Preface to the second edition

Recursive Methods

Much of this book is about how to use recursive methods to study macroeconomics. Recursive methods are very important in the analysis of dynamic systems in economics and other sciences. They originated after World War II in diverse literatures promoted by Wald (sequential analysis), Bellman (dynamic programming), and Kalman (Kalman filtering).

Dynamics

Dynamics studies sequences of vectors of random variables indexed by time, called *time series*. Time series are immense objects, with as many components as the number of variables times the number of time periods. A dynamic economic model characterizes and interprets the mutual covariation of all of these components in terms of the purposes and opportunities of economic agents. Agents *choose* components of the time series in light of their opinions about other components.

Recursive methods break a dynamic problem into pieces by forming a sequence of problems, each one posing a constrained choice between utility today and utility tomorrow. The idea is to find a way to describe the position of the system now, where it might be tomorrow, and how agents care now about where it is tomorrow. Thus, recursive methods study dynamics indirectly by characterizing a pair of *functions*: a transition function mapping the *state* of the model today into the state tomorrow, and another function mapping the state into the other endogenous variables of the model. The *state* is a vector of variables that characterizes the system's current position. Time series are generated from these objects by iterating the transition law.

Recursive approach

Recursive methods constitute a powerful approach to dynamic economics due to their described focus on a tradeoff between the current period's utility and a continuation value for utility in all future periods. As mentioned, the simplification arises from dealing with the evolution of state variables that capture the consequences of today's actions and events for all future periods, and in the case of uncertainty, for all possible realizations in those future periods. This is not only a powerful approach to characterizing and solving complicated problems, but it also helps us to develop intuition, conceptualize, and think about dynamic economics. Students often find that half of the job in understanding how a complex economic model works is done once they understand what the set of state variables is. Thereafter, the students are soon on their way to formulating optimization problems and transition equations. Only experience from solving practical problems fully conveys the power of the recursive approach. This book provides many applications.

Still another reason for learning about the recursive approach is the increased importance of numerical simulations in macroeconomics, and most computational algorithms rely on recursive methods. When such numerical simulations are called for in this book, we give some suggestions for how to proceed but without saying too much on numerical methods.[1]

Philosophy

This book mixes tools and sample applications. Our philosophy is to present the tools with enough technical sophistication for our applications, but little more. We aim to give readers a taste of the power of the methods and to direct them to sources where they can learn more.

Macroeconomic dynamics has become an immense field with diverse applications. We do not pretend to survey the field, only to sample it. We intend our sample to equip the reader to approach much of the field with confidence. Fortunately for us, there are several good recent books covering parts of the field that we neglect, for example, Aghion and Howitt (1998), Barro and Sala-i-Martin (1995), Blanchard and Fischer (1989), Cooley (1995), Farmer (1993), Azariadis

[1] Judd (1998) and Miranda and Fackler (2002) provide good treatments of numerical methods in economics.

(1993), Romer (1996), Altug and Labadie (1994), Walsh (1998), Cooper (1999), Adda and Cooper (2003), Pissarides (1990), and Woodford (2000). Stokey, Lucas, and Prescott (1989) and Bertsekas (1976) remain standard references for recursive methods in macroeconomics. Chapters 6 and Appendix A in this book revise material appearing in chapter 2 of Sargent (1987b).

Changes in the second edition

This edition contains seven new chapters and substantial revisions of important parts of about half of the original chapters. New to this edition are chapters 1, 11, 12, 18, 20, 21, and 23. The new chapters and the revisions cover exciting new topics. They widen and deepen the message that recursive methods are pervasive and powerful.

New chapters

Chapter 1 is an overview that discusses themes that unite many of the apparently diverse topics treated in this book. Because it ties together ideas that can be fully appreciated only after working through the material in the subsequent chapters, we were ambivalent about whether this chapter should be first or last. We have chosen to put this last chapter first because it tells our destination. The chapter emphasizes two ideas: (1) a consumption Euler equation that underlies many results in the literatures on consumption, asset pricing, and taxation; and (2) a set of recursive ways to represent contracts and decision rules that are history-dependent. These two ideas come together in the several chapters on recursive contracts that form Part V of this edition. In these chapters, contracts or government policies cope with enforcement and information problems by tampering with continuation utilities in ways that compromise the consumption Euler equation. How the designers of these contracts choose to disrupt the consumption Euler equation depends on detailed aspects of the environment that prevent the consumer from reallocating consumption across time in the way that the basic permanent income model takes for granted. These chapters on recursive contracts convey results that can help to formulate novel theories of consumption, investment, asset pricing, wealth dynamics, and taxation.

Our first edition lacked a self-contained account of the simple optimal growth model and some of its elementary uses in macroeconomics and public finance. Chapter 11 corrects that deficiency. It builds on Hall's 1971 paper by using the standard nonstochastic growth model to analyze the effects on equilibrium outcomes of alternative paths of flat rate taxes on consumption, income from capital, income from labor, and investment. The chapter provides many examples designed to familiarize the reader with the covariation of endogenous variables that are induced by both the transient (feedback) and anticipatory (feedforward) dynamics that are embedded in the growth model. To expose the structure of those dynamics, this chapter also describes alternative numerical methods for approximating equilibria of the growth model with distorting taxes and for evaluating the accuracy of the approximations.

Chapter 12 uses a stochastic version of the optimal growth model as a vehicle for describing how to construct a recursive competitive equilibrium when there are endogenous state variables. This chapter echoes a theme that recurs throughout this edition even more than it did in the first edition, namely, that discovering a convenient state variable is an art. This chapter extends an idea of chapter 8, itself an extensively revised version of chapter 7 of the first edition, namely, that a measure of household wealth is a key state variable both for achieving a recursive representation of an Arrow-Debreu equilibrium price system, and also for constructing a sequential equilibrium with trading each period in one-period Arrow securities. The reader who masters this chapter will know how to use the concept of a recursive competitive equilibrium and how to represent Arrow securities when there are endogenous state variables.

Chapter 18 reaps rewards from the powerful computational methods for linear quadratic dynamic programming that are discussed in chapter 5, a revision of chapter 4 of the first edition. Our new chapter 18 shows how to formulate and compute what are known as Stackelberg or Ramsey plans in linear economies. Ramsey plans assume a timing protocol that allows a Ramsey planner (or government) to commit, i.e., to choose once-and-for-all a complete state contingent plan of actions. Having the ability to commit allows the Ramsey planner to exploit the effects of its time t actions on time $t+\tau$ actions of private agents for all $\tau \geq 0$, where each of the private agents chooses sequentially. At one time, it was thought that problems of this type were not amenable recursive methods because they have the Ramsey planner choosing a history-dependent strategy. Indeed, one of the first rigorous accounts of the time inconsistency of a Ramsey plan focused on the failure of the Ramsey planner's problem to be recursive in the natural state variables (i.e., capital stocks and information variables). However, it turns out that the Ramsey planner's problem is recursive when the state is augmented by costate variables whose laws of motion are the Euler equations of private agents (or followers). In linear quadratic environments, this insight leads to computations that are minor but ingenious modifications of the classic linear-quadratic dynamic program that we present in chapter 5.

In addition to substantial new material, chapters 19 and 20 contain comprehensive revisions and reorganizations of material that had been in chapter 15 of the first edition. Chapter 19 describes three versions of a model in which a large number of villagers acquire imperfect insurance from a planner or money lender. The three environments differ in whether there is an enforcement problem or some type of information problem (unobserved endowments or perhaps both an unobserved endowments and an unobserved stock of saving). Important new material appears throughout this chapter, including an account of a version of Cole and Kocherlakota's (2001) model of unobserved private storage. In this model, the consumer's access to a private storage technology means that his consumption Euler inequality is among the implementability constraints that the contract design must respect. That Euler inequality severely limits the planner's ability to manipulate continuation values as a way to manage incentives. This chapter contains much new material that allows the reader to get inside the money-lender villager model and to compute optimal recursive contracts by hand in some cases.

Chapter 20 contains an account of a model that blends aspects of models of Thomas and Worrall (1988) and Kocherlakota (1996b). Chapter 15 of our first edition had an account of this model that followed Kocherlakota's account closely. In this edition, we have chosen instead to build on Thomas and Worrall's work because doing so allows us to avoid some technical difficulties attending Kocherlakota's formulation. Chapter 21 uses the theory of recursive contracts to describe two models of optimal experience-rated unemployment compensation. After presenting a version of Shavell and Weiss's (1979) model that was in chapter 15 of the first edition, it describes a version of Zhao's (2001) model of a "lifetime" incentive-insurance arrangement that imparts to unemployment compensation a feature like a "replacement ratio."

Chapter 23 contains two applications of recursive contracts to two topics in international trade. After presenting a revised version of an account of Atkeson's

(1991) model of international lending with both information and enforcement problems, it describes a version of Bond and Park's (2002) model of gradualism in trade agreements.

Revisions of other chapters

We have added significant amounts of material to a number of chapters, including chapters 2, 8, 15, and 16. Chapter 2 has a better treatment of laws of large numbers and two extended economic examples (a permanent income model of consumption and an arbitrage-free model of the term structure) that illustrate some of the time series techniques introduced in the chapter. Chapter 8 says much more about how to find a recursive structure within an Arrow-Debreu pure exchange economy than did its successor. Chapter 16 has an improved account of the supermartingale convergence theorem and how it underlies precautionary saving results. Chapter 15 adds an extended treatment of an optimal taxation problem in an economy in which there are incomplete markets. The supermartingale convergence theorem plays an important role in the analysis of this model. Finally, chapter 26 contains additional discussion of models in which lotteries are used to smooth nonconvexities facing a household and how such models compare with ones without lotteries.

Ideas

Beyond emphasizing recursive methods, the economics of this book revolves around several main ideas.

1. The competitive equilibrium model of a dynamic stochastic economy: This model contains complete markets, meaning that all commodities at different dates that are contingent on alternative random events can be traded in a market with a centralized clearing arrangement. In one version of the model, all trades occur at the beginning of time. In another, trading in one-period claims occurs sequentially. The model is a foundation for asset-pricing theory, growth theory, real business cycle theory, and normative public finance. There is no room for fiat money in the standard competitive equilibrium model, so we shall have to alter the model to let fiat money in.

2. A class of incomplete markets models with heterogeneous agents: The models arbitrarily restrict the types of assets that can be traded, thereby possibly igniting a precautionary motive for agents to hold those assets. Such models have been used to study the distribution of wealth and the evolution of an individual or family's wealth over time. One model in this class lets money in.

3. Several models of fiat money: We add a shopping time specification to a competitive equilibrium model to get a simple vehicle for explaining ten doctrines of monetary economics. These doctrines depend on the government's intertemporal budget constraint and the demand for fiat money, aspects that transcend many models. We also use Samuelson's overlapping generations model, Bewley's incomplete markets model, and Townsend's turnpike model to perform a variety of policy experiments.

4. Restrictions on government policy implied by the arithmetic of budget sets: Most of the ten monetary doctrines reflect properties of the government's budget constraint. Other important doctrines do too. These doctrines, known as Modigliani-Miller and Ricardian equivalence theorems, have a common structure. They embody an equivalence class of government policies that produce the same allocations. We display the structure of such theorems with an eye to finding the features whose absence causes them to fail, letting particular policies matter.

5. Ramsey taxation problem: What is the optimal tax structure when only distorting taxes are available? The primal approach to taxation recasts this question as a problem in which the choice variables are allocations rather than tax rates. Permissible allocations are those that satisfy resource constraints and implementability constraints, where the latter are budget constraints in which the consumer and firm first-order conditions are used to substitute out for prices and tax rates. We study labor and capital taxation, and examine the optimality of the inflation tax prescribed by the Friedman rule.

6. Social insurance with private information and enforcement problems: We use the recursive contracts approach to study a variety of problems in which a benevolent social insurer must balance providing insurance against providing proper incentives. Applications include the provision of unemployment

insurance and the design of loan contracts when the lender has an imperfect capacity to monitor the borrower.

7. Time consistency and reputational models of macroeconomics: We study how reputation can substitute for a government's ability to commit to a policy. The theory describes multiple systems of expectations about its behavior to which a government wants to conform. The theory has many applications, including implementing optimal taxation policies and making monetary policy in the presence of a temptation to inflate offered by a Phillips curve.

8. Search theory: Search theory makes some assumptions opposite to ones in the complete markets competitive equilibrium model. It imagines that there is no centralized place where exchanges can be made, or that there are not standardized commodities. Buyers and/or sellers have to devote effort to search for commodities or work opportunities, which arrive randomly. We describe the basic McCall search model and various applications. We also describe some equilibrium versions of the McCall model and compare them with search models of another type that postulates the existence of a matching function. A matching function takes job seekers and vacancies as inputs, and maps them into a number of successful matches.

Theory and evidence

Though this book aims to give the reader the tools to read about applications, we spend little time on empirical applications. However, the empirical failures of one model have been a main force prompting development of another model. Thus, the perceived empirical failures of the standard complete markets general equilibrium model stimulated the development of the incomplete markets and recursive contracts models. For example, the complete markets model forms a standard benchmark model or point of departure for theories and empirical work on consumption and asset pricing. The complete markets model has these empirical problems: (1) there is too much correlation between individual income and consumption growth in micro data (e.g., Cochrane, 1991 and Attanasio and Davis, 1995); (2) the equity premium is larger in the data than is implied by a representative agent asset-pricing model with reasonable risk-aversion parameter (e.g., Mehra and Prescott, 1985); and (3) the risk-free interest rate is too low relative to the observed aggregate rate of consumption growth (Weil, 1989). While there have been numerous attempts to explain these puzzles by altering the preferences in the standard complete markets model, there has also been work that abandons the complete markets assumption and replaces it with some version of either exogenously or endogenously incomplete markets. The Bewley models of chapters 16 and 17 are examples of exogenously incomplete markets. By ruling out complete markets, this model structure helps with empirical problems 1 and 3 above (e.g., see Huggett, 1993), but not much with problem 2. In chapter 19, we study some models that can be thought of as having endogenously incomplete markets. They can also explain puzzle 1 mentioned earlier in this paragraph; at this time it is not really known how far they take us toward solving problem 2, though Alvarez and Jermann (1999) report promise.

Micro foundations

This book is about micro foundations for macroeconomics. Browning, Hansen, and Heckman (2000) identify two possible justifications for putting microfoundations underneath macroeconomic models. The first is aesthetic and preempirical: models with micro foundations are by construction coherent and explicit. And because they contain descriptions of agents' purposes, they allow us to analyze policy interventions using standard methods of welfare economics. Lucas (1987) gives a distinct second reason: a model with micro foundations broadens the sources of empirical evidence that can be used to assign numerical values to the model's parameters. Lucas endorses Kydland and Prescott's (1982) procedure of borrowing parameter values from micro studies. Browning, Hansen, and Heckman (2000) describe some challenges to Lucas's recommendation for an empirical strategy. Most seriously, they point out that in many contexts the specifications underlying the microeconomic studies cited by a calibrator conflict with those of the macroeconomic model being "calibrated." It is typically not obvious how to transfer parameters from one data set and model specification to another data set, especially if the theoretical and econometric specification differs.

Although we take seriously the doubts about Lucas's justification for microeconomic foundations that Browning, Hansen and Heckman raise, we remain strongly attached to micro foundations. For us, it remains enough to appeal to the first justification mentioned, the coherence provided by micro foundations and the virtues that come from having the ability to "see the agents" in the artificial economy. We see Browning, Hansen, and Heckman as raising many legitimate questions about empirical strategies for implementing macro models with micro foundations. We don't think that the clock will soon be turned back to a time when macroeconomics was done without micro foundations.

Road map

An economic agent is a pair of objects: a utility function (to be maximized) and a set of available choices. Chapter 2 has no economic agents, while chapters 3 through 6 and chapter 16 each contain a single agent. The remaining chapters all have multiple agents, together with an equilibrium concept rendering their choices coherent.

Chapter 2 describes two basic models of a time series: a Markov chain and a linear first-order difference equation. In different ways, these models use the algebra of first-order difference equations to form tractable models of time series. Each model has its own notion of the state of a system. These time series models define essential objects in terms of which the choice problems of later chapters are formed and their solutions are represented.

Chapters 3, 4, and 5 introduce aspects of dynamic programming, including numerical dynamic programming. Chapter 3 describes the basic functional equation of dynamic programming, the Bellman equation, and several of its properties. Chapter 4 describes some numerical ways for solving dynamic programs, based on Markov chains. Chapter 5 describes linear quadratic dynamic programming and some uses and extensions of it, including how to use it to approximate solutions of problems that are not linear quadratic. This chapter also describes the Kalman filter, a useful recursive estimation technique that is mathematically equivalent to the linear quadratic dynamic programming problem.[2] Chapter 6 describes a classic two-action dynamic programming problem, the McCall search model, as well as Jovanovic's extension of it, a good exercise in using the Kalman filter.

[2] The equivalence is through duality, in the sense of mathematical programming.

While single agents appear in chapters 3 through 6, systems with multiple agents, whose environments and choices must be reconciled through markets, appear for the first time in chapters 7 and 8. Chapter 7 uses linear quadratic dynamic programming to introduce two important and related equilibrium concepts: rational expectations equilibrium and Markov perfect equilibrium. Each of these equilibrium concepts can be viewed as a fixed point in a space of beliefs about what other agents intend to do; and each is formulated using recursive methods. Chapter 8 introduces two notions of competitive equilibrium in dynamic stochastic pure exchange economies, then applies them to pricing various consumption streams.

Chapter 9 first introduces the overlapping generations model as a version of the general competitive model with a peculiar preference pattern. It then goes on to use a sequential formulation of equilibria to display how the overlapping generations model can be used to study issues in monetary and fiscal economics, including Social Security.

Chapter 10 compares an important aspect of an overlapping generations model with an infinitely lived agent model with a particular kind of incomplete market structure. This chapter is thus our first encounter with an incomplete markets model. The chapter analyzes the Ricardian equivalence theorem in two distinct but isomorphic settings: one a model with infinitely lived agents who face borrowing constraints, another with overlapping generations of two-period-lived agents with a bequest motive. We describe situations in which the timing of taxes does or does not matter, and explain how binding borrowing constraints in the infinite-lived model correspond to nonoperational bequest motives in the overlapping generations model.

Chapter 13 studies asset pricing and a host of practical doctrines associated with asset pricing, including Ricardian equivalence again and Modigliani-Miller theorems for private and government finance. Chapter 14 is about economic growth. It describes the basic growth model, and analyzes the key features of the specification of the technology that allows the model to exhibit balanced growth.

Chapter 15 studies competitive equilibria distorted by taxes and our first mechanism design problems, namely, ones that seek to find the optimal temporal pattern of distorting taxes. In a nonstochastic economy, the most startling finding is that the optimal tax rate on capital is zero in the long run.

Chapter 16 is about self-insurance. We study a single agent whose limited menu of assets gives him an incentive to self-insure by accumulating assets. We study a special case of what has sometimes been called the "savings problem," and analyze in detail the motive for self-insurance and the surprising implications it has for the agent's ultimate consumption and asset holdings. The type of agent studied in this chapter will be a component of the incomplete markets models to be studied in chapter 14.

Chapter 17 studies incomplete markets economies with heterogeneous agents and imperfect markets for sharing risks. The models of market incompleteness in this chapter come from simply ruling out markets in many assets, without motivating the absence of those asset markets from the physical structure of the economy. We must wait until chapter 19 for a study of some of the reasons that such markets may not exist.

The next chapters describe various manifestations of recursive contracts. Chapter 18 describes how linear quadratic dynamic programming can sometimes be used to compute recursive contracts. Chapter 19 describes models in the mechanism design tradition, work that starts to provide a foundation for incomplete assets markets, and that recovers specifications bearing an incomplete resemblance to the models of chapter 17. Chapter 19 is about the optimal provision of social insurance in the presence of information and enforcement problems. Relative to earlier chapters, chapter 19 escalates the sophistication with which recursive methods are applied, by utilizing promised values as state variables. Chapter 20 extends the analysis to a general equilibrium setting and draws out some implications for asset prices, among other things. Chapter 21 uses recursive contracts to design optimal unemployment insurance and worker compensation schemes.

Chapter 22 applies some of the same ideas to problems in "reputational macroeconomics," using promised values to formulate the notion of credibility. We study how a reputational mechanism can make policies sustainable even when the government lacks the commitment technology that was assumed to exist in the policy analysis of chapter 15. This reputational approach is later used in chapter 24 to assess whether or not the Friedman rule is a sustainable policy. Chapter 23 describes a model of gradualism of in trade policy that has some features in common with the first model of chapter 19.

Chapter 24 switches gears by adding money to a very simple competitive equilibrium model, in a most superficial way; the excuse for that superficial device is that it permits us to present and unify ten more or less well-known monetary doctrines. Chapter 25 presents a less superficial model of money, the turnpike model of Townsend, which is basically a special nonstochastic version of one of the models of chapter 17. The specialization allows us to focus on a variety of monetary doctrines.

Chapter 26 describes multiple agent models of search and matching. Except for a section on money in a search model, the focus is on labor markets as a central application of these theories. To bring out the economic forces at work in different frameworks, we examine the general equilibrium effects of layoff taxes.

Two appendixes collect various technical results on functional analysis and linear control and filtering.

Alternative uses of the book

We have used parts of this book to teach both first- and second-year courses in macroeconomics and monetary economics at the University of Chicago, Stanford University, New York University, and the Stockholm School of Economics. Here are some alternative plans for courses:

1. A one-semester first-year course: chapters 2–6, 8, 9, 10, and either chapter 13, 14, or 15.
2. A second-semester first-year course: add chapters 8, 12, 13, 14, 15, parts of 16 and 17, and all of 19.
3. A first course in monetary economics: chapters 9, 22, 23, 24, 25, and the last section of 26.
4. A second-year macroeconomics course: select from chapters 13–26.
5. A self-contained course about recursive contracts: chapters 18–23.

As an example, Sargent used the following structure for a one-quarter first-year course at the University of Chicago: for the first and last weeks of the quarter, students were asked to read the monograph by Lucas (1987). Students were "prohibited" from reading the monograph in the intervening weeks. During the middle eight weeks of the quarter, students read material from chapters 6 (about search theory); chapter 8 (about complete markets); chapters 9, 24, and 25 (about models of money); and a little bit of chapters 19, 20, and 21 (on social insurance with incentive constraints). The substantive theme of the course was the issues set out in a nontechnical way by Lucas (1987). However, to understand Lucas's arguments, it helps to know the tools and models studied in the middle weeks of the course. Those weeks also exposed students to a range of alternative models that could be used to measure Lucas's arguments against some of the criticisms made, for example, by Manuelli and Sargent (1988).

Another one-quarter course would assign Lucas's (1992) article on efficiency and distribution in the first and last weeks. In the intervening weeks of the course, assign chapters 16, 17, and 19.

As another example, Ljungqvist used the following material in a four-week segment on employment/unemployment in first-year macroeconomics at the Stockholm School of Economics. Labor market issues command a strong interest especially in Europe. Those issues help motivate studying the tools in chapters 6 and 26 (about search and matching models), and parts of 21 (on the optimal provision of unemployment compensation). On one level, both chapters 6 and 26 focus on labor markets as a central application of the theories presented, but on another level, the skills and understanding acquired in these chapters transcend the specific topic of labor market dynamics. For example, the thorough practice on formulating and solving dynamic programming problems in chapter 6 is generally useful to any student of economics, and the models of chapter 26 are an entry-pass to other heterogeneous-agent models like those in chapter 17. Further, an excellent way to motivate the study of recursive contracts in chapter 21 is to ask how unemployment compensation should optimally be provided in the presence of incentive problems.

Matlab programs

Various exercises and examples use Matlab programs. These programs are referred to in a special index at the end of the book. They can be downloaded via anonymous ftp from <ftp://zia.stanford.edu/pub/~sargent/webdocs/matlab>.

Answers to exercises

We have created a web site with additional exercises and answers to the exercises in the text. It is at <http://www.stanford.edu/~sargent>.

Notation

We use the symbol $\blacksquare$ to denote the conclusion of a proof. The editors of this book requested that where possible, brackets and braces be used in place of multiple parentheses to denote composite functions. Thus, the reader will often encounter $f[u(c)]$ to express the composite function $f \circ u$.

Brief history of the notion of the state

This book reflects progress economists have made in refining the notion of state so that more and more problems can be formulated recursively. The art in applying recursive methods is to find a convenient definition of the state. It is often not obvious what the state is, or even whether a finite-dimensional state *exists* (e.g., maybe the entire infinite history of the system is needed to characterize its current position). Extending the range of problems susceptible to recursive methods has been one of the major accomplishments of macroeconomic theory since 1970. In diverse contexts, this enterprise has been about discovering a convenient state and constructing a first-order difference equation to describe its motion. In models equivalent to single-agent control problems, state variables are either capital stocks or information variables that help predict the future.[3] In single-agent models of optimization in the presence of measurement errors, the true state vector is latent or "hidden" from the optimizer and the economist, and needs to be estimated. Here *beliefs* come to serve as the patent state. For example, in a Gaussian setting, the mathematical expectation and covariance matrix of the latent state vector, conditioned on the available history of observations, serves as the state. In authoring his celebrated filter, Kalman

[3] Any available variables that *Granger cause* variables impinging on the optimizer's objective function or constraints enter the state as information variables. See C.W.J. Granger (1969).

(1960) showed how an estimator of the hidden state could be constructed recursively by means of a difference equation that uses the current observables to update the estimator of last period's hidden state.[4] Muth (1960); Lucas (1972), Kareken, Muench, and Wallace (1973); Jovanovic (1979); and Jovanovic and Nyarko (1996) all used versions of the Kalman filter to study systems in which agents make decisions with imperfect observations about the state.

For a while, it seemed that some very important problems in macroeconomics could not be formulated recursively. Kydland and Prescott (1977) argued that it would be difficult to apply recursive methods to macroeconomic policy design problems, including two examples about taxation and a Phillips curve. As Kydland and Prescott formulated them, the problems were not recursive: the fact that the public's forecasts of the government's future decisions influence the public's current decisions made the government's problem simultaneous, not sequential. But soon Kydland and Prescott (1980) and Hansen, Epple, and Roberds (1985) proposed a recursive formulation of such problems by expanding the state of the economy to include a Lagrange multiplier or *costate* variable associated with the government's budget constraint. The costate variable acts as the marginal cost of keeping a promise made earlier by the government. Recently, Marcet and Marimon (1999) have extended and formalized a recursive version of such problems.

A significant breakthrough in the application of recursive methods was achieved by several researchers including Spear and Srivastava (1987); Thomas and Worrall (1988); and Abreu, Pearce, and Stacchetti (1990). They discovered a state variable for recursively formulating an infinitely repeated moral hazard problem. That problem requires the principal to track a history of outcomes and to use it to construct statistics for drawing inferences about the agent's actions. Problems involving self-enforcement of contracts and a government's reputation share this feature. A *continuation value* promised by the principal to the agent can summarize the history. Making the promised valued a state

[4] In competitive multiple-agent models in the presence of measurement errors, the dimension of the hidden state threatens to explode because beliefs about beliefs about ... naturally appear, a problem studied by Townsend (1983). This threat has been overcome through thoughtful and economical definitions of the state. For example, one way is to give up on seeking a purely "autoregressive" recursive structure and to include a moving average piece in the descriptor of beliefs. See Sargent (1991). Townsend's equilibria have the property that prices fully reveal the private information of diversely informed agents.

variable allows a recursive solution in terms of a function mapping the inherited promised value and random variables realized today into an action or allocation today and a promised value for tomorrow. The sequential nature of the solution allows us to recover history-dependent strategies just as we use a stochastic difference equation to find a "moving average" representation.[5]

It is now standard to use a continuation value as a state variable in models of credibility and dynamic incentives. We shall study several such models in this book, including ones for optimal unemployment insurance and for designing loan contracts that must overcome information and enforcement problems.

[5] Related ideas are used by Shavell and Weiss (1979); Abreu, Pearce, and Stacchetti (1986, 1990) in repeated games; and Green (1987) and Phelan and Townsend (1991) in dynamic mechanism design. Andrew Atkeson (1991) extended these ideas to study loans made by borrowers who cannot tell whether they are making consumption loans or investment loans.