1
Introduction

Models are the venue for expressing, comparing, and evaluating alternative ways of addressing important questions in economics. Applied econometricians are called upon to engage in these exercises using data and, often, formal methods whose properties are understood in decision-making contexts. This is true of work in other sciences as well.

There is a large literature on alternative formal approaches to these tasks, including both Bayesian and non-Bayesian methods. Formal approaches tend to take models as given, and the more formal the approach the more likely this is to be true. Whether the topic is inference, estimation, hypothesis testing, or specification testing, the formal treatment begins with a specific model. The same is also true of formal approaches to the tasks of prediction and policy evaluation.

Yet the ultimate success of these endeavors depends strongly on creativity, insight, and skill in the process of model creation. As the model builder entertains new ideas, casting off those that are not deemed promising and developing further those that are, he or she is engaged in a sophisticated process of learning. This process does not, typically, involve the specification of a great many models developed to the point of departure assumed in formal treatments in graduate courses, texts in econometrics and statistics, and journal articles. Discarding models that would ultimately be unsuccessful earlier rather than later in this process of learning improves the allocation of research time and talent.

Model development is inherently a task of learning under conditions of unstructured uncertainty. To assume that one's model fully accounts for the phenomenon under question is naive. A more defensible position is that of Box (1980): all models are wrong, but some are useful. To this it might be added that, with inspiration and perspiration, models can be improved. The process of information acquisition, learning, and behavior when objectives are well-specified in a utility or loss function is familiar ground in economics. In modeling the optimal behavior of economic agents in this situation the dominant paradigm is Bayesian learning, to the point that many in the profession are comfortable terming such behavior rational.

Can this paradigm be applied to model development? A number of obstacles suggest that the task may be demanding. First, the behavior of practicing econometricians regularly appears inconsistent with the Bayesian learning paradigm. In particular, the dominant statistical paradigm in econometrics has been frequentist and the inconsistencies of this approach with Bayesian inference and learning are well-known. Second, Bayesian model specification is more demanding than most non-Bayesian model specification, requiring prior distributions for inherently unobservable constructs like parameters, as well as for models themselves when multiple models are under consideration. Finally, whereas in academic treatments of Bayesian learning reality is fully specified, in applying economics to policy questions it is not even clear that the existence of a data-generating process has any epistemological standing at all.

The thesis of this monograph is that these objections can be met, and its essays are explorations of the prospects for more effective use of the Bayesian paradigm at the point where the investigator has much less information than is presumed in formal econometric approaches, be they Bayesian or non-Bayesian. At this point models are inherently incomplete: that is, they are lacking some aspect of a joint distribution over all parameters, latent variables, and models under consideration. Chapter 2 details more fully the concept of a complete model. It also establishes notation and serves as a primer on Bayesian econometrics.

Model incompleteness can take many forms, and the essays in this monograph take up three examples. Chapter 3 addresses the early steps of model construction—before the investigator has engaged the technical demands of formal inference or estimation and perhaps even before data have been collected. The emphasis in this chapter is on using formal Bayesian methods to compare and evaluate models. Model *comparison* at this stage amounts to prior predictive analysis, which was introduced to statistics by Box (1980) and emphasized in econometrics by Lancaster (2004) and Geweke (2005). These ideas are not new, but their potential for greatly accelerating effective research is not as yet well appreciated in the econometrics community. Model *evaluation* is the assessment of a specified model by absolute standards—a process in which economists regularly engage. The assertion, or conclusion, that a model is bad for a particular purpose is repeatedly heard in the economist's workday. But, as economists regularly point out, statements like this raise the question, bad compared with what? The final section of chapter 3 sets up an incomplete model as the basis of comparison implicit in such statements, and then extends the conventional apparatus of Bayesian model comparison to the complete model being evaluated and the incomplete model that is implicitly held as the standard. This treatment provides a fully Bayesian counterpoint to frequentist tests against an unspecified alternative, also known as pure significance testing.

No model is meant to specify all aspects of reality, even a sharply confined reality chosen for its relevance to a particular academic or policy question. This restriction can be straightforward for formal econometrics: for example, the stipulation that a regression model applies only over a specified range of values of the covariates, or that a structural model with several endogenous variables is intended only to provide the marginal distribution of a subset of these variables. But often the restriction is stronger. Chapter 4 takes up the case of structural models that are intended only to provide certain population moments as functions of the structural parameters of the model, a restriction that is especially common in dynamic stochastic general equilibrium models. The chapter shows that widely applied procedures, including conventional calibration, violate this restriction by taking the higher-order moments of the model literally in reaching conclusions about its structural parameters. It provides a constructive approach to this learning problem by treating explicitly the incompleteness of the structural model and then completing the model in a way that relies only on those aspects of the structural model intended to be taken literally. In the example used throughout the chapter this approach reverses widely held conclusions about the incompatibility of the U.S. equity premium with simple growth models.

Formal Bayesian methods provide a logically consistent and well-understood solution to the problem of using competing models with conflicting implications in a decision-making context. The critical element of this solution is the specification of prior model probabilities that sum to one. In so doing, the solution conditions on the process that actually generated the data being one of the models under consideration. Non-Bayesian methods that lead to rules for model choice also make the latter assumption. A widely observed characteristic of the formal Bayesian approach is that it often assigns posterior probability very close to unity for one of the models. The Bayesian solution then effectively amounts to model choice. This is not a problem for econometric theory, because in general the data-generating process is the one selected asymptotically. On the other hand, there is an evident conflict with reality: in reaching important decisions policymakers routinely wrestle with alternative models, leading to an apparent inconsistency of clear evidence with presumed rationality of the decision makers. The final chapter in this monograph steps back from the key assumption that reality lies somewhere in the space of models being considered. Replacing the assumption that the model space is completely specified with the alternative of a linear combination of predictive densities of future events that renders past events most probable, it shows that if the data-generating process is not among the group of models considered, then one will use several models. The weights given to these models will converge to positive limits asymptotically. The weights assigned in Bayesian model averaging are incorrect, under this alternative specification in which the model space is incomplete.