

COPYRIGHT NOTICE:

Kenneth J. Singleton: Empirical Dynamic Asset Pricing

is published by Princeton University Press and copyrighted, © 2006, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

Follow links for Class Use and other Permissions. For more information send email to: permissions@pupress.princeton.edu

1

Introduction

A DYNAMIC ASSET pricing model is refutable empirically if it restricts the joint distribution of the observable asset prices or returns under study. A wide variety of economic and statistical assumptions have been imposed to arrive at such testable restrictions, depending in part on the objectives and scope of a modeler's analysis. For instance, if the goal is to price a given cash-flow stream based on agents' optimal consumption and investment decisions, then a modeler typically needs a fully articulated specification of agents' preferences, the available production technologies, and the constraints under which agents optimize. On the other hand, if a modeler is concerned with the derivation of prices as discounted cash flows, subject only to the constraint that there be no "arbitrage" opportunities in the economy, then it may be sufficient to specify how the relevant discount factors depend on the underlying risk factors affecting security prices, along with the joint distribution of these factors.

An alternative, typically less ambitious, modeling objective is that of testing the restrictions implied by a particular "equilibrium" condition arising out of an agent's consumption/investment decision. Such tests can often proceed by specifying only portions of an agent's intertemporal portfolio problem and examining the implied restrictions on moments of subsets of variables in the model. With this narrower scope often comes some "robustness" to potential misspecification of components of the overall economy that are not directly of interest.

Yet a third case is one in which we do not have a well-developed theory for the joint distribution of prices and other variables and are simply attempting to learn about features of their joint behavior. This case arises, for example, when one finds evidence against a theory, is not sure about how to formulate a better-fitting, alternative theory, and, hence, is seeking a better understanding of the historical relations among key economic variables as guidance for future model construction.

As a practical matter, differences in model formulation and the decision to focus on a “preference-based” or “arbitrage-free” pricing model may also be influenced by the availability of data. A convenient feature of financial data is that it is sampled frequently, often daily and increasingly intraday as well. On the other hand, macroeconomic time series and other variables that may be viewed as determinants of asset prices may only be reported monthly or quarterly. For the purpose of studying the relation between asset prices and macroeconomic series, it is therefore necessary to formulate models and adopt econometric methods that accommodate these data limitations. In contrast, those attempting to understand the day-to-day movements in asset prices—traders or risk managers at financial institutions, for example—may wish to design models and select econometric methods that can be implemented with daily or intraday financial data alone.

Another important way in which data availability and model specification often interact is in the selection of the decision interval of economic agents. Though available data are sampled at discrete intervals of time—daily, weekly, and so on—it need not be the case that economic agents make their decisions at the same sampling frequency. Yet it is not uncommon for the available data, including their sampling frequency, to dictate a modeler’s assumption about the decision interval of the economic agents in the model. Almost exclusively, two cases are considered: *discrete-time* models typically match the sampling and decision intervals—monthly sampled data mean monthly decision intervals, and so on—whereas *continuous-time* models assume that agents make decisions continuously in time and then implications are derived for discretely sampled data. There is often no sound economic justification for either the coincidence of timing in discrete-time models, or the convenience of continuous decision making in continuous-time models. As we will see, analytic tractability is often a driving force behind these timing assumptions.

Both of these considerations (the degree to which a complete economic environment is specified and data limitations), as well as the computational complexity of solving and estimating a model, may affect the choice of estimation strategy and, hence, the econometric properties of the estimator of a dynamic pricing model. When a model provides a full characterization of the joint distribution of its variables, a historical sample is available, and fully exploiting this information in estimation is computationally feasible, then the resulting estimators are “fully efficient” in the sense of exploiting all of the model-implied restrictions on the joint distribution of asset prices. On the other hand, when any one of these conditions is not met, researchers typically resort, by choice or necessity, to making compromises on the degree of model complexity (the richness of the economic environment) or the computational complexity of the estimation strategy (which often means less econometric efficiency in estimation).

With these differences in modelers' objectives, practical constraints on model implementation, and computational considerations in mind, this book: (1) characterizes the nature of the restrictions on the joint distributions of asset returns and other economic variables implied by dynamic asset pricing models (DAPMs); (2) discusses the interplay between model formulation and the choice of econometric estimation strategy and analyzes the large-sample properties of the feasible estimators; and (3) summarizes the existing, and presents some new, empirical evidence on the fit of various DAPMs.

We briefly expand on the interplay between model formulation and econometric analysis to set the stage for the remainder of the book.

1.1. Model Implied Restrictions

Let \mathcal{P}_s denote the set of “payoffs” at date s that are to be priced at date t , for $s > t$, by an economic model (e.g., next period's cum-dividend stock price, cash flows on bonds, and so on),¹ and let $\pi_t : \mathcal{P}_s \rightarrow \mathbb{R}$ denote the pricing function, where \mathbb{R}^n denotes the n -dimensional Euclidean space. Most DAPMs maintain the assumption of no arbitrage opportunities on the set of securities being studied: for any $q_{t+1} \in \mathcal{P}_{t+1}$ for which $\Pr\{q_{t+1} \geq 0\} = 1$, $\Pr\{\pi_t(q_{t+1}) \leq 0\} \cap \{q_{t+1} > 0\} = 0$.² In other words, nonnegative payoffs at $t + 1$ that are positive with positive probability have positive prices at date t . A key insight underlying the construction of DAPMs is that the absence of arbitrage opportunities on a set of payoffs \mathcal{P}_s is essentially equivalent to the existence of a special payoff, a *pricing kernel* q_s^* , that is strictly positive ($\Pr\{q_s^* > 0\} = 1$) and represents the pricing function π_t as

$$\pi_t(q_s) = E[q_s q_s^* | \mathcal{I}_t], \quad (1.1)$$

for all $q_s \in \mathcal{P}_s$, where \mathcal{I}_t denotes the information set upon which expectations are conditioned in computing prices.³

¹ At this introductory level we remain vague about the precise characteristics of the payoffs investors trade. See Harrison and Kreps (1979), Hansen and Richard (1987), and subsequent chapters herein for formal definitions of payoff spaces.

² We let $\Pr\{\cdot\}$ denote the probability of the event in brackets.

³ The existence of a pricing kernel q^* that prices all payoffs according to (1.1) is equivalent to the assumption of no arbitrage opportunities when uncertainty is generated by discrete random variables (see, e.g., Duffie, 2001). More generally, when \mathcal{I}_t is generated by continuous random variables, additional structure must be imposed on the payoff space and pricing function π_t for this equivalence (e.g., Harrison and Kreps, 1979, and Hansen and Richard, 1987). For now, we focus on the pricing relation (1.1), treating it as being equivalent to the absence of arbitrage. A more formal development of pricing kernels and the properties of q^* is taken up in Chapter 8 using the framework set forth in Hansen and Richard (1987).

This result by itself does not imply testable restrictions on the prices of payoffs in \mathcal{P}_{t+1} , since the theorem does not lead directly to an empirically observable counterpart to the benchmark payoff. Rather, overidentifying restrictions are obtained by restricting the functional form of the pricing kernel q_s^* or the joint distribution of the elements of the *pricing environment* $(\mathcal{P}_s, q_s^*, \mathcal{I}_t)$. It is natural, therefore, to classify DAPMs according to the types of restrictions they impose on the distributions of the elements of $(\mathcal{P}_s, q_s^*, \mathcal{I}_t)$. We organize our discussions of models and the associated estimation strategies under four headings: preference-based DAPMs, arbitrage-free pricing models, “beta” representations of excess portfolio returns, and linear asset pricing relations. This classification of DAPMs is not mutually exclusive. Therefore, the organization of our subsequent discussions of specific models is also influenced in part by the choice of econometric methods typically used to study these models.

1.1.1. Preference-Based DAPMs

The approach to pricing that is most closely linked to an investor’s portfolio problem is that of the preference-based models that directly parameterize an agent’s intertemporal consumption and investment decision problem. Specifically, suppose that the economy being studied is comprised of a finite number of infinitely lived agents who have identical endowments, information, and preferences in an uncertain environment. Moreover, suppose that \mathcal{A}_t represents the agents’ information set and that the representative consumer ranks consumption sequences using a von Neumann-Morgenstern utility functional

$$E \left[\sum_{t=0}^{\infty} \beta^t U(c_t) \mid \mathcal{A}_0 \right]. \quad (1.2)$$

In (1.2), preferences are assumed to be time separable with period utility function U and the subjective discount factor $\beta \in (0, 1)$. If the representative agent can trade the assets with payoffs \mathcal{P}_s and their asset holdings are interior to the set of admissible portfolios, the prices of these payoffs in equilibrium are given by (Rubinstein, 1976; Lucas, 1978; Breeden, 1979)

$$\pi_t(q_s) = E[m_s^{s-t} q_s \mid \mathcal{A}_t], \quad (1.3)$$

where $m_s^{s-t} = \beta U'(c_s)/U'(c_t)$ is the intertemporal marginal rate of substitution of consumption (MRS) between dates t and s . For a given parameterization of the utility function $U(c_t)$, a preference-based DAPM allows the association of the pricing kernel q_s^* with m_s^{s-t} .

To compute the prices $\pi_t(q_s)$ requires a parametric assumption about the agent's utility function $U(c_t)$ and sufficient economic structure to determine the joint, conditional distribution of m_s^{s-t} and q_s . Given that prices are set as part of the determination of an equilibrium in goods and securities markets, a modeler interested in pricing must specify a variety of features of an economy outside of securities markets in order to undertake preference-based pricing. Furthermore, limitations on available data may be such that some of the theoretical constructs appearing in utility functions or budget constraints do not have readily available, observable counterparts. Indeed, data on individual consumption levels are not generally available, and aggregate consumption data are available only for certain categories of goods and, at best, only at a monthly sampling frequency.

For these reasons, studies of preference-based models have often focused on the more modest goal of attempting to evaluate whether, for a particular choice of utility function $U(c_t)$, (1.3) does in fact “price” the payoffs in \mathcal{P}_s . Given observations on a candidate m_s^{s-t} and data on asset returns $\mathcal{R}_s \equiv \{q_s \in \mathcal{P}_s : \pi_t(q_s) = 1\}$, (1.3) implies testable restrictions on the joint distribution of \mathcal{R}_s , m_t^{s-t} , and elements of \mathcal{A}_t . Namely, for each s -period return r_s , $E[m_s^{s-t}r_s - 1|\mathcal{A}_t] = 0$, for any $r_s \in \mathcal{R}_s$ (see, e.g., Hansen and Singleton, 1982). An immediate implication of this moment restriction is that $E[(m_s^{s-t}r_s - 1)x_t] = 0$, for any $x_t \in \mathcal{A}_t$.⁴ These unconditional moment restrictions can be used to construct method-of-moments estimators of the parameters governing m_s^{s-t} and to test whether or not m_s^{s-t} prices the securities with payoffs in \mathcal{P}_s . This illustrates the use of restrictions on the moments of certain functions of the observed data for estimation and inference, when complete knowledge of the joint distribution of these variables is not available.

An important feature of preference-based models of frictionless markets is that, assuming agents optimize and rationally use their available information \mathcal{A}_t in computing the expectation (1.3), there will be *no arbitrage opportunities* in equilibrium. That is, the absence of arbitrage opportunities is a consequence of the equilibrium price-setting process.

1.1.2. Arbitrage-Free Pricing Models

An alternative approach to pricing starts with the presumption of no arbitrage opportunities (i.e., this is not derived from equilibrium behavior). Using the principle of “no arbitrage” to develop pricing relations dates back at least to the key insights of Black and Scholes (1973), Merton (1973), Ross

⁴ This is an implication of the “law of iterated expectations,” which states that $E[y_s] = E[E(y_s|\mathcal{A}_t)]$, for any conditioning information set \mathcal{A}_t .

(1978), and Harrison and Kreps (1979). Central to this approach is the observation that, under weak regularity conditions, pricing can proceed “as if” agents are risk neutral. When time is measured continuously and agents can trade a default-free bond that matures an “instant” in the future and pays the (continuously compounded) rate of return r_t , discounting for risk-neutral pricing is done by the default-free “roll-over” return $e^{-\int_t^s r_u du}$. For example, if uncertainty about future prices and yields is generated by a continuous-time Markov process Y_t (so, in particular, the conditioning information set \mathcal{I}_t is generated by Y_t), then the price of the payoff q_s is given equivalently by

$$\pi_t(q_s) = E[q_s^* | Y_t] = E^{\mathbb{Q}}[e^{-\int_t^s r_u du} q_s | Y_t], \quad (1.4)$$

where $E_t^{\mathbb{Q}}$ denotes expectation with regard to the “risk-neutral” conditional distribution of Y . The term risk-neutral is applied because prices in (1.4) are expressed as the expected value of the payoff q_s *as if agents are neutral toward financial risks*.

As we will see more formally in subsequent chapters, the risk attitudes of investors are implicit in the exogenous specification of the pricing kernel q^* as a function of the state Y_t and, hence, in the change of probability measure underlying the risk-neutral representation (1.4). Leaving preferences and technology in the “background” and proceeding to parameterize the distribution of q^* directly facilitates the computation of security prices. The parameterization of $(\mathcal{P}_s, q_s^*, Y_t)$ is chosen so that the expectation in (1.4) can be solved, either analytically or through tractable numerical methods, for $\pi_t(q_s)$ as a function of Y_t : $\pi_t(q_s) = P(Y_t)$. This is facilitated by the adoption of continuous time (continuous trading), special structure on the conditional distribution of Y , and constraints on the dependence of q^* on Y so that the second expectation in (1.4) is easily computed. However, similarly tractable models are increasingly being developed for economies specified in discrete time and with discrete decision/trading intervals.

Importantly, though knowledge of the risk-neutral distribution of Y_t is sufficient for pricing through (1.4), this knowledge is typically not sufficient for econometric estimation. For the purpose of estimation using historical price or return information associated with the payoffs \mathcal{P}_s , we also need information about the distribution of Y under its *data-generating* or *actual* measure. What lie between the actual and risk-neutral distributions of Y are adjustments for the “market prices of risk”—terms that capture agents’ attitudes toward risk. It follows that, throughout this book, when discussing arbitrage-free pricing models, we typically find it necessary to specify the distributions of the state variables or risk factors under both measures.

If the conditional distribution of Y_t given Y_{t-1} is known (i.e., derivable from knowledge of the continuous-time specification of Y), then so typically is the conditional distribution of the observed market prices $\pi_t(q_s)$. The

completeness of the specification of the pricing relations (both the distribution of Y and the functional form of P_s) in this case implies that one can in principle use “fully efficient” maximum likelihood methods to estimate the unknown parameters of interest, say θ_0 . Moreover, this is feasible using market price data alone, even though the risk factors Y may be latent (unobserved) variables. This is a major strength of this modeling approach since, in terms of data requirements, one is constrained only by the availability of financial market data.

Key to this strategy for pricing is the presumption that the burden of computing $\pi_t(q_s) = P_s(Y_t)$ is low. For many specifications of the distribution of the state Y_t , the pricing relation $P_s(Y_t)$ must be determined by numerical methods. In this case, the computational burden of solving for P_s while simultaneously estimating θ_0 can be formidable, especially as the dimension of Y gets large. Have these considerations steered modelers to simpler data-generating processes (DGPs) for Y_t than they might otherwise have studied? Surely the answer is yes and one might reasonably be concerned that such compromises in the interest of computational tractability have introduced model misspecification.

We will see that, fortunately, in many cases there are alternative estimation strategies for studying arbitrage-free pricing relations that lessen the need for such compromises. In particular, one can often compute the moments of prices or returns implied by a pricing model, even though the model-implied likelihood function is unknown. In such cases, method-of-moments estimation is feasible. Early implementations of method-of-moments estimators typically sacrificed some econometric efficiency compared to the maximum likelihood estimator in order to achieve substantial computational simplification. More recently, however, various approximate maximum likelihood estimators have been developed that involve little or no loss in econometric efficiency, while preserving computational tractability.

1.1.3. Beta Representations of Excess Returns

One of the most celebrated and widely applied asset pricing models is the static capital-asset pricing model (CAPM), which expresses expected excess returns in terms of a security’s beta with a benchmark portfolio (Sharpe, 1964; Mossin, 1968). The traditional CAPM is static in the sense that agents are assumed to solve one-period optimization problems instead of multi-period utility maximization problems. Additionally, the CAPM beta pricing relation holds only under special assumptions about either the distributions of asset returns or agents’ preferences.

Nevertheless, the key insights of the CAPM carry over to richer stochastic environments in which agents optimize over multiple periods. There is

an analogous “single-beta” representation of expected returns based on the representation (1.1) of prices in terms of a pricing kernel q^* , what we refer to as an *intertemporal* CAPM or ICAPM.⁵ Specifically, setting $s = t + 1$, the benchmark return $r_{t+1}^* = q_{t+1}^*/\pi_t(q_{t+1}^*)$ satisfies⁶

$$E[r_{t+1}^*(r_{t+1} - r_{t+1}^*) | \mathcal{I}_t] = 0, \quad r_{t+1} \in \mathcal{R}_{t+1}. \quad (1.5)$$

Equation (1.5) has several important implications for the role of r_{t+1}^* in asset return relations, one of which is that r_{t+1}^* is a benchmark return for a single-beta representation of excess returns (see Chapter 11):

$$E[r_{j,t+1} | \mathcal{I}_t] - r_t^f = \beta_{jt} \left(E[r_{t+1}^* | \mathcal{I}_t] - r_t^f \right), \quad (1.6)$$

where

$$\beta_{jt} = \frac{\text{Cov}[r_{j,t+1}, r_{t+1}^* | \mathcal{I}_t]}{\text{Var}[r_{t+1}^* | \mathcal{I}_t]}, \quad (1.7)$$

and r_t^f is the interest rate on one-period riskless loans issued at date t . In words, the excess return on a security is proportional to the excess return on the benchmark portfolio, $E[r_{t+1}^* - r_t^f | \mathcal{I}_t]$, with factor of proportionality β_{jt} , for all securities j with returns in \mathcal{R}_{t+1} .

It turns out that the beta representation (1.6), together with the representation of r^f in terms of q_{t+1}^* ,⁷ constitute exactly the same information as the basic pricing relation (1.1). Given one, we can derive the other, and vice versa. At first glance, this may seem surprising given that econometric tests of beta representations of asset returns are often not linked to pricing kernels. The reason for this is that most econometric tests of expressions like (1.6) are in fact *not* tests of the joint restriction that $r_t^f = 1/E[q_{t+1}^* | \mathcal{I}_t]$ and r_{t+1}^* satisfies (1.6). Rather tests of the ICAPM are tests of whether a proposed candidate benchmark return r_{t+1}^β satisfies (1.6) alone, for a given information set \mathcal{I}_t . There are an infinite number of returns r_t^β that satisfy (1.6) (see Chapter 11). The return r_{t+1}^* , on the other hand, is the unique

⁵ By defining a benchmark return that is explicitly linked to the marginal rate of substitution, Breeden (1979) has shown how to obtain a single-beta representation of security returns that holds in continuous time. The following discussion is based on the analysis in Hansen and Richard (1987).

⁶ Hansen and Richard (1987) show that when the pricing function π_t is nontrivial, $\Pr\{\pi_t(q_{t+1}^*) = 0\} = 0$, so that r_{t+1}^* is a well-defined return. Substituting r^* into (1.1) gives $E[r_{t+1}^* r_{t+1} | \mathcal{I}_t] = \{E[q_{t+1}^{*2} | \mathcal{I}_t]\}^{-1}$, for all $r_{t+1} \in \mathcal{R}_{t+1}$. Since r_{t+1}^* is one such return, (1.5) follows.

⁷ The interest rate r_t^f can be expressed as $1/E[q_{t+1}^* | \mathcal{I}_t]$ by substituting the payoff $q_{t+1} = 1$ into (1.1) with $s = t + 1$.

return (within a set that is formally defined) satisfying (1.5). Thus, tests of single-beta ICAPMs are in fact tests of weaker restrictions on return distributions than tests of the pricing relation (1.1).

Focusing on a candidate benchmark return r_{t+1}^β and relation (1.6) (with r_{t+1}^β in place of r_{t+1}^*), once again the choices made regarding estimation and testing strategies typically involve trade-offs between the assumptions about return distributions and the robustness of the empirical analysis. Taken by itself, (1.6) is a restriction on the conditional first and second moments of returns. If one specifies a parametric family for the joint conditional distribution of the returns $r_{j,t+1}$ and r_{t+1}^β and the state Y_t , then estimation can proceed imposing the restriction (1.6). However, such tests may be compromised by misspecification of the higher moments of returns, even if the first two moments are correctly specified. There are alternative estimation strategies that exploit less information about the conditional distribution of returns and, in particular, that are based on the first two conditional moments for a given information set \mathcal{I}_t , of returns.

1.1.4. Linear Pricing Relations

Historically, much of the econometric analysis of DAPMs has focused on linear pricing relations. One important example of a linear DAPM is the version of the ICAPM obtained by assuming that β_{j_t} in (1.6) is constant (not state dependent), say β_j . Under this additional assumption, β_j is the familiar “beta” of the j th common stock from the *CAPM*, extended to allow both expected returns on stocks and the riskless interest rate to change over time. The mean of

$$u_{j,t+1} \equiv (r_{j,t+1} - r_t^f) - \beta_j (r_{t+1}^\beta - r_t^f) \quad (1.8)$$

conditioned on \mathcal{I}_t is zero for all admissible r_j . Therefore, the expression in (1.8) is uncorrelated with any variable in the information set \mathcal{I}_t ; $E[u_{j,t+1} x_t] = 0$, $x_t \in \mathcal{I}_t$. Estimators of the β_j and tests of (1.6) can be constructed based on these moment restrictions.

This example illustrates how additional assumptions about one feature of a model can make an analysis more robust to misspecification of other features. In this case, the assumption that β_j is constant permits estimation of β_j and testing of the null hypothesis (1.6) without having to fully specify the information set \mathcal{I}_t or the functional form of the conditional means of $r_{j,t+1}$ and r_{t+1}^β . All that is necessary is that the candidate elements x_t of \mathcal{I}_t used to construct moment restrictions are indeed in \mathcal{I}_t .⁸

⁸ We will see that this simplification does not obtain when the β_{j_t} are state dependent. Indeed, in the latter case, we might not even have readily identifiable benchmark returns r_{t+1}^β .

Another widely studied linear pricing relation was derived under the presumption that in a well-functioning—some say *informationally efficient*—market, holding-period returns on assets must be unpredictable (see, e.g., Fama, 1970). It is now well understood that, in fact, the optimal processing of information by market participants is not sufficient to ensure unpredictable returns. Rather, we should expect returns to evidence some predictability, either because agents are risk averse or as a result of the presence of a wide variety of market frictions.

Absent market frictions, then, one sufficient condition for returns to be unpredictable is that agents are risk neutral in the sense of having linear utility functions, $U(c_t) = u_0 + u_c c_t$. Then the MRS is $m_s^{s-t} = \beta^s$, where β is the subjective discount factor, and it follows immediately from (1.3) that

$$E[r_s | \mathcal{I}_t] = 1/\beta^s, \quad (1.9)$$

for an admissible return r_s . This, in turn, implies that r_s is unpredictable in the sense of having a constant conditional mean. The restrictions on returns implied by (1.9) are, in principle, easily tested under only minimal additional auxiliary assumptions about the distributions of returns. One simply checks to see whether $r_s - 1/\beta^s$ is uncorrelated with variables dated t or earlier that might be useful for forecasting future returns. However, as we discuss in depth in Chapter 9, there is an enormous literature examining this hypothesis. In spite of the simplicity of the restriction (1.9), whether or not it is true in financial markets remains an often debated question.

1.2. Econometric Estimation Strategies

While the specification of a DAPM logically precedes the selection of an estimation strategy for an empirical analysis, we begin Part I with an overview of econometric methods for analyzing DAPMs. Applications of these methods are then taken up in the context of the discussions of specific DAPMs. To set the stage for Part I, we start by viewing the model construction stage as leading to a family of models or pricing relations describing features of the distribution of an observed vector of variables z_t . This vector may include asset prices or returns, possibly other economic variables, as well as lagged values of these variables. Each model is indexed by a K -dimensional vector of parameters θ in an admissible parameter space $\Phi \in \mathbb{R}^K$. We introduce Φ

For instance, if \mathcal{I}_t is taken to be agents' information set \mathcal{A}_t , then the contents of \mathcal{I}_t may not be known to the econometrician. In this case the set of returns that satisfy (1.6) may also be unknown. It is of interest to ask then whether or not there are similar risk-return relations with moments conditioned on an observable subset of \mathcal{A}_t , say \mathcal{I}_t , for which benchmark returns satisfying an analogue to (1.6) are observable. This is among the questions addressed in Chapter 11.

because, for each of the DAPMs indexed by θ to be well defined, it may be necessary to constrain certain parameters to be larger than some minimum value (e.g., variances or risk aversion parameters), or DAPMs may imply that certain parameters are functionally related. The basic premise of an econometric analysis of a DAPM is that there is a unique $\theta_0 \in \Phi$ (a unique pricing relation) consistent with the population distribution of z . A primary objective of the econometric analysis is to construct an estimator of θ_0 .

More precisely, we view the selection of an *estimation strategy* for θ_0 as the choice of:

- A sample of size T on a vector z_t of observed variables, $\bar{z}_T \equiv (z_T, z_{T-1}, \dots, z_1)'$.
- An *admissible* parameter space $\Phi \subseteq \mathbb{R}^K$ that includes θ_0 .
- A K -vector of functions $\mathcal{D}(z_t; \theta)$ with the property that θ_0 is the unique element of Φ satisfying

$$E[\mathcal{D}(z_t; \theta_0)] = 0. \quad (1.10)$$

What ties an estimation strategy to the particular DAPM of interest is the requirement that θ_0 be the unique element of Φ that satisfies (1.10) for the chosen function \mathcal{D} . Thus, we view (1.10) as summarizing the implications of the DAPM that are being used directly in estimation. Note that, while the estimation strategy is premised on the economic theory of interest implying that (1.10) is satisfied, there is no presumption that this theory implies a unique \mathcal{D} that has mean zero at θ_0 . In fact, usually, there is an uncountable infinity of admissible choices of \mathcal{D} .

For many of the estimation strategies considered, \mathcal{D} can be reinterpreted as the first-order condition for maximizing a nonstochastic *population estimation objective* or *criterion* function $Q_0(\theta) : \Phi \rightarrow \mathbb{R}$. That is, at θ_0 ,

$$\frac{\partial Q_0}{\partial \theta}(\theta_0) = E[\mathcal{D}(z_t; \theta_0)] = 0. \quad (1.11)$$

Thus, we often view a choice of estimation strategy as a choice of criterion function Q_0 . For well-behaved Q_0 , there is always a θ^* that is the global maximum (or minimum, depending on the estimation strategy) of the criterion function Q_0 . Therefore, for Q_0 to be a sensible choice for the model at hand we require that θ^* be unique and equal to the population parameter vector of interest, θ_0 . A necessary step in verifying that $\theta^* = \theta_0$ is verifying that \mathcal{D} satisfies (1.10) at θ_0 .

So far we have focused on constraints on the population moments of z derived from a DAPM. To construct an estimator of θ_0 , we work with the sample counterpart of $Q_0(\theta)$, $Q_T(\theta)$, which is a known function of \bar{z}_T . (The subscript T is henceforth used to indicate dependence on the entire sample.)

The sample-dependent θ_T that minimizes $Q_T(\theta)$ over Φ is the *extremum* estimator of θ_0 . When the first-order condition to the population optimum problem takes the form (1.11), the corresponding first-order condition for the sample estimation problem is⁹

$$\frac{\partial Q_T}{\partial \theta}(\theta_T) = \frac{1}{T} \sum_{i=1}^T \mathcal{D}(z_i; \theta_T) = 0. \quad (1.12)$$

The sample relation (1.12) is obtained by replacing the population moment in (1.11) by its sample counterpart and choosing θ_T to satisfy these sample moment equations. Since, under regularity, sample means converge to their population counterparts [in particular, $Q_T(\cdot)$ converges to $Q_0(\cdot)$], we expect θ_T to converge to θ_0 (the parameter vector of interest and the unique minimizer of Q_0) as $T \rightarrow \infty$.

As noted previously, DAPMs often give rise to moment restrictions of the form (1.10) for more than one \mathcal{D} , in which case there are multiple feasible estimation strategies. Under regularity, all of these choices of \mathcal{D} have the property that the associated θ_T converge to θ_0 (they are *consistent* estimators of θ_0). Where they differ is in the variance-covariance matrices of the implied large-sample distributions of θ_T . One paradigm, then, for selecting among the feasible estimation strategies is to choose the \mathcal{D} that gives the most econometrically efficient estimator in the sense of having the smallest asymptotic variance matrix. Intuitively, the later estimator is the one that exploits the most information about the distribution of \bar{z}_T in estimating θ_0 .

Once a DAPM has been selected for study and an estimation strategy has been chosen, one is ready to proceed with an empirical study. At this stage, the econometrician/modeler is faced with several new challenges, including:

1. The choice of computational method to find a global optimum to $Q_T(\theta)$.
2. The choice of statistics and derivation of their large-sample properties for testing hypotheses of interest.
3. An assessment of the actual small-sample distributions of the test statistics and, thus, of the reliability of the chosen inference procedures.

The computational demands of maximizing Q_T can be formidable. When the methods used by a particular empirical study are known, we occasionally comment on the approach taken. However, an in-depth exploration of

⁹ In subsequent chapters we often find it convenient to define Q_T more generally as $1/T \sum_{i=1}^T \mathcal{D}_T(z_i; \theta_T) = 0$, where $\mathcal{D}_T(z_i; \theta)$ is chosen so that it converges (almost surely) to $\mathcal{D}(z_i; \theta)$, as $T \rightarrow \infty$, for all $\theta \in \Phi$.

alternative algorithms for finding the optimum of Q_T is beyond the scope of this book.

With regard to points (2) and (3), there are many approaches to testing hypotheses about the goodness-of-fit of a DAPM or the values of the parameters θ_0 . The criteria for selecting a test procedure (within the classical statistical paradigm) are virtually all based on large-sample considerations. In practice, however, the actual distributions of estimators in finite samples may be quite different than their large-sample counterparts. To a limited degree, Monte Carlo methods have been used to assess the small-sample properties of estimators θ_T . We often draw upon this literature, when available, in discussing the empirical evidence.

