Computational Macroeconomics for the Open Economy

G. C. Lim and Paul D. McNelis

The MIT Press
Cambridge, Massachusetts
London, England

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please e-mail special_sales@mitpress.mit.edu or write to Special Sales Department, The MIT Press, 55 Hayward Street, Cambridge, MA 02142.

This book was set in Palatino on 3B2 by Asco Typesetters, Hong Kong and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Lim, G. C. (Guay C.)
Computational macroeconomics for the open economy / G. C. Lim and Paul D. McNelis.
p. cm.
Includes bibliographical references and index.
ISBN 978-0-262-12306-8 (hbk. : alk. paper) 1. Econometric models. 2. Macroeconomics—Mathematical models. I. McNelis, Paul D. II. Title.
HB141.L54 2008
339.01′5195—dc22 2008011200

10 9 8 7 6 5 4 3 2 1

1 Introduction

The focus of this book is on a computational approach to the analysis of macroeconomic adjustments in an open or globalized economy. Specifying, calibrating, solving, and simulating a model for evaluating alternative policy rules can appear to be a cumbersome task. There are, of course, many different types of models to choose from, alternative views about likely parameter values, multiple approximation methods to try, and different options about simulation.

In this chapter we give a brief overview of the issues arising from the agenda we set for this book and the rationale for the structure of the book, the methodology adopted, and the economic experiments considered. Since the same solution method will be used throughout the book, to minimize repetitions, we provide more details here about the solution method, the approximating functions and the optimization algorithms used.

1.1 The Open Economy Setting

This book uses computational experiments to obtain insights about macroeconomic adjustments in the open economy setting. These analyses can then inform the design of policies such as the best inflation targeting program or the best tax regime.

Benigno and Woodford (2004) have pointed out, that too often monetary and fiscal policy rules have been discussed in isolation from each other, but they opt to work in a closed economy setting, within a linear quadratic framework to yield analytical closed form solutions for monetary and fiscal policy rules. In contrast, we adopt the open economy setting for our discussion of monetary and fiscal policies and abandon the quest for analytical results in favor of numerical approaches. In so doing, we also extend our discussion of policy issues to encompass inflation targeting and the problem of recurring deficits or surpluses in the fiscal and current-account deficits.

Incorporating the open economy setting, of course, raises issues about international trade and finance, external borrowing conditions and assumptions about "closing" the open economy. As Schmitt-Grohé and Uribe (2003) have pointed out, there are many alternative ways to do this, all of which involve further complications to the standard models used for monetary and fiscal policy analysis.

Discussions about monetary policy, by their very nature, involve assumptions about price stickiness. In the closed economy setting such stickiness can come about either in wage or price-setting behavior in monopolistically competitive markets. Once we move to an open economy environment, we face stickiness in the pricing of imported goods, and thus the case of incomplete pass-through of exchange-rate changes to the prices of imported goods.

The variety of shocks or exogenous forces affecting the economy also expands when we move to the open economy setting. In addition to the usual productivity changes driving a business cycle, there are terms of trade shocks, foreign interest rate developments, and global demand variables to consider. The open economy setting is much more exposed to varying types of shocks.

Discussions of optimal policy in the open economy, then, involve much more complexity than corresponding discussions in the closed economy setting. The models need to be closed, and there are different ways to do this (including the use of a two-country model). Furthermore a reasonable case can be made for "stickiness" in the pricing of imported goods, as well as in domestic price-setting behavior, which in turn involves both forward and backward-looking behavior in the imported-goods sector of the economy.

The models we use in this book are in the class of so-called open economy new neoclassical synthesis (NNS) models. Such models, as Goodfriend (2002) reminds us, incorporate classical features such as the real business cycle, as well as Keynesian features, such as monopolistically competitive firms and costly price adjustment. As Canzoneri, Cumby, and Diba (2004) note, such models have been routinely used to revisit the central issues of stabilization policy.

Different general equilibrium models can generate different effects, so it is essential to have a good strategy for developing a good dynamic stochastic general equilibrium (DSGE) model. As McCallum (2001) points out, it is desirable for a model to be consistent with both economic theory and empirical evidence, but this "dual requirement" is only a starting point for consideration of numerous issues. McCallum also points out that "depicting individuals as solving dynamic optimization problems," as is done in general equilibrium settings, is "useful in tending to reduce inconsistencies and forcing the modeler to think about the economy in a disciplined way" (McCallum 2001, p. 15). But adhering to dynamic general equilibrium models still leaves room for enormous differences, as the reader will see as the chapters unfold.

In this book we focus on variations of one prototype model of the open economy; complexity is introduced, by adding extra economic features, chapter by chapter. While there are many unresolved issues about macroeconomic adjustments and the conduct of policy in the open economy, the differing positions rest on specific assumptions in the models. Rather than review a myriad of conflicting positions based on differing models, we work with increasingly complex versions of the prototype model. The same productivity shock is considered in each case. However, to gain further insight, we also compare the dynamic responses of key variables to other shocks, such as exports and the terms of trade. The progressive addition of complexity highlights the contribution of each added economic feature and aids in the understanding of the economic results and the derived implications for policy rules in an open economy setting.

The model is calibrated rather than estimated—the recent development of estimation techniques for DSGE models deserves a separate book. However, the parameters are based on estimates which are widely accepted. Thus our model is not only completely based on underlying optimization decisions of economic agents, at the household, firm, and policy-making level, it is also meant to be reasonably realistic. To put this point another way, following Canova (2007), what is relevant for us is the extent to which our series of "false" models yield coherent explanations of interesting aspects of data, while maintaining highly stylized structures (Canova 2007, p. 251). Thus the models we use are widely shared, if not consensus, benchmarks of how to model an open economy for policy evaluation.

1.2 Solution Methods

DSGE models, no matter how simple, do not have closed form solutions except under very restrictive circumstances (e.g., logarithmic utility functions and full depreciation of capital). We have to use computational methods if we are going to find out how the models behave for a given set of initial conditions and parameter values. However, the results may differ, depending on the solution method. Moreover there is no benchmark exact solution for this model, against which we can compare the accuracy of alternative numerical methods.[1]

There are, of course, a variety of solution methods. Every practicing computational economist has a favorite solution method (or two). And even with a given solution method there are many different options, such as the functional form to use in any type of approximating function, or the way in which we measure the errors for finding accurate decision rules for the model's control variables. The selection of one method or another is as much a matter of taste as well as convenience, based on speed of convergence and the amount of time it takes to set up a computer program.

Briefly, there are two broad classes of solution methods: perturbation and projection methods. Both are widely used and have advantages and drawbacks. We can illustrate these differences with reference to the well-known example of an agent choosing a stream of consumption c_t that maximizes her utility function U, which then defines the capital k accumulation, given the production function f and productivity process z_t,

$$\max_{c_t} \sum_{t=1}^{\infty} \beta^t U(c_t), \tag{1.1}$$

$$k_{t+1} = f(z_t, k_t) - c_t, \tag{1.2}$$

$$z_t = \rho z_{t-1} + e_t, \qquad e_t \sim N(0, \sigma^2). \tag{1.3}$$

The first-order condition for the problem is

$$U'(c_t) = \beta U'(c_{t+1}) f'(k_{t+1}). \tag{1.4}$$

The system has one forward-looking variable for the evolution of c_t, and one state variable k_t that depends on the values of the forward-looking variable, c_t, and the previous period's values k_{t-1}. The key to solving the model is to find ways to represent functional forms ("decision rules")[2] for these controls, as these rules depend on the lagged values of the state variables. Once we do this, the system becomes fully recursive and the dynamic process is generated (given an initial value for k).

1.2.1 Perturbation Method

The first method—the perturbation method—involves a local approximation based on a Taylor expansion. For example, let $h(x_t)$ represent the decision rule (or policy function) for c_t based on the vector of state variables $x_t = [z_t, k_t]$ around the steady-state x_0:

$$h(x_t) = h(x_0) + h'(x_o)(x_t - x_0) + \frac{1}{2}h''(x_0)(x_t - x_0)^2 + \cdots .$$

Perturbation methods have been extensively analyzed by Schmidt-Grohé and Uribe (2004). The first-order perturbation approach (a first-order Taylor expansion around the steady state) is identical to the most widely used solution method for dynamic general equilibrium models, namely linearization or log linearization of the Euler equations around a steady state (for examples, see Uribe 2003). The linear model is then solved using the methods for forward-looking rational expectations such as those put forward by Blanchard and Kahn (1980) and later discussed by Sims (2001).

Part of the appeal of this approach lies with the fact that the solution algorithm is fast. The linearized system is quickly and efficiently solved by exploiting the fact that it can be expressed as a state-space system. Vaughan's method, popularized by Blanchard and Khan (1980), established the conditions for the existence and uniqueness of a rational expectations solution as well as providing the solution. Canova (2007) summarizes this method as essentially an eigenvalue–eigenvector decomposition on the matrix governing the dynamics of the system by dividing the roots into explosive and stable ones.

This first-order approach can be extended to higher order Taylor expansions. Moving from a first to a second-order approximation simply involves adding second-order terms linearly in the specification of the decision rules. Since the Taylor expansion has both forward-looking and backward-looking state variables, these methods also use the same Blanchard-Kahn (1980) method as the first-order approach. Collard and Julliard (2001a, b) offer first- and second-order perturbation methods in their DYNARE software system.

Log-linearization is an example of the "change of variable" method for a first-order perturbation method. Fernández-Villaverde and Rubio-Ramírez (2005) take this idea one step further within the context of the perturbation method. The essence of the Fernández-Villaverde and Rubio-Ramírez approach is to use a first or second-order perturbation method but with transformation of the variables in the decision rule from levels to power-functions. Just as a log-linear transformation is easily applied to the linear or first order perturbation representation, these power transformations may be done in the same way. The process simply involves iterating on a set of parameters for the power functions, in transforming the state variables, for minimizing the Euler equation errors. The final step is to back out the level of the series from the power transformations, once the best set of parameters is found. They argue that this method preserves the fast linear method for efficient solution while capturing model nonlinearities that would otherwise not be captured by the first-order perturbation method.

We note that the second-order method remains, like the first-order method, a local method. As such, as Fernandez-Villaverde (2006, p. 39) observes, it approximates the solution around the deterministic steady state and it is only valid within a specific radius of convergence. Overall, the perturbation method is especially useful when the dynamics of the model consists of small deviations from the steady-state values of the variables. It assumes that there are no asymmetries, no threshold effects, no types of precautionary behavior, and no big transitional changes in the economy. The perturbation methods are local approximations, in the sense that they assume that the shocks represent small deviations from the steady state.

While these methods are fast and easy to implement, they suffer from one important drawback: the shocks must be small.[3] First- and second-order perturbation methods go beyond linearization by making use of first- and second-order Taylor expansions of the Euler equations around the steady state. However, both linearization and perturbation methods leave out any possibility of asymmetric behavior widely observed in the adjustment of asset prices and other key macroeconomic variables. While this is fine for discussion of very small shocks, it is limiting for large or recurring disturbances.

1.2.2 Projection Methods and Accuracy Tests

This book applies the projection method to solve the DSGE models. The solution method seeks decision rules for c_t that are "rational" in that they satisfy the Euler equation (1.4) in a sufficiently robust way. It may be viewed intuitively as a computer analogue of the method of undetermined coefficients. The steps in the algorithm are as follows:

- Specify decision rules for the forward looking variables; for example, $\hat{c}_t = f(\Omega, x_t)$, where Ω are parameters, x_t are explanatory variables and f is an approximating function.

• Obtain the Euler error from the Euler equations

$$\epsilon_t = U'(\hat{c}_t) - \beta U'(\hat{c}_{t+1}) f'(k_{t+1}).$$

• Estimate Ω using various optimizing algorithm so that the Euler equation residuals, or the difference between the left- and right-hand sides of the Euler equation, is close to zero.

• Perform accuracy tests to check on the robustness of the results.

Approximating Functions For the example discussed here, the approximating function for consumption c_t, expressed as a function of the state variable known at time t, is

$$\hat{c}_t = \psi^c(\Omega^c, z_t, k_{t-1}). \tag{1.5}$$

The function ψ^c can be any approximating functions, and the decision variables are typically observations on the shocks and the state variable. In fact approximating functions are just flexible functional forms parameterized to minimize Euler equation errors that are well defined by a priori theoretical restrictions based on the optimizing behavior of the agents in the underlying the model.

Neural network (typically logistic) or the Chebychev orthogonal polynomial specifications are the two most common approximating functions used. The question facing the researcher here is one of robustness. First, given a relatively simple model, should one use a low-order Chebychev polynomial approximation or are there gains to using slightly higher order expansion for obtaining the decision rules for the forward-looking variable? Will the results change very much if we use a more complex Chebychev polynomial or a neural network alternative? Are there advantages to using a more complex approximating function, even if a less complex approximation does rather well? In other words, is the functional form of the decision rule robust with respect to the complexity of the model?

The question of using slightly more complex approximating functions, even when they may not be needed for simple models, illustrates a trade-off noted by Wolkenhauer (2001, p. ii): more complex approximations are often not specific or precise enough for a particular problem while simple approximations may not be general enough for more complex models. As a rule, the "discipline" of Occam's razor still applies: relatively simple and more transparent approximating functions are to be preferred over more complex and less transparent functions. Canova (2007) recommends starting with simple approximating functions such as a first- or second-order polynomial, and later checking the robustness of the solution with more complex functions.

In this book we use neural networks throughout. Sirakaya, Turnovsky, and Alemdar (2006) cite several reasons for using neural networks as approximating functions. First, as noted by Hornik, Stinchcombe, and White (1989), a sufficiently complex feedforward network can approximate any member of a class of functions to any degree of accuracy. Second, neural networks allow fewer parameters to be used to achieve the same degree of accuracy as orthogonal polynomials, which require an exponential increase in parameters. While the curse of dimensionality is still there, its "sting"—to borrow an expression from St. Paul, and expanded by Kenneth Judd[4]—is reduced. Third, neural networks, with logsigmoid functions, easily deliver control bounds on endogenous variables. Finally, such networks can be easily applied to models that admit bang-bang solutions [Sirakaya, Turnovsky, and Alemdar (2006): p. 3]. For all these reasons, neural networks can serve as a useful and readily available alternative or robustness check to the more commonly used Chebychev approximating functions.

While the outcomes of different approximating functions will not be identical since we cannot obtain closed form solutions for these models, we would like the results to be sufficiently robust, in terms of basic dynamic properties. In this book we also assess the performance of the function using accuracy tests. Before discussing these tests, we digress to present a brief overview of the neural network function.

Logistic Neural Networks Like orthogonal polynomial approximation methods, a logistic neural network relates a set of input variables to a set of one or more output variables, but the difference is that the neural network makes use of one or more hidden layers in which the input variables are squashed or transformed by a special function, known as a logistic or logsigmoid transformation. The following equations describe this form of approximation:

$$n_{j,t} = \omega_{j,0} + \sum_{i=1}^{i*} \omega_{j,i} x_{i,t}^{*}, \tag{1.6}$$

$$N_{j,t} = \frac{1}{1 + e^{-n_{j,t}}}, \tag{1.7}$$

$$y_t^* = \gamma_0 + \sum_{j=1}^{j^*} \gamma_j N_{j,t}. \tag{1.8}$$

Equation (1.6) describes a variable $n_{j,t}$ as a linear combination of a constant term $\omega_{j,0}$ and input variables observed at time t, $\{x_{i,t}\}$, $i = 1, \ldots, i^*$, with coefficient vector or set of "input weights" $\omega_{j,i}$, $i = 1, \ldots, i^*$. Equation (1.8) shows how this variable is squashed by the logistic function and becomes a neuron $N_{j,t}$ at time or observation t. The set of j^* neurons are then combined in a linear way with the coefficient vector $\{\gamma_j\}$, $j = 1, \ldots, j^*$, and taken with a constant term γ_0 to form the forecast $\hat{y}_t^*$ at time t.

This system is known as a feedforward network, and when coupled with the logsigmoid activation functions, it is also known as the multilayer perception (MLP) network. It is the basic workhorse of the neural network forecasting approach, in the sense that researchers usually start with this network as the first representative network alternative to the linear forecasting model. An important difference between neural network and orthogonal polynomial approximation is that the neural network approximation is not linear in parameters.

Optimizing Algorithm The parameters Ω_c are obtained by minimizing the squared residuals ϵ:[5]

$$\epsilon_t^c = U'(\hat{c}_t) - \beta U'(\hat{c}_{t+1}) f'(f(z_t, k_t) - \hat{c}_t). \tag{1.9}$$

To obtain the parameters, we use an algorithm similar to the parameterized expectations approach developed by Marcet (1988, 1992), and further developed in Den Haan and Marcet (1990, 1994) and in Marcet and Lorenzoni (1999). We solve for the parameters as a fixed-point problem. We make an initial guess of the parameter vector $[\Omega^c]$, draw a large sequence of shocks (e_t), and then generate time series for the endogenous variables of the model (c_t, k_t). We next iterate on the parameter set $[\Omega^c]$ to minimize a loss function $\mathscr{L}$ based on the Euler equation errors ϵ for a sufficiently large T.[6] We continue this process until convergence.

Note that the projection method does not require linearization, nor does it need the Blanchard-Khan algorithm. Instead, once expressions can be found for determining the forward-looking variables, the nonlinear model is solved for the other endogenous variables given the exogenously determined variables. A variety of optimization methods can be used to obtain the global optimum.[7] Fortunately optimization methods are becoming more effective for finding the global minima.

There are, however, drawbacks of this approach, as Canova (2005, p. 64) points out. One is that for more complex models, the iterations may take quite a bit of time for convergence. Fernández-Villaverde and Rubio-Ramírez (2006) also note that this is expensive in terms of computing time. We have found that with the right set of initial values the speed can be greatly reduced.

There is also the ever-present curse of dimensionality. The larger the number of state variables, the greater is the number of parameters needed to solve for the decision rules. There is no guarantee the Euler equation errors will diminish as the number of iterations grows when we deal with a very large number of parameters. The method relies on the sufficiency of the Euler equation errors. If the utility function is not strictly concave, for example, then the method may not give appropriate solutions. As Canova (2005) suggested, minimization of Euler equations may fail when there are large number of parameters or when there is a high degree of complexity or nonlinearity.

Heer and Maußner (2005) note another type of drawback of the approach. They point out that the Monte Carlo simulation will more likely generate data points near the steady-state values than far away from the steady state in the repeated simulations for estimating the parameter set $[\Omega_c]$ (Heer and Maußner 2005, p. 163). Fernández-Villaverde and Rubio-Ramírez (2006) have elaborated on this point. We want to weight the Euler equation errors by the percentage of time that the economy spends at those points. More to the point, we want to put more weight on the Euler equation errors where most of the action happens and less weight on the Euler equation errors that are not frequently realized. The problem, of course, is that we do not know the stationary distribution until we solve the model—that is, minimize the Euler equation errors.

That criticism is true, of course, if the innovations to the model represent small normally distributed disturbances around the steady-state equilibrium. If we simulate out for large sample, we are just staying close to the steady state. However, if we use, as Fernández-Villaverde (2005) suggests, either distributions with fat tails or with time-varying volatility, then the repeated simulations will be less likely to generate realizations concentrated near to the steady state. Similarly, if the process for the innovation distributions are realistic, based on well-accepted empirical results, then we are more than likely to stay in regions of the state space likely to be realized.

We have used normally distributed errors for most of this book, in order to show the effects of increasing model complexity and nonlinearity in the structural relations in the model. But we note that fat tails and volatility clustering are pervasive features of observed macroeconomic data, so there is no reason not to use wider classes of distributions for solving and simulating dynamic stochastic models. As Fernandez-Villaverde (2005) and Justiniano and Primiceri (2006) emphasize, there is no reason for a stochastic dynamic general equilibrium model not to have a richer structure than normal innovations. However, for the first-order perturbation approach, small normally distributed innovations are necessary. That is not the case for projection methods.

In summary, we work with one basic approach for solving models: the projection method, which is closely related to the Wright and Williams (1982, 1984, 1991) smoothing algorithm. We show that this method may be viewed as a computerized analogue of the method of undetermined coefficients commonly used to solve rational expectations models. With this method, as noted by Canova (2007), the approximation is globally valid as opposed to being valid only around a particular steady-state point as is the case for perturbation methods. The method is computationally more time-consuming than the perturbation method. But it has the advantage in that it is very useful for analyzing dynamics involving movements of key variables far away from their steady-state variables. And, of course, it allows us to incorporate asymmetries, threshold effects, and precautionary behavior. As Canova notes, the advantage of using this method is that the researcher or policy analyst can undertake experiments that are far away from the steady state, or involve more dramatic regime changes in the policy rule. Canova further notes two specific advantages of this approach: first, it can be used when inequality constraints are present, and second, it has a built-in mechanism to check if a candidate solution satisfies the optimality conditions of the model. These advantages are important when we take up open economy issues, such as constraints on foreign debt accumulation or the zero bound on nominal interest rates.

Another important reason for staying with the projection method is that it is a natural starting point for introducing learning on the part of the policy makers or on the part of the private decision makers in the model. Learning can be straightforwardly introduced and contrasted with the rational expectations when the setup comes from projection methods. Such learning represents stickiness in information in contrast to stickiness in price-setting behavior. As Orphanides and Williams (2002) put it, learning adds an additional layer of dynamic interactions between macroeconomic policies and economic outcomes.

Finally, Oveido (2005) argues, for us, convincingly, that the projection method is the appropriate approach to use for open economy models. The reason is that the net foreign asset position can deviate quite a bit from its steady-state value, since access to nearly frictionless world financial markets effectively separates saving from investment decisions. Since first- and second-order perturbation methods assume only small deviations of state variables from their steady-state variables, solutions based on these methods will overstate the volatility of macroeconomic aggregates.

Accuracy Tests To test the accuracy of stochastic simulation results, we have to work with the Euler equations. Since the model does not have any exact closed form solution against which we can benchmark numerical approximations, we have to use indirect measures of accuracy. Too often these accuracy checks are ignored when researchers present simulation results based on stochastic dynamic models. This is unfortunate, since the credibility of the results, even apart from matching key characteristics of observable data, rests on acceptable measures of computational accuracy as well as theoretical foundations. The accuracy tests used throughout the book are those due to Judd and Gaspar (1997) and to den Haan and Marcet (1994). They are based on the Euler equation errors.

Judd-Gaspar Statistic A natural way to start is to check to see if the Euler equations are satisfied, in the sense that the Euler equation errors are close to zero. Judd and Gaspar (1997) suggest transforming the Euler equation errors as follows:

$$JG_t^c = \frac{|\epsilon_t^c|}{C_t}; \tag{1.10}$$

that is they suggest checking the accuracy of the approximations by examining the absolute Euler equation errors relative to their respective forward looking variable. If the mean absolute values of the Euler equation errors, deflated by the forward-looking variable c_t, is 10^{-2}, Judd and Gaspar note that the Euler equation is accurate to within a penny per unit of consumption.

Den Haan-Marcet Statistic A drawback of the Judd and Gaspar criterion is that it is not based on any statistical distribution. It is purely a numerical method. At which point do the errors become statistically significant? For this reason we use another commonly used criterion, due to den Haan and Marcet (1994). This test is denoted $DM(m)$ and is defined as

$$DM(m) = T\mathbb{Q}'A^{-1}\mathbb{Q} \sim \chi^2(m), \tag{1.11}$$

$$\mathbb{Q} = \frac{1}{T}(\epsilon' x), \quad A = \frac{1}{T}\sum x_t x_t' \epsilon_t^2,$$

where the vector ϵ represents the vector of Euler equation errors, x is the instrument matrix with m columns. Under the null hypothesis of an accurate solution, $E(\epsilon' x) = 0$.

The authors recommend the following procedure for implementing this test: first, draw a sample of size T of den Haan and Marcet test of accuracy, with m degrees of freedom, repeatedly, say 500 times and calculate the DM statistics; second, compute the percentage of the DM statistics that is below the lower or above the upper 5 percent critical values of the $\chi^2(m)$ distribution. If these fractions are noticeably different from the expected 5 percent, then we have evidence for an inaccurate solution. They also recommend performing a "goodness-of-fit" type of test and to compare the empirical and theoretical cumulative density $\chi^2(m)$ function.

One of the goals of this book is to promote the reporting of accuracy statistics in computationally based research publications. We are no longer in the world of closed form solutions. However intuitively plausible the results of any research endeavor may be, it is important to know that they pass a minimum degree of computational accuracy.

1.3 Policy Goals, Welfare, and Scenarios

Whenever we discuss optimal policy, we have to specify the objectives of policy makers. Central banks, of course, have low inflation goals, and fiscal authorities may be concerned with fiscal sustainability. However, when we evaluate the overall performance of particular policy rules or stances of policy makers over the medium to long run, the overarching criterion for the performance of policy should be the welfare of households in the economy. By welfare, we mean an intertemporal index or measure of current and future consumption and leisure available to households.

Of course, policy is not made in a vacuum: the economy is subject to a variety of change, from external and internal sources, such as productivity, foreign interest rates, foreign demand, and terms of trade, all well beyond the control of any policy maker. So the measures of welfare, resulting from alternative rules for fiscal and monetary policy, also depend on factors beyond the scope of policy decisions. How can we evaluate the welfare consequences of specific policy rules when changes beyond the scope of policy are also taking place?

We make our case for computational approaches to policy evaluation precisely on this issue. With computational methods we can evaluate the distribution of welfare measures over a wide variety of realizations of shocks or exogenous changes affecting the economy, for different monetary and fiscal policy settings. We can specify a functional form for household utility and develop an intertemporal index, and compute this measure over a variety of policy settings. There is no need to substitute these direct welfare measures with quadratic loss functions or other ad hoc measures, since we are not linearizing the welfare function.

Moreover, whenever we discuss welfare, we present a histogram of welfare distributions. Given that any welfare index is based on realizations of one set of random shocks based on a given seed to a random number generator, it is important to know the dispersion of this welfare index for a wide set of realizations based on different seeds. We hope that this book will promote more widely the use of welfare distributions for assessing the payoff of different policy rules.

All chapters contain an alternative scenario or policy experiment, intended to motivate our readers to engage in computational experiments on their own. Many of the results come from one important difference between the open and closed economy setting. In the open economy consumers have access to international financial markets to smooth their consumption over time, when they face distortions in the domestic economy in the form of price or wage stickiness.

1.4 Plan of the Book

This book has eleven chapters. The goal of the computational experiments is to find robust conclusions regarding policy response to external and internal disturbances, under alternative assumptions about the structure of the economy and how agents react to new developments and policy change. We start with a very simple setting with no distortions or rigidities and gradually incorporate more distortions (e.g., in the form of price and wage stickiness, taxes, real rigidities in investment, financial frictions, and habit persistence in consumption).

Chapter 2 lays out the basic theoretical framework or model with fully flexible prices and with a simple Taylor rule for monetary policy. The model is closed by allowing for a debt elastic interest rate. We discuss how we calibrate the model and solve for the steady-state initial conditions of the model. Overall, we show that even this very simple framework involves forward-looking behavior and requires carefully constructed approximation methods for solution and simulation. Following the traditional literature, we show how the model can be solved for a given productivity shock with the projection method. We also present the results of the suggested accuracy checks. This chapter includes discussion about impulse-responses in response to a once-only shock as well as discussion of results from stochastic simulations resulting from recurring changes in productivity.

We believe that it is useful to consider simple flexible models because they are the benchmarks to evaluate welfare gains and loses of policy approaches under different types of rigidities and distortions. Consequently from the simulations we obtain benchmark welfare distributions under fully flexible prices for domestic and foreign goods, but bearing in mind that in these benchmark scenarios the monetary authority follows a simple Taylor rule aimed simply at inflation targets. The experiment conducted in this chapter is for the case of recurring changes in foreign demand. The results are compared with those obtained in response to changes in domestic productivity.

Chapter 3 takes up stickiness in domestic price setting. We examine how this form of stickiness reduces welfare, relative to the benchmark welfare distribution under fully flexible prices. We also explore more extensive Taylor rules responding not only to inflation targets but also to output gaps. The output gap is the difference between the actual level of output and the output which would occur in the flexible price economy. This chapter illustrates the effects of alternative policy targets.

The first few chapters were only concerned with monetary policy. In chapter 4 we analyze the welfare effects of alternative fiscal systems or tax bases, when there are recurring productivity shocks, for a given inflation-targeting monetary regime. We compare the case where the income tax rate is greater than the consumption tax rate with the reverse case where the income tax rate is less than the consumption tax rate.

The issue of domestic debt leads naturally to a consideration of the "twin" deficits in chapter 5. Here we let export demands react to the real exchange rate, and we explore the sensitivity of the relationship between the fiscal and current account deficits as the export elasticity of demand range from low to high for a productivity shock. Collectively, chapters 4 and 5 illustrate the sensitivity of results to alternative base case and alternative parameters.

Chapter 6 introduces capital accumulation into the basic models and considers the role of Tobin's Q in policy analysis. While the earlier chapters dealt with nominal stickiness associated with prices, this chapter is concerned with real rigidities and other types of distortions.

Chapter 7 expands the model to two sectors, which then allows us to broaden our scenario analysis to a consideration of a terms-of-trade shock. In particular, this chapter examines the case of productivity versus terms-of-trade shocks for an economy with a rich natural resource sector.

Chapter 8 introduces financial frictions by allowing for banking and financial frictions. This type of model is also called a limited participation model, since households are now restricted on the types of assets they can hold. In this chapter we compare the case of inflation targeting with a flexible exchange rate with the case of no inflation targeting with an effectively fixed exchange rate (which is akin to imported goods inflation targeting).

Chapter 9 is concerned with wage rigidities as a source of stickiness. Scenarios are simulated to explore how labor–leisure choices affect the outcomes of the productivity shock.

Chapter 10 introduces habit persistence into the consumption decision and considers the simulated results for two sets of comparisons: inflation targeting and no-inflation targeting, and productivity and terms-of-trade shocks.

The final chapter, chapter 11, makes use of the model with all of the bells and whistles and simulates a sudden stop as well as a large continuing capital inflow (and increasing external deficit) for an economy. Sudden stops have plagued emerging market economies in the last two decades, while the United States has experienced large and continuing external debt accumulation. This final chapter brings into sharp focus the advantages of using our nonlinear approximation algorithm for solving and simulating open economy stochastic dynamic models with sudden large shocks or increasing external debt levels. The aim of this chapter is to highlight, once again, the insights that can be obtained from simulating (nonlinear) DSGE models.

Of course, the order in which we have progressed, with increasing complexity—from the flexible price model, to sticky prices, to distortionary taxes, to capital accumulation, to sectoral production, to financial frictions, to sticky wages, to habit persistence—is a matter of taste. We are not suggesting that there is any deep evolutionary pattern in the ordering we have chosen, just that it follows roughly the development of the literature in open economy business-cycle analysis. Also as a final comment, we note that while we cover a range of topics familiar to students of open economy macroeconomics, this book is about methods for policy evaluation and not about policy evaluation itself.

Computational Exercises

At the end of chapters 2 through 10, we have added computational exercises. The MATLAB codes for the base flexible price model discussed in chapter 2 appears in the appendix at the end of the book.[8] This program estimates the decision rule coefficients as well as generates the impulse-response paths and the stochastic simulations for the model presented in chapter 2. As we move from chapter to chapter, the reader is invited to modify the codes from the base flexible price model to more complex extensions. Quite apart from programming to suit one's personal style and taste, we believe that the act of programming is an integral part of open economy macro research as it enhances the comprehension of the models and the simulated results.