y}) + \frac{h^2}{2}U''(\bar{Y}) + r_3(h) \qquad (6.7)$$

where r_3 is a remainder and r_3/h^2 tends to zero, as h tends to zero. If Y is now the random variable defined in equation (6.1), h is a random variable taking values $\pm\delta$ with equal probability, so the expected value $EU(Y)$ is found by taking the expectation of equation (6.7):

$$EU(\tilde{Y}) = U(\bar{Y}) + \tfrac{1}{2}\delta^2 U''(\bar{Y}) + Er_3(\delta). \qquad (6.8)$$

The certainty equivalent income defined in equations (6.2) and (6.3) can likewise be expressed in a Taylor series:

$$U(\hat{Y}) \equiv U(\bar{Y} - \rho) = U(\bar{Y}) - \rho U'(\bar{Y}) + r_2(\rho), \qquad (6.9)$$

where again r_2/ρ tends to zero with ρ. If δ is small, then the remainders can be ignored, and since by definition

$$EU(\tilde{Y}) = U(\hat{Y}),$$

it follows that the risk premium is approximately

$$\rho \approx -\tfrac{1}{2}\delta^2 \frac{U''(\bar{Y})}{U'(\bar{Y})} = \tfrac{1}{2}A \text{ Var } Y, \qquad (6.10)$$

so the absolute size of the risk premium is approximately equal to one-half the variance times the coefficient of absolute risk aversion. However, it is more usual (and more useful) to express the risk premium as a fraction of mean income:

$$\frac{\rho}{\bar{Y}} \approx -\tfrac{1}{2}\frac{\delta^2}{\bar{Y}^2}\frac{\bar{Y}U''(\bar{Y})}{U'(\bar{Y})} = \tfrac{1}{2}R\sigma_Y^2, \qquad (6.11)$$

where σ_Y is the coefficient of variation of income, and, like R, dimensionless. The relative risk premium is approximately equal to one-half the square of the coefficient of variation of income times the coefficient of relative risk aversion. In section 6.5, we shall ask how accurate these approximations are for more general distributions than the simple 2-point distribution of equation (6.1).

6.2.2 *Paramaterizing utility functions*

In the subsequent analysis, certain special utility functions are found to be extremely useful. These are:

(i) The utility function with constant relative risk aversion:

$$U(Y) = \frac{Y^{1-R}}{1-R}, \quad R \neq 1 \qquad (6.12)$$

$$U' = Y^{-R}$$

$$\frac{-YU''}{U'} = R.$$

A special case of this class of utility functions is the logarithmic or Bernoullian utility function,

$$U(Y) = \log Y : R = 1$$

which has unit relative risk aversion. For this class of utility functions the degree of relative risk aversion is equal to the elasticity of marginal utility, and the functions are usually described as constant elasticity utility functions. For these functions the *proportional* risk premium is independent of the level of wealth, Y.

(ii) The utility function with constant absolute risk aversion:

$$U(Y) = -ke^{-AY} \tag{6.13}$$

$$-U''/U' = A.$$

This class is often referred to as the exponential utility function. If the outcomes Y are normally distributed then the absolute risk premium is independent of the level of wealth – an implausible property which is nevertheless very useful in solving portfolio problems such as the choice of hedge in a futures market. Its properties are discussed more fully in section 6.5 and in Chapter 13.

(iii) The quadratic utility function:

$$U(Y) = -(a - bY)^2 \tag{6.14}$$

$$-U''/U' = \frac{1}{a/b - Y}.$$

The quadratic utility function has several attractive features – it has linear marginal utility (useful in solving dynamic buffer stock problems, as in Chapter 30) and it allows expected utility to be expressed in terms of the mean and variance of income alone. However, it has the obvious limitation that utility decreases with income beyond a certain level, a/b.

All three types of utility functions are members of a more general class which satisfy the equation

$$A(Y) = -U''/U' = \frac{1}{\alpha + \beta Y}. \tag{6.15}$$

Thus when $\alpha = 0$, $\beta = 1/R$ we have constant relative risk aversion; when $\beta = 0$, $\alpha = 1/A$, constant absolute risk aversion, and when $\beta = -1$, $\alpha = a/b$, the quadratic. Cass and Stiglitz (1970) describe the special properties of this wider class of functions.

The most important of these special properties as far as our later analysis is concerned is the linearity of asset demand functions with respect to wealth. Consider the case in which there is a risky asset yielding a return (per dollar invested) of $\tilde{r}$, and a safe asset with a return of s per dollar. Then if W_0 is the individual's initial wealth and Z is the amount invested in the risky asset, his wealth at the end of the period (when the risk has resolved) will be $\tilde{Y}$:

$$\tilde{Y} = \tilde{r}Z + s(W_0 - Z) + W_0. \tag{6.16}$$

The individual chooses Z to maximize expected utility, $EU(\tilde{Y})$, for which the first-order condition is

$$EU'(\tilde{Y})(\tilde{r} - s) = 0. \tag{6.17}$$

For the constant absolute risk aversion utility function of equation (6.13) this can be written as

$$kAe^{-A(1+s)W_0} \cdot \mathrm{E}(e^{-A(\tilde{r}-s)Z}(\tilde{r} - s)) = 0.$$

This has a solution Z^*, which satisfies

$$Ee^{-A(\tilde{r}-s)Z^*}(\tilde{r} - s) = 0 \tag{6.18}$$

and which is independent of wealth, W_0, though it does depend on the distribution of $\tilde{r} - s$. Later on, we shall be particularly interested in the special case in which r is normally distributed. In this case Y, which is a linear function of r, is also normally distributed. Expected utility can now be written as

$$EU = -kEe^{-A\tilde{Y}} = -k\exp\{-A\bar{Y} + \tfrac{1}{2}A^2E(\tilde{Y} - \bar{Y})^2\}. \tag{6.19}$$

(This result follows immediately from the definition of the moment-generating function of the normal distribution.) Consequently, maximizing expected utility of Y when Y is normally distributed is equivalent to maximizing

$$EY - \tfrac{1}{2}A \operatorname{Var} Y,$$

which is the maximand in the mean-variance model of asset demand discussed in section 6.5. In this particular case the solution Z^* must satisfy

$$Z^* = \frac{E(\tilde{r} - s)}{AE(\tilde{r} - s)(\tilde{r} - \bar{r})}. \tag{6.20}$$

Z is the dollar expenditure on the risky asset, and if its price is p per unit, then Z/p is the number of units purchased. Let X be the return per unit, so

$$r = X/p$$

$$\frac{Z^*}{p} = \frac{E(\tilde{X} - sp)}{AE(\tilde{X} - sp)(\tilde{X} - \bar{X})} = \frac{\bar{X} - sp}{AE(X - \bar{X})^2}. \tag{6.21}$$

Thus the demand functions for the asset are not only linear in wealth, they are also linear in price, p.

If the utility function exhibits constant relative risk aversion as in equation (6.12), then the solution to (6.17) is

$$Z^* = \alpha^* W_0$$

where α^* satisfies

$$EW_0^{-R}\{\alpha\tilde{r} + (1 - \alpha)s + 1\}^{-R}(\tilde{r} - s) = 0. \tag{6.22}$$

For the quadratic utility function of equation (6.14) Z^* is given by

$$Z^* = \beta + \alpha W_0 \tag{6.23}$$

where α and β solve

$$E[b(\tilde{r}-s)\beta + W_0 b\{\alpha(\tilde{r}-s)+(1+s)\} - a](\tilde{r}-s) = 0,$$

i.e.
$$\beta = \frac{a(\bar{r}-s)}{bE(\tilde{r}-s)^2}, \quad \alpha = -\frac{E(\tilde{r}-s)(1+s)}{E(\tilde{r}-s)^2}. \tag{6.24}$$

Thus, constant absolute risk aversion utility functions have a zero wealth elasticity of demand for the risky asset, constant relative risk aversion utility functions have a constant (positive) elasticity of demand, and quadratic utility functions have linear demand functions, with the demand for the risky asset decreasing with wealth.

It can be shown (see Stiglitz, 1970) that the only utility functions which always yield linear asset demand functions are those satisfying equation (6.15).

Although these parameterizations are extremely useful, it should be emphasized that they do have some special properties (such as, in the context of portfolio analysis, linearity of demand curves for assets as a function of wealth). In the context of poor countries, none of these utility functions adequately captures the large disutility associated with very low incomes (starvation). In such cases we would expect R to increase as income falls, or $dR/dY < 0$.

6.3 Measures of risk

One of the main objectives of this book is to analyse the effects of various stabilization schemes. Clearly, if we completely eliminated price or income instability, the new distribution of prices or income would be less variable than the old. However, for reasons which will become clearer, no stabilization scheme will ever *completely* eliminate risk. Accordingly, we are faced with the difficult task of comparing distributions, both of which are variable.

A natural solution to this problem which suggests itself is to look at some statistical measure of variability, like variance or the range of the distribution. Although this is a reasonable approach, and in many circumstances it may be the only practicable approach, there are certain limitations which need to be borne in mind. First, there are situations where the mean would remain the same, the variance be reduced, and yet expected utility be lowered. This can be seen diagrammatically in Fig. 6.2 where Y takes on three values, $Y_1 < Y_2 < Y_3$. The utility function is piecewise linear with a kink at $\hat{Y}$, where $Y_1 < \hat{Y} < Y_2$. In that case, any change in the distribution which keeps the mean, conditional on Y being greater than $\hat{Y}$, unchanged, leaves utility unaffected. In particular, there are many such changes which reduce variance. It is easy now to consider some further changes in the distribution which further reduce the variance, but which *lower* expected utility. Assume for instance we lower Y_1 a little and compensate

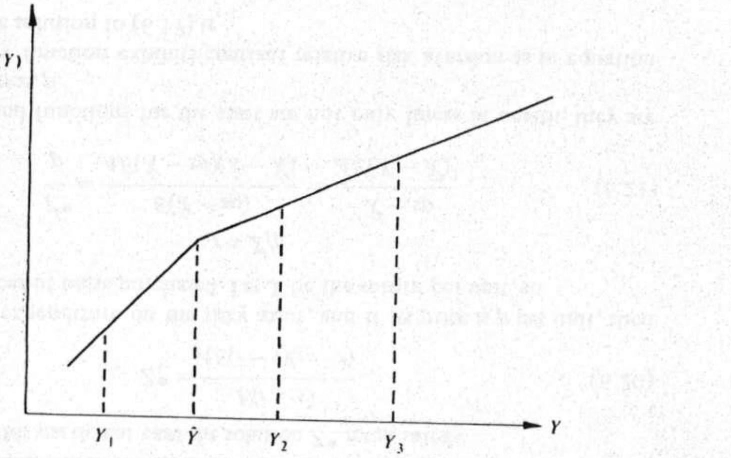

Fig. 6.2 Piecewise linear utility function

for it (to hold the mean constant) by a slightly increased Y_3, and also reduce the probability of either Y_1 or Y_3 occurring (to lower the variance). Clearly, this will make the individual worse off and lower expected utility.

Thus a reduction in variance, keeping mean constant, does not necessarily correspond to an increase in expected utility. Two questions naturally arise. Are there circumstances in which it does? The answer is 'yes', but they are very restrictive: we must either impose restrictions on the utility function or on the probability distribution function: (a) the utility function must be quadratic, or (b) the distribution function of *incomes* must be fully described by its mean and variance. The second condition appears to offer quite a wide range of applications, but the appearance is deceptive, as we show in section 6.5.

The second question which we can ask is: is there a way of ranking distributions which is valid, say, for *all* risk-averse individuals? The answer is 'yes', but we obtain only a partial ordering, that is, we cannot rank all distributions.

Intuitively, if we have two distributions for incomes, denoted by their distribution functions F and G, we can say that F is more variable than G if (i) F could have been derived from G by simply adding noise (that is, by adding an uncorrelated, purely random, term). (ii) F could have been generated from G by taking some probability weight from the centre of the distribution and putting it into two tails, so as to keep the mean constant, as depicted in Fig. 6.3 which shows the density function. ($f(x)$ is the density function if $f(z)dz$ is the probability of x lying between z and $z + dz$. The distribution function is then

$$F(x) = \int_{-\infty}^{x} f(z)dz,$$

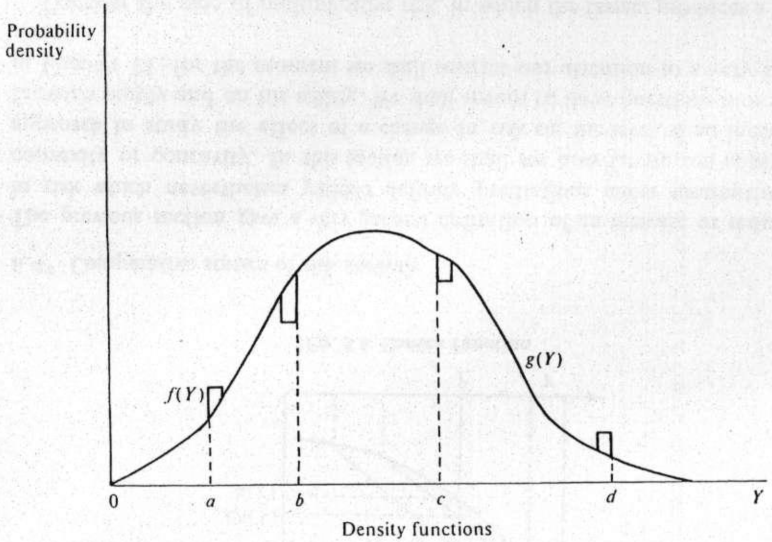

Fig. 6.3 Density functions

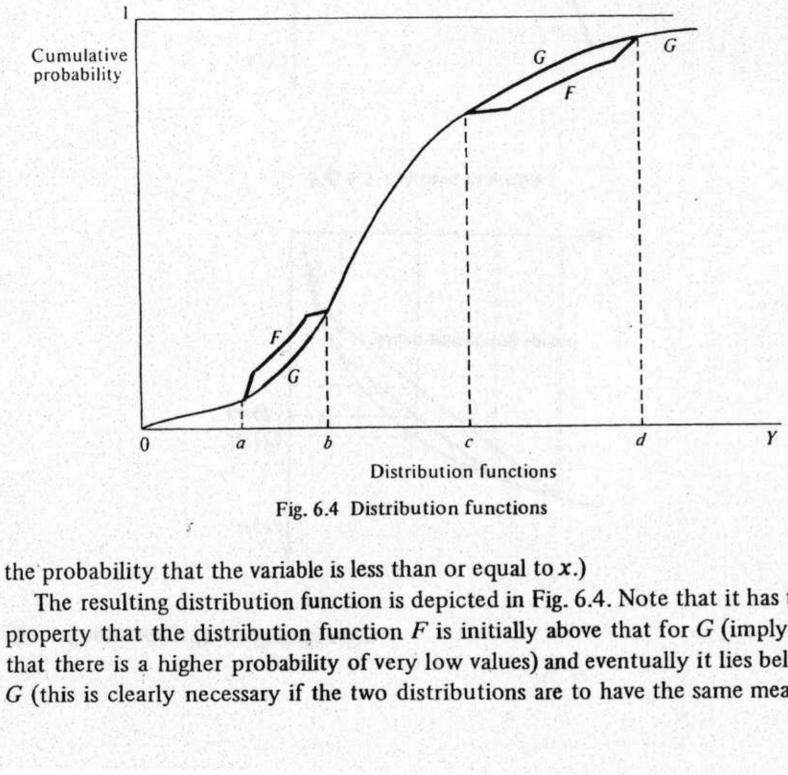

Fig. 6.4 Distribution functions

the probability that the variable is less than or equal to x.)

The resulting distribution function is depicted in Fig. 6.4. Note that it has the property that the distribution function F is initially above that for G (implying that there is a higher probability of very low values) and eventually it lies below G (this is clearly necessary if the two distributions are to have the same mean).

As a slight generalization to that, we can say that F is more variable than G if

$$\int_0^y F(Y)dY \geqslant \int_0^y G(Y)dY, \quad \text{for all } y \tag{6.25}$$

and

$$\int_0^\infty \{F(Y) - G(Y)\}dY = 0 \tag{6.26}$$

$$= [Y(F-G)]_0^\infty + \int_0^\infty Yg(Y)dY - \int_0^\infty Yf(Y)dY$$

(The second condition is simply that the two have the same mean, as the integration by parts confirms, since $[Y(F-G)]_0^\infty = 0$.)

Fortunately, as Rothschild and Stiglitz (1970) have shown, these different approaches are fully equivalent. If F could have been derived from G by simply adding noise, then equation (6.25) is always satisfied (and conversely), and all risk-averse individuals would prefer G to F (and conversely).

We have thus found a method of ranking distributions when the mean of the relevant variable is held constant. It is not, however, always apparent which is the appropriate variable whose mean is to be held constant. One of the important points we emphasize in Chapter 17 is that if we are comparing two price distributions it is not reasonable to hold the mean price constant, because this may not be feasible, but instead it is natural (and feasible) to hold constant the mean quantity sold. We have also found an unambiguous method of describing an increase in risk, for if we add a *mean-preserving spread* to a probability function g the resulting distribution f will be riskier. Moreover, just as a mean-preserving spread lowers expected utility for all concave utility functions, so that risk-averse individuals prefer the original distribution, as Figure 6.5 suggests, so mean-preserving spreads *increase* the expected value of convex functions as in Fig. 6.6. This is why the concept of a mean-preserving spread is so powerful in economics where many functions are known to be either convex or concave.

A closely related result which we shall make extensive use of is *Jensen's inequality*, which states that if $U(Y)$ is a concave function and $h(Y)$ is a convex function, then

$$EU(Y) \leqslant U(\bar{Y})$$
$$Eh(Y) \geqslant h(\bar{Y}).$$

These results follow directly from the definitions of concavity and convexity given in section 5.2 and are again intuitively clear from Figs. 6.5 and 6.6.

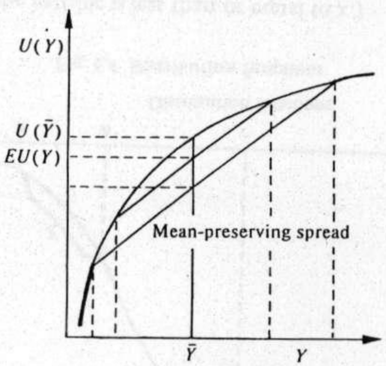

Fig. 6.5 Concave function

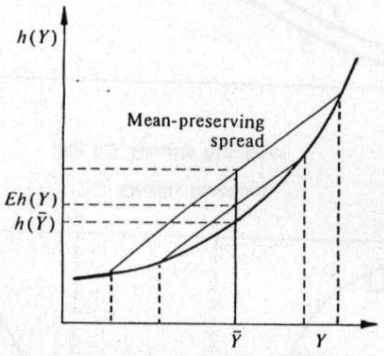

Fig. 6.6 Convex function

6.4* Comparative statics of risk analysis

The previous section gave a very general definition of an increase or reduction in risk which nevertheless yielded definite predictions under assumptions of convexity or concavity. In this section we shall see how far we can apply this approach to study the effect of a change in risk on the level of an individual farmer's supply and on his utility. We shall return to these questions more fully in Chapter 21; for the moment we shall restrict our attention to a very simple model.

Consider the case of multiplicative risk, in which the farmer produces a single crop. His revenue will be

$$\tilde{Y} = \tilde{p}\tilde{\theta}f(x). \tag{6.27}$$

If the only input of the farmer is his own labour, x, and if his utility is separable

in income and leisure, then the farmer will maximize

$$EU\{\tilde{p}\tilde{\theta}f(x)\} - wx, \tag{6.28}$$

where wx represents the disutility associated with labour. We shall make this assumption of separability repeatedly, since it greatly simplifies the analysis. It is equivalent to the assumption that income and leisure are on the borderline between being substitutes and complements. Since there is no clear empirical presumption either way, we thereby avoid having to deal with ambiguously signed cross-derivatives. The dedicated student can of course readily relax this assumption. For the moment, we also assume w is fixed (corresponding to constant marginal disutility of labour) but it is easy to show that this is not a critical assumption.

The utility maximizing farmer chooses x so that

$$EU'(\tilde{Y})\tilde{p}\tilde{\theta}f' = w. \tag{6.29}$$

Equation (6.29) can be solved for the optimal level of the input (effort or labour), x.

We now ask three fairly standard questions:

1. What is the effect of a change in risk on the supply of effort?
2. What is the effect of a change in risk on the level of expected utility?
3. What comparative statics results can be deduced from market stability conditions?

6.4.1 *Effect of risk on effort*

Consider the effect of a mean-preserving spread, i.e. a change in the distribution of $\tilde{p}\tilde{\theta}$ (and it is only the product of these with which he is concerned) which leaves $E\tilde{p}\tilde{\theta}$ unchanged.

For simplicity, let us define a new random variable, $\tilde{r}$, and rewrite equation (6.29) as

$$EU'\{\tilde{r}f(x)\}\tilde{r} = w/f'(x), \quad \tilde{r} \equiv \tilde{p}\tilde{\theta}. \tag{6.30}$$

Here $\tilde{r}$ is the random return to farming, which compounds the effect of price and output variability. Observe that $U\{\tilde{r}f(x)\}\tilde{r}$ can be viewed as a function of the random variable $\tilde{r}$. We know from our earlier discussion and Figs. 6.5 and 6.6 that a mean-preserving spread of a variable reduces the expected value of every concave function of that variable and, conversely, increases the expected value of every convex function.

To see whether this function is convex or concave in r, differentiate $rU'\{rf(x)\}$ twice with respect to r, obtaining

$$\frac{dU'r}{dr} = U' + U''Y$$

$$= U'(1 - R)$$

$$r\frac{d^2U'r}{dr^2} = U''Y(1 - R) - U'R'Y.$$

$U'r$ is convex or concave as this is positive or negative. Figure 6.7 shows the effect of an increase in the expected marginal return to greater effort on the equilibrium supply of effort, namely, to increase it from x^* to x^{**}. Originally, expected net utility W reaches a maximum at M, with a corresponding supply of effort x^*, and zero expected net *marginal* return to effort. If this increases to some positive level (the slope at N is positive), then the maximum must shift to the right, to P, and effort must be increased. Conversely if the return falls. Thus we obtain the result that the expected *marginal* return to greater

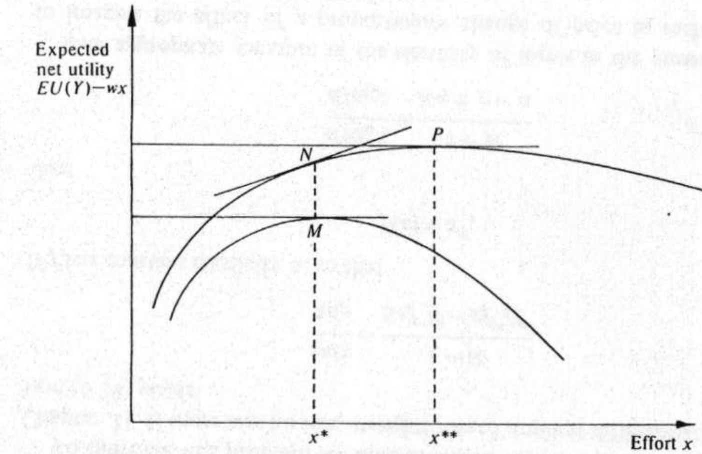

Fig. 6.7 Effect of increased risk on effort

effort is increased or decreased (and hence effort is increased or decreased for (6.30) to be satisfied again) as

$$R(1-R)+R'Y \lessgtr 0. \tag{6.31}$$

Certain special cases can now be identified:

(i) If $R'=0$ (constant relative risk aversion) effort is increased or decreased as $R \gtrless 1$. Individuals who are very risk averse increase their effort when risk is increased. They are worried, as it were, about the worst possible contingencies (e.g. starvation). When risk is increased, they have to work harder to avoid these extreme contingencies.

On the other hand, individuals who are less risk averse view the return to farming as lower; farming is a risky activity and risk is unattractive. Thus, they reduce their level of effort (output).

(ii) If, as seems likely, risk aversion is greater at low incomes, i.e. $R'<0$, it becomes more important to avoid low outcomes and thus an increase in risk is more likely to lead to increased output.

6.4.2 *Effect of risk on welfare*

Although the effect on output is ambiguous, the maximized value of expected utility is always decreased by a mean-preserving spread. For, denoting a small increase in risk by $d\theta$, we obtain

$$\frac{d(EU-wx)}{d\theta}=\frac{\partial(EU-wx)}{\partial x}\frac{dx}{d\theta}+\frac{\partial EU}{\partial\theta}. \tag{6.32}$$

The first term is zero, because the individual is assumed to choose x to maximize his expected utility. The second term is negative, because U is a concave function of r, and a mean-preserving spread always decreases the expected value of a concave function of the random variable.

Notice that the effect of a change in risk on the level of utility and on the level of marginal utility may be markedly different. They may even be of opposite sign. This has an important implication for the use of certainty equivalents of the kind referred to earlier. As we noted above there are two distinct notions of certainty equivalent that we might use in this context. The *action certainty equivalent* value of the return is the certain value of $\bar{r}$ which would lead the farmer to take precisely the same action (same level of effort) as the random $\tilde{r}$. The *utility certainty equivalent* value is the certain value of $\bar{r}$ which would leave the farmer with the same level of expected utility. These are distinctly different numbers (and indeed one may exceed the mean value of r and the other be less than it) and can clearly be affected in different ways by different policies. It is therefore preferable to work with the original utility functions rather than attempting to summarize their salient features in simple measures of certainty equivalence.

6.4.3 *Stability and comparative statics*

Samuelson's classic text (1947) identified two main sources of meaningful theorems in economics. The first set of theorems proceeds from the assumptions of maximizing behaviour of individuals, and may be exemplified by the Slutsky symmetry condition for demand theory (discussed in Chapter 8). The second set of fruitful theorems in comparative statics may be derived from the stability conditions of market interaction, using what Samuelson refers to as the *correspondence principle*, and set out in his Chapter IX. He illustrates this principle by analysing the stability of a single market in a way directly relevant for our present concern. The first, and perhaps most important, point to make is that market stability is a property of the disequilibrium adjustment mechanism. He contrasts Walrasian stability, in which price responds positively to excess demand, with Marshallian stability, in which supply responds positively to the excess of the demand price over the marginal cost, and with cobweb stability. We shall discuss these alternative concepts of stability in Chapter 23, but our present concern is whether any of these notions of stability permit deductions about comparative statics. As far as we know, very little work has been done applying the correspondence principle to risky markets, perhaps for good reasons. In a static riskless market, the relevant stability conditions are typically local and give information about the signs of various derivatives at the equilibrium point, which can then be used to derive comparative statics results about the directions of movement at this equilibrium point. In a risky market, the equilibrium typically depends on the shape of the various functions over the whole range of the probability distribution. Thus, for example, the supply of effort which solves equation (6.29) depends on the form of marginal utility over the whole range of values which income can take. Different distributions with the same mean (and variance) will give different results. Unless strong restrictions are placed on the functions (such as imposing constant relative risk aversion) little can be deduced about the local shape of these functions at equilibrium from the stability conditions, which depend on the average shapes of the functions over the range of the distribution.

To illustrate this problem we shall anticipate some of the stability analysis of Chapter 21. If there was no risk, straightforward implicit differentiation of equation (6.29) yields

$$\frac{pdx}{xdp}=\frac{1-R}{Rxf'/f-xf''/f'}. \tag{6.33}$$

If f has constant elasticity α, so that

$$f(x)=x^{\alpha},$$

then

$$\frac{d\log x}{d\log p}=\frac{1-R}{R\alpha+1-\alpha}. \tag{6.34}$$

One appropriate measure of the elasticity of input in the presence of risk is to imagine the effect of a proportionate change of price in each state of the world, so that the new price distribution is

$$\lambda\tilde{p}.$$

The elasticity of input to changes in the whole level of prices is then

$$\frac{\lambda dx}{xd\lambda}=\frac{EU'\cdot\tilde{p}\tilde{\theta}f'(1-R)}{EU'\cdot p\theta f'(R\alpha+1-\alpha)}=\frac{1-\bar{R}}{\bar{R}\alpha+1-\alpha}, \tag{6.35}$$

where

$$\bar{R}=\frac{EU'\tilde{p}\tilde{\theta}f'R}{EU'\tilde{p}\tilde{\theta}f'}.$$

Thus the *effective* degree of relative risk aversion is a weighted average value. The market stability conditions will impose restrictions on the value of $\bar{R}$ if

the market is to be stable, but unless R is assumed constant, this will not give useful information about the response of effort to risk, for example, which depends on R' as well: our knowledge of the sign of equation (6.31) will not be clarified by evidence on the average value of R.

6.5 Mean-variance analysis

It would obviously be convenient if we could describe attitudes to risk just in terms of the mean and variance of income, since these characteristics are simple to estimate and manipulate. The mean-variance model assumes that this is possible and has been extensively employed in the analysis of risk. Like most convenient models, it makes strong assumptions which prejudge the answers to various important questions and thus it is not suitable for proving general theorems. Provided its limitations are understood, it is very convenient for constructing examples and counter-examples, and we shall so use it in various places. Where possible, however, it is preferable to retain the more general utility framework, and appeal to general results on mean-preserving spreads on convex or concave functions. Even when it is necessary to approximate it is usually preferable to leave the Taylor series expansion until it is clear about which point to expand and how many terms to consider.

Given the simplicity and popularity of the mean-variance model it is worth discussing its main advantages and limitations in the analysis of risk. The first question to ask is under what assumptions on risk and the utility function is the mean-variance model valid? When, in other words, can we analyse the effects of risk on welfare and behaviour simply in terms of the mean and variance of income? Obviously, if the distribution of income is normal, this must be true, since the normal distribution is completely described by its mean and variance. If, therefore, the choices of the agent leave his income normally distributed, then all is well. This is particularly easy to see for utility functions with constant absolute risk aversion, for then (recalling the definition given in equation (6.13)) since

$$EU(Y)=-Ee^{-AY} \tag{6.36}$$

the expected utility can be found from the moment-generating function of the normal distribution; for if

$$Y=N(\bar{Y},V^2)$$

(i.e. Y is normally distributed with mean $\bar{Y}$, variance V^2, then

$$-E\exp(-AY)=-\exp\{-A(\bar{Y}-\tfrac{1}{2}AV^2)\}$$

and the utility certainty equivalent, $\hat{Y}$ is thus

$$\hat{Y}=\bar{Y}-\tfrac{1}{2}AV^2. \tag{6.37}$$

One important class of choices which satisfy this condition are portfolio choices in which the assets have jointly normally distributed returns. In the agricultural context if the income per acre from the *i*-th crop is $\tilde{r}_i$, and if the farmer allocates a fraction λ_i of each acre to crop i, then his total income per acre is

$$\tilde{Y} = \sum_{i=1}^{n} \lambda_i \tilde{r}_i \tag{6.38}$$

which has mean and variance

$$\bar{Y} = \sum \lambda_i \bar{r}_i$$

$$V^2 = \sum \lambda_i^2 \operatorname{Var}(r_i) + 2 \sum_{i \neq j} \lambda_i \lambda_j \operatorname{Cov}(r_i, r_j). \tag{6.39}$$

As the farmer varies his portfolio, that is, his farm plan, his income remains normally distributed, though its mean and standard deviation will depend on the choice of the fractions λ_i. Figure 6.8 plots the outcome of efficient portfolio choices (those which minimize V for given $\bar{Y}$), and the indifference curves associated with the utility function (that is, lines of constant expected utility).

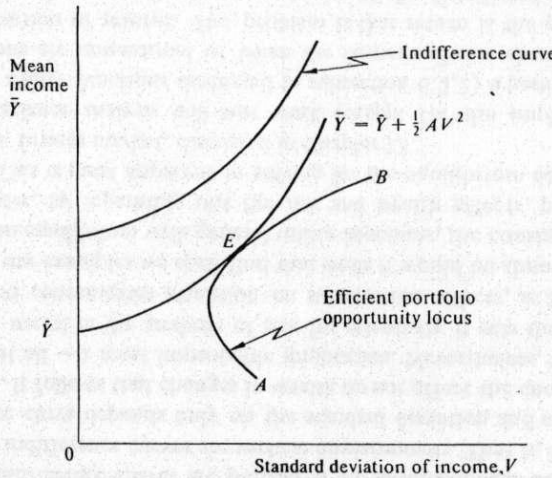

Fig. 6.8 Mean-variance portfolio choice

As shown the portfolio locus AEB is bowed to the left of the line AB, which merely requires that different crops are less than perfectly positively correlated. For example, if there were two crops of equal mean and variance, the portfolio mean would remain unchanged, but its variance would be

$$V^2 = \{\lambda^2 + 2\lambda(1-\lambda)\rho + (1-\lambda)^2\}\nu^2, \tag{6.40}$$

where ν is the variance of one crop and ρ is the correlation coefficient of the

returns. Minimizing the variance gives $\lambda = \frac{1}{2}$,

$$V^2 = \frac{1+\rho}{2}\nu^2 \tag{6.41}$$

which is less than the variance of each crop separately (if $\rho < 1$) and possibly much less, if ρ is negative.

If, in addition to the risky crops, the farmer has a perfectly safe crop which always yields the same return, r_s, then the mean return from allocating a fraction λ to the safe crop and $1 - \lambda$ to a single risky crop with return $\tilde{r}$, variance ν^2, is

$$\bar{Y} = \lambda r_s + (1-\lambda)\bar{r}$$

while the standard deviation is just

$$V = (1-\lambda)\nu.$$

This means that the opportunity locus is a straight line DE when plotted in mean standard deviation space, as in Figure 6.9. The point E represents the risky crop, which could as well be a risky portfolio of crops. For example, if there are many

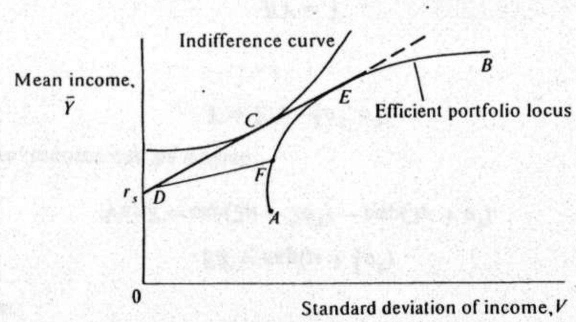

Fig. 6.9 Combination of risky and safe crops

risky crops with an opportunity locus AEB, then the point E is the tangent from D to the locus since any other point on the locus such as F, will yield less attractive final portfolio opportunities.

It follows that the farmer's decision can be made in two stages. First, given the returns on all the crops (including the safe crop), find the best crop pattern of risky crops as the point of tangency, E, from D to the locus of portfolio possibilities in Fig. 6.9. Second, decide on the amount of the risky crop pattern, which will be where the farmer's indifference curve touches DE at the point C. Only the second choice is affected by the farmer's attitude to risk so one can talk of *the* optimal crop plan for risky crops.

Now this portfolio separation theorem is very special and ceases to hold when the assumptions of mean-variance are relaxed. Nevertheless, the intuitive insight offered by mean-variance analysis — that it is the *covariance* of yields which are important determinants of the attractiveness of an asset or crop — is obviously a robust result which does not depend on the specific assumptions of mean-variance analysis.

The case of constant absolute risk aversion is special in a way which is obvious when its indifference curves are plotted in the mean-standard deviation Fig. 6.9 for all the indifference curves are vertical displacements. That is, the slope of the indifference curve depends only on the standard deviation and not on the mean of income. It follows that changes in wealth do not affect the choice of the risky portfolio at all – a most improbable implication. Nevertheless, this property is often very useful in the analysis of risk for essentially it sets the income effect to zero and concentrates attention on substitution effects, as risk changes. In several of our examples we shall find that while it would be almost impossible to solve for an equilibrium with general utility functions, the constant absolute risk aversion case, by separating out the risk and wealth effects, permits a simple solution. This is most apparent in solving for the equilibrium degree of speculation on the futures market, discussed in Chapter 13.

Mean-variance analysis will not work (except for the implausible case of quadratic utility functions discussed in subsection 6.2.2) where the probability distributions are non-normal or when the farmer's choice changes the *form* of the distribution of returns. The problem is that return is the product of price and quantity, less costs. Output cannot be normally distributed since that would imply some probability of negative output. Even if output were approximately normally distributed, and if this led to price being roughly normally distributed, their product would not be normal. One could cut through this problem by arguing that net returns *might* be roughly normally distributed to the extent that it would require considerable data to reject the hypothesis of normality. In many cases this is an adequate defence, but for some decisions the farmer will both be worried about extreme events (such as bankruptcy or starvation) and acutely aware that the risk is not symmetric or bell shaped. Indeed, his main decision problem might be to change the form of the distribution, to reduce the weight in the adverse tail. Most insurance schemes have this property and many actions can be thought of as insuring against adverse outcomes. For such decisions, mean-variance analysis may be seriously misleading.

There is one other special case which appears to offer the advantages of mean-variance analysis under a more plausible specification of risk and utility function, and that is the combination of constant relative risk aversion R, with a log-normal distribution of income (that is, the log of income is normally distributed).

6.5.1* *Properties of the log-normal distribution*

If a random variable X is log-normally distributed, with the mean of $\log X$ equal to μ, and the variance of $\log X$ equal to σ^2, then we write

$$X = \Lambda(\mu, \sigma^2) = e^Z, \qquad (6.42)$$

where Z is a random variable normally distributed with mean μ, variance σ^2:

$$Z = N(\mu, \sigma^2).$$

Its properties are detailed in Aitcheson and Brown (1957), and the expected value of powers of X can readily be found from the moment generating function of Z, defined as

$$M_Z(t) = E\exp(tZ).$$

Thus, since Z is normal

$$EX^\beta = M_Z(\beta) = \exp(\beta\mu + \tfrac{1}{2}\sigma^2\beta^2). \qquad (6.43)$$

In particular

$$EX = \exp(\mu + \tfrac{1}{2}\sigma^2)$$
$$\operatorname{Var} X = \exp(2\mu + 2\sigma^2) - \exp(2\mu + \sigma^2).$$

If, therefore, income can be written

$$Y = \bar{Y}\Lambda(-\tfrac{1}{2}\sigma^2, \sigma^2),$$

then

$$EY = \bar{Y}$$
$$\operatorname{Var} Y = \bar{Y}^2(\exp\sigma^2 - 1) \equiv \bar{Y}^2\sigma_y^2 \qquad (6.44)$$

But

$$\exp\sigma^2 = 1 + \sigma^2 + \sigma^4/2! + \dots$$

so the coefficient of variation of Y, σ_y, is approximately σ.

Equation (6.43) allows us to calculate the utility certainty equivalent income, $\hat{Y}$, directly for

$$\frac{\hat{Y}^{1-R}}{1-R} = \frac{EY^{1-R}}{1-R} = \frac{\bar{Y}^{1-R}}{1-R}\exp\left\{-R(1-R)\frac{\sigma^2}{2}\right\}$$

or, from (6.44)

$$\hat{Y} = \bar{Y}(1 + \sigma_y^2)^{-R/2}. \qquad (6.45)$$

This is an exact equation, to be compared to the approximate equation (6.11) above.

Log-normality is the natural assumption to make in many econometric specifications which assume constant elasticities and multiplicative risk, and it has the advantage that negative values are ruled out. Moreover, if output and price are log-normal, so will be revenue, while if output fluctuations *cause* the price fluctuations and demand is of constant elasticity, revenue will be log-normal, as required here. (These cases are discussed in more detail in section 6.6.) Unfortunately, we cannot use the previous portfolio analysis for crop choices in this case because the sum of two log-normal distributions is not log-normal.

Given the restricted conditions under which choices under risk can be accurately described in terms of means and variances alone, it is obviously important to ask whether it is possible to measure the impact of risk (and particularly changes in risk) *approximately* using just means and variances, and, if so under what conditions. The question is important for two reasons. In the first place, the concept of a mean-preserving spread is useful for qualitative analysis, but not for estimating the quantitative impact of changes in risk. Second, it is difficult enough to estimate the mean and variance of a probability distribution without having to investigate higher moments, and it would obviously simplify matters if these were not needed. The next section examines the conditions under which mean-variance analysis is approximately valid.

6.5.2 *Quadratic approximations and Taylor series expansions*

If the function $U(Y)$ is sufficiently differentiable then it can be expressed as a Taylor series:

$$U(Y) = U(\bar{Y}) + (Y-\bar{Y})U'(\bar{Y}) + \tfrac{1}{2}(Y-\bar{Y})^2U''(\bar{Y}) \qquad (6.46)$$

$$+ \ldots \frac{(Y-\bar{Y})^n}{n!}\{U^{(n)}(\bar{Y}) + \epsilon_n\}$$

where $\qquad \epsilon_n \to 0 \text{ as } Y \to \bar{Y}.$

ϵ_n can be thought of as a remainder, or an error, which depends on Y and n. Take the expectation of both sides to obtain

$$EU(Y) = U(\bar{Y}) + \tfrac{1}{2}U''(\bar{Y})E(Y-\bar{Y})^2 + \eta_2 \qquad (6.47)$$

$$\eta_2 = \tfrac{1}{2}E(Y-\bar{Y})^2\epsilon_2(Y).$$

This says that expected utility can be approximately expressed as a function of the mean and variance of income. We wish to know under what conditions is this approximation valid, i.e. when is the error η_2 small relative to the second term. It is tempting to suppose that as the variance of Y tends to zero, the error tends to zero, but it is easy to show that this is not necessarily true. For example, consider the family of distributions of Y in which

$$Y = \begin{cases} \bar{Y} + \pi^{-1/3} & \text{with probability } \pi \\ \bar{Y} - \pi^{-1/3} & \text{with probability } \pi \\ \bar{Y} & \text{with probability } 1 - 2\pi. \end{cases}$$

Then

$$E(Y-\bar{Y})^2 = 2\pi^{1/3} \to 0 \quad \text{as} \quad \pi \to 0$$

$$E(Y-\bar{Y})^3 = 0$$

$$E(Y-\bar{Y})^4 = 2\pi^{-1/3} \to \infty \quad \text{as} \quad \pi \to 0.$$

Thus, as the variance goes to zero, the fourth (and higher even moments) become infinitely large. It is easy to see what goes wrong, for it is not enough for the variance of a probability distribution to tend to zero for $Y - \bar{Y}$ to tend to zero, and hence it is not possible to argue that $\epsilon_2 \to 0$. We need a stronger notion of convergence, and Samuelson (1967, 1970) introduces the concept of 'compact' probabilities, such that as some specified parameter goes to zero, all the distributions converge on the certain outcome. Thus if the distribution function of Y can be written

$$F(Y) = P\left(\frac{Y-\bar{Y}}{V}\right), \qquad V^2 \equiv E(Y-\bar{Y})^2$$

for some given distribution function P, then as the variance of Y, V^2, tends to zero, the probability all piles up at $\bar{Y}$, and, assuming that P has finite moments, all the higher moments of Y will tend to zero. As Samuelson shows for such distributions, as the variance tends to zero, so the mean-variance approximation becomes progressively more accurate. In the counter example, as the variance went to zero, the form of the distribution was changing, and becoming *more* disperse, not less.

We shall therefore defend the use of Taylor series expansions on the grounds that they are valid for 'small risks', meaning not just that the variance is small, but the dispersion of the whole distribution is not 'too large'. However, it is usually advisable to delay taking approximations until as late as possible in the analysis to avoid the compounding of errors. For example, even when the quadratic approximation provides a good approximation to the level of expected utility, one must be careful in using mean-variance analysis for comparative statics analysis. That requires (for small variance) employing a quadratic approximation to the first-order condition, which may not be the same as the first-order condition derived from the quadratic approximation to the utility function.

As an illustration of the method of calculating expected values by Taylor series approximations consider

$$E\theta^\beta \quad \text{where} \quad \theta \equiv 1 + u, \quad Eu = 0, \quad Eu^2 = \sigma^2$$

$$\theta^\beta = (1+u)^\beta \simeq 1 + \beta u + \tfrac{1}{2}\beta(\beta-1)u^2,$$

hence

$$E\theta^\beta \simeq 1 + \tfrac{1}{2}\beta(\beta-1)\sigma^2. \qquad (6.48)$$

If in fact θ is log-normally distributed

$$\theta = \Lambda(-\tfrac{1}{2}s^2, s^2), \quad \sigma^2 = \exp s^2 - 1$$

then, from equation (6.43)

$$E\theta^\beta = \exp(-\tfrac{1}{2}\beta s^2 + \tfrac{1}{2}\beta^2 s^2)$$

$$E\theta^{\beta} = (1 + \sigma^2)^{\frac{1}{2}\beta(\beta-1)}. \tag{6.49}$$

For example, if $\beta = 3$, $\sigma = 0.3$, the error in the Taylor series approximation is 2 per cent, or 28 per cent of σ^2. If, however, $\beta = 1.5$, the error is only 1 per cent of σ^2. The accuracy of the approximation thus depends on the size of σ^2 and β.

6.6 Measuring the benefits of price and income stabilization

The most difficult part of the study of commodity price stabilization lies in measuring the responses by the various agents to the change in price instability. Farmers may change their allocation of inputs, their production techniques, their cropping plan, the amount of crop storage, and their dealings on futures markets. Consumers may also change their inventory policy, their dealings on futures markets, and, if they are intermediate producers, their production plans, though typically final consumers are not involved much in these activities. Much of the rest of the book is devoted to the study of these responses, but we can throw some immediate light on the central issue of the benefits of price stabilization if we are prepared to assume that there is *no* response to price stabilization. Obviously this is an extreme, and apparently unreasonable, assumption, but it can be defended, not only on the grounds that it is useful as a bench-mark. In the first place, it takes agents time to detect and measure changes in risk, for estimates of sample *variance* (to take one natural measure of risk) require large samples (or are themselves subject to wide sampling variation). Assuming no response is then equivalent to studying the short-run impact of the proposed stabilization scheme, to be compared with the long-run impact, after agents have fully adjusted to the new equilibrium. Second, we shall show that in some cases the long-run impact on producers is a simple fraction of the short-run impact. Finally, the short-run impact clarifies an important distinction between *transfer* benefits, which producers gain at the expense of consumers, or vice versa, and *efficiency* benefits, which represent net social gains. The long-run impact typically alters the transfer benefits, which are merely redistributive, without much affecting the efficiency benefits. For all these reasons, then, it is sensible to measure the producer benefits of stabilization schemes which change the variability of *incomes* (typically, by changing the variability of prices) without changing the level of inputs. Average output thus remains constant and we can ignore the (constant) disutility of effort. Moreover, we can apply some of the techniques developed in this chapter to obtain simple quantitative measures of the two types of producer benefits.

Suppose that initially a representative farmer has income $\tilde{Y}_0$ with mean $\bar{Y}$ and coefficient of variation σ_{y0}, and after stabilization this changes to $\tilde{Y}_1$ with mean $\bar{Y}_1$, coefficient of variation σ_{y1}. We wish to know what stabilization is worth to the farmer, that is, what sum of money, B, he would be willing to pay for the stabilization scheme to be introduced. This sum can be found by equating expected utility:

$$EU(\tilde{Y}_0) = EU(\tilde{Y}_1 - B). \tag{6.50}$$

Expand the left-hand side in a Taylor series:

$$\text{LHS} \cong U(\bar{Y}) + \tfrac{1}{2}E(\tilde{Y}_0 - \bar{Y})^2 U''(\bar{Y}) \tag{6.51}$$

and similar expand the right-hand side:

$$\text{RHS} \cong U(\bar{Y}) + (\Delta\bar{Y} - B)U'(\bar{Y}) + \tfrac{1}{2}E(\tilde{Y}_1 - \bar{Y} - B)^2 \cdot U''(\bar{Y}), \tag{6.52}$$

where Δ is the difference operator, so

$$\Delta\bar{Y} = \bar{Y}_1 - \bar{Y}_0.$$

Equating these two expansions and dividing by $\bar{Y}U'(\bar{Y})$ gives

$$\frac{B}{\bar{Y}} = \frac{\Delta\bar{Y}}{\bar{Y}} - \tfrac{1}{2}R\left\{\Delta\sigma_y^2 + \left(\frac{\Delta\bar{Y} - B}{\bar{Y}}\right)^2\right\}. \tag{6.53}$$

Evidently $(\Delta\bar{Y} - B)/\bar{Y}$ is of order σ_y^2 and its square can be ignored, given the accuracy of the approximation, so that

$$\frac{B}{\bar{Y}} = \frac{\Delta\bar{Y}}{\bar{Y}} - \tfrac{1}{2}R\Delta\sigma_y^2. \tag{6.54}$$

The first term is the *transfer* benefit, for the change in average income to the producer is matched by an equal change in average expenditure by consumers. The second term is the *efficiency* or *risk* benefit, the benefit from reducing costly risk to the farmer.

It is easy to quantify these benefits in particular cases. Suppose that the source of risk lies on the supply side, and that output is log-normally distributed. If demand is stable and has constant elasticity ϵ, so that

$$p = Q^{-1/\epsilon}, \quad Q = \bar{Q}\theta$$

with

$$\theta = \Lambda(-\tfrac{1}{2}\sigma^2, \sigma^2), \quad E\theta = 1, \text{Var}\,\theta \simeq \sigma^2,$$

then income is also log-normally distributed

$$Y = pQ = \bar{Q}^{1-1/\epsilon}\theta^{1-1/\epsilon}.$$

Expected values of powers of θ can be found directly from equation (6.43), for

$$E\theta^{\beta} = \exp\tfrac{1}{2}\beta(\beta-1)\sigma^2 \simeq 1 + \tfrac{1}{2}\beta(\beta-1)\sigma^2 \tag{6.55}$$

(since $\exp x \cong 1 + x$).

Initially, the CVs of prices and incomes are

$$\sigma_p = \frac{\sigma}{\epsilon}, \quad \sigma_y = \left(1 - \frac{1}{\epsilon}\right)\sigma.$$

Suppose that stabilization reduces the CV of prices to a fraction $1 - z$ of its original value, so that z measures the degree of stabilization, with $z = 1$ corresponding to perfect price stability:

$$p = \bar{Q}^{-(1/\epsilon)}\theta^{-(1-z)/\epsilon}, \quad \sigma_p = \frac{(1-z)\sigma}{\epsilon},$$

then the CV of income will fall to

$$\sigma_y = \left(1 - \frac{1-z}{\epsilon}\right)\sigma.$$

The transfer benefit B_T is

$$B_T = \frac{\Delta\bar{Y}}{\bar{Y}} = \frac{\{E\theta^{1-(1-z)/\epsilon} - E\theta^{1-1/\epsilon}\}}{E\theta^{1-1/\epsilon}}.$$

Since the numerator is of order σ^2, and the denominator is $1 + k\sigma^2$, where k is of order 1 the error in assuming that the denominator is exactly 1 is of order σ^4 and can be ignored. Hence, using the approximation of equation (6.55)

$$B_T = \tfrac{1}{2}\frac{z}{\epsilon}\left(1 - \frac{2-z}{\epsilon}\right)\sigma^2. \qquad (6.56)$$

Similarly, the risk benefit is

$$B_R = \tfrac{1}{2}R\frac{z}{\epsilon}\left(\frac{2-z}{\epsilon} - 2\right)\sigma^2 \qquad (6.57)$$

so the total benefit is

$$B = \tfrac{1}{2}\frac{z}{\epsilon}\left\{1 - 2R - \frac{2-z}{\epsilon}(1-R)\right\}\sigma^2, \quad 0 \leqslant z \leqslant 1. \qquad (6.58)$$

Evidently each part of the total benefits can be either negative or positive, depending on the magnitude of risk aversion, R, the elasticity of demand, ϵ, and the degree of stabilization, z.

More generally, if there is additional multiplicative randomness in demand, so that supply and demand are respectively

$$Q = \bar{Q}\theta$$

$$Q^d = p^{-\epsilon}\phi, E\phi = 1,$$

and if θ and ϕ are jointly log-normally distributed, then so will be price and quantity. Suppose the correlation coefficient of $\log p$ on $\log Q$ is r (typically negative) and that stabilization of degree z lowers the coefficient of variation to $1 - z$ of its original value σ_{po}.

The square of the CV of income (itself log-normally distributed) is approximately equal to the variance of $\log pQ$:

$$\mathrm{Var}(\log p + \log Q) = \sigma^2 + 2r(1-z)\sigma\sigma_{po} + (1-z)^2\sigma_{po}^2,$$

so that the risk benefit is now

$$B_R = \tfrac{1}{2}Rz\{(2-z)\sigma_{po}^2 + 2r\sigma\sigma_{po}\}. \qquad (6.59)$$

The transfer benefit is found by evaluating

$$B_T = \frac{\Delta\bar{Y}}{\bar{Y}} = \frac{EX_1X_2^{1-z} - EX_1X_2}{EX_1X_2}, \qquad (6.60)$$

where $X_1 = \theta, X_2 = (\phi/\theta)^{1/\epsilon}$ are jointly log-normally distributed. In the Appendix to Chapter 13 it is shown that

$$X_2 = \Lambda(\mu_2, \sigma_p^2); \quad \mu_2 = -\tfrac{1}{2}(\epsilon\sigma_p^2 + 2r\sigma\sigma_p),$$

while

$$X_1 = \Lambda(\mu_1, \sigma^2); \quad \mu_1 = -\tfrac{1}{2}\sigma^2.$$

Equation (13A11) demonstrates that

$$EX_1X_2^{1-z} - EX_1X_2 = -\tfrac{1}{2}z\{2\mu_2 + (2-z)\sigma_p^2 + 2r\sigma\sigma_p\}.$$

So, again approximating the denominator by unity, we obtain the same simple formula as before:

$$B_T = \tfrac{1}{2}z\{\epsilon - (2-z)\}\sigma_{po}^2. \qquad (6.61)$$

These formulae can be used to measure the short-run benefits of price stabilization as they are given in terms of the initial CV of output, σ, the initial CV of price, σ_{po}, the fractional reduction in the CV of price, $1 - z$, the correlation between price and supply, r, the elasticity of demand, ϵ, and the degree of relative risk aversion, R. Table 6.1 gives as an example the results of a very simple regression exercise reported in more detail in Chapter 20.

Table 6.1 *Benefits of stabilizing the cocoa price by 50 per cent*

Country	Elasticity[a]	r[a]	σ %	σ_p %	Risk benefit %	Transfer benefit %
Ghana	0.5	–0.74	21	31	1.2R	–2.4
Brazil	–0.3	0.33	24	31	3.0R	–4.3

[a] As defined, price elasticity and the correlation coefficient have opposite signs.

It is interesting to note the difference between countries for the same crop, and the difference in sign of the two components of the benefit.

Chapter 7

Empirical Measurements of Producers' Attitudes to Risk

The theory developed in Chapter 6 assumed that farmers acted as though they were maximizing expected utility. How reasonable is this assumption? What does the available evidence have to say about attitudes to risk? How important is risk in agriculture? In this chapter we first consider the theoretical problems with the expected utility hypothesis, and then discuss the results of an important recent experiment which provides evidence of farmers' attitudes to risk. In section 7.4 we examine the non-experimental empirical evidence, and, finally, conclude with some evidence on the magnitude of agricultural risk.

7.1 Distinguishing among alternative hypotheses

Much of the qualitative analysis contained in this book does not depend critically on the assumption employed throughout the work that individuals maximize their expected utility. Our analysis emphasized three points:

(i) a change in the price distribution may well increase the variability of income faced by individual farmers;

(ii) this change in risk may cause changes in the production decisions of producers; and possibly in the consumption decisions of consumers, and

(iii) there are important general equilibrium consequences of these changes.

These results would be true under virtually any theory which argued that producers (and consumers) are concerned with the riskiness of their income. Moreover, the expected utility hypothesis is consistent with a wide range of behaviour, e.g. an increase in risk of one crop could lead to the farmer growing more or less of the crop.

But for policy purposes, we often need more than just a qualitative analysis. We would like to know: (a) how important is risk in agriculture; and in particular (b) how large is the response (and in which direction) to a change in risk.

To answer the second question, in particular, to know how farmers would respond to the kinds of changes in risk induced by a commodity price stabilization programme, we need to be able to infer that the individual's behaviour towards this new risk situation will be similar to his behaviour towards earlier risky situations which he has faced.

This postulate of *consistency* – that individuals behave systematically when faced with risky situations – is thus the basic hypothesis underlying our analysis. The consistency hypothesis underlies all of the theory of consumer behaviour. But on *a priori* grounds it has somewhat less force in this context than in others:

if a consumer chooses oranges instead of apples and then finds that he dislikes the taste of oranges, he learns immediately that he has been mistaken. With risky choices it is more difficult to learn from experience, since it is not clear whether the consumer made the wrong choice, or whether he was merely unlucky. This hypothesis has been subjected to testing in a variety of laboratory experiments, where individuals are confronted with different gambles. Inconsistencies in their behaviour were often noted.

Parts of the apparent inconsistencies may be explained by the lack of familiarity of the individuals with the kinds of choices being faced; alternatively, because the pay-offs in the experimental situations are usually trivial, individuals may not take the experimental situation seriously. With larger pay-offs, one might expect to find more consistent behaviour. This in fact turns out to be the case, as the study we report later in section 7.3 suggests.

The expected utility theory, of course, implies more than just consistency. More generally, we can write utility as a function of income in each of the states of nature:

$$U = U\{Y(\theta_1), Y(\theta_2), Y(\theta_3), \ldots\}$$

where

$$Y(\theta_i) = \text{income in state } i.$$

The expected utility theory postulates that the utility function can be written in a particular form:

$$U = \sum_i U\{Y(\theta_i)\}\pi_i$$

where π_i is the probability that state i occurs and

$$\sum \pi_i = 1.$$

There are a variety of alternative axiomatic foundations for the expected utility hypothesis (see, for instance, Savage, 1954; Arrow, 1970; Luce and Suppes, 1965). The critical hypothesis is the 'compounding' axiom, that an individual is indifferent between a lottery of lotteries and a single lottery yielding the same outcomes with the same probabilities. For instance, consider a lottery that yields

$1 with probability 0.25

$2 with probability 0.5

$4 with probability 0.25

Now, consider an alternative situation, where the outcome of the first lottery is that the individual gets one of two lotteries, either lottery A, which yields

$1 with probability 0.5,

or

$2 with probability 0.5,

or lottery B which yields

$2 with probability 0.5,

or

$4 with probability 0.5.

If lotteries A and B are equally likely (the individual has a 50 : 50 chance of drawing lottery A or lottery B), then the expected utility hypothesis postulates that the individual ought to be indifferent between the original lottery and the lottery of lotteries.

This hypothesis too has been subjected to testing, and systematic behaviour which is not consistent with it has been uncovered in experimental situations. Part of the problem is that most individuals are not very good at estimating (and compounding) probabilities, particularly in unfamiliar circumstances. It is difficult to distinguish between the hypotheses that individuals make choices which are consistent with the expected utility hypothesis, given their subjective probabilities, but that there are systematic biases in the manner in which those subjective probabilities are formulated, and the hypothesis that individual behaviour is not consistent with the expected utility hypothesis. There has been considerable work (see, for instance, Tversky 1969) suggesting that there are systematic biases in the way in which individuals form their subjective probability estimates. For example, individuals act as though they systematically overestimate low objective probabilities (such as the chance of winning at football pools, or the chance of an aircraft accident, or small chances in experimental situations).

Some proponents of the expected utility theory argue that the theory ought to be viewed as a normative theory, how rational individuals ought to behave in the face of risk. Others go further and suggest that when individuals are made aware of their inconsistencies with the expected utility theory, they change their behaviour; that is, when individuals know how to calculate probabilities, they will, in fact, behave consistently with the expected utility hypothesis. Again, the results reported below lend some substance to that view: when individuals are faced with *serious* choices, involving large stakes, their behaviour appears to be more in conformity with what a proponent of the expected utility hypothesis might hold.

Most of this book is concerned with situations where, at least in the long run, we can speak of 'objective' (relative frequency) probabilities. When individuals appear to use probabilities which are different from these objective probabilities, we say that they have miscalculated these probabilities.

This raises some interesting and not completely resolved issues in evaluating policy changes. There is a widespread view that the individual welfare ought to be evaluated using *ex ante* expected utility, i.e. using the individual's own probabilities. There is another view, which in this context we think is more persuasive, that it is the average *experience* of individuals which is relevant, i.e. the probabilities we ought to use are the relative frequencies of the different events. In this view, then, we can distinguish between the effects of policies assuming that agents act rationally, and the effects given their likely misperceptions (differences between objective and subjective probabilities). For example, as we shall argue in Chapter 11, a price stabilization programme may eventually reduce the average degree of misperception by farmers of the appropriate certainty equivalent price, although it is also quite possible that the introduction of the scheme initially worsens farmers' perceptions and decisions. Occasionally some policies have been justified primarily on grounds of misperceptions: that it improves matters (only) when agents act irrationally. Such arguments need, however, to be used with caution: the gains may be transient and reversed once agents learn to improve their decision-making skills; and it may be easier to provide the necessary information to improve decisions than to adopt the policy.

7.1.1 *Risk-taking in a multi-period context*

The analysis so far has assumed that individuals live for one period, or at least look at each decision as if it were in isolation. This is a convenient simplification, but when it comes to empirical verification of the model it is clearly unsatisfactory.

Farmers typically have to make choices each year (and, indeed, during the course of a crop year). In the static (one-period) model, there is no need to distinguish between income and wealth, but in a sequential problem the two are quite different. Economists usually argue that utility is produced by consumption, not income, and if an individual has substantial wealth, fluctuations in his income, especially if they are random from one year to the next, should not seriously reduce his ability to consume and hence generate utility. If individuals have few liquid assets, and if they cannot easily lend or borrow, or store goods from one period to the next, then their consumption would be constrained by their income, and it would be legitimate to define utility in terms of current income. If they can transfer income from one period to another, and if they rationally choose a lifetime consumption plan, then the problem becomes more difficult, and attitudes to risk will depend not only on the shape of the utility function, but on initial wealth, the rate of interest, the degree of independence of successive risks, and future income possibilities, as discussed in Chapter 14.

The empirical importance of this distinction between income and wealth will become apparent below. This completes our survey of the theory of behaviour under risk and the next step is to confront the theory with the empirical evidence.

Ideally, we would like to find the answers to a number of related questions:

(i) Do farmers make consistent choices between risky alternatives in a way which can be described as maximizing an expected utility function defined on the outcomes and the objective probabilities? In short, does the Expected Utility Hypothesis describe their choices satisfactorily?

(ii) Do farmers respond similarly and consistently to changes in income and wealth so that their choices can be described by maximizing an expected utility function defined over income, Y, and wealth, W of the form $U(W + Y)$? This is the Asset Integration Hypothesis, and it implies that individuals realize that their utility depends on their ability to obtain consumption goods, which in turn depends on *all* the factors which influence their purchasing power, and not just the immediate consequences of the next decision. Of course, if W is not affected by their current choices, and if the individual has a well-defined utility function $U(W + Y)$, then we can define another, current utility function, $V(Y)$, defined on current income, by the relation

$$V(Y) \equiv U(W + Y).$$

The empirical problem is that ideally we wish to find the functional form of U (assuming that such a function can be found which describes choices), but typically all that we observe are changes in Y, so that we are only able to measure the shape of $V(Y)$, or, equivalently, to find out about the shape of U in the neighbourhood of the given level of wealth, W.

(iii) Do different farmers behave similarly to risky prospects, so that we can usefully talk of a representative farmer?

(iv) Is there a reasonably simple functional form of the utility function which describes attitudes to risk?

7.2 Empirical evidence

The empirical attempts to measure risk aversion fall into two categories. The direct method, due to von Neumann and Morgenstern (1947) and extensively employed by behavioural psychologists, consists in confronting the subject with choices between sure things and risky alternatives, or between different risky alternatives. These choices may be hypothetical or actual. In the first case the subject is asked to conduct a thought experiment, of the form 'which of the two alternatives *would* you choose if you had to choose between them?' Dillon and Scandizzo (1978) have used this approach for near-subsistence farmers in north-east Brazil, and Lin, Dean, and Moore (1974) have tried it on wealthy Californian farmers. The attraction of the method is that it is relatively cheap to conduct the experiments and it appears to allow the exact shape of the utility function to be traced out, but it runs the risk that there is little incentive for the subject to think carefully about his answer, since nothing is at stake. The better alternative is to offer actual choices, preferably comparable in size to the gains and losses of that aspect of economic activity which is under investigation. The cost of experimenting on wealthy Californian farmers would be prodigious, but Binswanger (1978a, b) has conducted extensive experiments in rural India where wage rates are very low, so that the choices offered to the subjects involved relatively large gains at a low US$ cost. The cost of the experiment was roughly $2500 in prize money and $2500 in other costs, while to replicate it in the US would probably have cost over $200 000. Binswanger was also able to compare the results of offering hypothetical choices with actual choices, and concluded that '***evidence on risk aversion from pure interviews is unreliable, non-replicable and misleading***, even if one is interested only in a distribution of risk aversion rather than reliable individual measurement' (Binswanger, 1978b, p. 45, emphasis in original). Given this finding, and since there are no other suitable experimental studies, we shall merely summarize Binswanger's results and refer the interested reader to the original for the definitive discussion of the problem.

The other, or indirect, approach involves deducing attitudes to risk by observing actual decisions, and using these observations to estimate the parameters in an explanatory model of the farmer's behaviour. This method has the advantage that a large number of observations can be collected in the course of a more general investigation into the determinants of production, but it depends crucially on how well specified the model is. A recent example is provided by Moscardi and de Janvry (1977), and is discussed in section 7.4.

First, however, we examine Binswanger's evidence to see how far it supports the expected utility hypothesis adopted in the previous chapter, and what, if anything, it tells us about the degree of risk aversion of poor farmers.

7.3 Experimental determination of attitudes to risk

Binswanger's experiment was performed with over 300 individuals randomly selected from six villages in semi-arid rural India, and consisted in playing a sequence of games with real and high pay-offs. The subjects were offered a choice between the eight alternatives described in the upper part of Table 7.1, after which a coin was tossed and the outcome paid. The monthly unskilled wage rate in this region was 60–80 Rs. and the modal individual wealth was

Table 7.1 *Results of Binswanger's experiments*

Choice	O	A	B	C	E	F	D*	D	
Reward: heads	50	45	40	30	10	0	35	20	
tails	50	95	120	150	190	200	125	160	
Greatest value of partial risk aversion, P	∞	7.5	1.74	0.82	0.32	0	inefficient		
	Cumulative frequency of choice %						frequency %		N obs.
Game level									
0.50 Rs.	1.7	7.6	36.1	56.3	71.4	89.9	10.1		119
5 Rs.	0.9	9.4	35.0	71.8	83.8	92.3	7.7		117
50 Rs.	2.5	7.6	42.4	82.2	89.0	90.7	9.3		118
500 Rs.[a]	2.5	16.1	67.8	96.6	96.6	97.5	2.5		118

[a]Hypothetical game

Source: Binswanger (1978)

about 10 000 Rs. (roughly US$1200). The game was played seven or eight times over a period of six weeks, starting with five games at the 0.5 Rs. level (where the rewards are as shown in Table 7.1, divided by 100). The rewards were then increased by 10 (a perfectly certain outcome of 5 Rs.), then increased again by 10, and finally, hypothetical choices at the 500 Rs. level were asked.

The number of respondents making each choice can be deduced from the data in Table 7.1, and the local shape of the individual's utility function can be inferred within limits on the assumption that the choice made yields higher expected utility than any other alternative.

As we remarked above, the observations only give us information about the local shape of the function U defined in total wealth $W + Y$, or, equivalently, about the shape of the function V defined on current outcomes, Y, where

$$V(Y) \equiv U(W + Y) \tag{7.1}$$

is defined for a particular level of wealth. The shape of U or V can be described in a number of ways, such as by the coefficient of absolute risk aversion, A:

$$A \equiv -\frac{U_{ww}}{U_w} = -\frac{V_{yy}}{V_y}. \tag{7.2}$$

However, as the size of the bets increased, it was found that the respondents tended to choose slightly, but not very much less risky alternatives, which implies that for the modal individual, at the 0.50 Rs. level the value of the absolute risk aversion centres about 1, but falls rapidly to about 0.002 at the 500 Rs. level. In short, A is not stable as the proportional size of the outcomes increases.

A measure of risk aversion which was found to be relatively stable as the size of bets increased was what Menezes and Hanson (1970) term the coefficient of partial risk aversion, P (called the size-of-risk aversion by Zeckhauser and Keeler, 1970), which is defined on the current outcome, Y:

$$P \equiv -Y\frac{V_{yy}}{V_y}. \tag{7.3}$$

The coefficient of partial risk aversion is related to the coefficient of relative risk aversion, R, which is defined on total wealth, as follows:

$$P = \left(\frac{Y}{W+Y}\right)R \quad \text{where } R = -(W+Y)\frac{U_{ww}}{U_w}. \tag{7.4}$$

Clearly, if R is thought to be independent of wealth (and in particular, the proportionate size of the risky outcome) then P will decrease with wealth and increase with Y.

The greatest value of partial risk aversion, P, for which the individual would choose an outcome rather than the next less risky alternative is given in Table 7.1 and found by solving the equation

$$y_h^{1-P} + y_t^{1-P} = x_h^{1-P} + x_t^{1-P},$$

where (y_h, y_t) are the rewards under heads or tails of the alternative, and (x_h, x_t) are the rewards for the next less risky alternative. Thus choice A is indifferent to B at $P = 1.74$ because

$$45^{-.74} + 95^{-.74} \simeq 40^{-.74} + 120^{-.74}.$$

If P were slightly greater than 1.74, A would be preferred to B.

7.3.1 *Interpreting the evidence*

The first, and basic, question is how far the results of the experiment allow us to distinguish between the expected utility approach and alternative theories, of which the most popular are those based on security motives, recently reviewed by Anderson (1979). In all these approaches the individual is assumed to have an overridingly important objective, either to minimize the probability of experiencing a shortfall below some critical minimum income level, or to maximize the income level below which income will fall only a specified proportion of the time.

Binswanger finds that his evidence rejects all such theories which make testable predictions. On the other hand, his evidence is consistent with the hypothesis that individuals maximize expected utilities defined on outcomes. The weak form of the Expected Utility Hypothesis predicts choices reasonably well, and it is noticeable that as the choices become larger the proportion of inefficient choices (i.e. choices inconsistent with a concave utility function) decrease.

The answers to questions (iii) and (iv) are that most farmers make similar choices (B or C) at different game levels, so that one can attach some meaning to the representative farmer, and his attitude to income risk is quite well approximated by the constant partial risk aversion utility function, with risk aversion, P, lying between 1.74 and 0.32.

As the size of potential gain increases, so the median degree of partial risk aversion increases. Table 7.1 shows that when the rewards are comparable to a day's wages (at 0.5 Rs.) only 36 per cent of the sample had risk aversion greater than 0.82, while at the highest level (500 Rs.) the percentage rose to 68 per cent. Nevertheless, the partial risk aversion is remarkably stable given that rewards increase by a factor of 1000. It is, however, difficult to accept the Asset Integration Hypothesis, for if it were to hold, it would imply a utility function for a representative farmer with initial wealth W_0, of the approximate form

$$U(W + Y) \doteq \frac{(W + Y - W_0)^{1-\rho}}{1-\rho}. \tag{7.5}$$

If the same function were to apply to farmers with different wealth, then the coefficient of partial risk aversion should decrease rapidly with wealth:

$$P(W) = \frac{\rho Y}{Y + W - W_0}.$$

However, Binswanger found a very small decline in partial risk aversion with increasing wealth, and one must either reject the hypothesis that different individuals have a similar utility function, given by equation (7.5), or reject the Asset Integration Hypothesis. Binswanger (1978b) argues persuasively for rejecting the latter. In particular, one can reject the hypothesis of constant *relative* risk aversion, R, which would require $W_0 = 0$ in equation (7.5), since for a modal individual with wealth of about 10 000 Rs. relative risk aversion *falls* from about 1000 at the lowest game level to between 1 and 2 at the highest game level. It appears, then, as if decisions are compartmentalized, so that the decision is not seen in the context of the individual's over-all asset position. Only current outcomes seem relevant to current choices. Why might this be? One possible explanation is that different attitudes to income and wealth correspond to differences in short- and long-run attitudes to risk, which could arise for two different reasons.

(i) In the short run, the individual has a large variety of commitments; thus, what an individual could have done (the enjoyment he could have received) from $x if he allocates all of it simultaneously may be markedly different from what he can do (the enjoyment he receives) when he first commits $\frac{1}{2}x$, in the belief that that is all he will have, and *then* is told he has an additional amount of $\frac{1}{2}x$ to spend.

(ii) Individuals' *perceptions* about the value of money (in particular increments in wealth) depend on their actual level of wealth. An individual with a wealth of $1000 is not likely to know how fast diminishing returns sets in, e.g. he may believe that he will be close to being satiated at a wealth of $10 000. But an individual with a wealth of $10 000 knows he is not satiated, but he could believe that satiation might set in at $100 000.

An alternative explanation is that individuals exhibit *bounded rationality*: that is, they limit the amount of data they consider in making decisions in order to simplify the choice.

From the positive point of view, it is obviously convenient that choices appear to depend only on current income risk, for it is then relatively simple to predict responses to risk and changes in risk without enquiring into the wealth position of the individual. However, from a normative point of view the situation is less satisfactory, for it is hard to believe that a small reduction in income risk would, over a long period of time, raise the average welfare of the farmer as much as is suggested by a utility function with constant partial risk aversion.

In most of the rest of the book, we shall be concerned with attitudes to current income risk, and where it is convenient to explore the implications of a particular choice of utility function we shall variously assume constant absolute risk aversion (when the distinction between income and wealth is irrelevant), or constant partial risk aversion, which, from now on, we shall refer to as constant relative risk aversion defined on income, and use the symbol R. In short, in most of the book, we shall follow convention, supported by the evidence, of confining attention to income risk.

The other main conclusions are quickly summarized. Few individuals have risk aversion much above 2, even at very high game levels (where the SD of outcomes is more than one-third average annual incomes; and therefore higher than risks experienced in agriculture). The results contrast sharply with those of Dillon and Scandizzo (1978) in which more than half the farmers typically had risk aversion greater than 3. Binswanger found that the interview technique employed by Dillon and Scandizzo gave very unreliable results when employed on this sample, which suggests caution in accepting such results at face value. Binswanger also correlated risk attitudes against a variety of personal characteristics, but found few clear-cut relationships. Past luck made individuals less risk averse, while wealth and schooling tended to reduce risk aversion, as, to a lesser extent, did salaried employment. Progressive farmers were slightly less risk averse than the average, but age, sex, family composition, and amount of land rented had negligible effect. In all cases the estimated relationship was weak, in the sense that massive changes in these characteristics are needed to change risk aversion substantially.

While the experimental results are consistent with expected utility-maximizing behaviour, they are not consistent with security-based theories of behaviour in which the agent is primarily concerned with achieving a subsistence level of income and avoiding disaster.

To conclude, most individuals are risk averse, but not very risk averse, and react to fluctuations in income rather than consolidating such changes into lifetime wealth. The coefficient of partial risk aversion typically increases from about 0.5 for small fluctuations in income (SD of about one month's wage) to about 1.2 for large fluctuations (SD about 50 per cent of annual income).

7.4 Other empirical studies of attitudes to risk

Most agricultural economists would agree that farmers' attitudes to risk are quantitatively important determinants of their decision-making, especially in less developed countries where risks are relatively larger, incomes lower, and risk-spreading options fewer. Despite this recognition, there are relatively few empirical studies of attitudes to risk and very few indeed which are at all satisfactory. The reason is that it is more difficult to identify the effects of risk on decision taking than almost any other factor (such as changes in prices), and it is difficult enough to study even the most straightforward influences on decision-making. A change in the relative prices of crops is in principle an objective, observable influence (at least, in a world of certainty, or if these prices are quoted on futures markets), while risk is difficult to observe and quantify and remains largely subjective. It takes a large number of observations on a random variable to establish even such comparatively crude measures of its underlying distribution as its mean and variance with any precision, let alone the exact form of the distribution.

Given this difficulty in conducting empirical investigations, it is obviously desirable to collect, compare, and assess as many different attempts as possible to see if any strong pattern emerges, and to learn from experience how best to conduct future investigations. Unfortunately, it is often difficult to interpret the published results because too little of the supporting evidence has been made available. In particular, research workers seem unfamiliar with the desirability of presenting their results in a dimension-free way. In demand studies most economists appreciate the value of providing *elasticities* of demand, evaluated typically at the sample mean, rather than just slope *coefficients* whose values depend on the units in which price and quantity are measured. Most economists presenting evidence on attitudes to risk seem unaware of the fact that most summary measures, such as the coefficient of absolute risk aversion, are not dimensionless, and so cannot be interpreted without a figure for mean income. Only the coefficients of relative and partial risk aversion, defined in equations (7.3) and (7.4), among the common measures, are dimensionless elasticities. Even for these measures, however, it is important to see whether risk aversion varies with income, so it remains useful to present data on income levels and wealth.

In this section we examine two examples of studies which attempt to infer attitudes to risk from observed behaviour in the presence of risk. In the first, Moscardi and de Janvry (1977) postulated a safety-first objective, in which households maximize the income level below which income will fall only a specified proportion (presumably low) of the time. That is, farmers maximize

$$EY - K\sqrt{(\text{Var } Y)}, \tag{7.4}$$

where

$$Y = p\theta f(\mathbf{x}) - \mathbf{w}\cdot\mathbf{x}, \quad E\theta = 1, \text{Var } \theta = \sigma^2, \tag{7.5}$$

and p is the output price, $\mathbf{x}$ is a vector of inputs, and $\mathbf{w}$ the vector of input prices. The solution to the problem is

$$1 - K\sigma = \frac{w_i x_i}{b_i p f(\mathbf{x})}; \quad b_i = \frac{x_i}{f}\frac{\partial f}{\partial x_i}, \tag{7.6}$$

where b_i is the imputed share of input i in average gross output, estimated by agronomists from observations on trial plots.

The results were a mean value of K of 1.12, standard error 0.61, and range between 0 and 2.0. Given Binswanger's rejection of safety-first rules, it is obviously interesting to ask how to interpret the data on the assumption that farmers maximize the expected utility of income

$$\underset{x}{\text{Max}}\ EU(Y) \tag{7.7}$$

which, given equation (7.5) yields

$$\frac{EU'(Y)\theta}{EU'(Y)} = \frac{w_i x_i}{b_i pf}. \tag{7.8}$$

If purchased inputs were negligible, or if riskless income matched the cost of these inputs, or if farmers were concerned with gross income and had constant (partial) relative risk aversion, R, then the left-hand side of equation (7.8) can be expressed as

$$\frac{E\theta^{1-R}}{E\theta^{-R}} \simeq 1 - R\sigma^2,$$

(where we assume that θ is log-normally distributed as $\Lambda(-\frac{1}{2}\sigma^2, \sigma^2)$ and use the expression of equation (6.43) to evaluate the expected values). If we knew the coefficient of variation of yield, σ, which the authors do not give, then we could estimate $R = K/\sigma$. Thus if, for example, $\sigma = 0.5$ (typical for wheat-growing areas in the US) then R would be about 2. Again, though, we are prevented from making full use of the data by the failure of the authors to present relevant information.

Moscardi and de Janvry actually used data on fertilizer inputs, a production function estimated from another, larger experiment, and collected additional socio-economic data from the forty-five farmers studied to relate risk attitudes to these socio-economic variables. As the authors point out, 'since risk aversion is measured as a residual . . . it tends to include other sources of discrepancy . . . such as, for example, imperfect market and agronomic information, restricted availability of financial capital and inputs, and high opportunity cost of family labour' (Moscardi and de Janvry, p. 711). To these sources of error one can add errors in specifying the production function and source of risk (as Cobb–Douglas and multiplicative), errors in the farmer's perception of the production function, errors in his perception of the relationship between price and output (discussed in more detail in 11.2.1), and the problem of identifying the sources of income risk (or income insurance). Since it is so hard for skilled econometricians to estimate the response to fertilizer, it is optimistic to assume that the *only* reason for under-supplying fertilizer is risk aversion (rather than, say, cautious learning behaviour).

In the second study, Schluter and Mount (1976) constructed a single time series of average yields and prices for six consecutive years for Surat District, India. Data from thirty-three unirrigated farms were then used to estimate the mean income and its mean absolute deviation (MAD) using the six-year data. A linear-programming (LP) model of the farm plan was built to find the trade-off between mean income and its MAD. The average MAD was 13 per cent, and the range was from 4 to 25 per cent. Income was defined as the gross value of farm output less purchased inputs, and if the randomness were roughly normal, then its coefficient of variation (CV) would be about 17 per cent, with a range from 5 to 32 per cent.

The LP model can be used to calculate the optimal farm plan for a representative farmer for differing levels of risk, and hence calculate the trade-off between mean income and its coefficient of variation (or more precisely, the mean absolute deviation.) Table 7.2 gives the mean income and coefficient of

Table 7.2 *Risk alternatives for representative unirrigated farm, India*

Income (Rs)	Coefficient of variation,[a] %	Relative risk aversion at which preferred
1530	11	3
1580	12	1.5
1630	14	1
1680	17	0.8
1730	20	0.5
1880	30	0.25
1950	40	0.05

[a]estimated as $\sqrt{(\pi/2)} \times$ MAD

Source: Schluter and Mount (1976)

variation (assumed to be $\sqrt{(\pi/2)} \times$ MAD) at different farm plans, the more risky involving more cash cropping and more fertilizer use. Since the table is for a representative farmer, and since actual farmers did not adopt risk-efficient farm plans (although they were typically within 5 per cent of the efficient frontier) it is difficult to deduce revealed attitudes to risk from the farm plans, since the data is not given in sufficient detail, but given that the average CV was 17 per cent, this corresponds to a coefficient of relative risk aversion of about 0.8.

The difficulty with this approach is that attitudes to risk are deduced from the frontier of an LP model, which will depend sensitively on the parameters of the model, the constraints imposed, and the magnitude of the risk. Many of the actual constraints facing farmers are too subtle to model in a simple LP, for example the seasonality of labour inputs and the extent to which they can be substituted, and one should therefore be cautious about deductions which depend on the shape of the estimated frontier.

7.5 The magnitude of income risk

The rough calculations of Chapter 6 showed that two factors influence decisions, the coefficient of risk aversion and the square of the coefficient of variation of income. How large is the CV of agricultural income? Although it is easy to measure the CV of agricultural yields, prices, and gross revenues, it is much harder to measure the relevant CV, which is that of income *net* of expenses, for the whole farm (and thus, typically, for a whole cropping pattern). Ideally, this should be measured from budget data over a run of years, but often this data is not available. Several writers have estimated the CV of net income from the underlying data on yield and price variability, the correlations between crop revenues, and a model which identifies the optimum farm plan for a given level of risk. In this section we give a few examples to suggest a range of plausible values; obviously the actual values will depend on a host of variables which vary from place to place. In Chapter 20 we provide more aggregative estimates of price, quantity, and revenue variability for some of the 'core' commodities for

which buffer stocks have been proposed, while here we consider a wider range of agricultural activities.

Girao, Tomek, and Mount (1974) provide some evidence from the combined records of both the farm business and household for fifty southern Minnesota farmers for the seven years 1963-9. In many cases rather longer time series were available, and the results were split into two groups – those with dairy farms (and relatively stable incomes) and those without (Table 7.3).

Table 7.3 *Coefficient of variation of income of Minnesota farmers*

Number of farms with CV	Dairy	No dairy
<0.3	16	3
0.3-0.4	4	6
0.4-0.5	2	6
>0.5	1	12
Average income 1965-9	\$9330	\$12 786

Source: Girao, Tomek, and Mount (1974, Table 2).

Evidently the twenty-seven farms without a dairy enterprise experience very unstable incomes, with the median CV between 40 and 50 per cent.

Heady (1952), in his classic textbook, gives a great variety of data on yield and price variability, and some evidence on Iowa livestock enterprises for the period 1918-49. The CV of income varies from 14 per cent for dairy enterprises, to 39 per cent for feeder lambs, with most examples around 34 per cent. Other crops was often very much more variable. Thus the CV of wheat yields cited by Heady (p. 457) ranges from 9 per cent (Montgomery, Penn.) to 92 per cent for Baca, Col., estimated for the period 1926-48. Those counties with more than 60 per cent of the area under wheat averaged over 50 per cent. Other crops, such as rice, appear even riskier (p. 456) though they are often combined in a diversified cropping pattern.

In underdeveloped countries, Roumasset (1976) has calculated the mean and standard deviation of Philippines rice farm incomes from the underlying data, and finds for rain-fed rice using traditional methods a CV of 20 per cent, rising to over 50 per cent for modern techniques (which yield a higher mean income). The modern techniques are superior to traditional methods for relative risk aversion parameters of 4.5 or less. For irrigated rice traditional methods yield a CV of 25 per cent, while modern techniques yield much higher mean income (almost double) with a CV of between 33 and 42 per cent.

Finally, we have the estimates of Schluter and Mount (1976) already given, where the average coefficient of variation of unirrigated Indian farms was about 17 per cent.

To put these various estimates in perspective, we can use the concept of a risk premium, or proportional risk premium, defined in equation (6.11). If the coefficient of relative risk aversion, R, is taken to be 2.0, and a representative CV of income is 33 per cent, then the risk premium is 10 per cent of income. If, on the other hand, the CV is 50 per cent, then the premium rises to 25 per cent of income, showing that high risk can be very costly.

It is apparent, then, that the welfare losses with which we are concerned, arising out of the risks facing farmers, are significant, and that policies which change these risks may have significant welfare consequences.

Chapter 8

Theory of Consumer Demand

8.1 Introduction

In Chapter 6 we showed that it was misleading to measure producers' profits under risk as the area between the price line and the supply curve. The same problem arises in measuring consumer surplus, the counterpart of producer surplus on the demand side. As before, the solution is the same: to develop a microeconomic theory of consumer behaviour and to model risk explicitly.

The purpose of this chapter is briefly to summarize those aspects of the theory of consumer demand which we shall need in the rest of the book, and to develop measures of the benefits of price stabilization. Readers who are interested in extending our analysis are strongly recommended to read Gorman's survey article (1976) which demonstrates the analytical simplification which a judicious choice of variables and functional form can provide. The two basic ideas which simplify the study of consumer demand are those of *duality* and *separability*. We have already seen that the principle of duality allows one to choose the appropriate independent variable with which to work. In production theory one can choose to work with quantities of inputs and outputs, using a production function, or in terms of prices and a profit function. With production, the focus is typically on the input-output choice, but for consumers it is prices that are often more important. We shall therefore derive the properties of the indirect utility function which is defined over prices.

Assumptions about separability impose structure on the problem, and are of central importance in the study of risk. Assumptions about attitudes to risk have strong implications for the separability of indirect utility functions, and vice versa. (See Stiglitz, 1969.)

Moreover, if utility functions are chosen to be separable in certain variables, then these variables will not interact directly. In many cases it is reasonable to impose such conditions explicitly, more often it affords considerable simplification in the analysis, and sometimes it is completely inappropriate as it assumes the answer to the main question. If the object is to study the impact of variable wheat yields on the price of cotton, then it will prejudge the issue to assume that wheat and cotton enter separably. In short, it is important to think carefully about the appropriateness of separability assumptions before starting the analysis and choosing functional forms. Ideally one would choose appropriate variables and the appropriate separability assumptions to suit the problem in question. Unfortunately, there is a difference between assuming that the direct utility function is separable, and assuming that the indirect utility function is separable, and in such cases one of the simplifications will typically have to be sacrificed. We shall return to this issue after deriving the duality results.

8.2* Duality, indirect utility functions, and the expenditure function

We start with the consumer's utility function, $U(\mathbf{q})$, defined over the vector of goods which he consumes, $\mathbf{q}$. This reflects his preferences, in the sense that if the consumer can freely choose between consumption bundles $\mathbf{q}^0$ and $\mathbf{q}^1$, and chooses, or reveals a preference for $\mathbf{q}^0$, then $U(\mathbf{q}^0) > U(\mathbf{q}^1)$. For most consumer theory, $U(\mathbf{q})$ can be thought of as an ordinal function, no more or less satisfactory than any monotonically increasing transform $\phi(U(\mathbf{q}))$, but this is not true if the consumer's choice under risk is to be described as one of maximizing expected utility as in the present context. In this case, $U(\mathbf{q})$ is defined up to an increasing linear transformation, i.e. if an individual's behaviour can be described *as if* he maximized $U(\mathbf{q})$, it could equally well be described as if he maximized $a + bU(\mathbf{q})$, $b > 0$. This amounts to requiring that $U(\mathbf{q})$ be a cardinal measure of satisfaction, while leaving the choice of origin and units of utility arbitrary. (It should be stressed that it is not necessary to represent choices as if the consumer had a cardinal utility function; merely convenient in that it allows the use of expected utilities. It is perfectly possible, though inconvenient, to define choices over outcomes and their probabilities without imposing additivity in the probabilities, see, e.g., Green, 1971, § 13.3.)

Just as it is convenient to derive cost and profit functions from the production function, so it is useful to derive exactly parallel expenditure and indirect utility functions from the direct utility function. The same duality relationships also hold, and are most readily derived from the expenditure function, which is defined as the minimum expenditure needed to achieve a given level of utility U at prices $\mathbf{p}$:

$$g(\mathbf{p}, U) = \underset{\mathbf{q}}{\text{Min}}\ \mathbf{p} \cdot \mathbf{q} \text{ such that } U(\mathbf{q}) \geqslant U. \tag{8.1}$$

(The dot product of the two vectors $\mathbf{p}$ and $\mathbf{q}$, written $\mathbf{p} \cdot \mathbf{q}$ or, more loosely, as $\mathbf{pq}$, is defined as the scalar

$$\mathbf{p} \cdot \mathbf{q} \equiv \sum p_i q_i,$$

where p_i is the i-th component of the n-vector $\mathbf{p}$. In the present context it is just the cost of the bundle of goods represented by the vector $\mathbf{q}$.)

The indirect utility function gives the level of utility achievable with lump-sum income I and prices p:

$$V(\mathbf{p}, I) = \underset{\mathbf{q}}{\text{Max}}\ U(\mathbf{q}) \text{ such that } \mathbf{p} \cdot \mathbf{q} \leqslant I. \tag{8.2}$$

(Lump-sum income is income that does not depend on any of the consumer's current consumption choices. Wage income is best treated as a lump-sum income equal to the wage rate, w, times twenty-four hours per day, part of which is spent on the purchase of the n-th consumption good, leisure, whose price $p_n = w$. In most of the book neither lump-sum nor wage income will depend on the variables under study, except in Chapter 25, where it becomes important to be precise about the definition of income.)

The relationship between the expenditure function and the indirect utility function is that for all levels of utility, U

$$V(\mathbf{p}, g(\mathbf{p}, U)) \equiv U. \tag{8.3}$$

The following are important properties of these two functions:

1. *The expenditure function is concave in prices*

This is equivalent, by definition, to

$$g(\lambda\mathbf{p}^0 + (1-\lambda)\mathbf{p}^1, U) \geqslant \lambda g(\mathbf{p}^0, U) + (1-\lambda)g(\mathbf{p}^1, U), \quad 0 \leqslant \lambda \leqslant 1, \tag{8.4}$$

for any pair of price vectors, $\mathbf{p}^0$ and $\mathbf{p}^1$, and utility level U. Define the weighted average price

$$\mathbf{p}^\lambda \equiv \lambda\mathbf{p}^0 + (1-\lambda)\mathbf{p}^1, \quad 0 \leqslant \lambda \leqslant 1, \tag{8.5}$$

and suppose that the cost-minimizing choice of consumption which achieves U at any price $\mathbf{p}^i$ is $\mathbf{q}^i$, so that,

$$g(\mathbf{p}^i, U) = \mathbf{p}^i \cdot \mathbf{q}^i \leqslant \mathbf{p}^i \cdot \mathbf{q}^j, \quad j \neq i.$$

The inequality follows because any other consumption bundle $\mathbf{q}^j$ which achieves U must cost more at prices $\mathbf{p}^i$, or else it would have been chosen instead. Then

$$\begin{aligned} g(\mathbf{p}^\lambda, U) &= \mathbf{p}^\lambda \cdot \mathbf{q}^\lambda = \lambda\mathbf{p}^0 \cdot \mathbf{q}^\lambda + (1-\lambda)\mathbf{p}^1 \cdot \mathbf{q}^\lambda, \\ &\geqslant \lambda\mathbf{p}^0 \cdot \mathbf{q}^0 + (1-\lambda)\mathbf{p}^1 \cdot \mathbf{q}^1, \end{aligned}$$

which is equivalent to equation (8.4), so the expenditure function is concave in prices. It follows that the average cost of achieving a given level of utility at random prices is less than the cost of achieving the same level of utility at prices stabilized at their mean (as another example of Jensen's inequality, see section 6.3).

2. *The indirect utility function is quasi-convex in prices and homogenous of degree zero*

A function $V(\mathbf{p}, I)$ is said to be quasi-convex in $\mathbf{p}$ for given I if for given U, the set of vectors satisfying

$$V(\mathbf{p}, I) \leqslant U$$

is convex. To prove quasi-convexity we need to show that

$$V(\mathbf{p}^i, I) \leqslant U \quad i = 0, 1$$

implies

$$V(\mathbf{p}^\lambda, I) \leqslant U$$

when $\mathbf{p}^\lambda$ is defined in equation (8.5). Suppose not, i.e.

$$V(\mathbf{p}^\lambda, I) = U(\mathbf{q}^\lambda) > U,$$

then

$$\mathbf{p}^\lambda \cdot \mathbf{q}^\lambda > \mathbf{p}^i \cdot \mathbf{q}^\lambda \quad i = 0, 1,$$

or else $\mathbf{q}^\lambda$ would have been chosen at prices $\mathbf{p}^i$. But this is impossible by the definition of $\mathbf{p}^\lambda$, hence V is quasi-convex in price. Homogeneity follows by noting that utility is unchanged by an equal proportional change in all prices and money income.

3. *Compensated demand* for good i is given by

$$q_i^c = \frac{\partial g}{\partial p_i} = q_i^c(\mathbf{p}, U). \tag{8.6}$$

This follows by noting first that the direct utility function is maximized at given prices, hence so is the Lagrangian

$$L = U(\mathbf{q}) + \lambda(I - \mathbf{p} \cdot \mathbf{q}),$$

whence

$$\frac{\partial U}{\partial q_j} = \lambda p_j. \tag{8.7}$$

Differentiate equation (8.1)

$$\frac{\partial g}{\partial p_i} = q_i + \sum_j p_j \frac{\partial q_j}{\partial p_i} = q_i + \frac{1}{\lambda} \sum_j \frac{\partial U}{\partial q_j} \cdot \frac{\partial q_j}{\partial p_i} \tag{8.8}$$

$$0 = \frac{dU}{dp_i} = \sum_j \frac{\partial U}{\partial q_j} \frac{\partial q_j}{\partial p_i}$$

and the second term of equation (8.8) vanishes, yielding the result.

The compensated demand schedule differs from the normal Marshallian or uncompensated demand schedule in that the consumer is kept at the same level of utility by money transfers as prices change, rather than holding money income constant and allowing utility to change. It is of central importance for welfare analysis, as the next result shows.

4. *The expenditure function is a direct measure of consumer surplus*
The consumer requires an amount of money equal to

$$CS = g(\mathbf{p}^1, U^0) - g(\mathbf{p}^0, U^0)$$

to compensate for moving from one set of prices $\mathbf{p}^0$ (and initial utility U^0) to another set of prices $\mathbf{p}^1$. This can be expressed as

$$CS = \int_{\mathbf{p}^0}^{\mathbf{p}^1} \sum_i \frac{\partial g}{\partial p_i} dp_i = \int_{\mathbf{p}^0}^{\mathbf{p}^1} \sum_i q_i^c dp_i. \tag{8.9}$$

If only one price varies, p_i, say, then this measure of consumer surplus is the area under the *compensated* demand curve. If several prices vary, then in general there is no simple geometric equivalent and it can be very misleading to measure areas separately, holding all other prices constant.

5. *Roy's identity* gives the uncompensated or Marshallian demands

$$q_i = -\frac{\partial V/\partial p_i}{\partial V/\partial I}. \tag{8.10}$$

Differentiate equation (8.3) holding U constant:

$$0 = \frac{dU}{dp_i} = \frac{\partial V}{\partial p_i} + \frac{\partial V}{\partial I}\frac{\partial g}{\partial p_i}$$

$$= \frac{\partial V}{\partial p_i} + q_i \frac{\partial V}{\partial I}$$

using equation (8.6). This gives equation (8.10) by rearrangement, and since demand is now a function of *income* and price we have found the uncompensated demands. Naturally these coincide at the point where income is sufficient to yield the reference level of utility specified in the compensated demands. In general they differ elsewhere as the next result shows.

6. *Slutsky's theorem*

$$\frac{\partial q_i}{\partial p_j} + q_j \frac{\partial q_i}{\partial I} = \frac{\partial q_j}{\partial p_i} + q_i \frac{\partial q_j}{\partial I}. \tag{8.11}$$

This follows if the expenditure function is twice continuously differentiable, for if we define

$$S_{ij} \equiv \frac{\partial^2 g}{\partial p_i \partial p_j}$$

then, since the compensated and uncompensated demands coincide for a given level of utility, equation (8.6) gives

$$q_i^c = \frac{\partial g}{\partial p_i} = q_i\{\mathbf{p}, g(\mathbf{p}, U)\}$$

$$S_{ij} = \frac{\partial}{\partial p_j} q_i\{\mathbf{p}, g(\mathbf{p}, U)\} = \frac{\partial q_i}{\partial p_j} + \frac{\partial q_i}{\partial I} \cdot \frac{\partial g}{\partial p_j}$$

$$S_{ij} = \frac{\partial q_i^c}{\partial p_j} = \frac{\partial q_i}{\partial p_j} + \frac{\partial q_i}{\partial I} \cdot q_j, \tag{8.12}$$

then clearly the Slutsky symmetry condition is satisfied, for the order of differentiation does not matter:

$$S_{ij} = S_{ji}.$$

This result shows that the uncompensated and compensated demand schedules will only coincide if income effects are zero.

Finally, notice the duality between the direct and indirect utility functions which emerges very clearly in equations (8.7) and (8.10), especially when it is realized that λ in equation (8.7), being the shadow price of the income constraint, is the marginal utility of income, $\partial V/\partial I$.

8.3* Attitudes towards risk

The results of the previous section depended only on the ordinal properties of the consumer's utility function. In other words, the consumer's behaviour can be described *as if* he maximized a utility function $U(\mathbf{q})$, and any other cardinalization of the utility function $\phi(U(\mathbf{q}))$ with $\phi' > 0$ would do as well. However, as was remarked in the Introduction it is convenient to describe the consumer's choices under risk as though he maximized expected utility, in which case the precise cardinalization of the utility function does matter. In Chapter 6, we discussed how the producer's behaviour could be analysed in the presence of risk. There, we showed how if he were averse to risk, we could describe his behaviour as if he maximized a concave utility function of income $EU(Y)$.

For consumers we need to distinguish between attitudes towards income variability at fixed prices, and attitudes towards price variability at fixed incomes. For both it is convenient to use the indirect utility function $V(\mathbf{p}, I)$. It is natural to hypothesize that individuals are averse to variations in income at fixed prices, i.e. $V_{II} < 0$. In the limiting case, where individuals are neutral to income variability (at every set of prices), we require the indirect utility function to be of the form

$$V(\mathbf{p}, I) = v(\mathbf{p}) + w(\mathbf{p})I,$$

where v is homogenous of degree zero and w is homogenous of degree -1 in prices (see Stiglitz, 1969).

On the other hand, there appears to be no natural restriction to impose on consumers' attitudes towards price variability. We have already pointed out that the analysis of changes in price dispersion, keeping mean price constant, was in general of dubious value. But even if we consider a mean-preserving increase in price risk, under perfectly reasonable conditions consumers' welfare may increase, decrease, or remain unchanged depending on whether utility is convex, concave, or linear in that price. To see which holds, we need to ascertain the sign of V_{pp}. Recall that from Roy's identity (equation (8.10))

$$V_p = -qV_I.$$

Hence, differentiating with respect to I

$$V_{pI} = -\frac{qV_I}{I}\cdot\left(\frac{Idq}{qdI} + \frac{IV_{II}}{V_I}\right) = \frac{qV_I}{I}\cdot(R^c - \eta), \tag{8.13}$$

where

$$\eta = \frac{d\log q}{d\log I}, \text{ the income elasticity of demand}$$

and

$$R^c = -\frac{IV_{II}}{V_I}, \text{ the consumer's relative risk aversion (to income variability).} \tag{8.14}$$

Differentiate Roy's identity with respect to p:

$$V_{pp} = \frac{qV_I}{p}\left(-\frac{pdq}{qdp}\right) - qV_{Ip}$$

$$V_{pp} = \frac{qV_I}{p}\{\epsilon - \beta(R^c - \eta)\}, \tag{8.15}$$

where

$$\epsilon = -\frac{d\log q}{d\log p}, \text{ the price elasticity of demand,}$$

and

$$\beta = \frac{pq}{I}, \text{ the expenditure share on the commodity.}$$

Although in general we would expect ϵ to be positive (as defined, i.e. demand schedules have a negative slope) and larger than $\beta(R^c - \eta)$, nevertheless this is not guaranteed, and indeed proposition 2 above only proved that V was quasi-convex in prices, not convex. Consequently, it is quite possible that consumers would prefer price variability to prices stabilized at their mean, and also possible that they would prefer price stability. However, they will prefer stable consumption to fluctuating consumption, since $U(\mathbf{q})$ is concave in $\mathbf{q}$.

Finally, it turns out that for many of the issues we shall be interested in, what is crucial is the effect of price variations on the marginal utility of income, i.e. the sign of

$$V_{Ip} = \frac{qV_I}{I}(R^c - \eta)$$

which, as is apparent, may be either greater or less than zero. For homothetic indifference maps (whose slopes are constant along rays through the origin, i.e. the income elasticity of demand is unity, $\eta = 1$) whether the marginal utility of income increases as price increases depends simply on whether the aversion to variability of income (R^c) is greater than 1.

If $V_{Ip} = 0$, price has no effect on marginal utility, and hence the indirect utility function must be separable:

$$V(\mathbf{p}, I) = v(p_1, p_2, \ldots, p_n) + w(I, p_2, \ldots, p_n),$$

where $p = p_1$, and both v and w are homogenous of degree zero. If all other prices $p_2, \ldots, p_n$, are constant, this can be abbreviated to

$$V(\mathrm{p}, I) = v(\mathrm{p}) + w(I),$$

in which case, by Roy's identity, the demand schedule is

$$q = -v'/w'$$

and the income elasticity is

$$\eta = -\frac{w''I}{w'}$$

and for this to be unity,

$$w(I) = a \log I + b. \tag{8.16}$$

Similarly, if the price elasticity is to be unity

$$V = -\alpha \log p + w(I). \tag{8.17}$$

These special parameterizations will play a key role in much of the analysis to follow, and will be considered further in the next section. What is important, however, to remember is that even if the individual is very averse to variations in income, he may not be so averse to variations in prices: the two are quite distinct and should not be confused.

8.4 The specification of demand risk and utility functions

In an ideal world if consumers acted as though they maximized expected utility and had consistent expectations and preferences, it would be possible to estimate econometrically the cardinal utility function describing their behaviour and identify the nature of the underlying risk. In practice, data are sparse, of poor quality, and relate to time periods over which expectations are unlikely to be stable. Moreover, most data refer to aggregate consumption, and there are few utility functions which generate individual demands which can be aggregated to give a total demand of the same functional form. One of the few utility functions which does aggregate is the Stone-Geary utility function

$$U = \sum_i \beta_i \log(q_i - c_i); \sum \beta_i = 1, \tag{8.18}$$

where c_i is to be understood as the minimum acceptable level of consumption of good i. This gives demands

$$q_i = c_i + \frac{\beta_i(I - \mathbf{p} \cdot \mathbf{c})}{p_i}, \tag{8.19}$$

which has linear Engel curves (relating q_i to income I) and expenditures $p_i q_i$ are linear functions of income, hence the system is often referred to as the linear expenditure system. The indirect utility function corresponding to the linear expenditure system is

$$V(p, I) = \log(I - \mathbf{p} \cdot \mathbf{c}) - \sum \beta_i \log p_i - \sum \beta_i \log \beta_i. \tag{8.20}$$

We have already encountered a special form of this function, for if $c = 0$, then relative risk aversion and all price and income elasticities are unity, and $V_{Ip} = 0$.

In general, one would not wish to make such restrictive assumptions, and so one should usually start with fairly general utility functions and impose appropriate, testable, empirical restrictions to explore their consequences – usually that elasticities are constant, or cross-elasticities are zero. This raises an important set of methodological questions, of which the most important is whether these empirical restrictions are consistent with any underlying utility function. We have already remarked that specific assumptions about attitudes to risk have strong implications for the structure of utility functions, and the converse is true, at least, once a particular cardinalization has been chosen (derived from, e.g., attitudes to income risk). For example, suppose we wished to impose the condition of constant income and price elasticities for two goods (in order to study the effect on one good of stabilizing another price). This suggests an indirect utility function which is some monotonic function of

$$V = v(\mathbf{p}_0) p_1^{1-\epsilon_1} p_2^{1-\epsilon_2} + \frac{1}{1-\eta}\left\{\frac{I}{g(\mathbf{p}_0)}\right\}^{1-\eta} + w(\mathbf{p}_0), \tag{8.21}$$

where $\mathbf{p}_0$ is a vector of all other prices, $v(\mathbf{p}_0)$ is homogenous of degree $\epsilon_1 + \epsilon_2 - 2$, g is homogenous of degree 1, and w is homogenous of degree 0. This form restricts the cross-elasticities of demand between the two goods, and if in addition the coefficient of relative risk aversion is to be constant, it must be equal to the income elasticity, η. Moreover, the aggregate demand function will not in general have constant income elasticity, making it difficult to argue that consumers in the aggregate behave as if they could be replaced by a single 'representative' consumer with a specific utility function.

We shall side-step all these problems by restricting our attention to small changes in prices or incomes so that it is reasonable to assume various parameters are locally constant without prejudging the general form of the utility function, and hence without necessarily imposing strong global restrictions on other parameters.

In order to estimate a demand system it is necessary to specify the form of demand risk. In Chapter 18 we shall demonstrate that different specifications lead to quite different predictions about the distributive impact of price stabilization, and so it is important to ask how reasonable the various alternatives are.

Since we are interested in studying the effects of the systematic variability in prices, we shall ignore non-systematic sources of variability, of which the most important are periodic changes in technology which affect the demand for goods. The invention of synthetic fabrics had an important effect on the demand for cotton, but is best viewed as a, perhaps intermittent, secular trend. Since invention occurs randomly, this secular trend will be random, making it very difficult to determine the level around which prices should be stabilized.

The two primary sources of systematic, or predictable, variability arise from variations in income (over the trade cycle) and variations in the prices of other commodities, perhaps caused by fluctuations in supply elsewhere.

The simplest econometric specification of a demand system is to assume constant elasticities:

$$Q_i = Ap_i^{-\epsilon_i} I^{\eta_i} \prod_{j \neq i} p_j^{\epsilon_{ij}}, \tag{8.22}$$

in which case the natural specification for demand risk is *multiplicative*, for either fluctuations in income or other prices will, on this specification, affect demand multiplicatively. The natural method for estimating this equation would be log-linear regression, and here too the errors are conveniently assumed multiplicative. (See Turnovsky (1976) for a defence of this assumption.)

The only serious alternative formulation is to suppose that risks are additive, which is consistent with a linear econometric specification. It is difficult to accept that demand is a linear function of income and prices, and that each of these has additive risk, as in

$$Q_i = \alpha_{i0} + \sum_{j \neq i} \alpha_{ij}(p_j + \tilde{u}_j) + \beta_i(I + \tilde{u}_i), \tag{8.23}$$

except, perhaps, as a local approximation. Although different risk specifications can affect the distributions of the benefits of price stabilization, the real difference lies not between additive and multiplicative risk, but rather whether it is the demand or inverse demand curve which is affected. The formulation

$$p = \theta g(Q, I)$$

(multiplicative shifts in the inverse demand curve, shown in Fig. 8.1a) leads to a very different (and implausible) kind of demand shift from the multiplicative demand shift shown in Fig. 8.1b, and adopted here:

$$Q = \theta D(p, I).$$

(Additive shifts are vertical or horizontal displacements of the demand schedule, and it makes little obvious difference whether it is the demand or inverse demand which is shifted.)

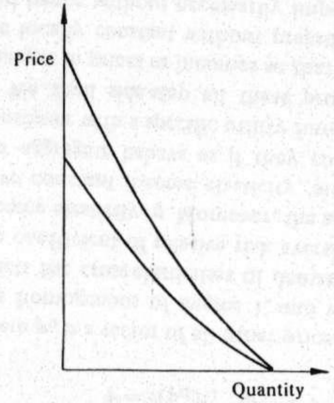

Fig. 8.1a Multiplicative inverse demand shifts

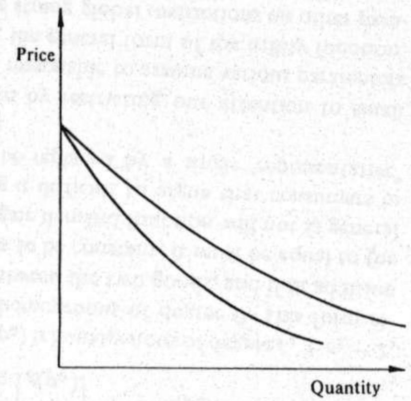

Fig. 8.1b Multiplicative demand shifts

Chapter 9*

Consumer Benefits of Price Stabilization

9.1 Introduction

The conventional measurement of the consumer benefits of price stabilization compares the Marshallian measure of consumer surplus (the area between the demand curve and the price) before and after the price is stabilized at its mean. This approach makes three errors.

(i) First, and least important, the Marshallian measure is only an approximation to the value of consumer surplus. In the first place, it only measures the cash value to the consumer accurately when the compensated and uncompensated demand curves coincide, as equation (8.9) showed. This requires a zero income elasticity of demand. Next, the value of the sum of money is only a good measure of value if the marginal utility of money remains constant, or $V_{Ip} = 0$ (see equation (8.13)). In some cases the two errors cancel (with logarithmic utility functions), but even when they do not the errors involved in using Marshallian measures are small, as Willig (1976) has shown.

(ii) The second, far more important error, is that it is rare for all other prices and incomes to be constant, especially in the context of risk analysis. If the demand *schedule* fluctuates, then this must be because of some more fundamental reason, such as some other price or income is varying. To use then a Marshallian measure as a means of calculating the benefits of, say, stabilizing this particular price, may be seriously misleading. To cite an example, suppose the demand for coffee fluctuates because the price of tea does. When the price of tea is low, the consumer substitutes tea for coffee, the demand for coffee drops, and the consumer is better off. However, the area under the coffee demand curve has fallen, and if used as a measure of the consumer's welfare would give exactly the wrong answer.

(iii) The third error lies in assuming that it is possible to stabilize a price at its mean. When dealing with price stabilization for producers we argued that in the short run inputs would remain fixed, in which case average supply would remain constant. We have already argued that unless demand schedules are linear, the average price will differ from the price of average supply. In the long run supply will adjust, and then, even with linear demand schedules, the price of average supply will change. The correct price around which to stabilize is the price at which average supply equals average demand.

We shall derive measures of the benefits of price stabilization for consumers which parallel those derived for producers in Chapter 6. However, since the formulae are more complex, we shall only calculate the benefits of completely stabilizing the price of the first commodity, leaving other prices

and income random, and leave the derivation of the benefits of partial stabilization or of simultaneously stabilizing several prices to the reader.

9.2 Efficiency and transfer benefits

The cash value to the representative consumer of completely eliminating the randomness in the first price, p_1, is B, the solution to the equation

$$EV(\tilde{p}_1, \tilde{p}_2, \ldots, \tilde{p}_n, \tilde{I}) = EV(\hat{p}_1, \tilde{p}_2, \ldots, \tilde{p}_n, \tilde{I} - B). \tag{9.1}$$

B is the amount of money which the consumer would give up in return for the price of the first commodity being stabilized at $\hat{p}_1$, leaving the average level of utility constant. The actual value of $\hat{p}_1$ will depend on the type of stabilization, but we shall assume that the mean output is held constant, so $\hat{p}_1$ is defined by the equation

$$\bar{q}_1 = ED_1(\tilde{p}_1, \tilde{p}_2, \ldots, \tilde{p}_n, \tilde{I}) = ED_1(\hat{p}_1, \tilde{p}_2, \ldots, \tilde{p}_n, \tilde{I}),$$

where D_1 is the demand for the first commodity, given by Roy's identity, equation (8.10). In general $\hat{p}_1$ will differ from the average price before stabilization, $\bar{p}_1$. Equation (9.1) can be solved, provided the coefficients of variations are small, by expanding each side in a Taylor series about $\bar{p}_1$, $\bar{I}$, as shown in the Appendix to this chapter, to yield:

$$B \cong -q_1(\bar{p})(\hat{p}_1 - \bar{p}_1) - \tfrac{1}{2}\bar{p}_1 q_1(\bar{p})\{\epsilon_1 \sigma_{p_1}^2 - 2 \sum_{i \neq 1} \epsilon_{i1} \rho(p_i, p_1)\sigma_{p_i}\sigma_{p_1} \tag{9.2}$$
$$+ 2(R - \eta_1)\rho(p_1, I)\sigma_{p_1}\sigma_I\},$$

where ϵ_1 is the own price elasticity of demand of good 1, defined to be positive for normal goods, ϵ_{ij} is the cross-price elasticity of demand defined to be positive for substitutes, σ_x is the coefficient of variation of x, and $\rho(x, y)$ is the correlation coefficient between variables x and y. All these terms are defined in the Appendix.

From now on, we shall drop the subscript 1, so that p and q are respectively the price and quantity of the first good. Although equation (9.2) is in a convenient form for calculating total consumer benefits, it does not distinguish between *transfer* benefits and the *efficiency* or *risk* benefits. As was pointed out in Chapter 6, the change in consumers' average expenditure is exactly matched by an equal and opposite change in the average revenue of producers (and/or the stabilizing authority), and as such it is a transfer and not a net social gain. The first term of equation (9.2) will be equal to the transfer benefit if supply is riskless, for then $q(\bar{p}) = q(\hat{p}) = \bar{q}$, which is constant, but otherwise it bears no simple relationship to any of the terms in the formula, for the simple reason that the formula was derived by expanding about mean price, $\bar{p}$, and not the price at mean quantity, $\hat{p} = p(\bar{q})$. The transfer benefit is

$$B_T = Epq - \hat{p}\bar{q}$$

or, writing $X(p) = pq(p)$, consumers' expenditure on the commodity,

$$B_T = EX(p) - X(\hat{p}). \qquad (9.3)$$

The first term of equation (9.2) is, dropping subscripts,

$$(\bar{p} - \hat{p})q(\bar{p}) = X(\bar{p}) - \hat{p}q(\bar{p}). \qquad (9.4)$$

Since we are interested in distinguishing between transfer and efficiency benefits, we shall use equation (9.3) to derive the transfer benefits, and obtain the efficiency benefits as the difference between the total benefits, given by equation (9.2), and the transfer benefits. The efficiency benefits will depend on the source of the instability, for this will affect the size and sign of the co-variance terms. If the source lies in the variability of other prices, the crucial determinant will be the cross-price elasticities of demand, while if the source is fluctuations in income, income elasticities will be important. Not much more can be said without making specific assumptions, and to these we now turn.

9.3 Income variability alone

One of the more important sources of demand fluctuations for primary commodities arises from the trade cycle, that is, from the variability of income in the importing countries. Consider for simplicity the risk benefit when supply does not vary and when all cross-price elasticities are zero. The price variability induced by income variability is readily calculated. Since

$$q = q(p, I)$$

is fixed,

$$\frac{dp}{dI} = -\frac{\partial q/\partial I}{\partial q/\partial p} = \frac{p}{I}\frac{\eta}{\epsilon}$$

so that the coefficient of variation is

$$\sigma_p = \frac{\eta}{\epsilon}\sigma_I.$$

Since quantities are constant, the first term of equation (9.2), given by equation (9.4), reduces to the transfer benefit of equation (9.3). With no other sources of risk, all prices will be perfectly positively correlated with income fluctuations, so that equation (9.2) reduces to

$$B = B_T - \tfrac{1}{2}X(\bar{p})\{\epsilon\sigma_p^2 + 2(R^c - \eta)\sigma_p\sigma_I\}. \qquad (9.5)$$

If we consider the special case of constant price and income elasticities discussed in section 8.4, equation (8.21):

$$q = p^{-\epsilon}I^{\eta} \quad I = \phi\bar{I},\ E\phi = 1, \qquad (9.6)$$

and if income risk is multiplicative, as in the natural specification, then

$$B_T = \bar{q}(\bar{p} - \hat{p}) = X(\bar{p})\{E\phi^{\eta/\epsilon} - (E\phi^{\eta})^{1/\epsilon}\}.$$

This can be evaluated using the methods of section 6.6 and is approximately

$$B_T = \tfrac{1}{2}X(\bar{p})\frac{\eta^2}{\epsilon^2}(1-\epsilon)\sigma_I^2 = \tfrac{1}{2}X(\bar{p})(1-\epsilon)\sigma_p^2. \qquad (9.7)$$

The efficiency benefit is the second term of equation (9.5):

$$B_E = -\tfrac{1}{2}X(\bar{p})\frac{\epsilon}{\eta}(2R^c - \eta)\sigma_p^2. \qquad (9.8)$$

For most agricultural commodities the income elasticity is small so that the efficiency benefit of price stabilization will usually be negative for consumers. The total benefit may also be negative, for

$$\frac{B}{X} = \tfrac{1}{2}\left(1 - \frac{2\epsilon R^c}{\eta}\right)\sigma_p^2. \qquad (9.9)$$

(Notice that since the difference between $X(\bar{p})$, $X(\hat{p})$, and $EX(p)$ is of order of magnitude σ_p^2, it is not necessary to distinguish between them in equations such as (9.9).)

9.4 Supply variability alone

Again suppose that the cross-elasticities of demand are sufficiently small that supply risk does not spill over into other markets. With multiplicative supply risk and constant elasticity of demand the transfer benefit is

$$B_T = EX(q) - X(\bar{q}) \cong X(\bar{q})(E\theta^{1-1/\epsilon} - 1),$$

where

$$\sigma_\theta = \sigma_q = \epsilon\sigma_p$$

Hence

$$B_T \cong \tfrac{1}{2}X(\bar{p})(1-\epsilon)\sigma_p^2. \qquad (9.10)$$

The total benefit is, from equation (9.2)

$$B = q(\bar{p})\bar{p}(E\theta^{-1/\epsilon} - 1) - \tfrac{1}{2}X\epsilon\sigma_p^2$$
$$B = \tfrac{1}{2}X\sigma_p^2, \qquad (9.11)$$

so that the efficiency benefit is

$$B_E = B - B_T = \tfrac{1}{2}X\epsilon\sigma_p^2. \qquad (9.12)$$

Here the efficiency benefit is unambiguously positive, as is the sum of the two terms. The same result can be derived directly from the indirect utility function by expanding about $\bar{q}$:

$$V(p) = V\{p(\bar{q})\} + (q - \bar{q})\frac{dV(\bar{q})}{dq} + \tfrac{1}{2}(q - \bar{q})^2\frac{d^2V}{dq^2}.$$

For constant elasticity demand curves

$$\frac{dV}{dq}=\frac{\partial V}{\partial p}\frac{dp}{dq}=\frac{\partial V}{\partial I}\frac{p}{\epsilon},$$

$$\frac{d^2V}{dq^2}\approx-\frac{\partial V}{\partial I}\frac{p}{q\epsilon^2},\ \text{if}\ \frac{\partial^2 V}{\partial p\,\partial I}\approx 0.$$

Hence the cash value of the total benefits from stabilization is approximately

$$B=\frac{V\{p(\bar{q})\}-EV(p)}{\partial V/\partial I}=\tfrac{1}{2}X\sigma_p^2. \tag{9.13}$$

9.5 Combining demand and supply variability

It would be convenient if the combined effect of income and supply variability were simply the sum of the two taken separately. We shall demonstrate that this is indeed the case under the assumptions made so far, that cross-price elasticities are zero, provided only that the income and supply variabilities are uncorrelated. In practice they are likely to be somewhat correlated for the following reason. Since the trade cycle has a period of between 4 and 8 years, this year's income is correlated with last year's income, which will be correlated with last year's price. Rational and myopic farmers alike will therefore adjust supply in the light of last year's price, introducing a correlation between income and supply. In our model, though, we are holding inputs constant, so this effect has been assumed away and the independence assumption seems reasonable.

If demands are independent, we have, by inverting the demand curve of equation (9.6):

$$p=q^{-1/\epsilon}I^{\eta/\epsilon}.$$

Hence if q and I are independent

$$\sigma_p^2=\frac{1}{\epsilon^2}\sigma_q^2+\left(\frac{\eta}{\epsilon}\right)^2\sigma_I^2, \tag{9.14}$$

which is just the sum of the supply and income effects separately. If the correlation coefficient between price and income is ρ, then $\rho^2=\lambda$ is the fraction of price variability caused by income variability:

$$\lambda\equiv\rho^2=\frac{(\eta/\epsilon)^2\sigma_I^2}{\sigma_p^2}. \tag{9.15}$$

Thus the terms in the braces of equation (9.2) are just the sum of the two impacts:

$$-\tfrac{1}{2}X\left\{\epsilon\sigma_p^2+2(R^c-\eta)\frac{\eta}{\epsilon}\sigma_I^2\right\}$$

$$=-\tfrac{1}{2}X\left\{\epsilon\frac{\sigma_q^2}{\epsilon^2}+(2R^c-\eta)\frac{\eta}{\epsilon}\sigma_I^2\right\} \tag{9.16}$$

from equation (9.14).

The transfer benefit is

$$B_T=Epq-\hat{p}\bar{q}.$$

If $q=\bar{q}\theta$, $I=\bar{I}\phi$, where θ and ϕ are independent, then

$$B_T\cong X(\bar{p})\{E\theta^{1-1/\epsilon}\phi^{\eta/\epsilon}-(E\phi^\eta)^{1/\epsilon}\}$$

$$\cong\tfrac{1}{2}X(\bar{p})(1-\epsilon)\left\{\frac{1}{\epsilon^2}\sigma_q^2+\left(\frac{\eta}{\epsilon}\right)^2\sigma_I^2\right\}$$

or, from equation (9.14)

$$B_T\cong\tfrac{1}{2}X(\bar{p})(1-\epsilon)\sigma_p^2 \tag{9.17}$$

as before.

The total benefit is the sum of the two terms of equation (9.2). The first term is

$$(\bar{p}-\hat{p})q(\bar{p})\approx X(\bar{p})\{E\theta^{-1/\epsilon}\phi^{\eta/\epsilon}-(E\phi^\eta)^{1/\epsilon}\}$$

$$=\tfrac{1}{2}X\left\{(1+\epsilon)\frac{\sigma_q^2}{\epsilon^2}+(1-\epsilon)\frac{\eta^2}{\epsilon^2}\sigma_I^2\right\}$$

$$=\tfrac{1}{2}X(1-\epsilon)\sigma_p^2+X\epsilon\frac{\sigma_q^2}{\epsilon^2}=B_T+X\epsilon\frac{\sigma_q^2}{\epsilon^2}.$$

Thus

$$B=B_T+\tfrac{1}{2}X\left\{\epsilon\left(\frac{\sigma_q^2}{\epsilon^2}\right)-(2R^c-\eta)\frac{\eta}{\epsilon}\sigma_I^2\right\}.$$

Hence, from equation (9.15)

$$B_E=\tfrac{1}{2}X\left\{\epsilon(1-\lambda)-\lambda\frac{\epsilon}{\eta}(2R^c-\eta)\right\}\sigma_p^2$$

or

$$B_E=\tfrac{1}{2}X\epsilon\left(1-\frac{2R^c\lambda}{\eta}\right)\sigma_p^2. \tag{9.18}$$

Unless consumers have very low risk aversion, or unless most of the price variability originates on the supply side, the consumer risk benefit will be negative (though small).

Total consumer benefit is the sum of equations (9.17) and (9.18):

$$B = \tfrac{1}{2}X\left(1 - \frac{2R^c\lambda\epsilon}{\eta}\right)\sigma_p^2, \tag{9.19}$$

where $\lambda = \rho^2$, the squared correlation coefficient of price on income. It therefore seems justified to treat the various sources of risk additively, provided that they are independent for the simple reason that the formulae contain only terms which are linear in expenditure (and hence additive) or linear in variance (also additive). Therefore it suffices to consider the remaining sources of price variability in isolation.

Appendix: Derivation of formulae

The left-hand side of equation (9.1) yields the following Taylor expansion

$$V(\mathbf{p}, I) = V(\bar{\mathbf{p}}, \bar{I}) + \sum_i \frac{\partial V}{\partial p_i}(p_i - \bar{p}_i) + \frac{\partial V}{\partial I}(I - \bar{I}) \tag{9.A1}$$
$$+ \tfrac{1}{2}\left\{\sum_i \sum_j \frac{\partial^2 V}{\partial p_i \partial p_j}(p_i - \bar{p}_i)(p_j - \bar{p}_j) + 2\sum_i \frac{\partial^2 V}{\partial I \partial p_i}(p_i - \bar{p}_i)(I - \bar{I})\right.$$
$$\left. + \frac{\partial^2 V}{\partial I^2}(I - \bar{I})^2\right\}.$$

Similarly, the right-hand side yields

$$V(\hat{p}_1, \ldots, I - B) = V(\bar{\mathbf{p}}, \bar{I}) + \sum_i \frac{\partial V}{\partial p_i}(p_i - \bar{p}_i) + \frac{\partial V}{\partial I}(I - B - \bar{I}) \tag{9.A2}$$
$$+ \tfrac{1}{2}\left\{\sum_i \sum_j \frac{\partial^2 V}{\partial p_i \partial p_j}(p_i - \bar{p}_i)(p_j - \bar{p}_j)\right.$$
$$\left. + 2\sum \frac{\partial^2 V}{\partial I \partial p_i}(p_i - \bar{p}_i)(I - B - \bar{I}) + \frac{\partial^2 V}{\partial I^2}(I - B - \bar{I})^2\right\},$$

where this time p_1 takes the non-random value $\hat{p}_1$. Take expectations of both equations, and note that $E(p_i - \bar{p}_i) = 0$, $E(I - \bar{I}) = 0$. Terms in B^2, $B(\hat{p}_1 - \bar{p}_1)$ and $(\hat{p}_1 - \bar{p}_1)^2$ can be ignored if $\hat{p}_1$ is sufficiently close to $\bar{p}_1$, since they are of order σ^4. Most of the cross-product terms cancel when the two expressions are equated, leaving

$$B\frac{\partial V}{\partial I} = -\tfrac{1}{2}\left\{\frac{\partial^2 V}{\partial p_1^2}\mathrm{Var}(p_1) + 2\sum_{i \neq 1}\frac{\partial^2 V}{\partial p_i \partial p_1}\mathrm{Cov}(p_i, p_1)\right. \tag{9.A3}$$
$$\left. + 2\frac{\partial^2 V}{\partial I \partial p_1}\mathrm{Cov}(I, p_1)\right\} + (\hat{p}_1 - \bar{p}_1)\frac{\partial V}{\partial p_1}.$$

The last term can be written, using Roy's identity, as

$$-(\hat{p}_1 - \bar{p}_1)q_1(\bar{\mathbf{p}}, \bar{I})V_I$$

and is discussed in the text. The remaining coefficients in equation (9.A3) can be further approximated as follows. We have from equation (8.13):

$$\frac{\partial^2 V}{\partial p_i \partial I} = \frac{\beta_i}{p_i}(R^c - \eta_i)V_I$$

where β_i is the expenditure share for the ith good.

Similarly, from Roy's identity

$$\frac{\partial^2 V}{\partial p_i \partial p_j} = \frac{\partial V}{\partial I}\frac{q_i}{p_j}\left\{-\frac{p_j \partial q_i}{q_i \partial p_j} - \beta_j(R^c - \eta_j)\right\}.$$

If we are discussing price stabilization for a limited range of agricultural goods, β_j is small (less than 1 per cent for those agricultural consumption goods which UNCTAD is concerned to stabilize). It follows that to a first approximation the second term in the braces can be ignored, leading to the simplification

$$\frac{\partial^2 V}{\partial p_i \partial p_j} = -\frac{\partial V}{\partial I}\frac{\beta_i I \epsilon_{ij}}{p_i p_j}; \quad \epsilon_{ij} = \frac{p_j \partial q_i}{q_i \partial p_j}. \tag{9.A4}$$

Here ϵ_{ij} is the cross-price elasticity of demand, positive for substitutes, negative for complements, and zero for independent goods. The uncompensated elasticity can be as accurately replaced by the symmetric compensated elasticity. Substituting these terms in the first part of equation (9.A3) gives

$$-\frac{\beta_1 I}{2}\left\{\epsilon_1 \sigma_{p_1}^2 - 2\sum_{i \neq 1} \epsilon_{i1}\rho(p_i, p_1)\sigma_{p_1}\sigma_{p_i} + 2(R^c - \eta_1)\rho(p_i, I)\sigma_{p_1}\sigma_I\right\} V_I \tag{9.A5}$$

where

$$\rho(x, y) = \frac{\text{Cov}(x, y)}{\sqrt{\text{Var}(x)\text{Var}(y)}}$$

is the correlation coefficient between x and y.

Chapter 10

Market Equilibrium with Rational Expectations

10.1 Introduction

In the previous six chapters we have outlined the theory of producer and consumer behaviour in the presence of risk. We now turn to the analysis of market equilibrium.

As we shall see, a critical determinant of the nature of the market equilibrium is the formation of expectations on the part of producers; if they could sell their crops forward, and if they could purchase crop insurance, then they would not need to form these expectations (any more than any other participant in the economy would; that is, they could completely *hedge* all their production risks; they *might* decide to speculate, but their speculative activity is quite distinct from their production activities). But for virtually all farmers complete hedging is not possible, and thus they must rely to some extent on expectations in making their production decisions.

This chapter begins with a general discussion of the problem of expectations formation. We then introduce the concept of rational expectations equilibria, i.e. equilibria in which the probability distribution of prices and outputs which individuals believe corresponds to the frequency distribution of prices and output. It has been argued that the rational expectations equilibrium is the natural equilibrium concept to employ in the presence of risk, and section 10.3 examines this argument. We discuss the limits on the usefulness of the rational expectations equilibrium concept and explain why we have employed it so extensively throughout this book.

Finally, in section 10.4 we show how, with simple parameterizations of the demand functions and the probability distributions, the rational expectations equilibrium may be completely described and analysed.

The next chapter is devoted to the analysis of market equilibria in which individuals do not have rational expectations.

10.2 The nature of expectations

Farmers typically draw up a farm plan at the beginning of the crop year before they know the weather, the crop yields, and the prices at which they will sell the crops. In Chapter 6 we assumed that there was an exogenous distribution of prices, and that the farmer knew the joint probability distribution of his output and the price. Armed with this knowledge, he was able to determine the optimum farm plan, that is, the plan which yields the highest expected utility.

In fact, the price that eventually clears the market is not exogenous, but depends on supply and demand. If we are to describe market equilibrium we must explain first how the price distribution is generated, and how it depends on supply decisions, and second how farmers form expectations about prices and so decide on the level of supply. If supply depends on expectations, it is tempting to try to capture this relationship in an equation such as the following:

$$p_t^e = \phi(p_{t-1}, p_{t-2}, \ldots; q_{t-1}, q_{t-2}, \ldots; s_{t-1}, \ldots; z_{t-1}, \ldots). \quad (10.1)$$

In this rather general formulation, the farmer forecasts the expected price at date t using his past observations on price, p, his own output, q, his perception of the states of the world, s (were prices high last year because of low supply or high demand?) and other relevant information, z (such as the level of exports of the commodity, the price of substitutes). In short, we imagine the farmer acting as a more or less sophisticated forecaster, given his prior beliefs about the way the economy works.

Before 1961 an econometrician attempting to model an agricultural market would have taken equation (10.1) and drastically simplified it to include only readily observed and obviously relevant variables such as past prices. We shall examine a typical example in the next chapter, but for the moment we are more concerned with whether this provides a satisfactory theoretical framework. In a seminal article in 1961, Muth argued convincingly that it does not. He observed that the behaviour of dynamic models was typically very sensitive to the specification of expectations, but these specifications were disturbingly *ad hoc*. Second, these naïve forecasting rules used information rather mechanically and therefore inefficiently. Many agencies, recognizing the importance of expectations, conduct regular surveys of expectations and intentions and publish them as leading economic indicators. These data show that averages of expectations in industry are typically more accurate than the forecasts of naïve models and as accurate as elaborate forecasting models. In other words, agents appear to use information more efficiently than these simple forecasting rules suggest.

The final objection is the most telling, for Muth pointed out that these forecasting rules did not explain how expectations were formed. Without such an explanation it is impossible to predict how expectations will adjust to a change in the amount of information or in the behaviour of the system. This is of central importance if we are interested in the introduction of price stabilization, for this will certainly affect the way the markets behave and the kind of information available to farmers.

Is there a more satisfactory theory of expectations which meets these objections and which is consistent with empirical evidence? Muth argued that there was and illustrated his argument in a simple model of agricultural supply and demand. He made two important simplifications – that supply and demand schedules were linear and that risk was additive, so that the certainty equivalent prices guiding supply decisions were equal to average prices. In Chapter 6 we argued that this was a very special case, and that in general the relationship between certainty equivalent prices and average prices depended on the sources of risk, the shape of the demand schedules, and the farmers' attitudes to risk. Moreover, the relationship between the certainty equivalent price and the average price would in general change if the economic environment changed. Since the main reason for developing a theory of expectations is to be able to deal with a change in the economic environment, it weakens the theory if we work in terms of an unchanged formula for the certainty equivalent price. For Muth's illustrative purposes, the simplification was well chosen to reveal the essence of his proposal, and his article is still an excellent introduction to the subject, but our concern is to develop a more general theory.

Instead of working in terms of a single certainty equivalent price this means that we should recognize that prices are uncertain and that farmers realize this. They therefore forecast a whole price distribution rather than a single-point estimate:

$$F^e(p_t) = \phi(p_{t-i}; q_{t-i}; s_{t-i}; z_{t-i}), \quad i = 1, 2, \ldots. \quad (10.2)$$

Here F^e is the *distribution function* of prices which the farmer expects, given his current information. Thus $F^e(p)$ is the probability that the price expected to rule at date t is no greater than p (with $F^e(0) = 0$, $F^e(\infty) = 1$). In Chapter 6 we showed how the choice of inputs could be systematically related to beliefs about the distribution of prices, as described by $F^e(p)$, and in particular to their variance and covariance with own production (and more subtle properties in the case of risk-averse producers). This choice of inputs, together with the underlying risk, generates a distribution of total output at date t, again represented by a distribution function:

$$G(Q_t) = \psi\{F^e(p_t)\}. \quad (10.3)$$

The demand for output may also be random, and depend on the state of nature, s. For the moment suppose that the commodity is not stocked, in which case demand is equal to current consumption, which will depend on the market clearing price, p_t, and possibly also s_t. The market clearing price today when supply is Q can thus be written

$$p_t = D_t(Q_t, s_t). \quad (10.4)$$

(If the commodity can be stocked, then part of the demand will be for addition to inventories, and will depend on expected future prices. A full analysis requires an intertemporal model of the kind set out in Part VII. As Muth shows, the notion of rational expectations equilibrium is not radically altered by such dynamic considerations.)

Finally, the distribution of supplies, $G(Q)$, together with the randomness in demand, will generate a distribution of market clearing prices, described by

$$F(p_t). \quad (10.5)$$

Figure 10.1 provides a simple geometric illustration for the case of non-random demand (and no storage) where the planned expected supply is $\bar{Q}$, distributed

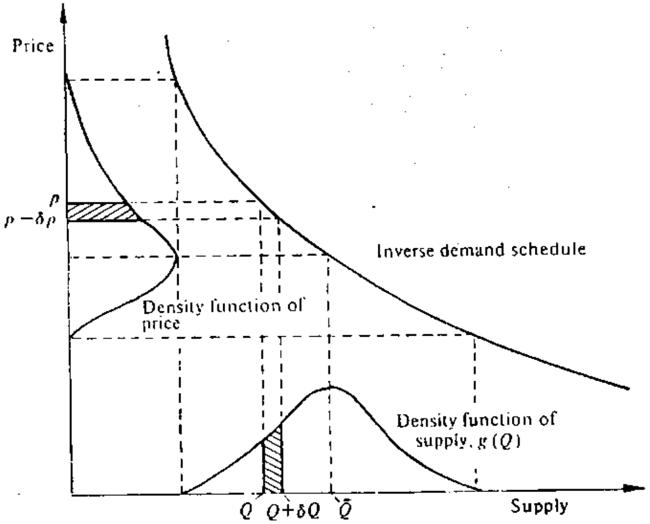

Fig. 10.1 Supply variability inducing price variability

as $G(Q)$. The density function of supplies, g, induces a density function of market clearing prices – the shaded areas are equal, and give the probability of supplies between Q and $Q + \delta Q$, prices between p and $p - \delta p$.

We have come full circle, for the actual distribution of prices will depend on the expected distribution of prices, and, of course, the structure of the system. It should now be clear that there is a natural description of expectations which meets the early objections of arbitrariness and informational inefficiency: a rational expectations equilibrium is a distribution of prices, $F^*(p)$, such that the actual price distribution is the same as the expected price distribution:

$$F^*(p) = F^e(p) = F(p).$$

The attraction of this concept is obvious, because it is the natural extension to a world of risk of the perfect foresight assumption so commonly used in traditional economic theory. If a unique rational expectations equilibrium exists (often a delicate question, though in our simple models such an equilibrium can be shown to exist) then it provides a logically consistent and satisfactory explanation of how expectations are related to information. It is the only specification of expectations which would not eventually be modified by observations, for if agents held different beliefs they would predict some outcomes as occurring more often than they would in fact, and others less often. Eventually, if the crop year were repeated often enough, they would collect enough observations to falsify their beliefs, and would, if rational, be forced to modify their expectations. To the extent that expectations were not rational, there would be scope for learning and the structure of the system would be changing. In practice,

as we shall comment later, things are more complicated. In a non-stationary environment it may not be easy for a farmer to tell whether his expectations are 'rational'.

Another way to characterize a rational expectations equilibrium is to suppose that farmers have some theory about the way their economic environment works, and that this theory is consistent with observations. They observe that price is low when their own supply is large, and they postulate a downward sloping *aggregate* demand schedule, although they also realize that they would sell as much as they like at the prevailing market price – the demand schedule facing an individual farmer is flat. For example, suppose that the output of farmer i depends on his inputs x^i and the state of the weather, s; then he realizes that total supply Q will also depend on s:

$$Q = Q(s) = \sum_i f^i(x^i, s). \qquad (10.6)$$

The market price will also depend on s, for, from equation (10.4):

$$p = p(s) = D(Q(s), s). \qquad (10.7)$$

Farmer i holds rational expectations if he can calculate profit in each state, $Y^i(x^i, s)$, and knows the probability of the occurrence of that state, $\pi(s)$, say:

$$Y^i(x^i, s) = p(s)f^i(x^i, s) - w \cdot x^i. \qquad (10.8)$$

He then chooses $x^i = x^{i*}$ to yield the highest expected utility. For example, if farmers were risk neutral, x^{i*} would satisfy

$$Ep(s)\frac{\partial f^i}{\partial x^i}(x^{i*}, s) = \sum_s \pi(s)p(s)\frac{\partial f^i}{\partial x^i} = w. \qquad (10.9)$$

A rational expectations equilibrium is one in which the price used in equation (10.9) simultaneously satisfies

$$p(s) = D\left(\sum f^i(x^{i*}, s), s\right) \qquad (10.10)$$

when all farmers choose x^{i*} according to equation (10.9).

The obvious objection to make against the realism of rational expectations is that it supposes an unreasonable degree of rationality. Organization theorists such as Simon (1959) argue that the assumption of rationality leads to theories which are unable to explain observed phenomena. Muth's reply is that dynamic economic models do not assume enough rationality, and therefore do not do as well at predicting as expectations surveys reveal actually happens. Moreover, although we have set out a general and therefore quite demanding description of rational expectations, in many cases the relevant information *can* be summarized in a simple statistic, a certainty equivalent. For example, when choosing his inputs using equation (10.9), if multiplicative risk is assumed, all that is required is that the farmer forecast expected gross returns per hectare; if the farmer is risk neutral, he need only forecast average returns. In this case, the rational expectations hypothesis is that he makes on average an unbiased forecast. (In the final section of the chapter we show how in certain circumstances farmers can compress the relevant information into a certainty equivalent price and so simplify their decision problem.) To conclude, the rational expectations hypothesis does not assert that predictions are accurate, but only that they cannot be improved without additional information. The claim is that agents do not waste scarce and valuable information. If they did, and if expectations were not moderately rational, then there would be opportunities for profitable speculation.

10.3 The efficiency of rational expectations equilibria

In most of the rest of the book we assume the rational expectations hypothesis whenever characterizing a market equilibrium, and it is important to understand why. We do not believe that the economy is always in a rational expectations equilibrium, and, indeed, we specifically argue that if some policy such as price stabilization is introduced, then there will be a lapse of time during which agents will be collecting information about the new environment during which they are unlikely to forecast accurately. Thus, we shall distinguish between the *impact* effect of a policy and its long-run effect, when the system has moved to a new rational expectations equilibrium. Our defence of the hypothesis is that it allows us to draw a clear distinction between two effects which a policy may have. If the economy starts and finishes in a rational expectations equilibrium, then the policy, by changing the equilibrium, changes (and, one would hope, improves) the efficiency of resource allocation. If the economy is not in rational expectations equilibrium, then the policy may, by changing the information available to agents, move them closer to the rational expectations equilibrium. The second effect is properly counted as a benefit of improved information, and it may be achievable in other ways, for example, by directly improving the information available to farmers. Put another way, if the only advantage of price stabilization was to improve farmers' forecasts of prices, then it might be much cheaper to produce crop forecasting services. On the other hand, it may be that it is difficult or costly to provide this information, in which case a price stabilization scheme might be attractive because it reduces the need for such information.

Moreover, if we did not assume rational expectations, then we would be forced to make some rather *ad hoc* assumption about expectations. The impact of any policy would then consist of a 'true' efficiency gain (or loss) and gain (or loss) of more closely approaching the rational expectations equilibrium. The magnitude of the second effect would depend sensitively on two *ad hoc* assumptions about the initial and final expectations. Different economists, attracted by different assumptions about expectations, might disagree extensively about the desirability of a policy, yet not realize that their disagreement resulted from such arbitrary assumptions. If the two impacts are clearly distinguished, then it should be easier to resolve differences and identify the key determinants of

policy success or failure. In the next chapter we shall demonstrate this distinction, and resolve a dispute which arises from this confusion.

Intuitively, a rational expectations equilibrium in a competitive economy would appear to be the best attainable equilibrium because it uses the available information about the future correctly, and, being competitive, does not allow producers to exploit their market power. If there were no uncertainty, rational expectations would correspond to perfect foresight, and the equilibrium would indeed be Pareto efficient. Unfortunately, however, with uncertainty this is no longer true (unless there is a complete set of insurance markets). Competitive rational expectations equilibria with incomplete markets can typically be improved by some form of intervention, as we show in two different contexts below. In Chapter 14 we discuss the nature of the bias in the choice of production technique under uncertainty, while in Chapter 23 we demonstrate the relative inefficiency of free trade in the presence of production risk.

Nevertheless, it must be admitted that the information required to improve the allocation of resources may be hard to come by; for the bias in the market allocation of resources depends on such features as the sign of the third derivative of the utility function, the source and nature of risk, and the heterogeneity of the population.

To conclude, the attraction of the rational expectations hypothesis is, first, that it can be applied to any dynamic system, and avoids the need for specific *ad hoc* assumptions about the formation of expectations. Second, it provides a natural bench-mark against which to measure deviations from rationality, or the potential gains from improved information and decision-making. Third, it allows us to draw the important distinction between changes in information which improve the efficiency of a given equilibrium, and movements between rational expectation equilibria (caused by, for example, a price stabilization programme). A similar and useful distinction is often made between movements towards the production possibility frontier (eliminating production inefficiency) and movements along the frontier (eliminating trade inefficiency).

10.4 Simple models of rational expectations with supply risk

In the case where the only source of disturbance to the market is supply variability, we can provide a fairly complete characterization of the rational expectations equilibrium. This will provide a basis for comparison of market equilibrium in which producers are 'ill informed' or 'irrational'. We have already illustrated the way in which supply variability is translated into price variability in Fig. 10.1, and note in passing that as shown the distribution of prices is skewed because the demand schedule is convex, so that the average price is above the price corresponding to average supply, $p(\bar{Q})$. This is another illustration of Jensen's inequality, discussed in section 6.3.

We can construct a rather simple model of supply risk to demonstrate the way in which a farmer could form rational expectations. Assume supply risk is multiplicative

$$q(s) = \theta(s)f(x), \quad E\theta = 1, \mathrm{Var}\,\theta = \sigma^2, \tag{10.11}$$

as in equation (5.12). If all farmers face the same, perfectly correlated risk and know the shape of the (static) demand curve, then a risk-neutral farmer would maximize:

$$\begin{aligned} EY &= f(x)E\theta p(\theta\bar{Q}) - wx \\ &= \hat{p}f(x) - wx \end{aligned} \tag{10.12}$$

where

$$\hat{p} = E\theta p(\theta\bar{Q}). \tag{10.13}$$

Again, $\hat{p}$ is the action certainty equivalent price defined and discussed in Chapter 5, and $\bar{Q}$ is average total supply, which the farmer can either predict (if he knows total planted area and that other farmers are the same) or learn. Moreover, $\hat{p}$ is proportional to average gross return per acre, and so in principle is easy to observe.

It is easy to extend the model to allow farmers to experience diverse multiplicative risk. If farmer i faces supply risk θ_i then this can be treated as the sum of two orthogonal components, the general supply risk θ, and the farmer's individual risk ν_i:

$$\theta_i = \theta + \nu_i, \quad E\nu_i = 0, E\theta\nu_i = 0.$$

Then

$$Q = \sum \theta\bar{q}_i + \sum \nu_i\bar{q}_i \simeq \theta\bar{Q}$$

since the individual uncorrelated risks will approximately cancel by the law of large numbers.

If the demand schedule is linear, so that

$$p = a - bQ, \tag{10.14}$$

then the certainty equivalent price is

$$\hat{p}_i = E(\theta + \nu_i)(a - b\theta\bar{Q}) = a - b\bar{Q}(1 + \sigma^2) = \bar{p} - b\bar{Q}\sigma^2, \tag{10.15}$$

which is the same for all farmers, and the same expression as equation (5.15). Notice that with a *linear* demand function the average price is the price of average output

$$Ep = E(a - bQ) = p(EQ),$$

but as Fig. 10.1 showed this is not generally true. If $\bar{\epsilon}$ is the elasticity of demand evaluated at the average price, then equation (10.15) can be written

$$\hat{p} = \bar{p}\left(1 - \frac{\sigma^2}{\bar{\epsilon}}\right) = \bar{p}(1 - \bar{\epsilon}\sigma_p^2), \tag{10.16}$$

where σ_p is the coefficient of variation of prices.

The only information required by the risk-neutral farmer for planning purposes is the average price, $\bar{p}$, the coefficient of variation of price, σ_p, and the elasticity, $\bar{\epsilon}$. The first two are directly observable from past data, but the last must be inferred from a knowledge of both price and output. Linear regression of price on his own output will give him an appropriate estimate of $\bar{\epsilon}$. Thus, given a *theory* of the determination of price (in this case that the demand schedule is linear, and that there is zero correlation between supply and demand risk), the farmer can determine the certainty equivalent price which on average maximizes his profits. Rational expectations on this view means acting on a correct theory which describes the market environment. If the theory is correct, it will not be falsified by observation.

In the case of linear demand curves the certainty equivalent price, $E\theta p(\theta\bar{Q})$, takes a particularly simple form, but even so its relation to the average price depends on the magnitude of risk and the shape of the demand function (as measured by $\bar{\epsilon}$). It is no longer possible to derive a supply function independent of the nature of risk and the demand function. The situation is very similar to the problem of deriving a supply function for a monopolist. *If* the elasticity of the demand curve facing a monopolist is constant at ϵ, then marginal revenue is $p(1 - 1/\epsilon)$, a constant fraction of price. Equating marginal revenue to marginal cost, MC, gives the following pseudo supply-curve:

$$p = \frac{\epsilon}{\epsilon - 1}\mathrm{MC}.$$

In the present case, since $\bar{p} = p(\bar{Q})$, average supply as a function of the price of average supply is given by

$$p(\bar{Q}) = \frac{\mathrm{MC}}{1 - \bar{\epsilon}\sigma_p^2}, \tag{10.17}$$

where MC is the marginal cost as a function of average supply, $\bar{q}_i$.

The formula for the certainty equivalent price given in equation (10.16) was derived for the special case of linear demand, but it remains approximately correct for other demand specifications, provided risk remains multiplicative and on the supply side. For example, if demand has constant elasticity, ϵ, and if θ is log-normally distributed (the natural econometric specification for multiplicative risk, as discussed in section 6.5)

$$\theta = \Lambda(-\tfrac{1}{2}\sigma^2, \sigma^2),$$

so that

$$E\theta = 1, \mathrm{Var}\,\theta = \exp\sigma^2 - 1 \simeq \sigma^2.$$

Then, from equation (10.13)

$$\hat{p} = Ep\theta = p(\bar{Q})E\theta^{1-1/\epsilon}. \tag{10.18}$$

The properties of the log-normal distribution (set out in equation (6.43)) allow us to evaluate this expression as

$$p(\bar{Q})\exp\left\{\tfrac{1}{2}\frac{1}{\epsilon}\left(\frac{1}{\epsilon}-1\right)\sigma^2\right\}.$$

Since

$$p = p(\bar{Q})\theta^{-1/\epsilon} = p(\bar{Q})\Lambda\left(\tfrac{1}{2}\frac{\sigma^2}{\epsilon}, \frac{\sigma^2}{\epsilon^2}\right)$$

the same technique gives the average price, $\bar{p}$

$$\bar{p} = p(\bar{Q})\exp\left\{\tfrac{1}{2}(1+\epsilon)\frac{\sigma^2}{\epsilon^2}\right\}. \tag{10.19}$$

Therefore

$$\hat{p} = \bar{p}\exp\left(-\frac{\sigma^2}{\epsilon}\right) \approx \bar{p}\left(1-\frac{\sigma^2}{\epsilon}\right) = \bar{p}(1-\epsilon\sigma_p^2). \tag{10.20}$$

The first part of equation (10.20) is exact, the second an approximation for small σ, giving (approximately) the same result as equation (10.16). (The same result can also be obtained without assuming θ to be log-normally distributed, provided σ is small, by taking Taylor series expansions, as in subsection 6.5.2.)

If, on the other hand, supply risk is (implausibly) *additive*:

$$q = f(x) + \tilde{u},\ E\tilde{u} = 0,\ E\tilde{u}^2 = \sigma^2, \tag{10.21}$$

then expected profits are

$$EY = E\tilde{p}\{f(x) + \tilde{u}\} - wx$$
$$= \bar{p}f(x) - wx + E\tilde{p}\cdot\tilde{u}$$

so

$$\hat{p} = \bar{p} \tag{10.22}$$

and the action certainty equivalent price is *equal* to the average price. (Note that the *utility* certainty equivalent price is *not* equal to the action certainty equivalent price, since $E\tilde{p}\tilde{u}$ is negative and risk therefore lowers average profits.) Thus whether risk is additive or multiplicative makes a difference to the calculation of the certainty equivalent price.

If the farmer is risk averse and maximizes an additively separable utility function as in equation (6.28):

$$EU\{p\theta f(x)\} - wx,$$

then the action certainty equivalent price, $\hat{p}$, is defined by the equation

$$\hat{p}U'(\hat{p}f) = Ep\theta U'(p\theta f). \tag{10.23}$$

If, as before, θ is log-normally distributed, demand is iso-elastic, and the utility function has constant relative risk aversion R (as in equation (6.12)), then

$$\hat{p} = \bar{p}\exp[-\{\epsilon + \tfrac{1}{2}R(\epsilon-1)^2\}\sigma_p^2], \tag{10.24}$$

which differs from equation (10.20) by the term in R. Provided the farmer knows his own attitude to risk he can calculate the appropriate certainty equivalent price without any additional information.

10.5 Concluding remarks

In this chapter we have defined the rational expectations equilibrium in the presence of risk, and in so doing have completed a logically consistent description of a competitive equilibrium in the presence of risk. We argue that the rational expectations equilibrium is the appropriate equilibrium concept to use in measuring the efficiency benefits of market intervention, as opposed to the informational benefits which might result if intervention improved the rationality of decision-making by making the market structure more transparent. In the next chapter we shall examine this distinction more carefully.