11

Learning and the Equilibrium Process: The Murky Mews

> [T]he fact that any satisfactory theory must be grounded in the theory of individual behavior has further consequences for some current work and for the way in which we view certain forms of analysis. To begin with, the stability problem is not satisfactorily solved by showing that there exists some adjustment processes which converge. However interesting certain adjustment processes may be, unless there is a reason to believe that they arise from the optimizing behavior of agents, they cannot be regarded as providing more than a computational algorithm for finding equilibria. Indeed, the situation here is worse than that involved in the ad hoc specialization of excess demand functions to achieve a stability proof. We know that such specialization *can* obtain under special circumstances. We often do not know that particular convergent processes are *ever* consistent with a sensible story about the behavior of individual agents.
>
> Franklin M. Fisher [1983, pp. 13–4]

> What we clearly need is some higher-level theory, which shows how rules ['of thumb'] are modified in the light of experience. Of necessity, this learning process would have to be non-Bayesian; but equally it could not be one of those rather depressing psychological theories of learning which imply that people never behave optimally however much experience they have. What this new learning theory will look like, I do not know; but there are rich rewards to be gained from it.
>
> John D. Hey [1983, p. 175]

Throughout our tour of questions about stability or disequilibria in this book we have repeatedly encountered questions concerning information and learning. We find disequilibrium theorists appreciating that a price-adjustment mechanism must include a role for learning and information, but they all seem to presume that there is only one way to learn. There are two problems with this presumption. One concerns why there is only one learning method to consider, and the other concerns the nature of the presumed learning method. As we have often noted, there is insufficient reason to claim that there is only one method let alone presume that the only learning method is the inductive method. Nevertheless, we still can understand why theorists (e.g. advocates of the Rational Expectations Hypothesis) presume the inductive learning method. Such a method always provides a strong natural connection between learning and information collection. The strong connection can be used to explain not only learning successes but also learning failures. Any failure to reach an equilibrium can be explained as the result of an insufficiency in the quantity or quality of available information.

The presumption that everyone learns inductively is difficult to justify even on its own terms. How does one learn to learn inductively? If we answer 'inductively', we have an infinite regress; and any other answer admits that learning involves something more than induction. So, there is no inductive proof that there is a necessary connection between the learning process and the accumulation of information. Relying exclusively on an inductive learning theory is self-contradictory! This is of critical importance for the recognition of a role for learning in the process of reaching an equilibrium or in the explanation for the absence of an equilibrium. If learning is still considered to be inductive and thus mechanically connected to the information collected, it will be virtually impossible to build a theory of stability or disequilibrium which is consistent with the requirements of methodological individualism. This is simply because inductive learning is considered an objectively 'rational process' that is so reliable that any rational individuals who collect the same information will reach the same conclusions. Learning in the usual neoclassical analysis is a universal process that is exogenously given.

These observations lead to some interesting questions for all equilibrium model-builders who are interested in the problems of stability. First, if there is no singular inductive learning method, is there a non-inductive method? Second, if there is more than one possible learning method, does this raise a choice-theoretic problem of how the individual chooses his or her learning method? Third, if there are many methods of learning, does the individual's choice of method affect the price-adjustment mechanism? And fourth, if one can choose one's learning method, can one choose the mechanism with which he or she will adjust the price? In this last chapter we will argue that there are

many methods of learning and that this fact must be recognized in any theory of the economy based on the behavior of autonomous individuals if individualism really matters.

1. Learning and Individualism

We wish to argue here that learning is a very individualistic activity and thus any commitment to methodological individualism requires a more fully developed theory of learning. Such a theory may require a reconsideration of methodological individualism itself. There is no good reason for why any two people facing exactly the same information will reach exactly the same conclusion. That is, there is no reason for why two individuals would learn the same thing from the same information set.

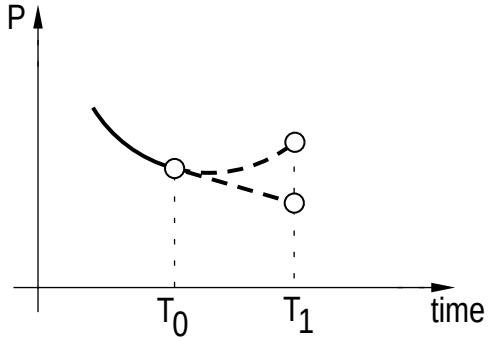

Figure 11.1. Expectations of future price

Consider the case of two individual consumers – see Figure 11.1 – observing, until time T_0, a falling price. The question, at time T_0, is whether to wait for the price to fall further or to buy now before it starts rising. One consumer might have the *a priori* view that prices cannot fall forever and must eventually rise, such that at time T_1 they will be higher. The other might have the *a priori* view that the price will continue to fall, such that it will be lower at time T_1. The former consumer will buy at time T_0 while the latter consumer will want to wait. Yet, at time T_0 the evidence of a falling price is the same for both consumers. The evidence is the same but the conclusions are different simply because the consumers have different *a priori* views of price dynamics in general. Without a reliable inductive learning method that would preclude the possibility of different *a priori* views, such a situation is not unlikely.

Situations as simple as this lie at the heart of disequilibrium macroeconomics. If everyone expects that all prices are going to fall further, there will be a significant deficiency of demand which yields a self-fulfilling expectation. Likewise, whenever one expects that prices are going to stop falling and start rising, one will find it wise to buy now rather than wait. If everyone expects prices to rise and all attempt to act accordingly, prices will be caused to rise by the sudden shift in the demand curve. The issue raised here is not just that expectations matter, but that any widespread agreement about expectations can have significant effects on price dynamics. If there really were only one learning method and it was entirely dependent on the available evidence, whenever everyone used the same evidence (such as in our simple example of a falling price) the expectations would be in widespread agreement.

If the stability of any neoclassical model depends on such widespread agreement to ensure equilibrium prices (as with rational expectations) recognizing that there is no single reliable learning process may mean that stability cannot be guaranteed, even when the available information is sufficient for inductive learning. Even worse, whenever the Keynesian models of persistent disequilibrium are based on a deficient demand caused by a widespread agreement concerning expectations, the absence of a reliable inductive learning method means that there is no sufficient reason for the persistence of the disequilibrium. The question raised here is why in the absence of a reliable inductive learning process would there ever be widespread agreement concerning expectations? Keynes seems to answer this question by saying we have three ways of forming our expectations.

> (1) We assume that the present is a much more serviceable guide to the future than a candid examination of past experience would show it to have been hitherto. In other words we largely ignore the prospect of future changes about the actual character of which we know nothing.
>
> (2) We assume that the *existing* state of opinion as expressed in prices and the character of existing output is based on a *correct* summing up of future prospects, so that we can accept it as such unless and until something new and relevant comes into the picture.
>
> (3) Knowing that our own individual judgment is worthless, we endeavor to fall back on the judgment of the rest of the world which is better informed. That is, we endeavor to conform with the behavior of the majority or the average. The psychology of a society of individuals each of whom is endeavoring to copy the

> others leads to what we may strictly term a *conventional* judgment.
>
> [Keynes, 1937, p. 214]

Now Keynes claims that everyone basing his or her technique of expectation formation on these three alternatives leads 'to sudden and violent changes'. That may be true in the long run, but in the short run it may be just the opposite. Let us apply Keynes' alternatives to our simple example. If prices have been falling, the first technique leads everyone to expect prices to continue to fall. Of course, this cannot go on forever, but how long does it take to get people to stop expecting prices to fall? The second technique does not make sense because falling (disequilibrium) prices may already imply an *incorrect* 'summing up of future prospects'. While the third technique begs the important question about why one individual is less able to form a judgment than the average individual, the short-run outcome is a very stable pattern of behavior, since everyone is following the same conventions – that is, the same 'rule of thumb'.

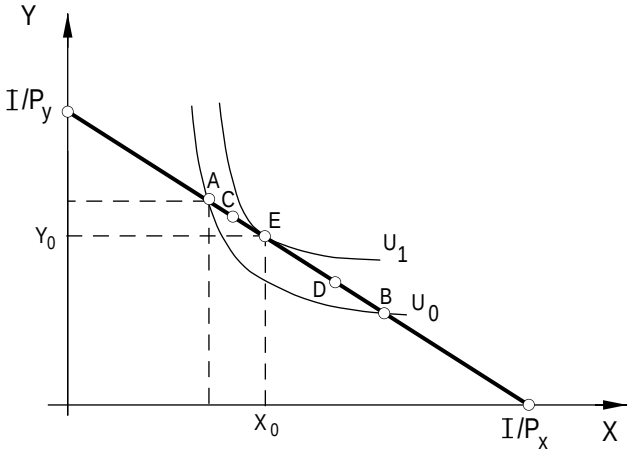

Figure 11.2. Choice theory

Without a single universal inductive learning method, the way each individual processes information must be explained if that information is to be relevant for the formation of expectations. In neoclassical economics, all explanations are matters of choice theory. How does the individual choose his or her learning technique? If there is a choice to be made, there must be many different techniques (Keynes gave us just three). To what extent does the choice of one technique over another imply a different pattern of behavior? If learning is to matter at all, differences between techniques must surely imply behavioral differences. Furthermore, if there are different learning techniques to choose from and if different techniques imply different patterns of demand or supply decisions, to what extent does the frequency distribution of those techniques over any given population affect the stability of the neoclassical equilibrium? And finally, if the distribution does matter, how do we explain it without violating the commitment to methodological individualism?

These questions, we think, form the agenda of any truly new microeconomics if such microeconomics is ever going to overcome the inadequacies of the numerous attempts to build models with stable equilibria or models with adequate explanations of persistent disequilibria. Let us consider each item on this proposed agenda.

2. Learning without Psychologism or Inductivism

Let us return to the paradigm of choice theory, the utility maximizing individual as illustrated in Figure 11.2. Our task is to reconsider how we would explain why an observed individual has purchased the quantities, X_0 and Y_0 represented by point *E*. We start from the usual explanation which says the individual *knows* his or her utility function or preference map and is given the income, *I*, and appropriate prices, P_x and P_y. The observed individual is claimed to have chosen the one point on the budget line where the slope of the budget line equals the slope of the indifference curve through the chosen point. Now, let us change our story. Let us say that while the prices are public knowledge, and the income is in the individual's pocket, the individual *does not know* his or her utility function. Again, the essential question is, why did the individual buy point *E* rather than any other point, such as point *A*? We could answer by claiming that the individual knows that *E* is better than *A*, but this begs the question of how the individual knows this. Did the individual learn this by trying all possible points? Unless all goods are restricted to discrete quantities (see Ch. 5), complete knowledge of the utility function is unlikely in a finite amount of time. There are just too many points to consider – even along the budget line. In the textbook versions which presume perfect divisibility, complete knowledge would be impossible because it would require an infinity of trials.

If the possibility of learning by exhaustive trial and error is effectively denied, what are our options? We could claim that the individual tried two points, *A* and *B*, and knew that they were not the optimum because

in each case the slope of the indifference curve was not equal to the slope of the budget line. This claim, however, would only beg the question of how the individual knows the slope of the indifference curve. Asking this question does not deny the individual's ability to compare points *A*, *B* and *E* once they have been purchased. Once purchased, any point will yield the utility indicated by the true, but incompletely known, utility function. Thus, point *E* is better than either *A* or *B*. But how does the individual know point *E* is the best of all the points between *A* and *B*? How does the individual know the best point is between *A* and *B*, even when he or she has learned that point *E* is better?

What is usually taken for granted is that the individual does not have to learn his or her utility function because it is a psychological given. This presupposition is much too convenient. We concede enough to psychology whenever we claim that the individual can compare two points *a posteriori* on the basis of the derived utilities. Claiming that the individual can compare points that have yet to be consumed goes too far. This is so even for the individual's perception of the slope of an indifference curve at one point, since in practical terms the slope amounts to the comparison of two points. To say otherwise brings up some difficult questions concerning the realism of infinitesimals and similar problems about the realism of calculus (see Ch. 5). Unfortunately, most neoclassical economists believe that a denial of psychology would be a denial of individualism. This belief, which we have called 'psychologism' (see Introduction), actually blocks the way to the neoclassical understanding of individual decision-making. While it may be possible to require that any neoclassical model exclude exogenous variables which are non-individualist and non-natural, the identification of the individualism with psychological states (we called such identification 'psychologistic individualism') reduces the role of the thinking individual to that of a simple mechanical link between his or her psychological state (e.g. tastes) and the optimum choice. There is neither autonomous thinking nor free will in this conception of the individual. We argued in Chapter 1 that the primary reason for building equilibrium models is that the concept of equilibrium allows individuals to make decisions freely (i.e. autonomously) yet it still permits us to explain the state of an entire economy. If our argument is correct, we need to avoid psychologism in future neoclassical models of stability or disequilibrium analysis.

The individual either learns his or her indifference map from experience alone or forms a conjecture about the map. Inductive learning, without the help of some sort of conjectures (Bayesian or otherwise), faces insurmountable problems in real time. If we are going to build realistic models of stability or disequilibrium analysis we must drop the reliance on Inductivism in the explanation of all expectation errors. We need instead to consider some form of autonomous conjectural knowledge. Unlike inductively based knowledge, where any insufficiency is supposedly due to problems with quality or quantity of information, conjectural knowledge has the potential of being wrong in many more ways. To deal with learning and expectations formation using conjectural knowledge we must come to grips with the many difficult questions of methodology [see Boland, 1978, 1979a].

Let us return to the simple paradigm of utility maximization as illustrated in Figure 11.2 where we continue to take prices and income as known givens. We want to continue considering an individual decision-maker who does not know *a priori* his or her indifference map even though there is a true map to be learned. That is, if one could try every point on the map, one's true map could be plotted (by connecting the points with the same level of utility). But since generally that is an impossible task in real time, one must form a conjecture about one's map. Each trip to the market is, then, a test of the individual's conjecture. This immediately raises a primary question of methodological importance. What will be the individual's response to a test which refutes his or her conjectured map? It is most important to note that when individuals have to base their demand or supply decisions on conjectures, whatever would have been considered a market equilibrium or disequilibrium is put into a different light. Even if the market clears today, unless the individuals are satisfied that their respective maps have been correctly conjectured, market clearance does not imply a (stable) equilibrium. If any individuals think they made a mistake even though the market cleared, there is no guarantee that the market will clear the next day.

3. Active Learning and Equilibrium Stability

For the most part, the assumption that decision-makers form conjectures does not dramatically affect our concept of a market equilibrium. Consider further our individual who does not know his or her indifference map and thus has formed a conjecture about his or her indifference map, has made the trip to the market, and is successful in purchasing the quantities planned. If the individual only has a conjecture about his or her indifference map and there have been relatively few trips to the market, how does the individual know that the chosen point is the one which maximizes utility. In simple terms, all that the individual has learned is the level of utility achieved for the chosen point, but he or she does not know whether that level is the best possible, since full knowledge of the map would at least require a very large number of trials.

Obviously there are many ways, each depending on the individual's method of learning from trial and error [see Boland, 1978, 1983a]. How much or what kind of evidence would it take to convince the individual that his or her conjecture is correct? How does an individual learn that his or her tastes have changed? This second question puts in doubt even states of equilibrium where all individuals are convinced that their conjectures are correct that day. If we admit that tastes change, how does the individual know they have changed without trying points which are not optimal according to the currently conjectured map? Both questions raise an important issue. The neoclassical equilibrium model presumes that every individual is choosing the point which maximizes utility, but here we are suggesting that from time to time the individual might deliberately choose a conjectured *sub-optimal* point to test either whether the currently conjectured map is true or whether tastes (and thus the map) have changed. If such perverse behavior is possible, what are the implications for neoclassical equilibrium models?

While such perverse behavior might at first seem to be devastating for the usual neoclassical maximization hypothesis, it has some constructive implications for stability analysis. To see this we need to consider again the conceivable market configurations illustrated in Figure 7.3. In Chapter 7 we said that when individuals facing a disequilibria can adjust the price in response to an insufficient demand or supply (i.e. Walrasian response behavior) and can also adjust the quantity in response to a difference between the demand price and the supply price (i.e. Marshallian behavior), the only configuration that has the possibility of ensuring stability is the market presented in most textbooks. Namely, only when the market is characterized by downward sloping demand curves and upward sloping supply curves, Figure 7.3(b), are we able to conceive of individuals facing a disequilibrium and making independent decisions that constitute stabilizing responses (i.e. convergence toward market clearance). If in Figure 7.3(b) the price were, for any reason, *not* the market-clearing price, the responses of the individuals facing the disequilibrium will (so long as they are small adjustments) always be in the right direction. Small adjustments will never be destabilizing, that is, never cause a greater discrepancy between demand and supply. This is obviously not so in the worst possible case, Figure 7.3(e), where both curves are sloping in the wrong direction; both ways of responding to a disequilibrium would make things worse.

Things are worse for anybody who wishes to argue that we should rely on the competitive market-system for social co-ordination [see Ch. 7 and Hayek, 1945/48]. Whenever people can opt for either Walrasian or Marshallian responses, it would be difficult to argue that prices are informative and thereby promote market equilibria unless the market appears like Figure 7.3(b) (or equivalently, Fig. 7.2). Whenever the market is otherwise, there is always the possibility that either type of response behavior could be destabilizing. In the worst case, Figure 7.3(e), unless demand equals supply when the market opens, either type of response behavior will cause the price to rise or fall at an increasing rate – that is, the market would virtually explode. A similar problem arises when the demand and supply curves slope in the same general direction. For example, if the market is configured like Figure 7.3(a) or (c), Walrasian behavior is stabilizing but Marshallian is destabilizing. Thus, whether the market is stable depends on which behavior dominates in the market. (We will leave this macroeconomic question for now.) Unless we have reasons to ensure that all markets are like Figure 7.3(b), we need to be very careful about recommending complete dependence on the market system as a means of organizing or coordinating society.

In the usual neoclassical model, if the price were accidentally set at the market-clearing price (perhaps with the help of an auctioneer), there would be no need to worry about the relative slopes of the demand and supply curves. When facing such a price, all demanders and suppliers would make the decisions that clear the market and the question of stability would not arise – even if the market were like Figure 7.3(e)! But this is only because in the usual neoclassical model all individuals are assumed to be making the correct decisions and to know that they are making the correct decisions. Any shift from equilibrium would, however, cause an explosive disequilibrium.

Consider now a modified neoclassical model where all decision-makers base their decisions on conjectured objective functions, as in our simple example of the utility maximizer. While the market may be clearing one day, there is no guarantee that it will clear the next day, even when the objective evidence is the same on both days. For example, a consumer may try a sub-optimal point the next day to test his or her conjectured map. Such a test will cause the demand curve to shift from the one which intersected the supply curve the day before.

If we allow people to test their conjectures *and* allow people to respond to disequilibria with either Walrasian or Marshallian behavior, the only type of market that will *guarantee* stability is the textbook's market with a downward sloping demand curve and an upward sloping supply curve, i.e. Figure 7.3(b). We must make such allowances if individuals are free to make any decision they wish. This means that the world outside our window, if it is truly populated by autonomous individuals as all neoclassical economists seem to think, must be a world like Figure 7.3(b) *since if any other configuration existed it would have exploded by now*.

This is a rather indirect argument for the textbook stability case and depends heavily on the predominance of the maximizing behavior usually assumed in neoclassical models. That is, only a small proportion of the market participants can be engaged in perverse testing of their conjectures, otherwise the usual concept of demand and supply curves would lose their meaning. If there is a small proportion of the market participants who are actively testing their conjectures everyday then the only stable market is the textbook market. These considerations emphasize the need to consider the macroeconomic question about which response behavior dominates the whole economy. Why should we expect that in configurations other than Figure 7.3(b) the dominant response will be stabilizing rather than destabilizing when people can behave in either way?

4. Macrofoundations of Microeconomics

Over the years there has been concern for providing microfoundations for equilibrium macroeconomic models – particularly in terms of the adequacy of explanations of disequilibrium phenomena [e.g. Phelps, 1970]. In many ways this may have missed the point. It may be argued that one thing we learned from Keynes is that we lack macrofoundations for equilibrium microeconomics [Boland, 1982a, pp. 79–94]. But surely the question of the stability of the whole market raises the macroeconomic questions noted at the end of the previous section. Which response behavior dominates must be explained using some sort of perspective on the economy as a whole which cannot be deduced from the behavior of individuals alone. Even if we allow for different types of rational response due to differing individual aims, we still must explain the macroeconomic distribution of those aims to ensure stability.

In Chapter 6 we considered a related issue. Hayek [1933/39] pointed out that a disequilibrium might require widespread expectational errors, thus begging the question of why so many people could be wrong. And as we noted above, Keynes [1937] raised questions of how individuals facing 'uncertainty' form expectations, and he answered that there are three different ways to form expectations. While he argued that all the noted ways were destabilizing in the long run, we argued that they may be stabilizing in the short run. Also, we could argue that his potentially destabilizing techniques were destabilizing only if they necessarily lead to false expectations. There is no reason why someone using one of Keynes' techniques of expectation formation cannot happen occasionally to form correct expectations. To think uncertainty necessarily leads to incorrect expectations is merely an expression of a belief in inductivism. Nevertheless, whether the use of any particular technique of forming expectations is stabilizing or destabilizing may depend on how widespread is its use. Unfortunately Keynes' only explanation of how widespread is the use of his techniques was entirely based on his 'psychological laws'.

Earlier in this chapter we saw how expectations can be self-fulfilling in the simple case of price dynamics in Figure 11.1. Clearly it matters whether everyone expects prices to continue falling, but what happens when expectations are mixed? This question is not addressed in the stability literature because of the tendency to think that there is only one technique of expectation formation. If we were to address this question in terms of contemporary stability analysis, we would have to ask why rational decision-makers would have different techniques. Surely, it might usually be said, there is only one technique and it is psychologically given. Obviously, we are continually pushing this line of questioning because we think that not only are there many techniques, but that they are not psychologically given. It is time now to present the problem of expectation formation in different terms.

5. Expectations and Conjectural Knowledge

Without reliable inductive learning methods and without psychologically given tastes and techniques, individuals make decisions on the basis of theories they conjecture to be true. There are many theories involved in any decision. The most simple theories are those about price expectations. For example, we could ask the individual why he or she thinks the price will rise or continue to fall between time T_0 and time T_1 in Figure 11.1. One individual may believe in the theory of inductive learning and say that the reason prices are expected to fall is that they have been observed to be doing nothing but falling for some time. Another individual may believe in an *a priori* theory that average prices are determined by real costs and that the daily price may oscillate about the average such that whenever the price falls for a while it will surely rise to restore the average. The question of widespread agreement over expectations then becomes a question of widespread agreement over either inductive learning or price oscillations. Perhaps a strong argument could be made that this is a question about the sociology of knowledge. It is certainly not a question about any differences between the information sets, as many of the stability theorists might think, since in our example the information set is the same for both individuals.

The idea of basing expectations on conjectural knowledge can easily

be extended to all decision-making processes involved in neoclassical economics [see Boland, 1978]. The individual decision-maker operates on the basis of a conjectured theory of learning, a conjectured indifference map, a theory of how to change the givens such as prices or incomes, and so on. The question of learning, which so many stability theorists are eager to consider, will thus have to be extended to how individuals learn these theories. If we insist that there is no universal learning method such as induction, and that theories are not psychologically given, there is no reason for why everyone should agree about any one operative theory, let alone all of them. To the extent that widespread agreement matters for either stability or persistent disequilibria, there is certainly a need to explain the extent of the agreement. Such an explanation would be an important beginning for a macrofoundation for a new microeconomics.

6. Towards a Generalized Methodological Individualism

We have argued from the beginning that the dominance of equilibrium model-building in neoclassical economics can be understood by seeing how the idea of an equilibrium allows social coordination of free-willed independent decision-makers. An equilibrium model is one designed to foster methodological individualist explanations of the economy. In such explanations only individuals make decisions, and they make them while being ultimately constrained only by the limits of Nature and guided by their own personal aims. Two individuals facing the same circumstances may make different decisions if they have different aims.

The major theoretical problem of neoclassical economics over the last twenty-five years is how to prove that methodological individualism and diverse aims are consistent with the possibility of an equilibrium. So far, the possibility of both diversity and methodological individualism can be shown only in the very special case of a long-run general equilibrium (see Ch. 1). Methodological-individualist consistency proofs can be provided if we preclude diversity, but only at the expense of our concept of individualism (see Ch. 3). We think it would be better to develop a general methodological individualism which would allow for both the diversity of individuals and the possibility of price systems where individuals can reach their aims simultaneously.

Many neoclassical economists are deceived by the attempts to provide proofs of consistency (i.e. ‘existence proofs’) because it is always thought that ‘rational’ decision-making is a psychological process and thus beyond question – the individual always knows what is best for him or her. The rational decision-making process that they have traditionally had in mind is nothing but a calculus-type maximization. It is questionable whether a rational decision-making process is psychologically given or is learned. But, since learning itself has traditionally been thought to be a psychological skill given to everyone, the only question left is whether everyone learns fast enough. The sole element of individuality here is the presumption that some people learn faster than others even though they all use the same method.

While most neoclassical model builders over the last fifty years have been satisfied to think that learning and maximizing are processes beyond question, some have noticed that inductive learning methods are not reliable whenever the available information is quantitatively or qualitatively inadequate (see Ch. 8). Some critics were quick to abandon equilibrium methods on the grounds that learning or maximization would require too much of any available information. Other model-builders have instead seen all this as an interesting puzzle to be solved. How can we assume the economy is in equilibrium when there are insufficient grounds to assume that the individual participants are capable of making decisions that are perfectly consistent with a state of equilibrium?

Some optimistic model-builders may still like to create a workable psychologistic version of methodological individualism that is consistent with inductive learning theory. Can we reject psychologism yet maintain inductivism? If rational decision-making is not psychologically given, perhaps it is learned. But do people learn to be inductive? The infinite regress here should be obvious. Alternatively, maybe we should admit that psychologically given learning methods are informationally insufficient in providing accurate expectations for correct rational decisions. Perhaps, then, individuals can be assumed to know how to deal with such insufficiency of the information basis for their expectations. The Rational Expectations Hypothesis was invented to close this circle. Expectations are rationally chosen like anything else – perhaps, it is said, the rational decision process is more like a combination of econometrics and ordinary calculus. The circle is thus closed without having to give up on the assumption that adequate psychological skills are exogenously given.

Some less optimistic model builders may see a different problem. If equilibria are to be the basis of explanation in economics, the possibility of the existence of a consistent equilibrium is a different issue than whether such an equilibrium can exist. The idea of an equilibrium is not captured by static properties of a single point such as the one where demand equals supply. It must also involve a process of reaching that point. In other words, the textbook idea of an unstable equilibrium is self-contradictory! Not only does the neoclassical idea of an equilibrium require an explanation of the process of reaching the equilibrium but,

most important, that explanation cannot violate the requirements of methodological individualism since this after all is why we are interested in the idea of an equilibrium (see Ch. 1).

The explanation of the process of reaching the equilibrium has been seen as a puzzle concerning how individuals learn independently to make decisions in a manner that is unintentionally stabilizing. It is unfortunate that the solution of this puzzle is still thought by many to require a fuller understanding of the psychology of learning. This will never do, simply because any success will be a denial of methodological individualism, understood as the requirement that all explanations must ultimately be in terms of individual choices, because by definition individuals do not choose their psychological givens. But even worse, the stability of an equilibrium based on expectations which individuals are said to have learned is either a false stability or a false individualism, since the stability is due only to a presumed theory of learning (i.e. where any two individuals facing the same information are thought to form the same expectations). So, the problem of equilibrium stability analysis must be seen as a problem, not only of satisfying methodological individualism by having only individuals make decisions, but of allowing the individuals to be autonomous. The problem of stability needs to be seen as that of how the equilibrium can be stable when individuals facing the same information are systematically forming different expectations.

The possibility of autonomous individualism does not necessarily lead to chaos or anarchy. The reason for stability of an equilibrium in a present day economy may just be that there is considerable homogeneity in the accepted views of learning and proper behavior – even though the homogeneity is not a psychological given. Many neoclassical theorists find pleasure in building equilibrium models which allow both diversity of individual aims and the possibility of a methodological-individualist explanation of prices. This is all too easy since diversity is always provided by the liberal assumption that everyone has different aims. Such diversity is not explained, it is just assumed in neoclassical models. Moreover, inventing new techniques of models with exogenous diversity does not constitute a new microeconomics.

The preoccupation with building equilibrium models with exogenously diverse tastes has overlooked a far more important intellectual challenge. The major theoretical task for neoclassical economics is not to explain why people make different choices when they are given the same information but why so many of them make the same choices when there is so much room to be different. Such homogeneity is endogenous rather than psychologically exogenous. The stability of any given society or economy may be due to individuals choosing to conform rather than being due to the independent individualism on which we like to base our neoclassical theories. A truly new microeconomics must not make individualism exogenous but make it a matter of choice. This surely involves the possibility of heterogeneous learning techniques. While allowing that disequilibrium responses, decision plan formats and diverse tastes are important, we cannot be satisfied with models that fail to provide a truly endogenous basis for the questions of stability or persistent disequlibria. The problem of stability rightfully occupies center stage for anyone interested in intellectually consistent equilibrium models.