

COPYRIGHT NOTICE:

Peter Bossaerts: The Paradox of Asset Pricing

is published by Princeton University Press and copyrighted, © 2005, 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.

For COURSE PACK and other PERMISSIONS, refer to entry on previous page. For more information, send e-mail to permissions@pupress.princeton.edu

CHAPTER I

Principles of Asset-Pricing Theory

1.1 Introduction

In this chapter, we will study the body of asset-pricing theory that is most appropriate to understanding the empirical tests that are reported later in this book. In particular, we focus on the discrete-time, stationary dynamic asset-pricing models that have been derived over the past thirty years or so. There are other classes of asset-pricing models, such as Merton's continuous-time model (see Merton [1973]), but these have been less important in empirical analysis. The models we will consider here capture the essence of all modern asset-pricing theory, namely, that financial markets equilibrate to the point that expected returns are determined solely by covariance with aggregate risk.

Such static models as the Capital Asset-Pricing Model (CAPM) can be considered special cases. We focus on the dynamic models because they potentially explain the (albeit small) predictability that one can find in historical securities return data. Moreover, returns are usually measured as tomorrow's price plus dividend divided by today's price. It is customary in static asset-pricing models to take these returns as given. That is, they are considered to be draws of some exogenous random-number generator. Yet returns are based on prices. In particular, they

involve the prices that will be set in tomorrow's markets. To understand how these prices are set, we need dynamic asset-pricing theory.

Dynamic asset-pricing theory builds on dynamic portfolio optimization, which itself is based on stochastic dynamic programming. The main mathematical technique that one uses to solve dynamic programming problems is the Bellman principle. We start out with a brief review of this principle. We will elaborate on a simplification in the optimality conditions that obtains in the portfolio optimization context. This simplification is of utmost importance in empirical analysis because it eliminates the need to solve for the key component of the Bellman principle, namely, the Bellman function.

At the heart of asset-pricing theory is the notion that portfolio optimizing agents meet in the marketplace, and that their demands interact to drive prices to an equilibrium. The theory then focuses on the properties of securities prices in the ensuing equilibrium. The predictions are very general and have been widely used in practical financial decision-making.

Although asset-pricing theory focuses on equilibrium, little attention has been paid to the question of how financial markets reach equilibrium. We shall refer to the latter throughout this book as the price discovery process. The chapter will end with a brief description of how one could go about modeling this process in a plausible and tractable way. The exercise leads to some interesting conclusions, opening up the way for much-needed further research.

1.2 Stochastic Dynamic Programming

We refrain from treating stochastic dynamic programming in its full generality. Instead, a specific problem is studied that encompasses most of the portfolio optimization problems that have been the basis of discrete-time, stationary dynamic asset-pricing models.

Let t index time: $t = 0, 1, 2, \dots, T$. (In much of what follows, we take $T = \infty$.) There is an $M \times 1$ *state vector* x_t and a $K \times 1$ *control vector* y_t . The state vector x_t summarizes the state of the investor's environment: wealth, parameters of the distribution of returns (if these change over time), the riskfree rate (if a riskfree asset is available), the history of consumption,

etc. The control vector y_t consists of investment decisions: how much does the investor choose to allocate to each available asset? We measure these quantities in dollars, so that the total amount allocated to all available assets is also our investor's gross investment, and consumption can be obtained from subtracting gross investment from wealth.

Our agent has a period- t -utility function denoted by $u_t(x_t, y_t)$, which we will assume to be time-invariant, so that we can drop the time subscript: $u(x_t, y_t)$. Notice that utility can be both a function of the state and the control variables. Because of our interpretation of state and control, this dual dependence is necessary even if utility is just a function of period consumption, because period consumption is not a separate control, but can be obtained by subtracting gross investment (the sum of the control variables) from wealth (one of the state variables), as we discussed in the previous paragraph.

The agent discounts future utility by using a subjective discount factor δ . We assume:

$$\delta \leq 1. \quad [1.1]$$

Let $N(\cdot)$ denote the discounted value of future utility:

$$N(x_0, x_1, \dots; y_0, y_1, \dots) = \sum_{t=0}^T \delta^t u(x_t, y_t). \quad [1.2]$$

The time-additivity seems restrictive, but the period-utility function $u(\cdot)$ depends on both the control and the state, and the latter may include the past history of consumption, so that period-utility may depend on past consumption, effectively linking utilities of different periods. This will be demonstrated later on, in the section on time nonseparable preferences.

R_{t+1} denotes the $N \times 1$ stochastic *shock vector*, with distribution function $F_t(R_{t+1}|x_t, y_t)$. The symbol R_{t+1} indicates how we interpret the shocks: they are the returns on the available assets (time $t+1$ prices plus dividends, divided by time t prices). We take the distribution function to be time-invariant, so that we can drop the time subscript: $F(R_{t+1}|x_t, y_t)$. This is without loss of generality: variation over time in any parameter of this distribution function can be captured by including the parameter in the state vector on which the distribution function still depends.

In portfolio optimization problems, the distribution of the shock will only be allowed to be a function of the state, and not of the control. In other words, we will not be able to control directly the distribution of the shock. Because the shock is interpreted as the securities returns, this restriction is effectively an assumption of *perfect competition*: the agent must take the distribution of returns as given. Thus, we will write $F(R_{t+1}|x_t)$. It may appear that our investor could impact the distribution of returns indirectly, through the effect of controls on wealth (one of the state variables). Later on, we will explicitly rule out such indirect control as well. In fact, this will lead to a vast simplification of the dynamic programming problem that our investor faces. Again, these restrictions are not ad hoc: they reflect the assumption of perfect competition.

The state moves (“transits”) from one state to another by means of the following *state-transition equation*, which we also take to be time-invariant (hence we drop the time index):

$$\begin{aligned} x_{t+1} &= g_t(x_t, y_t, R_{t+1}) \\ &= g(x_t, y_t, R_{t+1}). \end{aligned}$$

Notice that the distribution of the future state x_{t+1} , conditional on the past state x_t and the control y_t , is induced by the (conditional) distribution of R_{t+1} . Wealth is one of the state variables; hence, one of the state transition equations will merely capture how wealth changes as a function of returns and investment choice.

We want to maximize the expected discounted value of future utility:

$$\begin{aligned} &\max_{y_0, y_1, \dots} E \left[N(x_0, x_1, \dots; y_0, y_1, \dots) | x_0 \right] \\ &= \max_{y_0, y_1, \dots} E \left[\sum_{t=0}^T \delta^t u(x_t, y_t) | x_0 \right]. \end{aligned}$$

Hopefully the solution will be simple. In particular, we hope that the optimal policy y_t is a function only of the contemporaneous state x_t : $y_t = h_t(x_t)$, in which case one refers to a *Markovian strategy*. It is not obvious that such an assumption is correct. In fact, it is not even obvious whether the optimal policy will be measurable in the states: the optimal policy may involve randomization. But if indeed $y_t = h_t(x_t)$, then our assumptions imply that the joint process of the state and the control is Markovian (i.e., that its distribution conditional on the past depends only on the immediate past).

It would distract too much from the main purpose of this book if we were to elaborate on the issue of whether the optimal policy is Markovian. The reader is warned that the answer may be negative, and is referred for further information to the excellent treatments in Bertsekas and Shreve (1976) and Lucas and Stokey (1989).

If it exists and is measurable in x_t , solving for the optimal policy y_t is facilitated by the Bellman principle, which states that there exists a sequence of functions of the state only, $V_t(x_t)$, $t = 0, 1, 2, \dots$, called *value functions*, such that the optimal policies y_0, y_1, y_2, \dots , can be obtained by recursively solving the following sequence of optimization problems:

$$V_t(x_t) = \max_y \{ u(x_t, y) + \delta E [V_{t+1}(g(x_t, y, R_{t+1})) | x_t] \}. \quad [1.3]$$

The reader should verify that our assumptions so far imply that the expectation of functions of future information (state, policy, shock) will indeed be a function of the immediate past state and policy only.

Let us study the first-order conditions for the problem in (1.3). For this to make sense, we must take u as well as V_t ($t = 0, 1, 2, \dots$) to be continuously differentiable in their arguments. The former is a matter of assumption, the latter is not. Again, the issue would distract us here, but a discussion can be found in the aforementioned references.

The first-order conditions for control k are given in (1.4). Given x_t , the optimal strategy (point) is determined by:

$$\frac{\partial u(x_t, y)}{\partial y_k} + \delta E \left[\sum_{m=1}^M \frac{\partial V_{t+1}(g(x_t, y, R_{t+1}))}{\partial x_{m,t+1}} \frac{\partial g_m(x_t, y, R_{t+1})}{\partial y_k} \middle| x_t \right] = 0. \quad [1.4]$$

Varying x_t , we get a strategy function $y_t = h_t(x_t)$, which we hope is continuously differentiable. Again, this is not obvious, and must be proven.

Very often, it is indeed possible to prove that a continuously differentiable value function and Markovian strategies obtain, but computing them explicitly may be impossible. This certainly is annoying when we want to verify empirically an asset-pricing model that is based on dynamic portfolio optimization. Our inference error is already influenced by our ignorance of the key parameters that affect the distribution of the shocks (we may not even know this distribution at all, and have to rely on non-parametric procedures). We would like not to make things worse by introducing an additional error from numerical computation of the optimum portfolio strategies, if only because it is hard to control inference when

the error is partly statistical (sampling error) and partly deterministic (numerical error).

Fortunately, dynamic portfolio-allocation problems generally lead to a simplification that effectively circumvents these issues. As a matter of fact, we will have to compute neither the optimal policies nor the Bellman functions when testing the asset-pricing model that they induce. This obtains because of natural restrictions on the state transition equation, which, like the restriction that the distribution of the shocks (returns) be independent of the controls, are motivated by assuming perfect competition.

To see what happens in the abstract stochastic dynamic control problem, consider the envelope condition at $t + 1$:

$$\begin{aligned} & \frac{\partial V_{t+1}(x_{t+1})}{\partial x_{m,t+1}} \\ &= \frac{\partial u(x_{t+1}, h_{t+1}(x_{t+1}))}{\partial x_{m,t+1}} + \delta E \left[\sum_{l=1}^M \frac{\partial V_{t+2}(g(x_{t+1}, h_{t+1}(x_{t+1}), R_{t+2}))}{\partial x_{l,t+2}} \times \right. \\ & \quad \left. \frac{\partial g_l(x_{t+1}, h_{t+1}(x_{t+1}), R_{t+2})}{\partial x_{m,t+1}} \Big| x_{t+1} \right] \\ & \quad + \delta \frac{\partial E[V_{t+2}(g(x_{t+1}, h_{t+1}(x_{t+1}), R_{t+2})) | x_{t+1}]}{\partial x_{m,t+1}}, \end{aligned} \quad [1.5]$$

where the last term has to be interpreted as the derivative of the expected Bellman function w.r.t. the conditioning variable x_{t+1} .

Now consider the case where:

1. The first state does not influence any of the state transition equations:

$$\frac{\partial g_l(x_{t+1}, h_{t+1}(x_{t+1}), R_{t+2})}{\partial x_{l,t+1}} = 0, \quad [1.6]$$

where $l = 1, \dots, M$;

2. Neither does it influence the distribution of the future shock:

$$\frac{\partial F(R_{t+2} | x_{t+1})}{\partial x_{l,t+1}} = 0, \quad [1.7]$$

and hence,

$$\frac{\partial E[V_{t+2}(g(x_{t+1}, h_{t+1}(x_{t+1}), R_{t+2})) | x_{t+1}]}{\partial x_{l,t+1}} = 0;$$

3. None of the controls affects the state transition equations except for the first state:

$$\frac{\partial g_m}{\partial y_k} = 0, \quad [1.8]$$

where $m = 2, \dots, M$, $k = 1, \dots, K$.

By interpreting the first state variable as wealth, these amount to assumptions of competitive behavior on the part of the agent. The first assumption states that the agent cannot indirectly control her environment (through the impact of investment decisions—controls—on her wealth), except, of course, for her own wealth (the first state variable). The second states that she cannot indirectly control the distribution of asset returns (shocks). The third assumption states that our investor cannot directly control her environment (state vector, except wealth) either.

Combining (1.5) and (1.4) then implies that the first-order conditions simplify to:

$$- \delta E \left[\frac{\frac{\partial u(x_{t+1}, h_{t+1}(x_{t+1}))}{\partial x_{1,t+1}}}{\frac{\partial u(x_t, y)}{\partial y_k}} \frac{\partial g_1(x_t, y, R_{t+1})}{\partial y_k} \Big| x_t \right] = 1. \quad [1.9]$$

This is a remarkable simplification, because the first-order conditions do not directly involve the Bellman function. The result is quite surprising, because it says that the optimal control at t is entirely determined by the impact on the ratio of marginal utility at t and at $t + 1$, given the optimal policy at $t + 1$, as well as the effect on the state transition between t and $t + 1$. It is unnecessary to account for effects on utility beyond $t + 1$. In general, it is not sufficient to consider substitution effects between only t and $t + 1$, because changes in the time t policy may influence utility far beyond $t + 1$.

Because the expectation in (1.9) can be written as an integral, the equation is really an integral equation, which prompted economists to call it a *stochastic Euler equation*, a term borrowed from calculus of variation.

Economists insist on stationarity. That is, they would like the Bellman function and the optimal policies to be time-invariant functions of the state. This will be important in the empirical analysis, as explained in Chapter 2. Under stationarity, we can drop the time subscripts and merely

refer to the beginning-of-period state and policy as x and y , respectively, and to the end-of-period state and policy as x' and y' , respectively. Let R denote the period shock (return). The Bellman equation (1.3) becomes:

$$V(x) = \max_y \{u(x, y) + \delta E[V(g(x, y, R))|x]\}, \quad [1.10]$$

and the first-order conditions are now:

$$\frac{\partial u(x, y)}{\partial y_k} + \delta E \left[\sum_{m=1}^M \frac{\partial V(g(x, y, R))}{\partial x'_m} \frac{\partial g_m(x, y, R)}{\partial y_k} \middle| x \right] = 0. \quad [1.11]$$

The simplifying conditions are restated as follows.

1. The first state does not influence any of the state transition equations:

$$\frac{\partial g_l(x', h(x'), R')}{\partial x'_1} = 0,$$

where $l = 1, \dots, M$;

2. Neither does it influence the distribution of the future shock:

$$\frac{\partial F(R'|x')}{\partial x'_1} = 0,$$

and hence,

$$\frac{\partial E[V(g(x, h(x), R))|x']}{\partial x'_1} = 0;$$

3. None of the controls affects the state transition equations except for the first state:

$$\frac{\partial g_m}{\partial y_k} = 0,$$

where $m = 2, \dots, M$, $k = 1, \dots, K$. The stochastic Euler equations then simplify to:

$$\boxed{-\delta E \left[\frac{\frac{\partial u(x', h(x'))}{\partial x'_1}}{\frac{\partial u(x, y)}{\partial y_k}} \frac{\partial g_1(x, y, R)}{\partial y_k} \middle| x \right] = 1.} \quad [1.12]$$

1.3 Application to a Simple Investment-Consumption Problem

Let us now translate the abstract results of the previous section into the concrete problem of the simplest portfolio investment problem: how

much to consume and how much to invest? There is one good, which we refer to as the *dollar*. It can be either consumed or invested. In the latter case, it produces an uncertain return.

Although unorthodox, the most transparent translation defines the state and control (policy) variables as suggested in the previous section. In particular, let the first state variable x_1 be wealth (measured in terms of the single good) available for either consumption or investment at the beginning of the period (x'_1 will be available at the beginning of the next period). We interpret the control as investment, and the shock as return. Because there is only one asset, $N = K = 1$. Therefore, the policy variable y is the remaining wealth at the beginning of the period, after consumption has been subtracted. The shock variable R is the return on the investment over the period. We impose the simplifying assumption (1.7): the distribution of R_t is independent of x_1 . Since x_1 is the investor's wealth, the latter is really an assumption of competitive behavior, as already mentioned: the investor never has enough wealth to influence prices, and hence, the return distribution. The reader can immediately infer from this that the present model cannot be used to analyze investment choice in a strategic environment (e.g., a monopoly or oligopoly).

Our specification implies the following for the first transition equation:

$$x'_1 = g_1(x, y, R) = yR.$$

We need specify neither the nature of the remaining state variables nor their transition equations. These state variables only influence the outcome through their effect on the conditional distribution of the return R . But we do assume that their state transition equations are unaffected by investment, either directly or indirectly—assumptions (1.6) and (1.8). This effectively means that our investor cannot influence her environment through her actions, except for her own wealth. Again, these are assumptions about competitive behavior.

Let the utility be logarithmic, that is:

$$u(x, y) = \ln(x_1 - y).$$

To facilitate cross-reference to established formulae in the literature, let c denote consumption. Consumption is what remains after investing wealth: $c = x_1 - y$. Define:

$$\tilde{u}(c) = u(x, y).$$

Because we imposed (1.6)–(1.8), we can use the simplified first-order conditions in (1.12) to find the optimal consumption-investment policy (provided the outcome is stationary; otherwise, we have to refer to (1.9)). In this case, (1.12) reads:

$$\delta E \left[\frac{x_1 - y}{x'_1 - y'} R | x \right] = 1.$$

Let us guess an optimal policy where investment is proportional to the beginning-of-period wealth, that is:

$$y = \gamma x_1. \quad [1.13]$$

Of course, for stationarity, we should guess $y' = \gamma x'_1$ as well. The reader can easily verify that this policy does indeed satisfy the first-order conditions, with $\gamma = \delta$. Writing this in terms of consumption:

$$c = (1 - \delta)x_1,$$

that is, our investor consumes a fixed fraction of her wealth.

In fact, we have obtained a well-known result, namely, that an investor with logarithmic preferences invests in a *myopic* way: her investment policy does not change with the state of the world, except for her own wealth. There can be substantial predictability in returns (through the state variables x_2, \dots, x_M), but our investor keeps investing a fixed fraction of her wealth.

Notice that we did not have to solve for the Bellman function to find the optimal policy—a dramatic simplification of the problem.

1.4 A Nontrivial Portfolio Problem

We can readily extend the above to the case where there are multiple investment opportunities, called securities. The policy y is now a vector with N entries, namely, the number of dollars to be invested in each security. The return vector, R , has now N entries as well ($K = N$). The transition equation for the first state becomes:

$$x'_1 = g_1(x, y, R) = \sum_{n=1}^N y_n R_n. \quad [1.14]$$

Because we shall need it in the next chapter, we will consider general

one-period preferences. They are a function of consumption $c = x_1 - \sum_{n=1}^N y_n$ only:

$$\tilde{u}(c) = \tilde{u}(x_1 - \sum_{n=1}^N y_n) = u(x, y).$$

We make the same simplifying assumptions as in the previous section, and obtain the following stochastic Euler equations, written in terms of consumption. For $n = 1, \dots, N$,

$$\delta E \left[\frac{\frac{\partial \tilde{u}(c')}{\partial c'}}{\frac{\partial \tilde{u}(c)}{\partial c}} R_n | \mathcal{X} \right] = 1. \quad [1.15]$$

That is, the optimal portfolio is such that the (conditional) expectation of the marginal rate of substitution of consumption times the return on each asset equals δ^{-1} .²

Equation (1.15) will be crucial to understanding asset-pricing theory.

1.5 Portfolio Separation

In the static (one-period) case, theorists observed early on that the optimal portfolio strategy often involved *portfolio separation*, which means that the optimal portfolio for a variety (or even all) of risk-averse preferences can be obtained as a simple portfolio of a number of basic portfolios, referred to as mutual funds (see Ross [1978]). This facilitated the development of asset-pricing models. In particular, it led to the CAPM of Sharpe (1964), Lintner (1965), and Mossin (1966).

Given its importance in empirical research, we should examine the static case. Although we have been working in a dynamic context, it is fairly easily adjusted to accommodate the static models by considering the last period, $T - 1$ to T . To simplify notation, let the primed variables (x' , y') refer to time T , and those without a prime (x , y) refer to time $T - 1$. R denotes the return over $(T - 1, T)$.

2. The ratio of marginal utilities is traditionally referred to as the marginal rate of substitution. The term is suggested by the nonstochastic case, where the optimal rate of substitution of consumption over time or across goods is indeed given by the ratio of marginal utilities, as a consequence of the implicit function theorem. Any introductory textbook in economics establishes this relation.

At T , the investor consumes everything: $c' = x'_1$.

Consider *quadratic utility*:

$$\tilde{u}(c) = ac - \frac{b}{2}c^2,$$

where $a > 0$, $b > 0$. Marginal utility is a function of c :

$$\frac{\partial \tilde{u}(c)}{\partial c} = a - bc.$$

With quadratic utility, investors care only about the mean and variance of the return on their portfolio. For this reason, quadratic utility is often referred to as mean-variance preferences, and the optimal portfolio as the *mean-variance optimal*: it provides minimal return variance for its mean.

The optimality condition in (1.15) can be applied directly, producing:

$$\delta E[(a - bx'_1)R_n|x] = \lambda, \quad [1.16]$$

where

$$\lambda = a - bc.$$

(The second-order condition for optimality holds, because $b > 0$.)

To show how portfolio separation works, split the payoff on the optimal portfolio ($\sum_{n=1}^N y_n R_n$, or, equivalently, x'_1) into (1) a riskfree part with return R_F (assuming it exists) and (2) a portfolio of risky securities *only*, with return R_b :

$$x'_1 = \eta_{F,x} R_F + \eta_{b,x} R_b,$$

where $\eta_{F,x}$ is the dollar amount invested in the riskfree security, and $\eta_{b,x}$ is the dollar amount invested in risky securities (these quantities can vary with x , whence the subscript). Let us refer to the portfolio of risky securities only as the *benchmark portfolio*.

Now project the excess return on security n onto that of the benchmark portfolio:

$$R_n - R_F = \alpha_{n,x} + \beta_{n,x}(R_b - R_F) + \epsilon_n. \quad [1.17]$$

This is a conditional projection, which means that the error satisfies the following *conditional moment restrictions*:

$$E[\epsilon_n R_b|x] = E[\epsilon_n|x] = 0. \quad [1.18]$$

The conditioning justifies the subscript x on the intercept and slope, α_n and β_n , respectively. If time-variant, the riskfree rate R_F will be one of the remaining state variables x_2, \dots, x_M . Hence, (1.18) is equivalent to

$$E[\epsilon_n(R_b - R_F)|x] = E[\epsilon_n|x] = 0,$$

which is the usual definition of (conditional) projection.

Apply the optimality condition in (1.16) to R_n and R_F and take the difference:

$$\delta E[(a - bx'_1)(R_n - R_F)|x] = 0. \quad [1.19]$$

Next, apply (1.17):

$$\begin{aligned} & \delta E[(a - bx'_1)\alpha_{n,x}|x] + \delta\beta_{n,x}E[(a - bx'_1)(R_b - R_F)|x] \\ & + \delta E[(a - bx'_1)\epsilon_n|x] = 0. \end{aligned} \quad [1.20]$$

By construction (see (1.18)), the third term is zero:

$$E[(a - bx'_1)\epsilon_n|x] = aE[\epsilon_n|x] - \eta_{F,x}R_F bE[\epsilon_n|x] - \eta_{b,x}bE[\epsilon_n R_b|x] = 0.$$

The second term is zero, by the assumption that x'_1 is constructed optimally: (1.19) holds for each R_n , which means that it holds for any linear combination (portfolio) of returns, and, in particular, for R_b .

So, if R_b is to provide the return on the risky part of the optimal portfolio, it is necessary that, for all n :

$$\boxed{\alpha_{n,x} = 0.} \quad [1.21]$$

When (1.21) holds, the first term in (1.20) will be zero as well, as required by optimality.

The condition in (1.21) is *sufficient* as well: it can be used to construct an optimal portfolio for any person with quadratic utility (i.e., any choice of a and b): first determine for which benchmark portfolio the return R_b is such that $\alpha_{n,x} = 0$, for all n . Next, choose the weights $\eta_{F,x}$ and $\eta_{b,x}$ such that the second term in (1.20) is zero. The third term will be zero by construction.

We have obtained *portfolio separation*: the optimal portfolio for investors with quadratic utility can be obtained as a linear combination of the riskfree security and a benchmark portfolio that is the same for everyone. This also means that all investors effectively demand the same

portfolio of risky securities in the marketplace, a powerful result that can readily be exploited to get sharp asset-pricing results, as illustrated below.

It is interesting to work out the choice of $\eta_{b,x}$ that makes the resulting portfolio of the riskfree security and the benchmark portfolio optimal for a given investor. It is (the derivation is left as an exercise):

$$\eta_{b,x} = -\frac{E[R_b - R_F|x]}{\text{var}(R_b - R_F|x)} \frac{a - bE[x'_1|x]}{-b}. \quad [1.22]$$

That is, the optimal choice is minus the product of the reward to risk ratio times the ratio of the expected marginal utility of future wealth over the change in this expected marginal utility. The implications are intuitive: as the coefficient of risk aversion, b , increases, $\eta_{b,x}$ decreases (provided the reward to risk ratio is positive). This means that a more risk-averse person puts fewer dollars into risky investments.

An important caveat is in order. Portfolio separation will obviously not obtain if investors hold differing beliefs about the distribution of returns and states, because they would each compute different expectations, variances, and covariances. Their portfolio demands would reflect these differences in beliefs, and hence, cannot necessarily be described in terms of demand for a riskfree security and a benchmark portfolio, even if they all have quadratic preferences.

Two final remarks:

1. There is a way to obtain portfolio separation for all types of risk-averse preferences, not just quadratic preferences. We merely have to turn the linear projection conditions in (1.18) into *conditional mean independence*:

$$E[\epsilon_n | R_b - R_F, x] = 0. \quad [1.23]$$

(See Ross [1978].) In other words, the projection in (1.17) is a regression.³ If returns are (conditionally) normally distributed,

3. There is a subtle but important distinction between linear projection and regression that is not always made clear in introductory textbooks of statistics or econometrics. Projection is a mere mathematical exercise that is always possible (except for pathological cases): one computes the intercept and slope in such a way that the error is uncorrelated (i.e., orthogonal) to the variable on the right-hand side of the equation (i.e., the explanatory variable). Regression, in contrast, is a statistical exercise. One determines the function of the explanatory variable that provides the expectation of the variable on the left-hand side (i.e., the dependent variable). It is rare that this function is linear. If it is, then the

linear projection and regression coincide, which means that we automatically obtain portfolio separation, no matter what the investors' preferences are. That normally distributed returns give portfolio separation will be demonstrated explicitly in the last section of this chapter.

2. The conditional mean independence restriction in (1.23) is restrictive and may hold only if two or more benchmark portfolios are considered simultaneously. That is, one may need K benchmark portfolios ($K \geq 2$) with returns $R_{b,k}$, such that the error in

$$R_n - R_F = \alpha_{n,x} + \sum_{k=1}^K \beta_{k,n,x}(R_{b,k} - R_F) + \epsilon_n$$

satisfies the conditional mean independence restriction

$$E[\epsilon_n | R_{b,k} - R_F, k = 1, \dots, K; x] = 0.$$

If so, we will obtain $K + 1$ -fund portfolio separation: K benchmark portfolios will be needed to reconstruct an investor's optimal portfolio, in addition to the riskfree security.

1.6 Toward the First Asset-Pricing Model

Having explored portfolio choice, we are now ready to establish our first asset-pricing result. We already hinted at it in the previous section.

Let us assume that two-fund portfolio separation holds, that is, investors' optimal portfolios can be decomposed into a riskfree security and a benchmark portfolio of risky securities only. This would obtain if all investors use quadratic utility, or if returns on risky securities are jointly normally distributed. We also assume that investors hold common beliefs. The situation is thus vastly simplified: all investors demand the same portfolio of risky securities (the benchmark portfolio). This demand meets the supply in the marketplace. The supply of risky securities is called the *market portfolio*. For the market to be in equilibrium (i.e., for demand

coefficients can be found by projection, and the error will not only be uncorrelated with the explanatory variable, but also mean-independent. This means that the assumption of a linear regression function is a restriction on the data. A quite severe one, for that matter, but it does obtain when the explanatory and independent variables are jointly normal.

to match the supply) the market portfolio and the benchmark portfolio must coincide.

This implies, in particular, that *the market portfolio must be an optimal portfolio*, which means that it satisfies the same restrictions as the benchmark portfolio, namely, (1.21). Let R_M denote the return on the market portfolio. Project the excess return on all the assets onto that of the market portfolio:

$$R_n - R_F = \alpha_{n,x}^M + \beta_{n,x}^M (R_M - R_F) + \epsilon_n^M. \quad [1.24]$$

The error will have the following properties:

$$E[\epsilon_n^M R_M | x] = E[\epsilon_n^M | x] = 0.$$

In equilibrium:

$$\alpha_{n,x}^M = 0, \quad [1.25]$$

for all n .

This asset-pricing model has become known as the CAPM and was first derived by Sharpe (1964), Lintner (1965), and Mossin (1966). Taking expectations in (1.24), we can rewrite the condition in (1.25) in a more familiar form. For all n :

$$E[R_n - R_F | x] = \beta_{n,x}^M E[R_M - R_F | x]. \quad [1.26]$$

That is, the expected excess return on a security is proportional to its risk, as measured by the projection coefficient. The projection coefficient has become known as the security's *beta*.

There are two important remarks to be made about the developments so far. First, at the core of the CAPM is the notion of equilibrium. That is, the predictions that the theory makes about prices in a financial market rely on the belief that these markets somehow equilibrate. There is a school of thought in economics, the Neo-Austrian school, that rejects the very idea that markets equilibrate. We demonstrate later that equilibration, or *price discovery*, as we call it, is indeed far from a foregone conclusion.

Second, although (1.26) is referred to as an asset-pricing model, prices do not enter explicitly. They only enter implicitly, in that the return equals tomorrow's payoff divided by today's price. Equilibration, then, requires that the market search for the prices such that the return distributions for all the securities satisfy (1.26). There is an issue as to

whether there exist prices such that (1.26) can hold at all. That is, *equilibrium existence* is not a foregone conclusion. We will postpone discussion until the end of this chapter.

1.7 Consumption-Based Asset-Pricing Models

The argument that led to the asset-pricing model in the previous section is based on identification of a portfolio that must be optimal in equilibrium. A variation of this argument is to identify a consumption process that must be optimal in equilibrium. Let us investigate this now.

We again start with the first-order conditions, this time expressed directly in terms of consumption, namely 1.15:

$$\delta E \left[\frac{\frac{\partial \tilde{u}(c')}{\partial c'}}{\frac{\partial \tilde{u}(c)}{\partial c}} R_n | x \right] = 1. \quad [1.27]$$

This is a restriction at the individual level, prescribing how individual consumption must correlate with asset returns to be optimal. It is silent about market-wide phenomena, in particular, equilibrium.

It does become an equilibrium restriction, however, if it holds for *all* investors, that is, if all investors implement a consumption-investment policy that is optimal, and hence, satisfies (1.27).

It is hard to test such a proposition, for lack of data on individual consumption. One would like to work with aggregate data, which are more readily available. In particular, a restriction in terms of aggregate consumption is desirable.

We immediately conclude that if all investors are alike, aggregate consumption (both preferences and beliefs) must be optimal as well, for in that case, aggregate and private consumption are perfectly correlated (i.e., in equilibrium).

We thus obtain Lucas' consumption-based asset pricing model (Lucas [1978]), which states that the aggregate consumption at the beginning and end of each period, c_A and c'_A , respectively, must be such that

$$\delta E \left[\frac{\frac{\partial \tilde{u}(c'_A)}{\partial c'}}{\frac{\partial \tilde{u}(c_A)}{\partial c}} R_n | x \right] = 1, \quad [1.28]$$

for all n .

The assumption of identical investors is objectionable, but can readily be relaxed in two ways:

1. One could assume that financial markets are *complete*, which means that there are an equal number of securities and possible outcomes. This is equivalent to saying that all risk can be insured (even if not necessarily at a fair price). Arrow (1953) and Debreu (1959) have shown that in such a case: (i) markets equilibrate; and (ii) the equilibrium consumption processes are Pareto optimal, in the sense that they solve a dynamic economy-wide consumption-investment problem as in (1.3), with respect to a social welfare function that is of the same form. In other words, there exists a representative agent whose preferences are described by this social welfare function, and who finds the aggregate consumption process to be optimal given the returns provided in the financial markets. The asset-pricing restriction (1.28) would then hold for aggregate consumption (see also Constantinides [1982]).⁴
2. One could restrict attention to preferences that can be “aggregated,” in the sense that aggregate consumption and investment demands are as if determined by some aggregate investor. The idea is analogous to that of portfolio separation, where every investor essentially demands the same portfolio(s) of risky securities, but it extends to consumption as well. In this case, (1.28) should hold for aggregate consumption; otherwise the market is not in equilibrium (some investor, and hence, the aggregate investor, was not able to implement optimal consumption-investment plans). See Rubinstein (1974) for a list of preferences for which aggregation obtains.

It would distract us to elaborate on the technical aspects of either of the above relaxations on the assumption of identical investors. The interested reader should consult the references.

4. There seems to be little appreciation in the literature that this argument is problematic, because the welfare function will generally not be state-independent (as u in (1.28) is). The welfare function is really a weighted average of the utility function of all the agents in the economy, with the weights determined by marginal utility of wealth. Hence, the weights depend on the distribution of wealth, and the welfare function will as well. That is, the welfare function depends on state variables that capture the distribution of wealth.

The consumption-based model in (1.28) has the advantage that it is dynamic, in contrast to the CAPM, which is static. But the CAPM has the advantage that it provides equilibrium restrictions in terms of financial market data only. Instead, the consumption-based model is cast in terms of aggregate consumption. Although aggregate consumption data are available, they may not be as reliable as pure financial data. The frequent revision of older aggregate consumption statistics demonstrates their unreliability.

Rubinstein (1976) has derived a simple dynamic asset-pricing model, where he managed to substitute the market portfolio for aggregate consumption. We cover it in some detail here, because it features prominently in later chapters, both theoretically and empirically. (Rubinstein derived more general versions of the model than we discuss here; we focus on the simplest one, because of its pedagogical merits.)

We assume that investors are all alike (beliefs and preferences), and that they have logarithmic preferences, as in Section 1.3. In this case, optimal aggregate consumption is a fixed fraction of aggregate wealth. Letting x_A and x'_A denote aggregate wealth at the beginning and end of the period, respectively, this means:

$$c_A = (1 - \delta)x_A$$

and

$$c'_A = (1 - \delta)x'_A.$$

What is the future aggregate wealth? From the transition equation in (1.14), we can infer that it is determined by the payoff on the shares that the aggregate investor demands:

$$x'_A = \sum_{n=1}^N y_{A,n} R_n,$$

where $y_{A,n}$ denotes the aggregate demand for security n . At the same time, today's aggregate wealth equals the total amount invested (i.e., δx_A). Hence, the total *return* that is demanded equals $x'_A/(\delta x_A)$.

For the market to be in equilibrium, this demanded return must equal the total return available in the marketplace. Hence, it must be the return on the market portfolio, *if* we allow it to include the riskfree security (which we did not permit in the previous section). This implies:

$$\frac{x'_A}{\delta x_A} = R_M.$$

Plugging this in (1.28) generates the following equilibrium restrictions:

$$E \left[\frac{\delta x_A}{x'_A} R_n | x \right] = \boxed{E \left[\frac{1}{R_M} R_n | x \right] = 1}, \quad [1.29]$$

for all n . We will refer to this set of restrictions as Rubinstein's model.

There is a close relationship between the CAPM and Rubinstein's model, not surprisingly. A few additional assumptions bring us very close, but there is an important difference, which justifies paying some attention to the difference between CAPM and Rubinstein's model.

Assume, in particular, that R_M is conditionally lognormal with $E[\ln R_M | x] = \mu_{M,x}$, $\text{var}(\ln R_M | x) = \sigma_{M,x}^2$. Likewise, some individual asset returns are lognormal, say, assets $n = 1, \dots, N_1$, with $E[\ln R_n | x] = \mu_{n,x}$, $\text{var}(\ln R_n | x) = \sigma_{n,x}^2$. Not all assets can be lognormal, because the market portfolio is a linear combination of all the assets, and could not have a lognormally distributed return, because linear combinations of lognormal random variables are not lognormal.

The correlation between $\ln R_M$ and $\ln R_n$ ($n = 1, \dots, N_1$) is ρ_n .

Tedious algebra (see the Exercises) reveals the following:

$$\mu_{n,x} - \ln R_F = 2\tilde{\beta}_{x,n}^M (\mu_{M,x} - \ln R_F) - \frac{1}{2}\sigma_{n,x}^2. \quad [1.30]$$

The risk measure is:

$$\tilde{\beta}_{x,n}^M = \frac{\text{cov}(\ln R_n, \ln R_M | x)}{\text{var}(\ln R_M | x)}.$$

The correction term $-\frac{1}{2}\sigma_{n,x}^2$ is typical when lognormal random variables are transformed by the logarithmic function.

(1.30) is almost the CAPM; compare to (1.26). But the restriction obtains only for log returns, and even then is not formally the same. In particular, the market portfolio will not be optimal for an investor using quadratic utility (i.e., it will not be mean-variance optimal). But it is optimal for logarithmic preferences.

This version of Rubinstein's model also makes an interesting prediction about the expected log return of the market portfolio. As an exercise, the reader is asked to prove that:

$$\mu_{M,x} - \ln R_F = \sigma_{M,x}^2 - \frac{1}{2}\sigma_{M,x}^2. \quad [1.31]$$

(The term $-\frac{1}{2}\sigma_{M,x}^2$ is deliberately kept separate, because it is a standard correction for lognormal random variables.) That is, the average logarithmic return in excess of the log return on the riskfree asset is proportional to the asset's conditional variance. Hence, the risk premium on the market portfolio is determined by its variance. The higher the variance, the higher the marketwide risk premium. This implication is similar to the one found in Merton (1980) (which presents a continuous-time model).

1.8 Asset-Pricing Theory: The Bottom Line

Let us try to distill the common prediction of the asset-pricing models we have been studying. They all originate in the stochastic Euler equations of (1.15), and state that

$$E[AR_n|x] = 1, \quad [1.32]$$

where A measures aggregate risk. In the CAPM, A is a linear transformation of the return on the market portfolio. In Lucas' consumption-based model, A is the marginal rate of substitution of consumption in the beginning and end of a period. In the Rubinstein model, it equals the inverse return on the market portfolio, that is,

$$A = \frac{1}{R_M}. \quad [1.33]$$

We can express (1.32) in terms of covariances. Because it is always true that for two random variables Y and Z , $E[YZ|x] = \text{cov}(Y, Z|x) + E[Y|x]E[Z|x]$, we can restate (1.32) as follows:

$$E[R_n|x] - R_F = -\text{cov}\left(\frac{A}{E[A|x]}, R_n|x\right). \quad [1.34]$$

That is, mean excess returns are proportional to the covariance with aggregate risk.

Equation (1.34), then, is the central prediction of asset-pricing theory.

1.9 Arrow-Debreu Securities Pricing

It was mentioned before that one version of Lucas' model is the complete-markets model of Arrow (1953) and Debreu (1959).⁵ In it, a particular type of security plays an important role, namely, the *Arrow-Debreu security*, referred to as AD security. It pays one dollar in one state, and zero in all others. It may not be traded literally, but can be obtained by a portfolio of traded securities.

Consider two end-of-period states w and v . Let the beginning-of-period price of AD security w be $P_{x,w}$. Let $P_{x,v}$ denote the beginning-of-period price of AD security v . The return on the former, R_w , equals:

$$R_w = \begin{cases} \frac{1}{P_{x,w}} & \text{if state } w \text{ occurs,} \\ 0 & \text{otherwise.} \end{cases}$$

Likewise,

$$R_v = \begin{cases} \frac{1}{P_{x,v}} & \text{if state } v \text{ occurs,} \\ 0 & \text{otherwise.} \end{cases}$$

Let $\pi_{x,w}$ and $\pi_{x,v}$ denote the conditional probabilities of state w and v , respectively.

Apply Lucas' model (1.28) to obtain:

$$P_{x,w} = \delta \pi_{x,w} \frac{\frac{\partial \bar{u}(c'_{A,w})}{\partial c'}}{\frac{\partial \bar{u}(c_A)}{\partial c}},$$

where $c'_{A,w}$ denotes the aggregate end-of-period consumption in state w . Likewise,

$$P_{x,v} = \delta \pi_{x,v} \frac{\frac{\partial \bar{u}(c'_{A,v})}{\partial c'}}{\frac{\partial \bar{u}(c_A)}{\partial c}}.$$

Assume states w and v are equally likely (i.e., $\pi_{x,w} = \pi_{x,v}$). Taking the ratio of the two state prices produces:

5. The two models are not nested: Lucas' model is stationary, whereas Arrow and Debreu's model obtains in a nonstationary world as well; Lucas' model may obtain in an incomplete market (i.e., a market where not all risk can be insured); the Arrow-Debreu model requires completeness. Most importantly, the Arrow-Debreu model does not require that investors hold common beliefs; they may disagree, but not to the point that one investor thinks a state is impossible whereas another one thinks that it is possible.

$$\frac{P_{x,w}}{P_{x,v}} = \frac{\frac{\partial \tilde{u}(c'_{A,w})}{\partial c'}}{\frac{\partial \tilde{u}(c'_{A,v})}{\partial c'}}. \quad [1.35]$$

That is, the ratio of the AD securities prices for two equally likely states is given by the ratio of the marginal utilities of aggregate consumption.

This has an important implication. Because marginal utility is decreasing (reflecting risk aversion), states with lower aggregate consumption will command a higher price. Loosely speaking, insurance for states with low aggregate consumption is relatively expensive.

1.10 Roll's Critique

The CAPM as well as Rubinstein's model are examples of a class of models that we could best describe as *portfolio-based asset-pricing models*. They identify a particular portfolio that must be optimal for the market to be in equilibrium. In both models, the market portfolio must be optimal. In the CAPM, the market portfolio is mean-variance optimal, and includes only risky assets. In Rubinstein's model, the market portfolio is optimal for logarithmic preferences, and must include the supply of riskfree securities.

If we cannot observe the portfolio that portfolio-based asset-pricing models predict to be optimal, the theory is without empirical content. Using a proxy portfolio will not do. For an optimal portfolio always exists (absent arbitrage opportunities), and hence, we could by chance choose a proxy that happened to be optimal (i.e., that satisfies the restrictions of asset-pricing theory).

This is the core argument of the *Roll critique* (Roll [1977]). In the context of the CAPM, Roll demonstrated that high correlation between the return on the proxy and the market portfolio is no indication that there is much to be learned from the properties of the proxy about the mean-variance optimality of the market portfolio.

In particular, let us suppose that, for some benchmark portfolio with return R_b , we find that $\alpha_{n,x} = 0$ for all n in

$$R_n - R_F = \alpha_{n,x} + \beta_{n,x}(R_b - R_F) + \epsilon_n.$$

This is not a test of the CAPM, but only an indication that the benchmark portfolio is mean-variance optimal. At best, it is a test that no arbitrage opportunities exist, for otherwise there would not be a mean-variance optimal portfolio.

The existence of such a portfolio can be exploited, however, to summarize the data. If we can find a benchmark portfolio (or combination of benchmark portfolios) that is mean-variance optimal, we can use a security's beta to determine its expected excess return:

$$E[R_n - R_F|x] = \beta_{n,x}E[R_b - R_F|x]. \quad [1.36]$$

If the same portfolio (or set of portfolios) is found to be optimal across markets and over time, we conclude that there is a regularity in the data that can be used to predict expected excess returns in the future or in cross-section. Again, this is not evidence that financial markets equilibrate according to asset-pricing theory, but it is an interesting empirical fact that provides useful summary information about how markets price securities.

One wonders whether this is how recent work in empirical asset pricing has to be interpreted, because benchmark portfolios are being used that bear little relationship with the theory, and yet are found to be mean-variance optimal. The prototype is Fama and French (1996); a recent survey is Cochrane (1999). When multiple benchmark portfolios are found to explain the cross-section of expected excess returns, the outcome is called a *multifactor asset-pricing model*. This term is objectionable: unless there is a theoretical reason why the factor portfolios “work,” it should not be referred to as an asset-pricing model.

1.11 Time Nonseparable Preferences

Early in this chapter, it was mentioned that our framework could accommodate time nonseparabilities, even if, from (1.2), our preferences look purely separable. In particular, we can make utility depend on past consumption levels.

Consider the following utility function, evaluated at a state and control of the future (variables with a prime):

$$u(x', y') = \frac{(x'_1 - y' + \lambda x'_s)^{\gamma+1}}{\gamma + 1}, \quad [1.37]$$

where x'_1 denotes future wealth before consumption (as before), y' denotes future wealth after consumption (the control), and x'_s is a new state variable with transition equation:

$$x'_s = x_1 - y;$$

λ is a parameter. Effectively, future-period utility is a function of (1) future consumption $x'_1 - y'$, and (2) prior-period consumption $x_1 - y$. This is a simple way of overcoming time separability in standard preferences. When $\lambda < 0$, consumption is intertemporally complementary, and one refers to this situation as habit persistence. A continuous-time version of this class of preferences was studied, among others, in Constantinides (1990) and Sundaresan (1989).

To simplify notation, let

$$z = x_1 - y + \lambda x_s.$$

The future equivalent carries a prime:

$$z' = x'_1 - y' + \lambda x'_s.$$

We use double primes when referring to two periods in the future:

$$z'' = x''_1 - y'' + \lambda x''_s.$$

As an exercise, the reader is asked to write down the resulting Bellman function for a standard portfolio problem, derive the first-order conditions, and to use the envelope condition to conclude that, for optimality,

$$-z^\gamma + \delta E \left[\left((z')^\gamma + \delta \lambda E[(z'')^\gamma | x'] \right) R_n | x \right] - \delta E[\lambda (z')^\gamma | x] = 0. \quad [*]$$

Under Lucas' assumption of a representative agent, these optimality conditions generate the following asset-pricing model. We will use the subscript A to indicate aggregate variables (i.e., those based on aggregate consumption):

$$E[AR_n | x] = 1,$$

with

$$A = \delta \frac{(z'_A)^\gamma + \delta \lambda E[(z''_A)^\gamma | x']}{z_A^\gamma + \delta \lambda E[(z'_A)^\gamma | x]}. \quad [1.38]$$

1.12 Existence of Equilibrium

We defined equilibrium in a rather casual way as a restriction on moments of the joint distribution of return and aggregate risk (see (1.34)) such that demand for securities equals their supply. But, in the context of the CAPM, we already hinted that existence of equilibrium and equilibration (price discovery) are not foregone conclusions, and hence, deserve some attention. We will first discuss existence of equilibrium.

One difficulty with proving existence of equilibrium is that prices are only implicit in the asset-pricing restrictions. In fact, returns are non-linear functions of beginning-of-period prices, so that a set of prices may not exist such that (1.34) holds.

There is a graver difficulty. In a dynamic context, end-of-period payoffs are the sum of dividends plus the prices at the beginning of the next period, and hence, are endogenous to the model. We can take prices at the beginning of the next period to be equilibrium prices, but we do have to formulate how investors form beliefs about the distribution of future prices, or at least hypothesize what these beliefs are.

This issue did not come up in the complete-markets world of Arrow and Debreu. In it, all risks are supposed to be insurable at all times. Thus all possibilities are already priced at time zero by means of a straightforward Walrasian equilibrium,⁶ and investors can simply read the future pricing of risks in the pattern of prices of AD securities with different maturities (see Debreu [1959]).

In the general context, the standard solution due to Radner makes a stronger assumption about agents' predictive capabilities. It is hypothesized that investors can list tomorrow's states, and, for each state, write down what securities prices would be. In this case the investors know the

6. A Walrasian equilibrium is defined to be the set of prices such that, if investors submit their securities demands for those prices, total demand equals total supply.

mapping from states to prices. This hypothesis has become known as *rational expectations* (RE; see Radner [1972]).

In part, RE originates in the reasonable assumption that investors should be at least as good as the average economist, and hence, should be capable of working out equilibrium prices in any given state. The latter was first suggested by Muth (1961). But to work out equilibrium prices in future states requires an enormous amount of structural knowledge about the economy that even the best economist does not have.

It will be important for the remainder of this book that the reader distinguish between RE and unbiased predictions. To have RE merely makes the strong assumption that investors know what prices would obtain in each possible state. The beliefs that they hold about the chances that each state occurs may still be wrong, and hence, investors make biased forecasts. It should be equally clear, however, that RE and the assumption of unbiased predictions are not entirely unrelated.

Many do not distinguish between RE and unbiased predictions, assuming that investors know the future distribution of prices, effectively stating that they know: (1) the mapping from states to prices; and (2) the distribution of states. In other words, investors' beliefs are right. The assumption is motivated by the concern that stationarity of equilibrium return distributions will not obtain if investors' beliefs can be wrong. If investors realize that they are wrong, they learn, and hence, the state vector (x_t) should include their beliefs. Hopefully, beliefs converge. If so, the state vector does not constitute a stationary (i.e., time-invariant) process. Yet, stationarity facilitates empirical research, as will be discussed in the next chapter.

Even if investors' beliefs are correct, the existence of stationary equilibria is not a foregone conclusion. By correct beliefs we mean that the probability measure with which investors assess uncertainty coincides with the probability measure with which states are factually drawn. This is the assumption that Lucas made in deriving his asset-pricing model, discussed in earlier sections. Lucas referred to it as rational expectations, but, to distinguish it from Radner's notion, we will refer to it as *Lucas-RE*.

To study the existence of equilibrium in Lucas' framework, we obviously cannot work with returns, and must write the equilibrium restrictions explicitly in terms of prices. This also means, among other things, that the stochastic shocks now cannot be identified as returns. Instead,

only the dividends that enter the return can be considered as shocks. Let D'_n denote the end-of-period dividend of security n .⁷ D'_n will be the n th shock. The existence of equilibrium becomes a question of the existence of a *fixed point* in a mapping, as is now shown.

Decompose the return into prices and dividends:

$$R_n = \frac{P'_n + D'_n}{P_n}.$$

The question is: does there exist a function ϕ_n mapping the state into the positive part of the real line, such that:

$$P'_n = \phi_n(x'),$$

$$P_n = \phi_n(x),$$

and the equilibrium condition (1.28) holds? Equivalently, does there exist a mapping ϕ such that:

$$\delta E \left[\frac{\frac{\partial \bar{u}(c'_A)}{\partial c'}}{\frac{\partial \bar{u}(c_A)}{\partial c}} (\phi_n(x') + D_n) \middle| x \right] = \phi_n(x), \quad [1.39]$$

for all n . In more abstract notation: does there exist a fixed point⁸ to the function G , mapping functions into functions, where:

$$G(\phi_n)(\cdot) = \delta E \left[\frac{\frac{\partial \bar{u}(c'_A)}{\partial c'}}{\frac{\partial \bar{u}(c_A)}{\partial c}} (\phi_n(x') + D_n) \middle| \cdot \right]?$$

Under particular assumptions, Lucas demonstrates that the answer is affirmative. This is comforting, because it means that our asset-pricing theory makes sense after all.

I.13 Price Discovery

In addition to the existence of equilibrium, one wonders how financial markets equilibrate. In general equilibrium theory, this question has traditionally been cast in terms of the *stability* of equilibrium: if an economy is kicked off its equilibrium, will it return to it? There is an extensive literature on this in the context of the Arrow-Debreu model (see, e.g., Negishi

7. We could also insist that the dividends add up to the end-of-period consumption ($\sum_{n=1}^N D'_n = c'_A$), but this is not necessary.

8. The fixed point is the value ϕ_n such that $G(\phi_n)(x) = \phi_n(x)$, for all x .

[1962], Arrow and Hahn [1971]). But we have just concluded that our asset-pricing models require a more sophisticated notion of equilibrium than the Walrasian equilibrium that is used in Arrow and Debreu's world, namely, RE.

Most of the analysis to date has focused on how a market would learn the mapping from states to prices that is a crucial component of the RE equilibrium. Nice examples are Jordan (1985) and Marcet and Sargent (1989). Little attention is paid, however, to the question of how financial markets equilibrate given that this mapping is known, perhaps because the analyses are done using static models.

We will see that there may be problems here even within a simple CAPM example.

We need a model of price discovery. We have been assuming competitive behavior so far (i.e., investors cannot control their environment beyond their own consumption and investment policies). We will continue to make the assumption, rather than delve into strategic (i.e., game-theoretic) models of price discovery. Game-theoretic models have become popular in an area referred to as market microstructure theory.⁹

The prototype *competitive price-discovery model* is the Walrasian *tatonnement*. An auctioneer calls out a set of prices for all securities, investors reveal their demand at these prices, and the auctioneer collects the demands. If the total excess demand for a security is nonzero, the auctioneer calls out new prices, and a new round of demand revelation starts. The process stops if all markets clear. In the adjustment between rounds, the auctioneer could simply change the price of a security in the direc-

9. Market microstructure models, starting with Kyle (1985) and Glosten and Milgrom (1985), study how a market discovers an asset's "liquidation value," and hence, can also be considered equilibrium price-discovery models if the liquidation value is interpreted as equilibrium price. But these models are also equilibrium models, and beg the question of how their equilibrium is discovered. That is, the issue of equilibration is pushed back just one step. Many other aspects of market microstructure models make them less obvious candidates as models of price discovery of the equilibria in asset-pricing theory: because they are game-theoretic, market microstructure models make far stronger assumptions about common knowledge; categorization of traders into the classes of market makers, noise traders, and discretionary liquidity traders is absent from asset-pricing theory; risk neutrality is often assumed, whereas risk aversion is at the core of asset-pricing theory; etc. Finally, most market microstructure models focus on information aggregation—an issue that is largely ignored in the asset-pricing theory that has been used to explain historical data.

tion of the excess demand: if it is negative, raise the price; if it is positive, lower the price. We will use this simple adjustment process.

In the standard tatonnement process, there is no trade until all markets clear. A more realistic model of price discovery would have trade take place in each market, perhaps by randomly assigning demand to the suppliers when excess demand is positive, or randomly assigning supply to demanders when excess demand is negative. Models of price discovery with intermediate trading are referred to as *nontatonnement*. In the simple CAPM example given below, however, intermediate trading will be shown not to affect the evolution of prices. Hence, the distinction between tatonnement and nontatonnement is inconsequential as far as the evolution of prices is concerned.

The tatonnement process need not be taken literally, of course. There need not be an auctioneer, no rounds, etc. The mathematics will reflect only the essential parts: (1) investors cannot affect the adjustment process (competition); (2) adjustment works through prices, with prices changing as a function of demand pressure. Everything else is immaterial.

To focus on price discovery, we will concentrate on the static (one-period) case, which means that investors consume their entire end-of-period wealth. In addition, we will exclude consumption in the beginning of the period. This way, the model will be ready to be used to analyze some of the experimental results we will talk about in Chapter 4.

With the exception of the previous section, we have been working with returns rather than prices. An analysis of price discovery, however, requires that one explicitly state the problem in terms of prices. This means that we will have to redefine a few variables. We will also drop unnecessary conditioning, which would clutter the notation. Hence, the state variable x will be one-dimensional, and measure only an investor's wealth. On the other hand, we will have to explicitly consider differences among investors. We will introduce the index j to identify Q types of investors, distinguishable in terms of their risk aversion: $j = 1, \dots, Q$. (It will not be necessary to distinguish investors according to their wealth, because we are going to assume constant absolute risk aversion.)

We will generate the CAPM here by assuming normally distributed returns. In this case, preferences trade-off means against variance (the only two parameters in the normal distribution). Moreover, we will as-

sume constant absolute risk aversion, which implies that the trade-off between mean and variance will be independent of wealth.

As before, we will assume that a riskfree security exists. To facilitate notation, however, we will count $N + 1$ securities, the first one (security $n = 0$) being the riskfree security, and the remaining securities, $n = 1, \dots, N$, being risky.

The price of the riskfree security will be used as numeraire at the beginning of the period, and it is assumed to pay one dollar at the end of the period. Investor j demands h_j units of the riskfree security. In terms of the earlier notation (with the addition of the subscript j , to identify the investor):

$$y_{j,0} = h_j.$$

Investor j demands $z_{j,n}$ units of each of the N risky securities. Security n carries a price of P_n . The latter is a normalized price, expressed in terms of the price of the riskfree security (rather than beginning-of-period dollars). Hence, in terms of the earlier notation:

$$y_{j,n} = z_{j,n}P_n,$$

where $n = 1, \dots, N$.

It will be easier to switch to vector notation. In particular, let z_j denote the vector of units of the risky securities that investor j demands: $z_j = [z_{j,n}]_{n=1,\dots,N}$. Similarly, let P denote the vector of the beginning-of-period prices of the n risky securities: $P = [P_n]_{n=1,\dots,N}$.

Let D_n denote the dollar payoff of security n at the end of the period. From what we said before, $D_0 = 1$. For the risky securities, let D denote the vector of their payoffs ($D = [D_n]_{n=1,\dots,N}$). Let μ denote the vector of expected payoffs ($\mu = [\mu_n]_{n=1,\dots,N}$) and let Σ denote the covariance matrix,

$$\Sigma = [\text{cov}(D_w, D_v)]_{w=1,\dots,N;v=1,\dots,N}.$$

We assume that all investors agree on the means, variances, and covariances (i.e., beliefs are homogeneous).

Investor j 's end-of-period utility, \hat{u}_j , will be a function only of h_j and z_j . In particular,

$$\hat{u}_j(h_j, z_j) = h_j + z_j^T \mu - \frac{b_j}{2} z_j^T \Sigma z_j, \quad [1.40]$$

where T denotes transpose. The parameter b_j measures investor j 's risk aversion (trade-off between mean and variance). It is assumed to be independent of wealth (i.e., the investor has constant absolute risk aversion).

At the beginning of the period there is no consumption, but investor j is endowed with a certain number h_j^0 of riskfree securities, and a vector z_j^0 of risky securities. Therefore, her budget constraint becomes:

$$h_j + P^T z_j = h_j^0 + P^T z_j^0. \quad [1.41]$$

We can ignore discounting, because this is a static model.

Standard maximization of the utility (1.40) subject to the budget constraint in (1.41) generates the following first-order conditions:

$$z_j = \frac{1}{b_j} \Sigma^{-1}(\mu - \lambda_j P),$$

where λ_j is the Lagrange multiplier, which equals 1, from the first-order condition of the demand for the riskfree security. This means:

$$z_j = \frac{1}{b_j} \Sigma^{-1}(\mu - P). \quad [1.42]$$

Incidentally, notice that we have here an example of portfolio separation. In (1.42), only the scalar b_j is specific to an individual. That is, up to a scaling constant, demand for risky securities is the same for everybody, namely,

$$\Sigma^{-1}(\mu - P).$$

The latter can be written as the *weights of a benchmark portfolio*, by normalization:

$$\Sigma^{-1}(\mu - P) / \hat{1}^T \Sigma^{-1}(\mu - P),$$

where $\hat{1}$ denotes a vector of ones.

With an equal number of investors of each type, total excess demand is just the sum of $h_j - h_j^0$ (riskfree security) and $z_j - z_j^0$ (risky securities). In equilibrium, prices must be such that:

$$\sum_{j=1}^Q (h_j - h_j^0) = 0,$$

and

$$\sum_{j=1}^Q (z_j - z_j^0) = 0.$$

The latter can be solved for prices (the former will then automatically be satisfied), which generates:

$$P = \mu - \frac{1}{\sum_{j=1}^Q \frac{1}{b_j}} \Sigma z^0, \quad [1.43]$$

where z^0 denotes the vector of total initial holdings, that is,

$$z^0 = \sum_{j=1}^Q z_j^0.$$

Although it may not be immediately clear, the equilibrium in (1.43) is the CAPM. (In an exercise, the reader is asked to verify the link between (1.43) and (1.26).) As an aside, note that existence is not an issue here, and that the equilibrium is unique.

Now let us introduce tatonnement. Let r measure time in the tatonnement process. Rounds are held at regular intervals Δr , at times r_i ($i = 1, 2, \dots$). The price that is called in round i is P_{r_i} , and the revealed demand for riskfree securities is h_{j,r_i} ; that for risky securities is z_{j,r_i} . Between rounds, the auctioneer adjusts the prices in proportion to the excess demand:

$$\Delta P_r = P_{r_i} - P_{r_{i-1}} = \left[\sum_{j=1}^Q z_{j,r_{i-1}} - z^0 \right] \Delta r.$$

Plug in the formula for individual demands (1.43) to obtain:

$$\Delta P_r = \left[\sum_{j=1}^Q \frac{1}{b_j} \Sigma^{-1} (\mu - P_{r_{i-1}}) - z^0 \right] \Delta r. \quad [1.44]$$

To facilitate the analysis, let the interval between two tatonnement rounds become smaller and smaller: $\Delta r \rightarrow 0$. In the limit, we get continuous adjustment. Let dP/dr denote the vector of derivatives of the price path with respect to tatonnement time. Simple algebra (see exercises) reveals:

$$\frac{dP}{dr} + \left[\sum_{j=1}^Q \frac{1}{b_j} \right] \Sigma^{-1} P = K, \quad [1.45]$$

where the vector of constants K equals:

$$K = \left[\sum_{j=1}^Q \frac{1}{b_j} \right] \Sigma^{-1} \mu - z^0.$$

This is a standard system of ordinary differential equations.

In the exercises, the reader is asked to verify the following:

1. The unique stationary point of the system is the equilibrium in (1.43).
2. The system is globally stable (i.e., solutions converge to the stationary point from anywhere in the price space).

These points provide an answer to the question we set out to tackle, namely, how can markets find the CAPM equilibrium? We proposed that markets adjust as in our tatonnement model, and confirmed that, if so, they would discover equilibrium from anywhere in the price space.

There are at least two aspects of our tatonnement adjustment process that can be criticized. First, no transactions take place until the process has converged. In most financial markets (including the experimental markets of Chapter 3), transactions take place continuously, whether the market is in or out of equilibrium. That is, realistic price adjustment is nontatonnement.

Consequently, we should ask how robust our findings are to intermediate trade. They are unaffected, because of our assumption of constant absolute risk aversion, which implies that the aggregate excess demand, and hence, the price adjustment is independent of investors' wealth (see (1.44)). In other words, no matter how much investors gain or lose in the off-equilibrium transactions, it does not influence their demand in subsequent rounds.

Another way to see this is to notice that the equilibrium prices do not depend on the distribution of initial endowments of risky securities, but only on the total supply, z^0 (see (1.43)).

A second potential criticism concerns the adjustment of the price of the riskfree security. It is implicitly assumed in the standard tatonnement process that the riskfree security does not adjust independently to its own excess demand. This is because the entire analysis is carried out in terms of relative prices: the prices of the risky securities are expressed in terms of that of the riskfree security, and these relative prices adjust only to excess demand in their own market (see (1.44)).

Walrasian equilibrium restricts only relative prices, and our equilibrium is Walrasian. Therefore, it is standard practice to normalize the prices in terms of the price of some numeraire (in our case: the risk-

free security).¹⁰ One then carries out tatonnement analysis, as if the normalization that is justified in equilibrium remains possible out of equilibrium. It may, but it does imply a special adjustment process, because it does not allow the price of the numeraire to change separately as a function of excess demand in its own market.

The simplest way to force the excess demand in the numeraire to affect the price discovery process is to subtract it from the adjustment equation for each relative price. In our case, (1.44) changes to:

$$\Delta P_r = \left[\sum_{j=1}^Q \frac{1}{b_j} \Sigma^{-1} (\mu - P_{r_{j-1}}) - z^0 - \left(\sum_{j=1}^Q h_{j,r_{j-1}} - h^0 \right) \hat{1} \right] \Delta r, \quad [1.46]$$

where $\hat{1}$ denotes a vector of ones, and h^0 is the aggregate supply of riskfree securities ($h^0 = \sum_{j=1}^Q h_j^0$). If the excess demand for the riskfree security is positive, its absolute price increases, and hence, the relative prices of the risky securities decrease. The opposite is true if the excess demand for the riskfree security is negative.

In the limit, the price adjustment process in (1.46) becomes a system of ordinary differential equations:

$$\begin{aligned} \frac{dP}{dr} + \left[\sum_{j=1}^Q \frac{1}{b_j} \right] \Sigma^{-1} P - \left(\left[\sum_{j=1}^Q \frac{1}{b_j} \right] [\Sigma^{-1} \mu] - z^0 \right)^T P \hat{1} \\ + \left[\sum_{j=1}^Q \frac{1}{b_j} \right] P^T \Sigma^{-1} P \hat{1} = K, \quad [1.47] \end{aligned}$$

where the vector of constants K equals:

$$K = \left[\sum_{j=1}^Q \frac{1}{b_j} \right] \Sigma^{-1} \mu - z^0.$$

This is Abel's system of ordinary differential equations of the first kind.

The system in (1.47) has several interesting properties. We single out two of them, relevant for our analysis. Exercises allow the reader to verify these:

10. Incidentally, in a RE equilibrium, one cannot readily normalize prices because absolute prices may carry information that is lost in the normalization.

1. The system can have several stationary points; only one of them is the CAPM equilibrium in (1.43).
2. The system is not globally stable, and the CAPM equilibrium may or may not be locally stable. The latter means that even if prices start out close to the equilibrium, they may never reach it.

These conclusions are troubling. They demonstrate that price discovery in the CAPM environment may be problematic if the price in the market for the riskfree security adjusts separately to its own excess demand, rather than through the excess demand in the risky securities, as in standard tatonnement analysis.

The latter demonstrates that the issue of price discovery in financial markets deserves far more attention than it has received so far. This section provides a framework for thinking about price discovery by staying within the competitive paradigm, rather than moving into strategic issues, which would undoubtedly complicate things.

Exercises

1. Prove (1.5), which is an application of the envelope theorem.
2. Prove that $\gamma = \delta$ in (1.13). (δ is defined in (1.1).)
3. Prove (1.22).
4. Demonstrate that portfolio separation obtains for general risk-averse preferences (i.e., $d\tilde{u}(c)/dc > 0$, $d^2\tilde{u}(c)/dc^2 < 0$) when (1.23) holds.
5. Prove that (1.21) holds in projections of excess returns on individual securities onto that of the market portfolio when the latter minimizes variance of return for a given mean (excess return).
6. Derive (1.30) and (1.31). (*Hint:* If $\ln X$ is normally distributed with mean μ and variance σ^2 , then X is lognormal with mean $e^{\mu + \frac{1}{2}\sigma^2}$ and variance $e^{2\mu + \sigma^2}(e^{\sigma^2} - 1)$.)
7. Derive the Bellman function for a standard portfolio problem using the nonseparable utility function in (1.37), write down the first-order conditions, and use the envelope condition to prove (*).
8. Show that (1.26) and (1.43) are identical.

9. Prove that (1.45) has a unique stationary point, namely (1.43), and that the system is stable at this point.
10. Prove that (1.47) has multiple stationary points (how many exactly?), and evaluate stability at each point. Is (1.43) a stationary point, and, if so, is it always stable?