© Lawrence A. Boland

3

Psychologism vs Disequilibrium Models

it really is assumed ... that what you see when you look out the window is an economy in ordinary general equilibrium....

This view has obvious (and intended) affinities to nineteenth-century economic thought, Say's Law, and all that. Like that tradition, the new equilibrium school faces a basic problem: how can it account for the 'obvious' large-scale divergences from equilibrium that we think we see, especially in prolonged depressions?

Robert Solow [1979, p. 341]

Every explanation of economic crises must include the assumption that entrepreneurs have committed errors. But the mere fact that entrepreneurs do make errors can hardly be regarded as a sufficient explanation of crises. Erroneous dispositions which lead to losses all round will appear probable only if we can show why entrepreneurs should all simultaneously make mistakes in the same direction. The explanation that this is just due to a kind of psychological infection or that for any other reason most entrepreneurs should commit the same avoidable errors of judgement does not carry much conviction. It seems, however, more likely that they may all be equally misled by following guides or symptoms which as a rule prove reliable. Or, speaking more concretely, it may be that the prices existing when they make their decisions and on which they had to base their views about the future have created expectations which must necessarily be disappointed. In this case we might have to distinguish between what we may call justified errors, caused by the price system, and sheer errors about the course of external events.

Friedrich Hayek [1933/39]

The Current Problem-Situation of Neoclassical Theory

We have been arguing that all neoclassical research programs are based on a hidden agenda consisting of two main items. One is the acceptance of the need to deal with the so-called Problem of Induction either directly or, what is more common, indirectly by dealing with its variant, the Problem of Conventions. The other item is the requirement of methodological individualism – that every explanation must assume that only individuals make decisions. However, we also pointed out that at present neoclassical theory is based on a reductive version of methodological individualism – specifically, one which identifies the individuals with their exogenous psychological states (such as their given utility functions). The strict reliance on the reductive version – that is, on psychologistic individualism – always presents a general problem of explanation which we have called the problem of simple psychologistic individualism: if everyone is governed by the same laws of psychology, then there is no psychological basis for individuality. We noted that neoclassical theory provides a solution to this problem by restricting the laws of psychology to only that which specifies that everyone faces diminishing marginal utility (or its equivalent). This solution allows people to have different utility functions; hence it provides a means of solving the more general explanatory problem of methodological individualism as well as the Problem of Conventions. The reductive methodological individualism of neoclassical economics accepts only models which exclude all exogenous variables except psychological states and natural givens.

We thus note that the two agenda items are not independent, as the latter one is sustained partly because it supports the former. That is, it would be difficult for most neoclassical economists to give up their reliance on psychologistic individualism and their solution to the problem of simple psychologism because that would entail the lack of a means of dealing with the Problem of Conventions. Furthermore, this difficulty is compounded by the fact that most neoclassical economists take the Problem of Conventions for granted; hence it is difficult to see that there may be a need to deal with any of these problems.

It is not our job to form a final judgement or 'methodological appraisal' of the existing neoclassical research program, as some might wish [e.g., Blaug, 1980]. Instead, our purpose is to establish a clear understanding of what neoclassical economics is rather than to determine what some philosophers think it should be. We also wish to examine what practicing economists think economics should be while recognizing that they may not all share the same view. Although some of what we shall argue is critical of certain aspects of some neoclassical models, we do not intend to present a destructive criticism of neoclassical economics. If we did intend a destructive criticism at this stage, we would immediately set about eliminating the Problem of Induction, which would appear to eliminate the more impressive aspects of neoclassical theory. But our purpose is quite different. In particular we wish to understand the methodological nature and purpose of exisiting neoclassical research programs which accept the items on the hidden agenda, and thus for our purposes we have to take the two items as givens. Besides, as we shall argue, *although the hidden agenda is necessary for the explanation of neoclassical methodology, it is not necessary for neoclassical economic theories*.

When we say that we take the hidden agenda items as givens it does not mean that we are thereby accepting them. On the contrary, we shall be arguing that, paradoxically, it is the acceptance of the two items which gives rise both to the many theoretical problems that avant-garde economists find fascinating and to the obstacles which block their solution.

We will begin by considering the neoclassical research programs which are based on the acceptance of the neoclassical solution to the problem of simple psychologistic individualism – e.g., those which are based on the view that the only psychological law affecting all individuals is the negative slope of their marginal utility curves and that individuality is provided by there being an infinity of utility curves satisfying this requirement. Although we have argued before that psychologistic individualism is accepted by some 'conservative' methodologists because it is viewed as one way of providing the 'atoms' for a long-run inductive science, we shall postpone consideration of the direct effects of the Problem of Induction until Chapter 4. Instead, we will examine the short-run theoretical consequences of attempting to deal with the Problem of Conventions within the confines of the problem of simple psychologism.

The Price System and Psychologistic Individualism

By saying that neoclassical economics is based on psychologistic individualism, we are saying specifically that neoclassical theories or analyses must permit only two types of exogenous variable: natural constraints and psychological states. Of particular concern is the psychologistic individualist requirement that no social institution that appears in our explanations must be allowed to play the role of an exogenous given. For reference, let us call this particular requirement the 'problem of social institutions'. Specifically:

> The *problem of social institutions* is: how do we assure that every institution which is introduced as a given in the short-run (or partial

equilibrium) version of a model can be explained in terms that include only the exogenous variables permitted in the long-run (or general equilibrium) version of the model?

The Price system as a social institution

To begin, let us consider what Schumpeter noted: 'prices are obviously social phenomena' [1909]. What does this mean? Does this mean that the price system is an exogenous social institution? Or is it merely an epiphenomenon of the psychological states of the individuals in society? If, as required by psychologistic individualism, it is only an epiphenomenon, then it must be true that (1) the actual price system can be influenced by each individual, and (2) in our explanation of the price system, the value of all prices can be determined only by reference to all exogenous variables, namely, the natural givens and the psychological states of *all* individuals.

Condition 1 is the basis for some of the interesting theoretical questions raised by Kenneth Arrow [1959]. Specifically, in what circumstances is it *possible* for all individuals to be influencing the determination of the market price, yet at the same time for no one individual alone to determine the actual price and thereby deny the influence of all other individuals? Arrow argued that condition 1 is satisfied only when the market is in equilibrium. When the market is in equilibrium all individuals influence the price by their participation on the demand side or the supply side, since if they withdrew from either side, the price would change. Also, in equilibrium no individual can force the price to change to a specific price other than the value of the equilibrium price, since any effort to do so would cause a disequilibrium. We see then that according to the neoclassical (i.e., psychologistic) view of the market determination of actual prices, the individual's 'influence' on the price level is only indirect or 'unintended'. Given this, all that is required for the logical adequacy of this view is that the states of all the permitted exogenous variables do indeed entail a determinant price system (i.e., all markets are potentially stable). This is the requirement that for any model of the price system there must exist a solution for the values of all the prices – i.e., we must be able to provide a so-called 'existence proof' [see Boland, 1970a].

Over the last fifty years the mathematical requirements for any existence proof have been explored and analyzed in excessive detail, to the boredom of most economists. Such proofs are no longer the basis of research programs in economic theory, although it could be argued that the existing proofs still require too much of the real world [Boland, 1975]. There nevertheless remains an unanswered part of Arrow's argument. What happens to the methodology of psychologistic individualism when the market is not in equilibrium? Arrow argues that in order to explain the determination of prices, we violate the requirements of psychologistic individualism, since in order to get the price to return to the equilibrium, at least one individual (i.e., the bidder) must set the price, and that means that at least one individual is determining the price? This observation of Arrow has led to two schools of thought. One argues that we need a theory of 'disequilibrium trading' [e.g. Clower, 1965]; the other argues that we need a neoclassical theory of the individual bidder [Gordon and Hynes, 1970]. Neither school has been completely successful, thus Arrow's challenge still stands.

We shall argue here that the theoretical puzzles based on condition 1 are the direct consequence of the acceptance of the methodological requirements of psychologistic individualism and, in particular, of condition 2 – the requirement that all institutions must be endogenous. It will be easy to see that in any case the problem of social institutions has immediate consequences for our concept of the price system as a social system which *endogenously* coordinates all individuals, and in which they are all presumed to be engaged in independent, rational decision-making.

The concept of a social institution is not often specified in economic models, although a few writers have discussed some of the important attributes of institutions from the perspective of neoclassical models [Buchanan and Tullock, 1962; Davis and North, 1971]. An obvious example might be the legislated 'ground rules' that define property rights. Generally, it is accepted that institutions exist potentially to constrain all individual decision-makers. But what makes institutional constraints important is that they are not natually given but have themselves been created by other individuals acting in concert. That any institution may effectively constrain only one individual is not the issue here; rather, the issue is that its existence is dependent on the activities of many individuals, including any individuals who may be constrained by the institution.

In order to distinguish institutions from individual actions in the most general terms, we are saying that any constraint, the establishment of which requires the implicit or explicit participation of many individuals, is in some sense an institution. For this reason, some economists might consider a system of all market-determined prices to be an institution whose function is to provide the decision-maker with a 'summary of information about the production possibilities, resource availabilities and preferences of all other decision-makers' [Koopmans 1957, p. 53]. In the neoclassical market everyone faces the same price, and in this way the price is an institution with which individuals' social behavior can be coordinated.

Individualism vs. coordination

Methodologically speaking, the neoclassical theorist cannot rest until it is shown that the nature of any institution is what it is only because people have directly or indirectly chosen that it should be what it is [Boland, 1979b]. Recognizing the price system as an institution is interesting in this sense because the price system serves a dual purpose. On the one hand, it has to be *responsive* (no matter how small the degree) to every individual's psychological state and, on the other, it has to be a relatively *stable* signal indicating to every individual decision-maker the wishes of every other individual in society. How can a social institution serve such a dual purpose? How can something be both volatile (i.e., responsive) and stable?

To answer these questions, we need to understand the neoclassical conception of market-equilibrium prices (i.e., socially coordinated prices). Specifically, we need to understand how prices would have to be determined in a neoclassical model in a manner consistent with psychologistic individualism. As Arrow argued, in a consistent neoclassical model, prices are determinant only when the influences of *all* individuals are, in some non-accidental way, *in balance*. But, as suggested by Koopmans, if prices are to fulfill the requirements of a social institution, not only must we assure the possibility of such a balance being a stable institution, but that balance must *also* be an *equilibrium*. That is, any accidental disturbance of the balance will be corrected without the extraordinary influence of any one individual or institution. Any going price will be the one price at which the influence of individuals (through their willingness to demand or supply some of that good) are in balance. What the existence of an equilibrium implies is the following. The going (observable) price of a good is not an accidental price. It is not accidental because had it been higher or lower there would have been *reasons* for it to return to the balancing price [see further, Boland, 1977b].

Unfortunately, an existence proof does not usually provide behavioral reasons for the occurence of an equilibrium. All of the mathematical studies concerning existence proofs have only assured the existence of a possible *balance* (for example, as a solution of a system of simultaneous equations representing demand and supply functions). Nothing more was intended [Hahn, 1973]. The question still may be open as to whether a potential balance is also an equilibrium. To many, the distinction between a balance and an equilibrium may not appear to be very significant because in economics textbooks the concept of equilibrium is often confused with that of a balance. But the distinction is essential to an understanding of avant-garde theoretical research programs.

To understand this distinction, consider a coin balanced on its edge. If it is tilted slightly to either side, its physical position will permanently change. That is, there is no reason for the coin to bring itself back into the original upright balance. Textbooks would say that this is an unstable equilibrium, but, more correctly, it is an unstable balance – the concept of an unstable equilbrium is self-contradictory. Similarly, textbooks would say that an equality between supply and demand is an equilibrium when actually it is only a balance. For an equilbrium, more is required.

If there are reasons why a balance is not accidental (e.g., if it is the result of competition), then those reasons imply that the balance is stable, i.e., it is an equilibrium. In neoclassical theory this is of particular importance to the concept of equilibrium prices. The equilibrium price can be thought to be determined by the reasons which guarantee that demand and supply will be in balance, because these are the reasons which guarantee the existence of a stable balance. But to accommodate psychologistic individualism, the reasons must be related to individuals' psychological states.

Equilibrium and psychologism

Now let us consider how the acceptance of psychologistic individualism affects the neoclassical concept of a market equilibrium. The psychologistic individualist explanation of equilibrium goes as follows. On the basis of a posited relationship between the quantity demanded of a good and its going market price, and of a posited relationship between the quantity supplied of that good and the *same* going price, the equilibrium price will of necessity be the one price that brings into balance these two quantities as an unintended consequence of competition. What that price will be depends on the two posited relationships. Generally, if either relationship changes, the result will be a new equilibrium price. It is the sensitivity of the demanded (or supplied) quantity's relationship to the price which assures the responsiveness of the equilibrium price to changes in individual psychological states. In neoclassical theory the relationship is merely a consequence of maximization based on given utility functions. Changes in the psychological state of any individual must have some effect on the equilibrium price (even though, when there are many individuals, the result may appear to be negligible) if psychologistic individualism is to be maintained.

Formally, one part of the psychologistic individualist explanation is easy to provide. As we have just stated, the relationship between the quantity demanded (or supplied) by any individual is asserted to be the consequence of maximization (i.e., rational decision-making). One aspect of maximization is that its consequences are usually determinant (and non-arbitrary); that is, it leads to a unique quantity for any given

price (for the given budget). If we are given the going price, then theoretically we can calculate the quantity demanded and supplied for any good. If there is a discrepancy between the quantities demanded and supplied, then it must be the case that at least one individual is not maximizing? For example, for the usual case in which the demand curve is negatively sloped and the supply curve is positively sloped, if the price were lower than that which would clear the market, then the transacted quantity must be less than the quantity which would maximize every individual's utility (at that price). In other words, whenever the market is not in equilibrium not all individuals can be maximizing according to their psychologically given utility functions (recognizing that a profit function is also the result of a psychological desire to maximize profits or to survive [Alchian, 1950]).

Equilibria and incentives

This point needs to be stressed, since it is the center of the methodological problem facing many theorists today. *The neoclassical theory that all individuals are maximizers can be true only if all markets are in equilibrium*. For an equilibrium to exist, there cannot be any incentive for individuals to change their behavior, that is, change their choices. If an incentive does exist, then we would have to explain why it has not been pursued. If the individual chose not to pursue it, it could not have been an incentive. If the individual is in any way constrained from pursuing it, then additional constraints must be included among the exogenous variables of our explanation. This leads to the central theoretical problem of neoclassical economics today. How can there be any disequilibria? Would anyone choose a state of disequilibrium?

Before discussing the methodological problems of disequilibria, let us be clear about the more elementary relationship between equilibria and incentives. Basically, a true equilibrium says that all possible gains from trade or from adjustments to behavior have been exhausted. If there were possible gains available, then there would be reasons for change. In a state of market equilibrium there cannot exist any incentives for change. This does not mean that there are no constraints, but only that all operative constraints must be beyond choice.

Psychologism and the World of Adam Smith

Self-interest vs. social optima

For the purposes of psychologistic individualism, it is essential that all incentives be individualist and not social. Consider, for example, the profits of the firm in a competitive economy. We say that excess profits must in the long run be zero. Perhaps from the social point of view this may be desirable, but to be consistent with psychologistic individualism we must not allow social objectives to be imposed on individual firms. Even from the social point of view, profit itself is not necessarily interesting. As Schumpeter pointed out, 'as the rise and decay of industrial fortunes is *the* essential fact about the social structure of capitalist society, both the emergence of what is, in any single instance, an essentially temporary gain, and the elimination of it by the working of the competitive mechanism, obviously are more than "frictional" phenomena, as is that process of underselling by which its achievements result in higher real incomes all round' [1928, pp. 380-1]. The point of Adam Smith's classical vision is that the pursuit of private interest (i.e., 'selfishness') *unintentionally* produces a social good. It does this only in a world of competition where profits are indirectly eliminated. Zero profits are an 'unintended consequence' of the combination of profit maximization and competition.

The only individualist incentive we use to explain the behavior of a firm in the short run is the maximization of profits – the difference between total revenue and cost. When costs equal revenue, average cost must equal average revenue. For a profit maximizing firm, however, average cost and revenue is irrelevant *with regard to maximization*. As is well known today, it is marginal revenue and marginal cost that matter.

The firm can respond to its incentives (possible improvements in its profits) in two different ways. Primarily, it can internally and independently alter its output to adjust its costs and revenue. If it is maximizing its profits, then, of course, marginal revenue must just equal marginal cost. But also, if it is maximizing profits, any increase in output must produce a situation in which marginal cost exceeds marginal revenue. So long as the firm is not incurring losses, there is no incentive for it to change its output. Secondarily, it can also deal with its situation by altering its external situation – but only if there exist other possibilities. If it is making losses (even though it may be minimizing them), it can do nothing to change its situation unless there exist better alternatives. But such contingencies are beyond its control in a competitive economy. Either a losing firm eventually quits or it switches to another industry. Its decision is a private matter.

The assumption of profit maximization, then, only assures that marginal cost equals marginal revenue. Which of the individualist variables assure the attainment of zero profits? If some firms in one industry are making profits while firms in another are making losses, then there is an incentive for the losing firms to change industries. In doing so the firms entering the profit-making industry only drive the market price down or reduce every other firm's share of the market and, either way, reduce everyone's profits. Even so, the existence of profits is

an incentive for individual entrepreneurs to enter. The incentive to enter disappears only when profits are driven down to zero – however unintentionally.

So the individual's pursuit of self-interested profits internally and externally eventually leads to zero profits. But zero profit combined with profit maximization does not necessarily mean that the social goal of optimum resource allocation has been served. Or does it? It does whenever all maximizers are also price-takers.

Social optima as forced, unintended consequences

The allocation of resources is optimum only if there is no possibility of reducing their utilization without reducing outputs. Traditionally, this is illustrated by a U-shaped average cost curve which represents all the possible levels of the cost of the resource used per unit of output. If average cost can be decreased by producing more, then the current output is not efficiently produced. Maximum efficiency in this sense then occurs only at the level of output where average cost is minimum. This is the key to connecting the individual's concerns to the social objective. It is also the key to understanding the role of natural constraints.

Since the individual maximizer is only concerned with marginal values, we need to note an elementary point: the behavior of the average is not independent of its relationship to the margin. Specifically, to cause the average to fall, the margin must be below the average. Similarly, if the average rises, it can only be because the margin is above the average. Thus, with this elementary point in mind, we see that whenever the average is at a true minimum and thus temporarily unchanging with respect to output, the margin equals the average. So, in order to get the firm to use its resources efficiently, we need only have the firm produce where the marginal cost equals the average cost.

But profit maximization, our individualist incentive, only assures the equality of marginal cost with marginal revenue. Similarly, reducing profits to zero only assures that average cost equals average revenue. There is nothing here to bring average and marginal cost into equality. Now here is where the idea that firms (and buyers) are price-takers becomes crucial. If a firm is a price-taker – that is, the price is given and does not change in response to the firm's behavior (such as when there are very many small firms or prices are externally fixed) – then marginal revenue will necessarily equal the average revenue (which is the price). In this special case, if the individual firm's profit is maximized, the price (marginal revenue) will equal its marginal cost. If the individual firm inadvertently causes profits to be reduced to zero, the price (average revenue) will equal its average cost. Thus, indirectly we obtain the desired efficiency; the firm's marginal cost will equal its average cost *without the deliberate action by any individual* in that regard!

Do we need Adam Smith's World?

What has to be seen is that in an individualist world of price-takers (that is, where no one individual or small group can set the price for everyone else), it is the combination of zero profits and profit maximization that implies the achievement of an optimum allocation of resources – maximization of efficiency, so to speak. While we can see that an individual firm might wish to maximize profits, no individual would want to eliminate profits. This is both the perversity and the beauty of Adam Smith's world. No individual has to have zero profits as his or her goal. It is the free pursuit of private interest which, unintentionally, is *sufficient* for the provision of the social good.

Adam Smith's world of greed and virtue

If one examined only the sufficient conditions for an economy which is in a state of psychologistic individualist equilibrium, the beauty of Adam Smith's world would be lost. Surely any entrepreneurs who took a broad perspective (or a course in the principles of economics) would see that the outcome of any promotion of free competition must lead to the situation in which, without further changes in the natural constraints, everyone ends up making no profits beyond the costs of production. With this realization, it is easy to see why some critics of neoclassical theory [e.g., Robinson and Eatwell, 1973] might argue that the real incentive for any entrepreneurs is to restrict competition or eliminate their competitors so as to create so-called monopoly profits. Although they are correct about a realistic world of governmental regulations and patents, this view completely misses the point of Adam Smith's unregulated world.

What Smith's world provides is an incentive for entrepreneurs to alter their technological constraints [see further, Schumpeter, 1942/50]. For example, if we are all in a state of long-run equilibrium – in which, everyone's supply price just covers his or her production costs – one way to get ahead is to improve one's technology of production in order to reduce the costs and thereby create short-run profits. The profits will be only short-run gains because in Adam Smith's world, where there are no patents, no marketing boards, etc., other producers will attempt to duplicate the cost-reducing technologies, and in this way everyone (i.e., each consumer) benefits from the original entrepreneur's 'greediness' – so long as free competition prevails. This is how 'virtue' is unintentionally the outcome of greed or 'selfishness' in Adam Smith's world.

Freedom vs. necessity

We have noted that Adam Smith's world is concerned with the sufficiency of free competition when combined with rational decision-making for the achievement of a social optimum. We wish to point out that some neoclassical theorists have also been concerned with its necessity. Hayek [1937/48] specifically wished to show that other world-views (e.g., 'collectivism') were not sufficient. His argument was that exogenous social institutions were informationally inefficient. Specifically, in Adam Smith's world the individual only needs to know his or her own situation (tastes, prices, income, and the location of the market). In the contrasting liberal socialist world, for example, where a central authority may plan the workings of the entire economy, made up of individual but constrained decision-makers, the central authority would have to know the same information but for all individuals!

The primary message of Hayek's view is that if one realizes that all decisions require information and one assumes that the objective of every socialist economy is the achievment of a social optimum, then, if one adopts both psychologistic individualism and Hayek's view that only individuals know what is best for them, the determination of the social optimum must depend on the psychological states of all individuals. Hayek asserts that there is no way a socialist central planner could ever be able to calculate the social optimum in order to implement policies to reach it. What he presumes is that in a psychologistic individualist world there are private facts that affect each individual's view of what is best for him or her. Such private facts are by definition beyond the acquisition of any central planner, yet they are necessary for the calculation of the social optimum. Thus, with Hayek's view, we can see that, *given psychologism*, Adam Smith's world is a necessity, as all other world views would give a role to an exogenous institution which would necessarily have insufficient knowledge to formulate adequate policies.

One has to admit that Adam Smith's view is magnificent, almost magic. But there is no magic here, only simple arithmetic. What is magnificent is the total reliance on individual decision-making. No social institution would seem to be necessary. The final outcome is the result only of the actions of individuals. But there may seem to be a paradox here. The key element to yielding the optimum (beyond maximization) is the inability of one individual firm or consumer to affect the price; that is, competition must be perfect. Nevertheless, individuals are not powerless, since they are allowed to make their personal contribution to supply or demand. The end result is both a social optimum and an equilibrium. All this can exist without any recourse to either non-natural or non-individualist variables or constraints. The only assumption in this neoclassical vision of Adam Smith's world is that every individual is an independent (i.e., self-interested) maximizer. If one could show that the existing world is possibly an instance of an Adam Smith world in long-run equilibrium, then one would have proven the logical feasibility of a psychologistic individualist research program. But what about the assumption that all decision-makers are price-takers? And how do we know when the world is in long-run equilibrium?

Psychologism and Long-run Equilibria

Equilibrium vs. imperfect competition

Theorists either explain why something exists or they explain it away [cf. Agassi, 1977; Solow, 1979]. For those theorists bound by psychologistic individualism, disequilibria must be explained away. In the absence of constraints, neoclassical theory would argue that an equilibrium must exist, since without constraints universal maximization is entailed. If there is a disequilibrium, then it follows that there must be an operative constraint somewhere. Thus, for psychologistic individualism, it must be shown that what appears to be a disequilibrium is really a chosen event or the consequence of a natural constraint. This is because the only allowable exogenous (i.e., non-chosen) variables are natural constraints and psychological states.

The concept of equilibrium is a contingent proposition. There is a disequilibrium only if there are unexploited gains that *can* be obtained. Is it always a question of assessing the (transactions) cost of obtaining the gains? But, possibly more important, how do we measure the potential gains? Too often the alleged gains are an illusion caused by comparing the existing state with an ideal state. As Coase put it,

> very little analysis is required to show that an ideal world is better than a state of laissez faire, unless the definitions of a state of laissez faire and an ideal world happen to be the same. But the whole discussion is largely irrelevant for questions of economic policy since whatever we may have in mind as our ideal world, it is clear that we have not yet discovered how to get to it from where we are.... [1960, section 10]

The question here is whether the state of *laissez-faire* can be one in which there is imperfect competition. The approach offered by Coase allows us to argue that the ideal world is the one with perfect competition – that is, the one where the achievement of private goals indirectly assures the achievement of social goals. However, it may cost us too much to have that much competition. Imperfect competition may be the realistic *laissez-faire* optimum.

If this approach is taken in order to explain away the disequilibia (relative to the ideal world), then we would need to show that the resulting *laissez-faire* equilibrium (i.e., the imperfectly competitive equilibrium) is the result only of individuals' pursuits of their private interests. The question for us then is: how can an imperfectly competitive equilibrium be seen as a social optimum? To be an equilibrium, there should not be any possibility of an improvement, that is, there should not be any incentive. This is assured only if everyone is maximizing with respect to every variable of choice.

This is where the old problem of increasing returns comes in [cf. Sraffa, 1926]. The textbook diagram of an imperfectly competitive equilibrium clearly shows the firm's profit maximizing output to be at a level where the average cost curve is negatively sloped (i.e., to the left of the bottom of the U-shaped average cost curve). If the output is at such a level, it is possible to reduce the average cost (hence the average use of resource inputs) per unit of output – that is, to reduce the ratio of inputs to outputs. If that ratio can be reduced, then its inverse – the returns for each unit of input – can be increased. This possibility is sometimes called a situation of increasing returns. Any situation in which there are increasing returns would seem to indicate the possibility of reducing costs, which would benefit everyone in society. From the perspective of society, increasing returns imply disequilibria, since the existing potential cost reduction is an unexploited incentive, hence increasing returns imply that we have not yet reached a social optimum. Yet an imperfectly competitive *equilibrium* appears necessarily to entail increasing returns.

Although this is an elementary point of price theory, we must treat it with care. Let us then look again at imperfect competition from the perspective of the individual decision-maker who is supposed to pursue profit maximization. If a firm is an imperfect competitor, then by definition it cannot be a price-taker, since its output decisions affect the price. Whenever the price varies with the level of output, marginal revenue is not equal to the price. Furthermore, since it is always assumed that the demand curves are downward-sloping, marginal revenue is always less than average revenue. Now, keeping this in mind, recognize that profit maximization implies the equality of marginal revenue with marginal cost. If we also recognize that a competitive equilibrium painted in any color implies the absence of excess profits (over the cost of producing the chosen level of output) – the absence of incentives for new firms to enter the competition – then the price must equal average cost. Putting all these implications together means that profit maximization with competitively imposed long-run zero profits does not entail the lowest possible average cost. In particular, since marginal revenue is below the price and since profit maximization means that marginal revenue equals marginal cost, then necessarily marginal cost is below average cost – which means that average cost must be falling (i.e., there are increasing returns) whenever there is an imperfectly competitive *equilibrium*.

The conjunction of these implications forms the textbook picture of an imperfectly competitive equilibrium – a tangency between the demand curve and the average cost curve – and since the demand curve is negatively sloped, so must the average cost curve be downward-sloping. If the average cost curve is necessarily falling for a competitor facing the usual downward-sloping demand curves, then *in equilibrium* there must exist the possibility of reducing average cost further. Thus, whenever there is an imperfectly competitive equilibrium, there appears to be a necessary conflict between the individual decision-maker's optimum (profit maximization) and what might be society's optimum (minimizing average cost).

How can the imperfectly competitive equilibrium ever represent a social optimum? If the individual firms' average cost could be reduced, society would benefit, since the available resources could seem to be made to produce more output for the same input. Thus, the possibility of social benefits (reduced average cost) coexists with the absence of any incentive for the producer to change its behavior, since profits are both maximum and zero. But, on the other hand, if every producer is maximizing profits and profits are zero (and the demand curve reflects utility maximization by all individuals), how can there be any disequilibrium? The common view of an imperfectly competitive equilibrium as a social disequilibrium may be only an illusion created by comparing it to an unrealistic ideal world that nobody really wants. If any imperfectly competitive equilibrum is a *laissez-faire* equilibrium (i.e., the consequence of everyone's pursuit of profit or utility maximization), then there is no disequilibrium (unexploited gains) in the real world.

Imperfect competition vs. psychologistic individualism

Let us consider the implications for possible theories of the imperfectly competitive firm – the firm which either is not a price-taker or has such a large share of the market that its output decisions do affect the price. The general question is: in the long run, when ultimately the firm's profits are driven down to zero but are still maximized for its non-negligible share of the market, can the firm really be considered to be in equilibrium? Following the works of Robinson [1933] and Chamberlin [1934], most textbook theories say yes. But, unlike the perfectly competitive world where anything goes, the imperfectly competitive world seems to be based on an arbitrary institutional assumption that restricts competition. That is, the nature of the market situation has been

© Lawrence A. Boland

exogenously given. Unless the degree of non-perfect competition is explained, it may be an unacceptable given in our explanation. Does this mean that one cannot complete a psychologistic individualist program if one attempts to develop a theory of an imperfect competitor in equilibrium? Or does this merely mean that an imperfectly competitive equilibrium is an illusion and thus that the imperfectly competitive firm is doomed to perpetual disequilibrium? Can an imperfectly competitive firm ever be in equilibrium and thus be explainable?

In order to explain how there can ever be an imperfectly competitive equilibrium, we only need to explain why the possibility of internally reducing average cost would be ignored. The explanation is that if average cost and average revenue are equal and if average revenue will fall faster (rise more slowly) than average cost, then there would be no incentive to reduce cost further. What does this explanation say about the 'apparent' increasing returns? It says that they never really existed or, more generally, that the assessment of costs and benefits is misleading.

This raises an interesting theoretical question. How do we know there are increasing returns? What is the source of the increasing returns? So as to avoid repeating all of the volumes of articles devoted to the puzzle presented by the concept of increasing returns, let us bluntly state the analytical case concerning the existence of increasing returns for a given production function, say *f*, where *f* is defined as

$$\text{Output} = f\,(\text{labor, capital}). \qquad [1]$$

If we were to double both factors and the result is that the output more than doubles, then we would have a case of 'increasing returns'. But how is it possible for there to be increasing returns? If the doubling process has merely meant building an identical plant next door, what is the source of the increase in output beyond the level of the original plant? Either the source is external or the production function has been misspecified, since there must be some third factor which has been more than doubled to account for the increased output. These two possibilities are really the same thing. Some constraint was not stated in the original production function. It should have been,

$$\text{Output} = f\,(\text{labor, capital}, X) \qquad [2]$$

where X is the missing factor. As Harvey Leibenstein would probably say, there could only have been the possibility of increasing returns because one of the factors (namely, X) was previously being used inefficently; that is, the optimum quantity of X was not chosen. Stated another way, there could be the possibility of increasing returns only because the original plant was not maximizing profits with respect to *all* inputs.

Any attempt to explain the existence of increasing returns only brings into question the true nature of the production function [Samuleson, 1947/65, p. 84]. If everything is variable, then exact duplication is possible; hence no production functions can exhibit increasing returns. If increasing returns are possible, then there must be something constraining the variability of one or more of the factors so as to create the possibility of improving efficiency. But if there is something constraining the factors, then there is something which should have been included in the specification of the production function, that is, a missing factor. If it is not included, then we have the methodological problem prescribed by psychologistic individualism. Any non-natural, non-individualist constraints must be explained away.

Explaining disequilibria away

Attempts to explain imperfectly competitive firms raises the key dilemma facing modern theorists. On the one hand, if one is to fulfill the commitment to the psychologistic individualist program, then there cannot be any unexplained non-natural, non-individualist constraints. That is, there cannot be any disequilibria, since a disequilibrium is only possible because something is constraining the attainment of an equilibrium by constraining universal maximization. On the other hand, if imperfect competition exists, then there is something which is constraining competition, and thus something is left unexplained. Even worse, some may say that an imperfectly competitive equilibrium is still a disequilibrium in terms of perfect competition. Only in a perfectly competitive equilibrium is it possible to fulfill all of the requirements of a psychologistic individualist research program.

The key question here is the following. If we accept that a realistic concept of the existence of disequilibrium implies the existence of an endogenous constraint, do we also have to accept the reverse, namely, that the existence of an endogenous constraint implies the existence of a disequilbrium? If one considers the reverse, then the way is open to explaining away disequilbria. One can simply deny the existence of permanent (long-run) endogenous constraints. And if disequilibria can be explained away, then psychologistic individualism will be a feasible research program.

One way disequilibria are explained away is to show that all non-natural constraints are matters of choice. Thus what appears to be a disequilibrium is really an equilibrium, as there are no real possibilities of improvement [e.g., Coase, 1960]; if there were, they would have been

pursued. This way may not appeal to everyone, since this is really an indirect argument that in some way assumes what it is supposedly proving. There is another way which, while more mysterious, is at least direct. It argues that the formal transaction prices do not reflect the actual prices. The actual price is the sum of the formal price and the average cost borne by the *buyer*. For example, many people will wait in a queue to save money at a price-reduction sale. Those who do not wish to wait may go elsewhere and pay a higher price [e.g., DeVany, 1976].

This 'invisible-price' approach can go a long way toward explaining why some may see increasing returns when there really are none. The actual average cost curve may be minimum at the output level corresponding to the textbook's imperfectly competitive long-run equilibrium. The actual demand curve may be perfectly elastic, since all reductions in prices are compensated by offsetting increases in transactions costs. If this is the case, then the formal imperfectly competitive long-run equilibrium is actually a perfectly competitive equilibrium? Even more important is the consequence that the price any individual pays is no longer a social institution. Every individual's actual price is specific to that individual's psychological state concerning willingness to wait. This invisible-price approach gives new meaning to Hayek's view of the impossibility of a successful social planner.

Psychologism in the Short Run

Although in the long run *we* may all be dead, in any long-run equilibrium psychologism and Inductivism live. It is easy to see that psychologism is not jeopardized if we can adopt a view of the world where everything is in long-run equilibrium. Does this mean that if one wishes to build more realistic short-run models, one must abandon the psychologistic individualist research program in favor of a more complicated disequilibrium approach?

The consensus among avant-garde theorists today gives a negative answer to this last question. That is, there seems to be agreement that a realistic short-run neoclassical theory must involve disequilibria that cannot be explained away, yet the requirement of psychologistic individualism must be retained. We would argue that, methodologically, this is self-contradictory – disequilibria imply the existence of non-natural and non-individualist givens, while psychologistic individualism implies only individualist or natural givens. The problem facing contemporary theorists is to find a way either to explain the existence of disequilibria even though individuals are seen to be free to follow only their self-interest or to explain disequilibria away. And so,

> theorists solve the problem by depending primarily on expectational errors as the prime source of divergences from full equilibrium. Economic agents optimize subject to what they perceive to be their circumstances.... Agents have to form expectations about ... unknown or imperfectly known circumstances. One necessary part of the definition of equilibrium in this kind of world is that those expectations be confirmed, at least in some reasonable statistical sense. The way is now open to explain major departures from equilibrium as mainly the result of unusually large and/or unusually prolonged expectational errors. [Solow, 1979, p. 341]

In the next chapter we shall examine this 'expectational errors' approach to short-run disequilibria. We shall argue that as a solution to the methodological problem of disequilibria it is an illusion, as it is based on the acceptance of Inductivism.

© Lawrence A. Boland

4

Rational Expectations and Theories of Knowledge

> Expectations, since they are informed predictions of future events, are essentially the same as the predictions of the relevant economic theory....
>
> The [rational expectations] hypothesis ...[is] that expectations of firms ... tend to be distributed, *for the same information set*, about the prediction of the theory....
>
> John Muth [1961, p. 316, emphasis added]

> At the logical level, Benjamin Friedman has called attention to the omission from MREH [macro rational expectations hypothesis] of an explicit learning model, and has suggested that, as a result, it can only be interpreted as a description not of short-run but of long-run equilibrium in which no agent would wish to recontract. But then the implications of MREH are clearly far from startling, and their policy relevance is almost nil. At the institutional level, Stanley Fischer has shown that the mere recognition of long-term contracts is sufficient to generate wage rigidity and a substantial scope for stabilization policies. But the most glaring flaw of MREH is its inconsistency with the evidence: if it were valid, deviations of unemployment from the natural rate would be small and transitory – in which case *The General Theory* would not have been written.... Sargent (1976) has attempted to remedy this fatal flaw by hypothesizing that the persistent and large fluctuations in unemployment reflect merely corresponding swings in the natural rate itself. In other words, what happened in the United States in the 1930s was a severe attack of contagious laziness! I can say that, despite Sargent's ingenuity, neither I nor, I expect, most others at least of the nonmonetarists' persuasion are quite ready yet to turn over the field of economic fluctuations to the social psychologist!
>
> Franco Modigliani [1977, p. 6]

What is really surprising about rational expectations models is that they employ a 500 year-old theory of knowledge and at the same time ignore the 200 year-old refutation of that theory! It is also surprising that although 'expectations' are now considered a central concern of many mainstream theorists, there is virtually no discussion of the theories of knowledge which must support any concept of rational expectations. This is the case even for the critics of rational expectations. Contrary to Benjamin Friedman's view, all models that employ a Rational Expectations Hypothesis do have a theory of learning. Their theory of learning is not discussed simply because it is presumed that everyone understands and accepts it, since there are no recognized alternatives. Specifically, all such models implicitly presume an Inductivist theory of learning. Rational expectations are nothing more than the standard Conventionalist response to the realization that true induction would require an infinity of both time and information.

Knowledge, Expectations and Equilibrium

Knowledge and long-run equilibria

It was noted in Chapter 3 that one of the necessary conditions for (general) market equilibrium is that all demanders and suppliers are maximizing, which implies also that all potential gains from trade are being exploited. It follows then that equilibrium of all markets entails the successful acquisition of *adequate* knowledge for the purposes of maximization.

Now just what constitutes 'adequate knowledge for the purposes of maximization'? Recall Solow's comment, quoted above, that disequilibria are being explained away today by referring to expectational errors. Specifically, he noted that optimization requires the formation of expectations about 'prices that rule in the future, as well as other facts about the future that cannot now be known'. If an individual or firm is ever going to be successful in maximizing utility or profit, the expectations must be correct. Expectational errors lead to failures to maximize. This leads us to ask what constitutes 'expectational errors'? And more perversely, can one ever expect to be able to avoid 'expectational errors'?

Knowledge and learning in the short run

Years ago Hayek argued that since the individual's acquisition of the (true) knowledge of his or her circumstances, the givens, is essential for any (stable) equilibrium, in order to explain how the economy changes over time we must be able to explain how individuals acquire their knowledge [1937/48, p. 47]. Hayek's concern was that there was no

(inductive) way to show how any individual could ever acquire true knowledge. Hayek pessimistically confessed his inability to offer an explanation for even one individual's knowledge acquisition process [pp. 47-8]. Explicitly, he admitted that he could not specify 'assertions about causal connections, about how experience creates knowledge' [p. 47]; implicitly, he was merely admitting the impossiblity of an inductive proof.

Today, neoclassical theorists are more optimistic. Their optimism is based on the acceptance of Conventionalism (and an ignorance of such things as the paradox of confirmation which we discussed in Chapter 1). Mainstream theorists today do not require that any individual decision-maker have absolutely true knowledge because they would readily admit that inductive proofs have always been impossible. Instead, it would be argued that nobody's knowledge is ever absolutely true but only 'true' according to some degree of probability. Thus, a more moderate view of knowledge would be asserted. Today many theorists would argue that absolutely true knowledge has a probability of 1.00 and that a realistic view of knowledge would say the actual knowledge of any individual or group of individuals has a probability of less than 1.00. Of course, the closer the probability is to 1.00 the better is the knowledge [Malinvaud, 1966]. Given this view of knowledge, it could be argued that learning takes place whenever the probability of one's knowledge is increased – for example, whenever the degree of confirmation has increased.

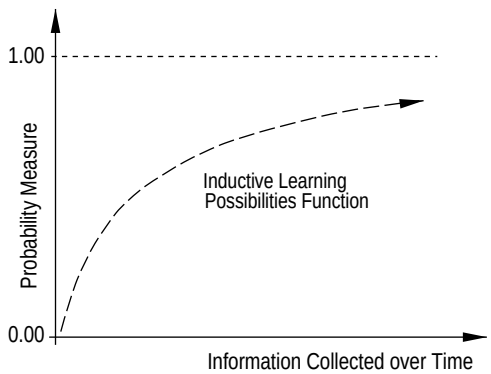

Figure 4.1 *Inductivist Learning*

Let us illustrate this view of learning and knowledge with a simple diagram, Figure 4.1, in which the curve indicates that the inferred probability of one's knowledge increases as information is collected. We shall call such a curve the 'inductive learning possibilities function'.

It is a very short further step to argue that the probablity of the truth of one's knowledge is like a utility function and that learning is only a matter of maximizing the probability. It is precisely this step that has been taken in the formation of the concept of rational expectations.

The Conventionalist theory of learning

The basis of virtually every neoclassical model that involves the recognition of limited or 'imperfect' knowledge is a Conventionalist theory of knowledge and learning – a theory which is merely a short-run version of the old Inductivist theory of knowledge and learning. When we say every neoclassical model we are including in this claim all models of rational expectations and efficient markets, as well as the theories of imperfect information and uncertainty.

Let us review the Inductivist theory of knowledge and learning which we examined in Chapter 1. Briefly stated, this old theory said that individuals learn by collecting (objectively provable) facts and when they have enough of them they are able to induce the true theory which would explain the phenomena encompassed by those facts. Inductivism, as we have said, presumes that such an inductive process is indeed possible. For any specific case the only question at issue is whether enough facts have been collected, or possibly whether the quality of those facts is adequate, or both.

Now the Conventionalist theory of learning, which we also discussed in Chapter 1, recognized that there really is no way to collect enough facts to prove absolutely the truth of any explanation. Instead, the best we can do is to maximize the quantity of facts collected or improve their quality (which sometimes turns out to be the same thing as collecting more facts). One learns either by improving the empirical support for one's theory or by finding a better theory. Switching to another theory would be considered a case of long-run learning, and improving the support of one's present theory would be considered short-run learning.

The important point to be realized here is that the Conventionalist theory of learning is merely a version of the Inductivist theory. The difference is only that absolute proofs (i.e., where probability equals 1.00) are no longer required. Learning, in a sense, has been quantified. Either one learns directly by collecting *more* information (i.e., information about additional variables deemed to be relevant), or one learns indirectly by collecting *more* secondary facts to improve the estimates contained in the present set of information.

© Lawrence A. Boland

The Conventionalist theory of information and knowledge

This leads to the Conventionalist theory of information. At any point in time, the current knowledge is a specific 'information set' (a collection of empirical hypotheses which summarize all information bits, i.e., all data) used to derive propositions about relevant decision variables [Hirshleifer and Riley, 1979]. The quality of the information in the information set is reflected in specific probabilities of the truth of those hypotheses, and thus the quality of information is a direct function of the quantity of evidence available to support the required information set. Thus 'more information' can mean two different things: either more supporting evidence for the hypotheses included in the present information set or an expansion of the information set to include additional variables and hypotheses.

From the perspective of Conventionalism, knowledge is merely a name which refers to the information set of the decision-maker. Knowledge can never differ from, or be anything more than, what is contained in the information set, which in turn is nothing more than what can be inferred from the available empirical evidence and logical analyses relating to the evidence. The more information we have, the better is our knowledge. For those who speak in terms of probabilities, the more information we have, the higher will be the probability that our knowledge is true because either there is more evidence or there are more ways to relate the evidence. There are some who draw an analogy between information and resources and thereby speak of information being efficient [see Hirshleifer and Riley, 1979]. That is, one information set can be more efficient than another if it provides a higher probability, given the same or a smaller quantity of supporting evidence.

We have illustrated this theory graphically in Figure 4.2, in which the curve in Figure 4.1 has been modified by placing the quantity of information for a specific information set (e.g., for a given model of the economy) on the horizontal axis and the probability of our knowledge being true on the vertical. The Conventionalist theory of knowledge posits a specific relationship between these two, which is a monotonically increasing function that increases at a decreasing rate such that it approaches *asymptotically* the horizontal line drawn through the point representing a probability of 1.00. To illustrate the possibility of there being two different information sets (or models), there are two 'inductive learning possibilities functions', one for each information set. Each curve preserves two essential properties of the learning theory presumed in Conventionalism. Each curve is monotonically increasing in order to represent the inductive learning; and the Conventionalist denial of the possibility of inductive proofs is represented by the curve being asymptotic to the horizontal line through 1.00, which means that the probability can never reach a maximum for any finite amount of information.

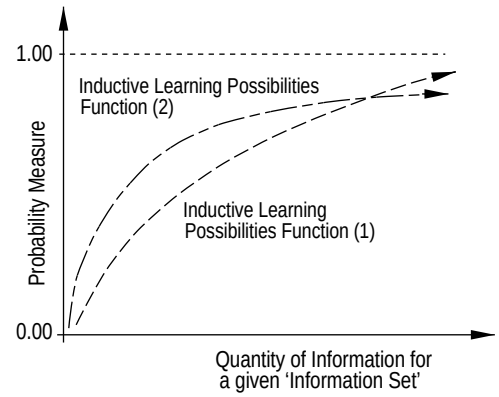

Figure 4.2 *Conventionalist learning*

The Economics of Knowledge and Information

Once one has put the relationship between the quantity of information and the probability of the truth of knowledge into the context of an input-output relation, the way is open to apply economic analysis to the status and acquisition of knowledge. We will suspend our disbelief in such a relationship for now and instead examine the 'economics of information acquisition'.

Given the conception of an informational input-output relationship, we next draw an analogy between it and the utility function of demand theory, in which probability plays the role of utility and information plays the role of the consumed good. Now this is the key: when we say the consumer maximizes his or her utility we do not mean that the absolute maximum is reached but only the highest level of utility that the consumer can afford. The same will apply to the neoclassical conception of a 'rational' acquisition of knowledge. We will see that the rational expectation hypothesis does not assume that the absolute maximum for the probability of one's knowledge is ever obtained. What is assumed is

that the 'learner' must assess the benefits and costs of increasing the probability.

Unlike consumer theory, in which absolute utility maxima may be allowed for finite quantities of any good (i.e., 'bliss points'), this Conventionalist theory of learning specifically denies a maximum probability in the real world. Probability 1.00 is reached *only* with an infinity of information which would require an infinity of time. Thus the marginal productivity of information, so to speak, is always positive, although it approaches zero as the size of the information set grows larger.

Perhaps for philosophers there is never enough information, but for practical economists the benefits of more information may be an insignificant increase in the probability of one's theory, while the cost of the new information may be quite significant. This then is the economics of information. The 'rational' learner has assessed the costs and benefits of seeking one more bit of information (or one more alternative model to consider) and has stopped acquiring information when the extra benefits no longer outweigh their extra cost. The obvious fact that information is always costly, then, adds support to this Conventionalist theory of knowledge and information. In effect, Conventionalism is self-supporting!

Let us illustrate the economics of information with another diagram. Consider Figure 4.3, in which the cost of information is represented by a fixed dollar multiple of the quantity of information collected, and the benefit of information is represented by a fixed dollar multiple of the probability. The economics of information is simply that the optimum amount of information will be obtained when the marginal benefits (the slope of the benefits curve) just equal the marginal cost (the slope of the cost curve). That is, the optimum amount of information collected means only that quantity of information for which the net benefits have been maximized. This also means that as long as there are positive costs for each bit of information collected, the optimum amount of information will not support a probability of 1.00. Economists might argue that even if induction were logically possible it might not be economical!

It should be noted that it is quite easy to make this Conventionalist theory of knowledge part of a psychologistic individualist program. All that we would have to do is to add the additional interpretation that the probability in question is a 'subjective probability' and then the analogy between probablity and utility would be almost redundant. It must be added, though, that this additional interpretation is an unnecessary frill. It only helps some theorists feel that they are not far from fulfilling the requirement of a psychologistic individualist research program.

The Role of the Rational Expectations Hypothesis

The point of all this is that expectations are rational if they are inductively based on the 'best' available information set. The expectations will not usually be absolutely true for the simple reason that to make them so, even if it were logically possible, would cost far too much. We can now state the Rational Expectations Hypothesis that has been receiving so much attention in recent years. Those models employing the Rational Expectations Hypothesis assume merely that every decision-maker has acquired information only to the point that its acquisition is economical. In effect the Rational Expectations Hypothesis is a straightforward application of the maximization hypothesis to knowledge acquisition in a real-world setting where opportunity costs matter – along the lines of Figure 4.3.

Now, let us consider how the Rational Expectations Hypothesis is used in neoclassical research programs based on the hidden agenda of the Problem of Induction (or of Conventions) and psychologistic individualism. For our purposes, we only need to show how the agenda items are served by the Rational Expectations Hypothesis.

Rational expectations and the Problem of Conventions

Whether any given assumption serves as one of the main items on the hidden agenda is only a matter of logical consistency and adequacy. Consider the first item on the agenda – either the Problem of Induction itself or some variant in the form of the Problem of Conventions. On the

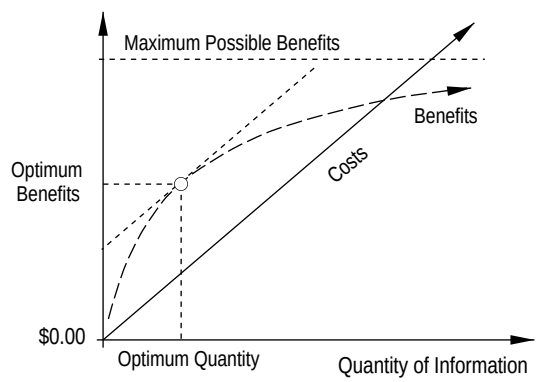

Figure 4.3 *Rational knowledge acquisition*

basis of what we have been saying here, it is relatively easy to see that the Rational Expectations Hypothesis is quite compatible with the Problem of Induction, although no solution is offered. It is the Problem of Conventions that is solved. In one sense, the Rational Expectations Hypothesis is the only logically adequate solution to the Problem of Conventions (although the solution appears to border on the realm of Instrumentalism [see Boland, 1979a]).

The Rational Expectations Hypothesis solves the Problem of Conventions by saying that one should choose the model or theory which maximizes the net benefits of the present information set. This is analogous to maximizing short-run profits given the current capital. Furthermore, rather than providing a solution to the Problem of Induction, it provides a good reason for not requiring inductive proofs, since, even if logically possible, they would be too expensive. The basis of the rational expectations version of Conventionalism might thus be called 'partial' induction. And in this sense it could be argued that traditional Inductivism is a special case based on what might be called 'extreme' induction – which assumes that successful induction implies an absolute proof. A more moderate view of induction might be satisfied with a large quantity of supporting facts even though they do not constitute an absolute proof [e.g., Rotwein, 1980].

Rational expectations and individualism

The major beneficiary of the Rational Expectations Hypothesis is the other main item on the neoclassical hidden agenda, psychologistic individualism. To the extent that individual decision-makers must form some expectations about the market in order to make ordinary decisions, the maintenance of any stable balance or equilibrium necessitates a minimal adequacy of those expectations – in the sense that the benefits of better expectations would not outweigh their cost. For any model which is to be consistent with the hidden agenda of neoclassical economics, the Rational Expectations Hypothesis is supposed to provide the minimal adequacy of the decision-makers expectations .

Although the individual decision-maker is not assumed to hold absolutely true expectations, the Rational Expectations Hypothesis would at least provide that the actual expectations are consistent with any state of equilibrium. Specifically, if the Rational Expectations Hypothesis holds, then there cannot be a possibility of gains to be made by collecting more information. Thus, as long as the remainder of the model implies an equilibrium, the state of the information set will not be a destabilizing influence. In effect, the Rational Expectations Hypothesis repairs the older versions of neoclassical economics which had to presume perfect knowledge in order to assure the existence of an equilibrium.

If everyone is created equal when it comes to knowledge acquisition, then for general equilibrium the Rational Expectations Hypothesis presumes that everyone has gathered the optimum amount of information. Furthermore, if all actions depend on available knowledge, then the decision-makers' knowledge can be revealed by their actions in the same sense that their preferences are revealed by their choices. We note further that such an assumption of equilibrium precludes privileged (or 'insider') information, as any attempt to benefit from secret information must reveal that information in the market for everyone else to see (or infer). This is the basis for some versions of the so-called 'efficient markets hypothesis' which in this sense is closely related to the Rational Expectations Hypothesis. In either case, the idea of an equilibrium – the absence of any affordable further gains – also precludes any gains from the further collection of information.

In terms of our discussion of psychologistic individualism, what the Rational Expectations Hypothesis provides is a 'naturalization' of a potential source of non-individualist or non-natural exogenous variables. It does this by presuming the existence of a natural given which we have labeled the 'inductive learning possibilities function'. We illustrated this in Figures 4.1, 2 and 3. It has the natural property of being everywhere monotonically increasing. Even in those versions of the Rational Expectations Hypothesis which identify subjective probabilities as the end result of learning, the process of learning is constrained by the Nature-given inductive learning possibilities function. The only endogenous variable introduced is the extent of the information collection and that variable is made a matter of choice, like everything else which is not Nature-given.

A Critique of the Critics of Rational Expectations

It is interesting that although there are many critics of the Rational Expectations Hypothesis, none of the leading critics has noted its dependence on induction or inductive learning. Instead, the critics are concerned with the fact that the Rational Expectations Hypothesis is employed in models which deny any effective governmental intervention such as that implied by Keynesian stabilization policies. What is alleged is that the use of the Rational Expectations Hypothesis yields models of the economy in which it is always possible for individuals to gain by 'outguessing' the government. That is, by including the information collected by observing the government's behavior over time, any individuals can induce the true (*ceteris paribus*) outcome of the governmental policy and, if it is in their interest, alter their behavior to change the final outcome in their favor.

The rational expectations learning theory

Given that the critics do not deny the Inductivist underpinnings of all current versions of neoclassical economics, it is difficult to see how they can argue that there is anything logically wrong with equilibrium models which employ the Rational Expectations Hypothesis. Even Benjamin Friedman's criticism [1979] (and Herbert Simon's [1979]) that there is no discussion of an explicit learning theory to back up the acquisition of the expectations falls short of the mark because, as we noted above, there is a learning theory built into the hidden agenda. Thus, no additional learning theory is logically necessary.

Comparative advantage of individualism

Those critics who favor governmental interventions, such as stabilization policies, imply that such policies are necessary because individuals' expectations are often wrong. But does this mean that somehow the government can know more than any individual? On what basis can the critic argue that the government can form better expectations? We suspect that the answer to the latter question is that the government in these modern days of the computer and high-speed communications can gather and process much more information. This is an effective argument if one believes in induction. If there is no inductive logic, then it is just as easy for governments to make mistakes in forming expectations as it is for any individual. If induction is denied, then arguments for stabization policies are just as weak as the typical Rational Expectations model.

Adaptive learning vs. rational expectations

Some critics of Rational Expectations models argue that in the long-run the Rational Expectations Hypothesis is equivalent to the more elementary 'adaptive expectations hypothesis' [B. Friedman, 1979]. That this claim is supposed to be a criticism suggests that there is something wrong with adaptive expectations. What is wrong is that adaptive expectations explanations are *ad hoc*. Adaptive expectations are formed merely by trial and error; each subsequent prediction (expectation) is adjusted as indicated by the sign and magnitude of the previous error. For example, to predict the price, P, one could adjust the prediction according to the difference between the previous prediction (P_E) and the observed price (P_O) as follows:

$$dP_t / dt = h(P_O - P_E) \qquad [1]$$

The usual version of adaptive expectations simply assumes that the parameter h is a fixed constant. In effect, Benjamin Friedman argued that if h is appropriately adjusted, then, in the long-run, rational expectations are the same as adaptive expectations. But, if that is the case, what is wrong with adaptive expectations?

Let us offer a different view. Rather than assuming that h adjusts, the equivalence can be obtained by assuming that the 'true model' being estimated is constant. In these circumstances and given Conventionalism, learning by trial and error is constrained by the inductive learning possibilities functions, as illustrated in Figure 4.4. The parameter h determines the speed of adjustment, but as long as the number of trials is not infinite the probability will still not be 1.00. Trials continue until the probability of one's estimate is 'sufficiently close' to 1.00. That is, referring to Figure 4.4, once the probability is within the distance e of the maximum, the estimate is considered a sufficient approximation.

We argued above that the only question of learning in such models is abou how to determine the quantity of information (or trials) needed. Rational expectations are determined economically. Adaptive expectations are formed arbitrarily in the following sense. The determining factor in this Conventionalist version of adaptive expectations is the 'arbitrary' error factor, e. If e is chosen to conform to standard statistical test criteria, is the number of trials necessarily arbitrary? A proponent of Rational Expectations Models could argue that unless e is chosen so as to indicate the same optimum number of trials as would be determined on the basis of recognizing the cost of the trials, the adaptive learning model will not be the same in the long run.

What is more significant for our purposes is the recognition that the adaptive learning model illustrated here may accurately portray a Conventionalist theory of learning. Yet, *if* one accepts Conventionalism, the rational expectations approach is definitely superior, since it is not arbitrary.

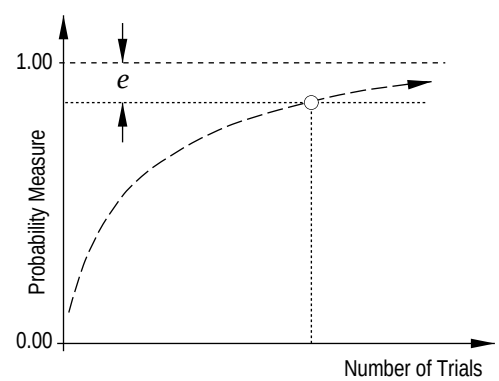

Figure 4.4 *Adaptive Conventionalist learning*