Chapter 1
Expectations and the Learning Approach

1.1 Expectations in Macroeconomics

Modern economic theory recognizes that the central difference between economics and natural sciences lies in the forward-looking decisions made by economic agents. In every segment of macroeconomics expectations play a key role. In consumption theory the paradigm life-cycle and permanent income approaches stress the role of expected future incomes. In investment decisions present-value calculations are conditional on expected future prices and sales. Asset prices (equity prices, interest rates, and exchange rates) clearly depend on expected future prices. Many other examples can be given.

Contemporary macroeconomics gives due weight to the role of expectations. A central aspect is that expectations influence the time path of the economy, and one might reasonably hypothesize that the time path of the economy influences expectations. The current standard methodology for modeling expectations is to assume *rational expectations* (RE), which is in fact an equilibrium in this two-sided relationship. Formally, in dynamic stochastic models, RE is usually defined as the mathematical conditional expectation of the relevant variables. The expectations are conditioned on all of the information available to the decision makers. For reasons that are well known, and which we will later explain, RE implicitly makes some rather strong assumptions.

Rational expectations modeling has been the latest step in a very long line of dynamic theories which have emphasized the role of expectations. The earliest references to economic expectations or forecasts date to the ancient Greek philosophers. In Politics (1259a), Aristotle recounts an anecdote concerning the pre-Socratic philosopher Thales of Miletus (c. 636–c. 546 B.C.). Forecasting one winter that there would be a great olive harvest in the coming year, Thales placed deposits for the use of all the olive presses in Chios and Miletus. He then made a large amount of money letting out the presses at high rates when the harvest time arrived.[1] Stories illustrating the importance of expectations in economic decision making can also be found in the Old Testament. In Genesis 41–47 we are told that Joseph (on behalf of the Pharaoh) took actions to store grain from years of good harvest in advance of years in which he forecasted famine. He was then able to sell the stored grains back during the famine years, eventually trading for livestock when the farmers' money ran out.[2]

Systematic economic theories or analyses in which expectations play a major role began as early as Henry Thornton's treatment of paper credit, published in 1802, and Émile Cheysson's 1887 formulation of a framework which had features of the "cobweb" cycle.[3] There is also some discussion of the role of expectations by the Classical Economists, but while they were interested in dynamic issues such as capital accumulation and growth, their method of analysis was essentially static. The economy was thought to be in a stationary state which can be seen as a sequence of static equilibria. A part of this interpretation was the notion of perfect foresight, so that expectations were equated with actual outcomes. This downplayed the significance of expectations.

Alfred Marshall extended the classical approach to incorporate the distinction between the short and the long run. He did not have a full dynamic theory, but he is credited with the notion of "static expectations" of prices. The first explicit analysis of stability in the cobweb model was made by Ezekiel (1938). Hicks (1939) is considered to be the key systematic exposition of the temporary equilibrium approach, initiated by the Stockholm school, in which expectations

[1] In giving this story, as well as another about a Sicilian who bought up all the iron from the iron mines, Aristotle also emphasized the advantage of creating a monopoly.

[2] The forecasting methods used in these stories provide an interesting contrast with those analyzed in this book. Thales is said to have relied on his skill in the stars, and Joseph's forecasts were based on the divine interpretation of dreams.

[3] This is pointed out in Schumpeter (1954, pp. 720 and 842, respectively). Hebert (1973) discusses Cheysson's formulation. The bibliographical references are Cheysson (1887) and Thornton (1939).

of future prices influence demands and supplies in a general equilibrium context.[4] Finally, Muth (1961) was the first to formulate explicitly the notion of rational expectations and did so in the context of the cobweb model.[5]

In macroeconomic contexts the importance of the state of long-term expectations of prospective yields for investment and asset prices was emphasized by Keynes in his General Theory.[6] Keynes emphasized the central role of expectations for the determination of investment, output, and employment. However, he often stressed the subjective basis for the state of confidence and did not provide an explicit model of how expectations are formed.[7] In the 1950s and 1960s expectations were introduced into almost every area of macroeconomics, including consumption, investment, money demand, and inflation. Typically, expectations were mechanically incorporated in macroeconomic modeling using adaptive expectations or related lag schemes. Rational expectations then made the decisive appearance in macroeconomics in the papers of Lucas (1972) and Sargent (1973).[8]

We will now illustrate some of these ways of modeling expectations with the aid of two well-known models. The first one is the cobweb model, though it may be noted that a version of the Lucas (1973) macroeconomic model is formally identical to it. The second is the well-known Cagan model of inflation (see Cagan, 1956). Some other models can be put in the same form, in particular versions of the present-value model of asset pricing.

These two examples are chosen for their familiarity and simplicity. This book will analyze a large number of macroeconomic models, including linear as well as nonlinear expectations models and univariate as well as multivariate models. Recent developments in modeling expectations have gone beyond rational expectations in specifying learning mechanisms which describe the evolution of expectation rules over time. The aim of this book is to develop systematically this new view of expectations formation and its implications for macroeconomic theory.

[4]Lindahl (1939) is perhaps the clearest discussion of the approach of the Stockholm school. Hicks (1965) has a discussion of the methods of dynamic analysis in the context of capital accumulation and growth. However, Hicks does not consider rational expectations.

[5]Sargent (1993) cites Hurwicz (1946) for the first use of the term “rational expectations.”

[6]See Keynes (1936, Chapter 12).

[7]Some passages, particularly in Keynes (1937), suggest that attempting to forecast very distant future events can almost overwhelm rational calculation. For a forceful presentation of this view, see Loasby (1976, Chapter 9).

[8]Most of the early literature on rational expectations is collected in the volumes Lucas and Sargent (1981) and Lucas (1981).

1.2 Two Examples

1.2.1 The Cobweb Model

Consider a single competitive market in which there is a time lag in production. Demand is assumed to depend negatively on the prevailing market price

$$d_t = m_I - m_p p_t + v_{1t},$$

while supply depends positively on the expected price

$$s_t = r_I + r_p p_t^e + v_{2t},$$

where m_p, $r_p > 0$ and m_I and r_I denote the intercepts. We have introduced shocks to both demand and supply. v_{1t} and v_{2t} are unobserved white noise random variables. The interpretation of the supply function is that there is a one-period production lag, so that supply decisions for period t must be based on information available at time $t-1$. We will typically make the representative agent assumption that all agents have the same expectation, but at some points of the book we explicitly take up the issue of heterogeneous expectations. In the preceding equation p_t^e can be interpreted as the average expectation across firms.

We assume that markets clear, so that $s_t = d_t$. The reduced form for this model is

$$p_t = \mu + \alpha p_t^e + \eta_t, \tag{1.1}$$

where $\mu = (m_I - r_I)/m_p$ and $\alpha = -r_p/m_p$. Note that $\alpha < 0$. $\eta_t = (v_{1t} - v_{2t})/m_p$ so that we can write $\eta_t \sim \text{iid}(0, \sigma_\eta^2)$. Equation (1.1) is an example of a temporary equilibrium relationship in which the current price depends on price expectations.

The well-known Lucas (1973) aggregate supply model can be put in the same form. Suppose that aggregate output is given by

$$q_t = \bar{q} + \pi\left(p_t - p_t^e\right) + \zeta_t,$$

where $\pi > 0$, while aggregate demand is given by the quantity theory equation

$$m_t + v_t = p_t + q_t,$$

where v_t is a velocity shock. Here all variables are in logarithmic form. Finally, assume that money supply is random around a constant mean

$$m_t = \bar{m} + u_t.$$

Here u_t, v_t, and ζ_t are white noise shocks. The reduced form for this model is

$$p_t = (1+\pi)^{-1}(\bar{m} - \bar{q}) + \pi(1+\pi)^{-1}p_t^e + (1+\pi)^{-1}(u_t + v_t - \zeta_t).$$

This equation is precisely of the same form as equation (1.1) with $\alpha = \pi(1+\pi)^{-1}$ and $\eta_t = (1+\pi)^{-1}(u_t + v_t - \zeta_t)$. Note that in this example $0 < \alpha < 1$.

Our formulation of the cobweb model has been made very simple for illustrative purposes. It can be readily generalized, e.g., to incorporate observable exogenous variables. This will be done in later chapters.

1.2.2 The Cagan Model

In a simple version of the Cagan model of inflation, the demand for money depends linearly on expected inflation,

$$m_t - p_t = -\psi\left(p_{t+1}^e - p_t\right), \qquad \psi > 0,$$

where m_t is the log of the money supply at time t, p_t is the log of the price level at time t, and p_{t+1}^e denotes the expectation of p_{t+1} formed in time t. We assume that m_t is iid with a constant mean. Solving for p_t, we get

$$p_t = \alpha p_{t+1}^e + \beta m_t, \tag{1.2}$$

where $\alpha = \psi(1+\psi)^{-1}$ and $\beta = (1+\psi)^{-1}$.

The basic model of asset pricing under risk neutrality takes the same form. Under suitable assumptions all assets earn the expected rate of return $1+r$, where $r > 0$ is the real net interest rate, assumed constant. If an asset pays dividend d_t at the beginning of period t, then its price p_t at t is given by $p_t = (1+r)^{-1}p_{t+1}^e + d_t$.[9] This is clearly of the same form as equation (1.2).

1.3 Classical Models of Expectation Formation

The reduced forms (1.1) and (1.2) of the preceding examples clearly illustrate the central role of expectations. Indeed, both of them show how the current

[9] See, e.g., Blanchard and Fischer (1989, pp. 215–216).

market-clearing price depends on expected prices. These reduced forms thus describe a temporary equilibrium which is conditioned by the expectations. Developments since the Stockholm School, Keynes, and Hicks can be seen as different theories of expectations formation, i.e., how to close the model so that it constitutes a fully specified dynamic theory. We now briefly describe some of the most widely used schemes with the aid of the examples.

Naive or static expectations were widely used in the early literature. In the context of the cobweb model they take the form of

$$p_t^e = p_{t-1}.$$

Once this is substituted into equation (1.1), one obtains $p_t = \mu + \alpha p_{t-1} + \eta_t$, which is a stochastic process known as an AR(1) process. In the early literature there were no random shocks, yielding a simple difference equation $p_t = \mu + \alpha p_{t-1}$. This immediately led to the question whether the generated sequence of prices converged to the stationary state over time. The convergence condition is, of course, $|\alpha| < 1$. Whether this is satisfied depends on the relative slopes of the demand and supply curves.[10] In the stochastic case this condition determines whether the price converges to a stationary stochastic process.

The origins of the *adaptive expectations* hypothesis can be traced back to Irving Fisher (see Fisher, 1930). It was formally introduced in the 1950s by Cagan (1956), Friedman (1957), and Nerlove (1958). In terms of the price level the hypothesis takes the form

$$p_t^e = p_{t-1}^e + \lambda\left(p_{t-1} - p_{t-1}^e\right),$$

and in the context of the cobweb model one obtains the system

$$p_t^e = (1 - \lambda(1 - \alpha))\, p_{t-1}^e + \lambda\mu + \lambda\eta_{t-1}.$$

This is again an AR(1) process, now in the expectations p_t^e, which can be analyzed for stability or stationarity in the usual way.

Note that adaptive expectations can also be written in the form

$$p_t^e = \lambda \sum_{i=0}^{\infty} (1 - \lambda)^i\, p_{t-1-i},$$

which is a distributed lag with exponentially declining weights. Besides adaptive expectations, other distributed lag formulations were used in the litera-

[10] In the Lucas model the condition is automatically satisfied.

ture to allow for extrapolative or regressive elements. Adaptive expectations played a prominent role in macroeconomics in the 1960s and 1970s. For example, inflation expectations were often modeled adaptively in the analysis of the expectations-augmented Phillips curve.

The *rational expectations* revolution begins with the observations that adaptive expectations, or any other fixed-weight distributed lag formula, may provide poor forecasts in certain contexts and that better forecast rules may be readily available. The optimal forecast method will in fact depend on the stochastic process which is followed by the variable being forecast, and as can be seen from our examples this implies an interdependency between the forecasting method and the economic model which must be solved explicitly. On this approach we write

$$p_t^e = E_{t-1} p_t \quad \text{and} \quad p_{t+1}^e = E_t p_{t+1}$$

for the cobweb example and for the Cagan model, respectively. Here $E_{t-1} p_t$ denotes the mathematical (statistical) expectation of p_t conditional on variables observable at time $t - 1$ (including past data) and similarly $E_t p_{t+1}$ denotes the expectation of p_{t+1} conditional on information at time t.

We emphasize that rational expectations is in fact an equilibrium concept. The actual stochastic process followed by prices depends on the forecast rules used by agents, so that the optimal choice of the forecast rule by any agent is conditional on the choices of others. An RE equilibrium imposes the consistency condition that each agent's choice is a best response to the choices by others. In the simplest models we have representative agents and these choices are identical.

For the cobweb model we now have $p_t = \mu + \alpha E_{t-1} p_t + \eta_t$. Taking conditional expectations E_{t-1} of both sides yields $E_{t-1} p_t = \mu + \alpha E_{t-1} p_t$, so that expectations are given by $E_{t-1} p_t = (1 - \alpha)^{-1}\mu$ and we have

$$p_t = (1 - \alpha)^{-1}\mu + \eta_t.$$

(This step implicitly imposes the consistency condition described in the previous paragraph.) This is the unique way to form expectations which are "rational" in the model (1.1).

Similarly, in the Cagan model we have $p_t = \alpha E_t p_{t+1} + \beta m_t$ and if m_t is iid with mean $\bar{m}$, a rational expectations solution is $E_t p_{t+1} = (1 - \alpha)^{-1}\beta\bar{m}$ and

$$p_t = (1 - \alpha)^{-1}\alpha\beta\bar{m} + \beta m_t.$$

For this model there are in fact other rational expectations solutions, a point we will temporarily put aside but which we will discuss at length later in the book.

Two related observations should be made. First, under rational expectations the appropriate way to form expectations depends on the stochastic process followed by the exogenous variables η_t or m_t. If these are not iid processes, then the rational expectations will themselves be random variables, and they often form a complicated stochastic process. Second, it is apparent from our examples that neither static nor adaptive expectations are in general rational. Static or adaptive expectations will be "rational" only in certain special cases.

The rational expectations hypothesis became widely used in the 1970s and 1980s and it is now the benchmark paradigm in macroeconomics. In the 1990s, approaches incorporating learning behavior in expectation formation have been increasingly studied.

Paralleling the rational expectations modeling, there was further work refining the temporary equilibrium approach in general equilibrium theory. Much of this work focused on the existence of a temporary equilibrium for given expectation functions. However, the dynamics of sequences of temporary equilibria were also studied and this latter work is conceptually connected to the learning approach analyzed in this book.[11] The temporary equilibrium modeling was primarily developed using nonstochastic models, whereas the approach taken in this book emphasizes that economies are subject to random shocks.

1.4 Learning: The New View of Expectations

The rational expectations approach presupposes that economic agents have a great deal of knowledge about the economy. Even in our simple examples, in which expectations are constant, computing these constants requires the full knowledge of the structure of the model, the values of the parameters, and that the random shock is iid.[12] In empirical work economists, who postulate rational expectations, do not themselves know the parameter values and must estimate them econometrically. It appears more natural to assume that the agents in the economy face the same limitations on knowledge about the economy. This suggests that a more plausible view of rationality is that the agents act like statisti-

[11] Many of the key papers on temporary equilibrium are collected in Grandmont (1988). A recent paper in this tradition, focusing on learning in a nonstochastic context, is Grandmont (1998).

[12] The strong assumptions required in the rational expectations hypothesis were widely discussed in the late 1970s and early 1980s; see, e.g., Blume, Bray, and Easley (1982), Frydman and Phelps (1983), and the references therein. Arrow (1986) has a good discussion of these issues.

cians or econometricians when doing the forecasting about the future state of the economy. This insight is the starting point of the adaptive learning approach to modeling expectations formation. This viewpoint introduces a specific form of "bounded rationality" to macroeconomics as discussed in Sargent (1993, Chapter 2).

More precisely, this viewpoint is called adaptive learning, since agents adjust their forecast rule as new data becomes available over time. There are alternative approaches to modeling learning, and we will explain their main features in Chapter 15. However, adaptive learning is the central focus of the book.

Taking this approach immediately raises the question of its relationship to rational expectations. It turns out that in many cases learning can provide at least an asymptotic justification for the RE hypothesis. For example, in the cobweb model with unobserved iid shocks, if agents estimate a constant expected value by computing the sample mean from past prices, one can show that expectations will converge over time to the RE value. This property turns out to be quite general for the cobweb-type models, provided agents use the appropriate econometric functional form. If the model includes exogenous observable variables or lagged endogenous variables, the agents will need to run regressions in the same way that an econometrician would.[13]

Another major advantage of the learning approach arises in connection with the issue of multiple equilibria. We have briefly alluded to the possibility that under the RE hypothesis the solution will not always be unique. To see this we consider a variation of the Cagan model $p_t = \alpha E_t p_{t+1} + \beta m_t$, where now money supply is assumed to follow a feedback rule $m_t = \bar{m} + \xi p_{t-1} + u_t$. This leads to the equation

$$p_t = \beta\bar{m} + \alpha E_t p_{t+1} + \beta\xi p_{t-1} + \beta u_t.$$

It can be shown that for many parameter values this equation yields two RE solutions of the form

$$p_t = k_1 + k_2 p_{t-1} + k_3 u_t, \tag{1.3}$$

where the k_i depend on the original parameters $\alpha, \beta, \xi, \bar{m}$. In some cases both of these solutions are even stochastically stationary.

[13] Bray (1982) was the first to provide a result showing convergence to rational expectations in a model in which expectations influence the economy and agents use an econometric procedure to update their expectations over time. Friedman (1979) and Taylor (1975) considered expectations which are formed using econometric procedures, but in contexts where expectations do not influence the economy. The final section in Chapter 2 provides a guide to the literature on learning.

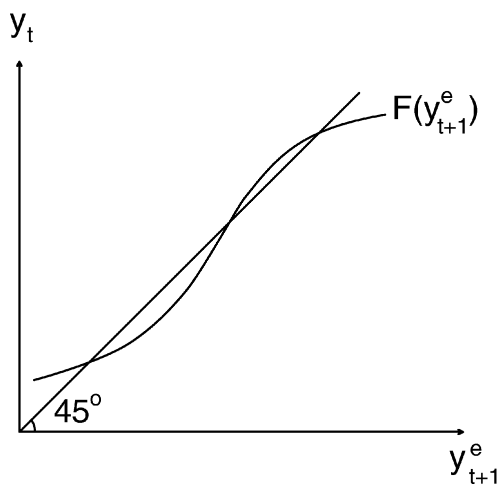

Figure 1.1.

For rational expectations this is a conundrum. Which solution should we and the agents choose? In contrast, in the adaptive learning approach it is supposed that agents start with initial estimates of the parameters of a stochastic process for p_t taking the same functional form as equation (1.3) and revise their estimates, following standard econometric procedures, as new data points are generated. This provides a fully specified dynamical system. For the case at hand it can be shown that only one of the RE solutions can emerge in the long run. Throughout the book the multiplicity issue will recur frequently, and we will pay full attention to this role of adaptive learning as a selection criterion.[14]

In nonlinear models this issue of multiplicity of RE solutions has been frequently encountered. Many nonlinear models can be put in the general form

$$y_t = F\left(y^e_{t+1}\right),$$

where random shocks have here been left out for simplicity. (Note that this is simply a nonlinear generalization of the Cagan model.) Suppose that the graph of $F(\cdot)$ has the S-shape shown in Figure 1.1. The multiple steady states $\bar{y} = F(\bar{y})$ occur at the intersection of the graph and the 45-degree line. We will later give an example in which y refers to output and the low steady states represent coor-

[14]Alternative selection criteria have been advanced. The existence of multiple equilibria makes clear the need to go in some way beyond rational expectations.

dination failures. Under learning, a number of interesting questions arise. Which of the steady states are stable under adaptive learning? Are there statistical learning rules for which there can be rational or nearly rational fluctuations between the steady states? We will treat these and other issues for nonlinear models, allowing also for random shocks.

Finally, the transition under learning to rational expectations may itself be of interest. The process of learning adds dynamics which are not present under strict rationality and they may be of empirical importance. In the cases we just described, these dynamics disappear asymptotically. However, there are various situations in which one can expect learning dynamics to remain important over time. As an example, if the economy undergoes structural shifts from time to time, then agents will need periodically to relearn the relevant stochastic processes. Moreover, if agents know that they are misspecifying a model which undergoes recurrent shifts, they may allow for this in their learning in a way which leads to persistent learning dynamics.

1.5 Statistical Approach to Learning

As already discussed, the approach taken in this book views economic agents as behaving like statisticians or econometricians when they make forecasts of future prices and other economic variables needed in their decision making. As an illustration, consider again the cobweb model (1.1).

Assume that agents believe that the stochastic process for the market price takes the form $p_t = \text{constant} + \text{noise}$, i.e., the same functional form as the RE solution. The sample mean is the standard way for estimating an unknown constant, and in this example the estimate is also the forecast for the price. Thus, suppose that agents' expectations are given by

$$p_t^e = \frac{1}{t}\sum_{i=0}^{t-1} p_i .$$

Combining this with equation (1.1) leads to a fully specified stochastic dynamical system. It can be shown that the system under learning converges to the RE solution if $\alpha < 1$. This result holds, too, for the basic Cagan model (1.2) with iid shocks.

It is easy to think of generalizations. If the economic model incorporates exogenous or lagged endogenous variables, it is natural for the agents to estimate the parameters of the perceived process for the relevant variables by means of least squares. As an illustration, suppose that an observable exogenous variable w_{t-1} is introduced into the cobweb model, so that equation (1.1) takes the form

$$p_t = \mu + \alpha p_t^e + \delta w_{t-1} + \eta_t. \tag{1.4}$$

It would now be natural to forecast the price as a linear function of the observable w_{t-1}. In fact, the unique RE solution is of this form.[15] Under learning, agents would forecast according to

$$p_t^e = a_{t-1} + b_{t-1} w_{t-1}, \tag{1.5}$$

where a_{t-1} and b_{t-1} are parameter estimates obtained by a least squares regression of p_t on w_{t-1} and an intercept.[16]

This way of modeling expectations formation has two major parts. First, the economy is taken to be in a temporary equilibrium in which the current state of the economy depends on expectations. Second, the statistical approach to learning makes the forecast functions and the estimation of their parameters fully explicit. A novel feature of this situation is that the expectations and forecast functions influence future data points. Mathematically, this *self-referential* feature makes these systems nonstandard. Analyzing their dynamics is not trivial and requires special techniques. An overview of those techniques is in the next chapter, and they are presented more formally in later chapters.

1.6 A General Framework

The examples described in the previous sections can be placed in a more general framework. As already noted, the approach taken in this book is *adaptive* in the sense that expectation rules are revised over time in response to observed data. We use the phrase “adaptive learning” to contrast the approach with both “eductive learning” and “rational learning.” In eductive approaches agents engage in a process of reasoning and the learning takes place in logical or notional time. The central question is whether coordination on an REE (rational expectations equilibrium) can be attained by a mental process of reasoning based on

[15]The unique REE is $p_t = \bar{a} + \bar{b} w_{t-1} + \eta_t$, where $\bar{a} = (1-\alpha)^{-1}\mu$ and $\bar{b} = (1-\alpha)^{-1}\delta$.

[16]Bray and Savin (1986) and Fourgeaud, Gourieroux, and Pradel (1986) analyzed cobweb and Cagan models for learning.

common knowledge assumptions.[17] Rational learning takes place in real time,[18] but retains the rational expectations equilibrium assumptions, at each point in time, which we do not want to impose *a priori*. The adaptive learning approach instead assumes that agents possess a form of bounded rationality, which may, however, approach rational expectations over time.

To describe our general framework, let y_t be a vector of variables that agents need to forecast and let y_t^e denote the expectations formed by the agents. y_t could include future values of variables of interest as well as unknown current values. (If agents are heterogeneous in the sense that they have differing expectations, then one can treat this by letting $y_t^e(k)$ denote the expectations of agents k. One would then need to examine the evolution of $y_t^e(k)$ for each agent. For simplicity, we continue the discussion under the assumption of homogeneous expectations.) If the optimal actions of agents depend on the second or higher moments as well as the mean of certain variables, then this can be treated by including powers of these variables in y_t. Similarly, expectations of nonlinear functions of several variables may also be included as components of y_t. Thus, at this stage our framework is very general.

Suppose that agents, when they are making their forecasts y_t^e, have observations on a vector of variables X_t. X_t might include a finite number of lags of some or all components of y_t, and could also include lagged values of y_t^e as well as other exogenous and endogenous observables. The forecasts y_t^e are assumed to be a function of the observables so that

$$y_t^e = \Psi(X_t, \theta_{t-1}),$$

where θ_{t-1} is a vector of parameters that may evolve over time. Inclusion of the parameter θ_{t-1} is a crucial aspect of the adaptive learning approach, as we will discuss. However, the framework so far is broad enough to include static expectations, adaptive expectations, and rational expectations as special cases with appropriate fixed values of θ.[19]

Under the statistical approach to learning, the forecast rule $\Psi(X_t, \theta)$ is based on an econometric model specification, i.e., on a perceived law of mo-

[17] We briefly discuss the eductive approach in Section 15.4 of Chapter 15. There are close links between the stability analyses of eductive and adaptive learning. For a forceful presentation of the eductive viewpoint see Guesnerie (1999).

[18] An example of fully rational learning is Townsend (1978). See the section on the discussion of the literature in Chapter 2 for further references.

[19] Under rational expectations $y_t^e = \Psi(X_t) \equiv E(y_t \mid X_t)$, the mathematical conditional expectation of y_t given X_t, provided $E(y_t \mid I_t) = E(y_t \mid X_t)$, where $I_t = \{X_t, X_{t-1}, \ldots, X_0\}$ is the information set at time t. This will often hold if X_t is chosen appropriately.

tion for the variables of interest, and the vector θ represents unknown parameters which must be estimated statistically in order to implement the forecast rule. As an example, in the cobweb model the forecast rule (1.5) for $y_t^e = p_t^e$ is a linear function of the observables, where X_t includes the variables 1 and w_t and θ_{t-1} includes the parameters a_{t-1} and b_{t-1}. The forecasting framework is completed by specifying a rule for estimating θ and updating the estimates over time as data is accumulated. We will assume that this takes a recursive form $\theta_t = \mathcal{G}(t, \theta_{t-1}, X_t)$. It is convenient to write this in the equivalent form[20]

$$\theta_t = \theta_{t-1} + \gamma_t Q(t, \theta_{t-1}, X_t),$$

where γ_t is a given deterministic "gain" sequence which governs how responsive estimate revisions are to new data.[21] Recursive estimators are sometimes called "on-line," in contrast to "off-line" estimators in which θ_t could depend on the full history $X_1, \ldots X_t$. However, as we will see, many standard statistical estimators such as least squares can be rewritten in recursive form. A simple special case is that the sample mean $a_t = t^{-1} \sum_{i=1}^{t} p_i$, for $t = 1, 2, 3, \ldots,$ can be written in recursive form as $a_t = a_{t-1} + t^{-1}(p_t - a_{t-1})$, where $a_1 = p_1$. Thus, while not completely general, our formulation remains quite general. The recursive version of least squares estimation will be developed in Chapter 2.

The system as a whole is specified once the dynamic process governing the state variables X_t is described. Since the model is self-referential, the dependence of key variables on expectations, manifest in the cobweb model via equations (1.4) and (1.5), will be reflected either in the specification of the process followed by X_t or in the specification of the updating equation $Q(\cdot)$ for the parameter estimates θ_t, or both. This self-referential aspect is what prevents us from analyzing the resulting stochastic dynamic systems using standard econometric tools.

We have thus arrived at a stochastic dynamic system in which economic variables depend through the forecasts of agents on the agents' estimates of key parameters and those parameters are updated over time in response to the evolution of the variables themselves. Analyzing the evolution of this stochastic dynamic system over time is the heart of the adaptive learning approach to expectations formation and the subject of this book.

[20]As we will see in the next chapter, in order to make possible a recursive formulation, θ_t must often include auxiliary parameters in addition to the parameters of interest.

[21]The importance of a recursive formulation for obtaining general adaptive learning results was stressed in Marcet and Sargent (1989c).

1.7 Overview of the Book

In the remainder of Part I we describe the adaptive learning approach to expectations formation in some detail and illustrate it with numerous examples from the recent macroeconomic literature. The analysis in Part I is presented in simplified terms to show the key features and applicability of the approach. The level of exposition is aimed at graduate students and other economists with some familiarity of standard macroeconomic theory. In other parts of the book the style of analysis is rigorous and requires familiarity with techniques presented in Part II. However, Chapters 8, 11, 13, and 15 should be by and large accessible after reading Part I.

Chapter 2 spells out the details of the approach in the context of the cobweb model with agents updating their forecast parameters using recursive least squares. This model, in which there is a unique rational expectations equilibrium (REE), is particularly convenient for introducing the technical framework. We show explicitly how to represent the model under learning as a stochastic recursive algorithm (SRA), how to approximate the system with an associated ordinary differential equation (ODE), and how the asymptotic stability of the REE under learning hinges on a stability condition called "expectational stability" or "E-stability." The discussion of the techniques in this chapter is introductory, emphasizing the heuristic aspects. Chapter 3 shows how some simple variations can lead to interesting further results. In particular, we explore the implications of modifying or misspecifying the recursive least squares learning rule. In that chapter we also discuss a simple form of adaptive learning which can be used in nonstochastic models. The standard coordination failure model is used as an illustration of learning in a nonstochastic context.

Chapter 4, the last chapter of Part I, shows how to use the techniques to study adaptive learning in a wide range of frequently encountered models. The examples include the standard overlapping generations model, the Ramsey optimal growth model, simple linear stochastic macro models, the Diamond growth model, and a model with increasing social returns. We give examples of convergence to REEs, and illustrate the possibility of REEs which are not stable under adaptive learning. In this chapter we also provide an illustration of convergence to a "sunspot" solution, i.e., to a solution which depends on extraneous variables because agents learn to coordinate their expectations on these variables.

Part II provides a systematic treatment of the technical tools required for the analysis of SRAs. Chapter 5 provides a summary of standard results on economic dynamics, with an emphasis on stability results. Topics include difference and differential equations, both deterministic and stochastic, as well as a number of specialized results which will be needed. Chapter 6 presents a formal statement of the key technical results on the stability of SRAs which makes possible the systematic study of adaptive learning in macroeconomic models with expectations. Separate local stability, global stability, and instability results are given. As an illustration we obtain the stability conditions for convergence to the unique REE in the multivariate cobweb model. Chapter 7 presents some additional technical results, including speed of convergence and asymptotic approximations for constant-gain algorithms.

In Parts III and IV we apply the techniques systematically to linear and nonlinear economic models. In these parts we continue to focus on the issue of the conditions under which adaptive learning converges to REE. A major emphasis of these two parts is the possibility of multiple equilibria. As we have already stressed, macroeconomic models in which the state of the economy depends on expectations have the potential for "self-fulfilling prophecies," taking the form of multiple REEs. Local stability under learning provides a natural selection principle for assessing these equilibria.

Part III is a systematic treatment of linear models. Chapters 8 and 9 examine univariate linear models, covering many of the standard workhorses of macroeconomics. Part III begins with a full treatment of several special cases in which the full set of REEs can be readily listed. The solutions take the form of one or more minimal state variable (MSV) solutions and one or more continua of ARMA solutions, possibly depending on "sunspot" variables. In some cases there is a unique REE which is nonexplosive, but examples with multiple stationary solutions do arise. In looking at the local stability of these solutions under least squares learning, we emphasize the role of the E-stability conditions, and we distinguish "weak E-stability conditions," which govern local stability when the REE is correctly specified, and "strong E-stability" conditions, which are relevant when the perceived law of motion estimated by the agents overparameterizes the REE solution of interest. We show how to use the tools of Part II to prove that these conditions govern local stability and instability, under least squares learning, for certain classes of solutions, and we provide supporting numerical simulations for other cases where formal proofs are not available.

Economic examples covered in Chapters 8 and 9 include the Sargent–Wallace "ad hoc" model, Taylor's real balance and overlapping contracts models, the Cagan inflation model, asset price models, investment under uncertainty, Muth's inventory model, and a version of Dornbusch's exchange rate model. Recent dynamic general equilibrium models, such as the Real Business Cycle (RBC) model, are inherently multivariate, as are conventional large-scale macroeconometric models. Although the RBC model is nonlinear, a good approximation is often given by linearized versions. In Chapter 10 we take up multivariate linear expectations models. We show how our techniques to assess the local stability of REE under least squares learning can be extended in a straightforward way to the multivariate setting and we present the E-stability conditions for REE in multivariate linear models. This chapter discusses both "regular" cases, for which there is a unique stationary REE, and "irregular cases" with multiple stationary REE. Economic examples include an IS-LM-new Phillips curve model, the RBC model, and the Farmer–Guo irregular model.

Part IV turns to nonlinear models. From the viewpoint of formal macroeconomic theory these models give rise to the possibility of multiple steady-state REEs, as well as (in nonstochastic models) perfect-foresight cycles and stochastic equilibria which depend on an extraneous variable, a "sunspot." The possibility of multiple steady-state REEs in nonstochastic models is considered in Part I but is more systematically discussed in Chapter 11. This chapter also considers solutions to nonlinear models subject to white noise intrinsic shocks, for example due to random technology, preference, or policy shocks, and we show the existence of "noisy steady states" for nonlinear models with small white noise shocks. We obtain the E-stability conditions and show that they govern local stability under adaptive learning for steady states and noisy steady states. In addition to the overlapping generations (OG) models, with and without stochastic shocks, our examples include the increasing social returns model, the hyperinflation (seignorage) model, and the Evans–Honkapohja–Romer model of growth cycles.

Chapter 12 continues the systematic treatment of nonlinear models. Perfect-foresight cycles can arise in nonlinear models such as the OG model.[22] Chapter 12 shows the possibility of "noisy cycles" in nonlinear models with white noise shocks. We then obtain the E-stability conditions for (perfect-foresight or noisy) cycles and show that these govern the local stability of these cycles under adaptive learning. We derive both weak E-stability conditions, in which the perceived law of motion held by the agents correctly specifies the order of the cycle under consideration, and strong E-stability conditions, which are required for stability when the agents overparameterize the order of the cycle. Sunspot equilibria were originally discovered to exist in nonlinear models, taking the form of Markov chains.[23] Chapter 12 also obtains corresponding weak and strong E-stability conditions and shows that these govern the local stability of stationary sunspot equilibria (and noisy stationary sunspot equilibria) under adaptive

[22] See Grandmont (1985).

[23] See Shell (1977), Azariadis (1981), Azariadis and Guesnerie (1982), and Cass and Shell (1983). The possibility of convergence to sunspots under adaptive learning was shown by Woodford (1990).

learning. Particular attention is paid to the E-stability conditions for Markov sunspot solutions which are close to REE cycles or pairs of steady states.

Part V returns to general issues in adaptive learning. We have been modeling the economic agents as making forecasts in the same way as econometricians. This is a weakening of the rational expectations assumption, but one which would appear reasonable since, after all, economists themselves use econometrics as the principal tool for forecasting. As with all bounded rationality assumptions, one can consider further strengthening or weakening of the degree of rationality. The emphasis of much of the book is on the possibility that econometric learning will asymptotically converge to fully rational expectations. Indeed, we have advocated local stability under adaptive learning as a selection criterion when multiple REEs exist. Chapters 13 and 14 consider the possibility that natural econometric learning rules may fail to converge fully to REEs even in the limit.

In Chapter 13 we consider the implications of agents using a misspecified model. When the perceived law of motion estimated by the agents does not nest the REE under consideration, convergence of learning to that REE is, of course, impossible. This does not, however, preclude convergence of the estimators. We give several examples in which underparameterized learning converges to a restricted perceptions equilibrium under least squares learning. This equilibrium, though not rational, may be optimal given the restricted class of perceived laws of motion entertained by the forecasters. In the cobweb model the stability condition is unaffected, but in other examples misspecification can affect the stability condition as well as the asymptotic point of convergence. This chapter and the next also discuss the model of misspecified learning by monetary policy makers recently set forth in Sargent (1999).

If agents have a misspecified model, they may, however, be aware of the possible misspecification and make allowances for this in their learning rule. One way this can be done is to choose the "gain" sequence, which measures the sensitivity of estimates to new data points, so that it is bounded above zero asymptotically. This is in contrast to standard statistical procedures, like least squares, in which the gain shrinks to zero over time as more data is accumulated. Such nondecreasing or constant-gain estimators have the disadvantage in correctly specified models that estimators fluctuate randomly in the limit, precluding convergence to full rationality. But they have the advantage, in some kinds of misspecified models, of being able to track an economic structure which is evolving in some unknown way. Chapter 14 discusses the implications of constant-gain learning in the context of the cobweb model, the increasing social returns model, and Sargent's inflation model. In some cases, dramatic and new persistent learning dynamics can arise because of the incomplete learning.

Chapter 15 contains a discussion of extensions, alternatives, and new approaches to adaptive learning that have been recently employed. Genetic algorithms, classifier systems, and neural networks are alternative forecasting methods available from the computational intelligence literature. We also discuss eductive approaches, as well as extensions which permit agents to use nonparametric methods or to weigh the costs and benefits of improving forecasts. The chapter ends with a discussion of experimental work and recent empirical applications.

Chapter 16 concludes the book with a perspective on what has been achieved and points out some issues for further research.