Chapter 1

Background and Overview

> Science is facts; just as houses are made of stones, so is science made of facts; but a pile of stones is not a house and a collection of facts is not necessarily science.
>
> —Henri Poincaré

1.1 Background

The seminal contribution of Kydland and Prescott (1982) marked a sea change in the way macroeconomists conduct empirical research. Under the empirical paradigm that remained predominant at the time, the focus was either on purely statistical (or reduced-form) characterizations of macroeconomic behavior, or on systems-of-equations models that ignored both general-equilibrium considerations and forward-looking behavior on the part of purposeful decision makers. But the powerful criticism of this approach set forth by Lucas (1976), and the methodological contributions of, for example, Sims (1972) and Hansen and Sargent (1980), sparked a transition to a new empirical paradigm. In this transitional stage, the formal imposition of theoretical discipline on reduced-form characterizations became established. The source of this discipline was a class of models that have come to be known as dynamic stochastic general equilibrium (DSGE) models. The imposition of discipline most typically took the form of "cross-equation restrictions," under which the stochastic behavior of a set of exogenous variables, coupled with forward-looking behavior on the part of economic decision makers, yield implications for the endogenous stochastic behavior of variables determined by the decision makers. Nevertheless, the imposition of such restrictions was indirect, and reduced-form specifications continued to serve as the focal point of empirical research.

Kydland and Prescott turned this emphasis on its head. As a legacy of their work, DSGE models no longer serve as indirect sources of theoretical discipline to be imposed upon statistical specifications. Instead, they serve directly as the foundation upon which empirical work may be conducted. The methodologies used to implement DSGE models as foundational empirical models have evolved over time and vary considerably. The same is true of the statistical formality with which this work is conducted. But despite the characteristic heterogeneity of methods used in pursuing contemporary empirical macroeconomic research, the influence of Kydland and Prescott remains evident today.

This book details the use of DSGE models as foundations upon which empirical work may be conducted. It is intended primarily as an instructional guide for graduate students and practitioners, and so contains a distinct how-to perspective throughout. The methodologies it presents are organized roughly following the chronological evolution of the empirical literature in macroeconomics that has emerged following the work of Kydland and Prescott; thus it also serves as a reference guide. Throughout, the methodologies are demonstrated using applications to three benchmark models: a real-business-cycle model (fashioned after King, Plosser, and Rebelo, 1988); a monetary model featuring monopolistically competitive firms (fashioned after Ireland, 2004a); and an asset-pricing model (fashioned after Lucas, 1978).

1.2 Overview

The empirical tools outlined in the text share a common foundation: a system of nonlinear expectational difference equations representing a DSGE model. The solution to a given system takes the form of a collection of policy functions that map a subset of the variables featured in the model—state variables—into choices for the remaining subset of variables—control variables; the choices represent optimal behavior on the part of the decision makers featured in the model, subject to constraints implied by the state variables. Policy functions, coupled with laws of motion for the state variables, collectively represent the solution to a given DSGE model; the solution is in the form of a state-space representation.

Policy functions can be calculated analytically for only a narrow set of models, and thus must typically be approximated numerically. The text presents several alternative methodologies for achieving both model approximation and the subsequent empirical implementation of the corresponding state-space representation. Empirical implementation is demonstrated in applications involving parameter estimation, assessments of fit and model comparison, forecasting, policy analysis, and measurement of unobservable facets of aggregate economic activity (e.g., measurement of productivity shocks).

This book is divided into five parts. Part I contains three chapters. Following this brief introduction, chapter 2 establishes the consistent set of notation used throughout the text, and presents an important preliminary step—linearization—of the model solution stage. Chapter 3 then introduces the three benchmark models that serve as examples throughout the text.

Part II contains two chapters devoted to model-solution techniques. Chapter 4 describes solution techniques applicable to linear approximations of systems of expectational difference equations. Chapter 5 then describes three alternative classes of methods for obtaining nonlinear model approximations: projection, iteration, and perturbation. Although nonlinear approximations are necessarily more accurate than their linear counterparts, linear approximations are nevertheless valuable, as they often serve as the foundation upon which nonlinear representations are constructed.

Part III is devoted to data preparation and representation, and contains three chapters. Chapter 6 presents two important preliminary steps often needed for priming data for empirical analysis: removing trends and isolating cycles. The purpose of these steps is to align what is being measured in the data with what is being modeled by the theory. For example, the separation of trend from cycle is necessary in confronting trending data with models of business cycle activity.

Chapter 7 presents tools used to summarize properties of the data. First, two important reduced-form models are introduced: autoregressive moving average models for individual time series, and vector autoregressive models for sets of time series. These models provide flexible characterizations of the data that can be used as a means of calculating a wide range of important summary statistics. Next, a collection of popular summary statistics (along with algorithms available for calculating them) are introduced. These statistics often serve as targets for estimating the parameters of structural models, and as benchmarks for judging their empirical performance. Empirical analyses involving collections of summary statistics are broadly categorized as limited-information analyses.

Chapter 8 concludes Part III with an overview of state-space representations. Such representations characterize the evolution over time of two types of variables: those that are observed (possibly subject to measurement error) by the empirical analyst, and those that are not. Conditional on the parameterization of a given model, and the evolution of the observable variables through a given point in time, evolution of the unobservable variables up to the same point in time may be inferred through a process known as filtering. As a by-product of the filtering process, the likelihood function of the parameterized model can be derived. Chapter 8 first demonstrates the mapping of a given DSGE model into a generic state-space representation. It then provides a general overview of likelihood evaluation and filtering, and then characterizes the Kalman filter, which achieves likelihood evaluation and filtering analytically for the special case in which stochastic innovations and measurement errors are Normally distributed, and model representations are linear. Chapter 8 concludes by characterizing a special class of reduced-form state-space representations: dynamic factor models.

With the exception of the linear/Normal class of state-space representations, likelihood evaluation and filtering requires the calculation of integrals that cannot be computed analytically, but instead must be approximated numerically. Part IV contains two chapters that characterize numerical integration techniques, or Monte Carlo methods. Chapter 9 presents four methods, all in the context of generic integration problems: direct Monte Carlo integration; model simulation; importance sampling; and Markov Chain Monte Carlo. Chapter 10 then characterizes the application of numerical integration techniques to the specific objective of likelihood evaluation and filtering in applications involving state-space representations. The particle filter is first introduced as a foundational method for achieving these objectives. Next, concerns regarding numerical efficiency are introduced, with an eye towards achieving efficiency enhancements through the process of adaption. The EIS filter is then introduced generically as a specific means of achieving such enhancements. Finally, the specific application of the EIS filter to state-space representations associated with DSGE models is outlined in detail.

Taking as input the representation of a given DSGE model in a state-space framework, Part V concludes the text with four chapters devoted to alternative methods for achieving empirical implementation. Specifically, chapters 11 through 14 present the following empirical methodologies: calibration, methods of moments, maximum likelihood, and Bayesian inference. Each chapter contains a general presentation of the methodology, and then presents applications to the benchmark models in pursuit of alternative empirical objectives. (For an alternative textbook presentation of these methods, see Canova, 2007.)

Chapter 11 presents the most basic empirical methodology covered in the text: the calibration exercise, as pioneered by Kydland and Prescott (1982). Original applications of this exercise sought to determine whether models designed and parameterized to provide an empirically relevant account of long-term growth were also capable of accounting for the nature of short-term fluctuations that characterize business-cycle fluctuations, summarized using collections of sample statistics measured in the data. More generally, implementation begins with the identification of a set of empirical measurements that serve as constraints on the parameterization of the model under investigation: parameters are chosen to ensure that the model can successfully account for these measurements. (It is often the case that certain parameters must also satisfy additional a priori considerations.) Next, implications of the duly parameterized model for an additional set of statistical measurements are compared with their empirical counterparts to judge whether the model is capable of providing a successful account of these additional features of the data. A challenge associated with this methodology arises in judging success, because this second-stage comparison is made in the absence of a formal statistical foundation.

The method-of-moments techniques presented in chapter 12 serve as one way to address problems arising from the statistical informality associated with calibration exercises. Motivation for their implementation stems from the fact that there is statistical uncertainty associated with the set of empirical measurements that serve as constraints in the parameterization stage of a calibration exercise. For example, a sample mean has an associated sample standard error. Thus there is also statistical uncertainty associated with model parameterizations derived from mappings onto empirical measurements (referred to generally as statistical moments). Method-of-moment estimation techniques account for this uncertainty formally: the parameterizations they yield are interpretable as estimates, featuring classical statistical characteristics. Moreover, if the number of empirical targets used in obtaining parameter estimates exceeds the number of parameters being estimated (i.e., if the model in question is overidentified), the estimation stage also yields objective goodness-of-fit measures that can be used to judge the model's empirical performance. Prominent examples of moment-matching techniques include the generalized and simulated methods of moments (GMM and SMM), and indirect-inference methods.

Moment-matching techniques share a common trait: they are based on a subset of information available in the data (the targeted measurements selected in the estimation stage). An attractive feature of these methodologies is that they may be implemented in the absence of explicit assumptions regarding the underlying distributions that govern the stochastic behavior of the variables featured in the model. A drawback is that decisions regarding the moments chosen in the estimation stage are often arbitrary, and results (e.g., regarding fit) can be sensitive to particular choices. Chapters 13 and 14 present full-information counterparts to these methodologies: likelihood-based analyses. Given a distributional assumption regarding sources of stochastic behavior in a given model, chapter 13 details how the full range of empirical implications of the model may be assessed via maximum-likelihood analysis, facilitated either by use of the Kalman filter or a numerical counterpart. Parameter estimates and model evaluation are facilitated in a straightforward way using maximum-likelihood techniques. Moreover, given model estimates, the implied behavior of unobservable variables present in the model (e.g., productivity shocks) may be inferred as a by-product of the estimation stage.

A distinct advantage in working directly with structural models is that, unlike their reduced-form counterparts, one often has clear a priori guidance concerning their parameterization. For example, specifications of subjective annual discount rates that exceed 10% may be dismissed out of hand as implausible. This motivates the subject of chapter 14, which details the adoption of a Bayesian perspective in bringing full-information procedures to bear in working with structural models. From the Bayesian perspective, a priori views on model parameterization may be incorporated formally in the empirical analysis, in the form of a prior distribution. Coupled with the associated likelihood function via Bayes' Rule, the corresponding posterior distribution may be derived; this conveys information regarding the relative likelihood of alternative parameterizations of the model, conditional on the specified prior and observed data. In turn, conditional statements regarding the empirical performance of the model relative to competing alternatives, the implied behavior of unobservable variables present in the model, and likely future trajectories of model variables may also be derived.

In the spirit of reducing barriers to entry into the field, we have developed a textbook website that contains the data sets that serve as examples throughout the text, as well as computer code used to execute the methodologies we present. The code is in the form of procedures written in the GAUSS programming language. Instructions for executing the procedures are provided within the individual files. The website address is http://www.pitt.edu/~dejong/text.htm. References to procedures available via this site are provided throughout this book. In addition, a host of freeware is available throughout the Internet. In searching for code, good starting points include the collection housed by Christian Zimmerman in his Quantitative Macroeconomics web page, and the collection of programs that comprise DYNARE: http://dge.repec.org/, http://www.dynare.org/.

Much of the code provided at our website reflects the modification of code developed by others, and we have attempted to indicate this explicitly whenever possible. Beyond this attempt, we express our gratitude to the many generous programmers who have made their code available for public use. Finally, we thank Ryan Davis and Lauren Donahoe for assistance with updating the data sets reported in the first edition of the text.