

Asset Pricing

Asset Pricing

John H. Cochrane

Princeton University Press
Princeton and Oxford

Copyright © 2001 by Princeton University Press
Published by Princeton University Press, 41 William Street,
Princeton, New Jersey 08540
In the United Kingdom: Princeton University Press, 3 Market Place, Woodstock,
Oxfordshire OX20 1SY

All Rights Reserved

Library of Congress Cataloging-in-Publication Data

0-691-07498-4 CIP to come

This book has been composed in Baskerville.

Printed on acid-free paper. ☺

www.pup.princeton.edu

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Acknowledgments

This book owes an enormous intellectual debt to Lars Hansen and Gene Fama. Most of the ideas in the book developed from long discussions with each of them, and from trying to make sense of what each was saying in the language of the other. I am also grateful to all my colleagues in Finance and Economics at the University of Chicago, and to George Constantinides especially, for many discussions about the ideas in this book. I thank all of the many people who have taken the time to give me comments, and especially George Constantinides, Andrea Eisfeldt, Gene Fama, Wayne Ferson, Michael Johannes, Owen Lamont, Anthony Lynch, Dan Nelson, Monika Piazzesi, Alberto Pozzolo, Michael Roberts, Juha Seppala, Mike Stutzer, Pierto Veronesi, an anonymous reviewer, and several generations of Ph.D. students at the University of Chicago. Peter Dougherty was instrumental in shepherding this book to publication. I thank the National Science Foundation and the Graduate School of Business for research support.

You will find typos, corrections, and additional materials, as they develop, on my website <http://gsbwww.uchicago.edu/fac/john.cochrane/research/Papers>

To Sally, Eric, Gene and Lydia

Contents

Acknowledgments	v
Preface	xiii
Part I. Asset Pricing Theory	3
1 Consumption-Based Model and Overview	5
1.1 Basic Pricing Equation	6
1.2 Marginal Rate of Substitution/Stochastic Discount Factor	8
1.3 Prices, Payoffs, and Notation	10
1.4 Classic Issues in Finance	12
1.5 Discount Factors in Continuous Time	28
Problems	33
2 Applying the Basic Model	37
2.1 Assumptions and Applicability	37
2.2 General Equilibrium	39
2.3 Consumption-Based Model in Practice	44
2.4 Alternative Asset Pricing Models: Overview	46
Problems	48
3 Contingent Claims Markets	51
3.1 Contingent Claims	51
3.2 Risk-Neutral Probabilities	53
3.3 Investors Again	54
3.4 Risk Sharing	56
3.5 State Diagram and Price Function	58
4 The Discount Factor	63
4.1 Law of One Price and Existence of a Discount Factor . . .	64

4.2	No Arbitrage and Positive Discount Factors	69
4.3	An Alternative Formula, and x^* in Continuous Time	74
	Problems	77
5	Mean-Variance Frontier and Beta Representations	79
5.1	Expected Return-Beta Representations	80
5.2	Mean-Variance Frontier: Intuition and Lagrangian Characterization	83
5.3	An Orthogonal Characterization of the Mean-Variance Frontier	86
5.4	Spanning the Mean-Variance Frontier	91
5.5	A Compilation of Properties of R^* , R^{e*} , and x^*	92
5.6	Mean-Variance Frontiers for m : The Hansen–Jagannathan Bounds	95
	Problems	100
6	Relation between Discount Factors, Betas, and Mean-Variance Frontiers	101
6.1	From Discount Factors to Beta Representations	102
6.2	From Mean-Variance Frontier to a Discount Factor and Beta Representation	105
6.3	Factor Models and Discount Factors	108
6.4	Discount Factors and Beta Models to Mean-Variance Frontier	112
6.5	Three Risk-Free Rate Analogues	113
6.6	Mean-Variance Special Cases with No Risk-Free Rate	119
	Problems	122
7	Implications of Existence and Equivalence Theorems	123
8	Conditioning Information	133
8.1	Scaled Payoffs	134
8.2	Sufficiency of Adding Scaled Returns	136
8.3	Conditional and Unconditional Models	138
8.4	Scaled Factors: A Partial Solution	146
8.5	Summary	148
	Problems	148
9	Factor Pricing Models	149
9.1	Capital Asset Pricing Model (CAPM)	152
9.2	Intertemporal Capital Asset Pricing Model (ICAPM)	166
9.3	Comments on the CAPM and ICAPM	168

9.4	Arbitrage Pricing Theory (APT)	173
9.5	APT vs. ICAPM	183
	Problems	184
Part II.	Estimating and Evaluating Asset Pricing Models	185
10	GMM in Explicit Discount Factor Models	189
10.1	The Recipe	190
10.2	Interpreting the GMM Procedure	192
10.3	Applying GMM	196
11	GMM: General Formulas and Applications	201
11.1	General GMM Formulas	202
11.2	Testing Moments	206
11.3	Standard Errors of Anything by Delta Method	207
11.4	Using GMM for Regressions	208
11.5	Prespecified Weighting Matrices and Moment Conditions	210
11.6	Estimating on One Group of Moments, Testing on Another	219
11.7	Estimating the Spectral Density Matrix	220
	Problems	228
12	Regression-Based Tests of Linear Factor Models	229
12.1	Time-Series Regressions	230
12.2	Cross-Sectional Regressions	235
12.3	Fama–MacBeth Procedure	244
	Problems	251
13	GMM for Linear Factor Models in Discount Factor Form	253
13.1	GMM on the Pricing Errors Gives a Cross-Sectional Regression	253
13.2	The Case of Excess Returns	256
13.3	Horse Races	258
13.4	Testing for Characteristics	259
13.5	Testing for Priced Factors: Lambdas or b 's?	260
	Problems	264
14	Maximum Likelihood	265
14.1	Maximum Likelihood	265
14.2	ML is GMM on the Scores	268
14.3	When Factors are Returns, ML Prescribes a Time-Series Regression	270

14.4	When Factors are Not Excess Returns, Regression ML Prescribes a Cross-Sectional	273
	Problems	275
15	Time Series, Cross-Section, and GMM/DF Tests of Linear Factor Models	277
15.1	Three Approaches to the CAPM in Size Portfolios	278
15.2	Monte Carlo and Bootstrap	285
16	Which Method?	291
Part III.	Bonds and Options	307
17	Option Pricing	311
17.1	Background	311
17.2	Black–Scholes Formula	318
	Problems	324
18	Option Pricing without Perfect Replication	325
18.1	On the Edges of Arbitrage	325
18.2	One-Period Good-Deal Bounds	327
18.3	Multiple Periods and Continuous Time	334
18.4	Extensions, Other Approaches, and Bibliography	344
	Problems	346
19	Term Structure of Interest Rates	347
19.1	Definitions and Notation	347
19.2	Yield Curve and Expectations Hypothesis	352
19.3	Term Structure Models—A Discrete-Time Introduction	355
19.4	Continuous-Time Term Structure Models	360
19.5	Three Linear Term Structure Models	366
19.6	Bibliography and Comments	377
	Problems	380
Part IV.	Empirical Survey	383
20	Expected Returns in the Time Series and Cross Section	387
20.1	Time-Series Predictability	389
20.2	The Cross Section: CAPM and Multifactor Models	434
20.3	Summary and Interpretation	448
	Problems	453

Preface

Asset pricing theory tries to understand the prices or values of claims to uncertain payments. A low price implies a high rate of return, so one can also think of the theory as explaining why some assets pay higher average returns than others.

To value an asset, we have to account for the *delay* and for the *risk* of its payments. The effects of time are not too difficult to work out. However, corrections for risk are much more important determinants of many assets' values. For example, over the last 50 years U.S. stocks have given a real return of about 9% on average. Of this, only about 1% is due to interest rates; the remaining 8% is a premium earned for holding risk. *Uncertainty*, or *corrections for risk* make asset pricing interesting and challenging.

Asset pricing theory shares the positive versus normative tension present in the rest of economics. Does it describe the way the world *does* work, or the way the world *should* work? We observe the prices or returns of many assets. We can use the theory positively, to try to understand why prices or returns are what they are. If the world does not obey a model's predictions, we can decide that the model needs improvement. However, we can also decide that the *world* is wrong, that some assets are "mis-priced" and present trading opportunities for the shrewd investor. This latter use of asset pricing theory accounts for much of its popularity and practical application. Also, and perhaps most importantly, the prices of many assets or claims to uncertain cash flows are not observed, such as potential public or private investment projects, new financial securities, buyout prospects, and complex derivatives. We can apply the theory to establish what the prices of these claims *should* be as well; the answers are important guides to public and private decisions.

Asset pricing theory all stems from one simple concept, presented in the first page of the first chapter of this book: price equals expected discounted payoff. The rest is elaboration, special cases, and a closet

full of tricks that make the central equation useful for one or another application.

There are two polar approaches to this elaboration. I call them *absolute pricing* and *relative pricing*. In *absolute pricing*, we price each asset by reference to its exposure to fundamental sources of macroeconomic risk. The consumption-based and general equilibrium models are the purest examples of this approach. The absolute approach is most common in academic settings, in which we use asset pricing theory positively to give an economic explanation for why prices are what they are, or in order to predict how prices might change if policy or economic structure changed.

In *relative pricing*, we ask a less ambitious question. We ask what we can learn about an asset's value *given* the prices of some other assets. We do not ask where the prices of the other assets came from, and we use as little information about fundamental risk factors as possible. Black–Scholes option pricing is the classic example of this approach. While limited in scope, this approach offers precision in many applications.

Asset pricing problems are solved by judiciously choosing how much absolute and how much relative pricing one will do, depending on the assets in question and the purpose of the calculation. Almost no problems are solved by the pure extremes. For example, the CAPM and its successor factor models are paradigms of the absolute approach. Yet in applications, they price assets “relative” to the market or other risk factors, without answering what determines the market or factor risk premia and betas. The latter are treated as free parameters. On the other end of the spectrum, even the most practical financial engineering questions usually involve assumptions beyond pure lack of arbitrage, assumptions about equilibrium “market prices of risk.”

The central and unfinished task of absolute asset pricing is to understand and measure the sources of aggregate or macroeconomic risk that drive asset prices. Of course, this is also the central question of macroeconomics, and this is a particularly exciting time for researchers who want to answer these fundamental questions in macroeconomics and finance. A lot of empirical work has documented tantalizing stylized facts and links between macroeconomics and finance. For example, expected returns vary across time and across assets in ways that are linked to macroeconomic variables, or variables that also forecast macroeconomic events; a wide class of models suggests that a “recession” or “financial distress” factor lies behind many asset prices. Yet theory lags behind; we do not yet have a well-described model that explains these interesting correlations.

In turn, I think that what we are learning about finance must feed back on macroeconomics. To take a simple example, we have learned that the risk premium on stocks—the expected stock return less interest rates—is much larger than the interest rate, and varies a good deal

more than interest rates. This means that attempts to line investment up with interest rates are pretty hopeless—most variation in the cost of capital comes from the varying risk premium. Similarly, we have learned that some measure of risk aversion must be quite high, or people would all borrow like crazy to buy stocks. Most macroeconomics pursues small deviations about perfect-foresight equilibria, but the large equity premium means that volatility is a first-order effect, not a second-order effect. Standard macroeconomic models predict that people really do not care much about business cycles (Lucas [1987]). Asset prices reveal that they do—that they forego substantial return premia to avoid assets that fall in recessions. This fact ought to tell us something about recessions!

This book advocates a discount factor/generalized method of moments view of asset pricing theory and associated empirical procedures. I summarize asset pricing by two equations:

$$p_t = E(m_{t+1}x_{t+1}),$$

$$m_{t+1} = f(\text{data, parameters}),$$

where p_t = asset price, x_{t+1} = asset payoff, m_{t+1} = stochastic discount factor.

The major advantages of the discount factor/moment condition approach are its simplicity and universality. Where once there were three apparently different theories for stocks, bonds, and options, now we see each as special cases of the same theory. The common language also allows us to use insights from each field of application in other fields.

This approach allows us to conveniently separate the step of specifying economic assumptions of the model (second equation) from the step of deciding which kind of empirical representation to pursue or understand. For a given model—choice of $f(\cdot)$ —we will see how the first equation can lead to predictions stated in terms of returns, price-dividend ratios, expected return-beta representations, moment conditions, continuous versus discrete-time implications, and so forth. The ability to translate between such representations is also very helpful in digesting the results of empirical work, which uses a number of apparently distinct but fundamentally connected representations.

Thinking in terms of discount factors often turns out to be much simpler than thinking in terms of portfolios. For example, it is easier to insist that there is a positive discount factor than to check that every possible portfolio that dominates every other portfolio has a larger price, and the long arguments over the APT stated in terms of portfolios are easy to digest when stated in terms of discount factors.

The discount factor approach is also associated with a state-space geometry in place of the usual mean-variance geometry, and this book emphasizes the state-space intuition behind many classic results.

For these reasons, the discount factor language and the associated state-space geometry are common in academic research and high-tech practice. They are not yet common in textbooks, and that is the niche that this book tries to fill.

I also diverge from the usual order of presentation. Most books are structured following the history of thought: portfolio theory, mean-variance frontiers, spanning theorems, CAPM, ICAPM, APT, option pricing, and finally consumption-based model. Contingent claims are an esoteric extension of option pricing theory. I go the other way around: contingent claims and the consumption-based model are the basic and simplest models around; the others are specializations. Just because they were discovered in the opposite order is no reason to present them that way.

I also try to unify the treatment of empirical methods. A wide variety of methods are popular, including time-series and cross-sectional regressions, and methods based on generalized method of moments (GMM) and maximum likelihood. However, in the end all of these apparently different approaches do the same thing: they pick free parameters of the model to make it fit best, which usually means to minimize pricing errors; and they evaluate the model by examining how big those pricing errors are.

As with the theory, I do not attempt an encyclopedic compilation of empirical procedures. The literature on econometric methods contains lots of methods and special cases (likelihood ratio analogues of common Wald tests; cases with and without risk-free assets and when factors do and do not span the mean-variance frontier, etc.) that are seldom used in practice. I try to focus on the basic ideas and on methods that are actually used in practice.

The accent in this book is on understanding statements of theory, and working with that theory to applications, rather than rigorous or general proofs. Also, I skip very lightly over many parts of asset pricing theory that have faded from current applications, although they occupied large amounts of the attention in the past. Some examples are portfolio separation theorems, properties of various distributions, or asymptotic APT. While portfolio theory is still interesting and useful, it is no longer a cornerstone of pricing. Rather than use portfolio theory to find a demand curve for assets, which intersected with a supply curve gives prices, we now go to prices directly. One can then find optimal portfolios, but it is a side issue for the asset pricing question.

My presentation is consciously informal. I like to see an idea in its simplest form and learn to use it before going back and understanding all the foundations of the ideas. I have organized the book for similarly minded readers. If you are hungry for more formal definitions and background, keep going, they usually show up later on.

Again, my organizing principle is that everything can be traced back to specializations of the basic pricing equation $p = E(mx)$. Therefore, after reading the first chapter, one can pretty much skip around and read topics in as much depth or order as one likes. Each major subject always starts back at the same pricing equation.

The target audience for this book is economics and finance Ph.D. students, advanced MBA students, or professionals with similar background. I hope the book will also be useful to fellow researchers and finance professionals, by clarifying, relating, and simplifying the set of tools we have all learned in a hodgepodge manner. I presume some exposure to undergraduate economics and statistics. A reader should have seen a utility function, a random variable, a standard error, and a time series, should have some basic linear algebra and calculus, and should have solved a maximum problem by setting derivatives to zero. The hurdles in asset pricing are really conceptual rather than mathematical.

Asset Pricing

PART I
Asset Pricing Theory

1

Consumption-Based Model and Overview

An investor must decide how much to save and how much to consume, and what portfolio of assets to hold. The most basic pricing equation comes from the first-order condition for that decision. The marginal utility loss of consuming a little less today and buying a little more of the asset should equal the marginal utility gain of consuming a little more of the asset's payoff in the future. If the price and payoff do not satisfy this relation, the investor should buy more or less of the asset. It follows that the asset's price should equal the expected discounted value of the asset's payoff, using the investor's marginal utility to discount the payoff. With this simple idea, I present many classic issues in finance.

Interest rates are related to expected marginal utility growth, and hence to the expected path of consumption. In a time of high real interest rates, it makes sense to save, buy bonds, and then consume more tomorrow. Therefore, high real interest rates should be associated with an expectation of growing consumption.

Most importantly, risk corrections to asset prices should be driven by the covariance of asset payoffs with marginal utility and hence by the covariance of asset payoffs with consumption. Other things equal, an asset that does badly in states of nature like a recession, in which the investor feels poor and is consuming little, is less desirable than an asset that does badly in states of nature like a boom in which the investor feels wealthy and is consuming a great deal. The former asset will sell for a lower price; its price will reflect a discount for its "riskiness," and this riskiness depends on a *covariance*, not a *variance*.

Marginal utility, not consumption, is the fundamental measure of how you feel. Most of the theory of asset pricing is about how to go from marginal utility to observable indicators. Consumption is low when marginal utility is high, of course, so consumption may be a useful indicator. Consumption is also low and marginal utility is high when the

investor's other assets have done poorly; thus we may expect that prices are low for assets that covary positively with a large index such as the market portfolio. This is a Capital Asset Pricing Model. We will see a wide variety of additional indicators for marginal utility, things against which to compute a covariance in order to predict the risk-adjustment for prices.

1.1 Basic Pricing Equation

An investor's first-order conditions give the basic consumption-based model,

$$p_t = E_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} x_{t+1} \right].$$

Our basic objective is to figure out the value of any stream of uncertain cash flows. I start with an apparently simple case, which turns out to capture very general situations.

Let us find the value at time t of a *payoff* x_{t+1} . If you buy a stock today, the payoff next period is the stock price plus dividend, $x_{t+1} = p_{t+1} + d_{t+1}$. x_{t+1} is a random variable: an investor does not know exactly how much he will get from his investment, but he can assess the probability of various possible outcomes. Do not confuse the *payoff* x_{t+1} with the *profit* or *return*; x_{t+1} is the value of the investment at time $t + 1$, without subtracting or dividing by the cost of the investment.

We find the value of this payoff by asking what it is worth to a typical investor. To do this, we need a convenient mathematical formalism to capture what an investor wants. We model investors by a *utility function* defined over current and future values of consumption,

$$U(c_t, c_{t+1}) = u(c_t) + \beta E_t[u(c_{t+1})],$$

where c_t denotes consumption at date t . We often use a convenient power utility form,

$$u(c_t) = \frac{1}{1-\gamma} c_t^{1-\gamma}.$$

The limit as $\gamma \rightarrow 1$ is

$$u(c) = \ln(c).$$

The utility function captures the fundamental desire for more *consumption*, rather than posit a desire for intermediate objectives such as mean and variance of portfolio returns. Consumption c_{t+1} is also random; the investor does not know his wealth tomorrow, and hence how much he will decide to consume tomorrow. The period utility function $u(\cdot)$ is increasing, reflecting a desire for more consumption, and concave, reflecting the declining marginal value of additional consumption. The last bite is never as satisfying as the first.

This formalism captures investors' impatience and their aversion to risk, so we can quantitatively correct for the risk and delay of cash flows. Discounting the future by β captures impatience, and β is called the *subjective discount factor*. The curvature of the utility function generates aversion to risk and to intertemporal substitution: The investor prefers a consumption stream that is steady over time and across states of nature.

Now, assume that the investor can freely buy or sell as much of the payoff x_{t+1} as he wishes, at a price p_t . How much will he buy or sell? To find the answer, denote by e the original consumption level (if the investor bought none of the asset), and denote by ξ the amount of the asset he chooses to buy. Then, his problem is

$$\max_{\{\xi\}} u(c_t) + E_t[\beta u(c_{t+1})] \quad s.t.$$

$$c_t = e_t - p_t \xi,$$

$$c_{t+1} = e_{t+1} + x_{t+1} \xi.$$

Substituting the constraints into the objective, and setting the derivative with respect to ξ equal to zero, we obtain the first-order condition for an optimal consumption and portfolio choice,

$$p_t u'(c_t) = E_t[\beta u'(c_{t+1}) x_{t+1}], \quad (1.1)$$

or

$$p_t = E_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} x_{t+1} \right]. \quad (1.2)$$

The investor buys more or less of the asset until this first-order condition holds.

Equation (1.1) expresses the standard marginal condition for an optimum: $p_t u'(c_t)$ is the loss in utility if the investor buys another unit of the asset; $E_t[\beta u'(c_{t+1}) x_{t+1}]$ is the increase in (discounted, expected) utility he obtains from the extra payoff at $t+1$. The investor continues to buy or sell the asset until the marginal loss equals the marginal gain.

Equation (1.2) is *the* central asset pricing formula. Given the payoff x_{t+1} and given the investor's consumption choice c_t, c_{t+1} , it tells you what market price p_t to expect. Its economic content is simply the first-order conditions for optimal consumption and portfolio formation. Most of the theory of asset pricing just consists of specializations and manipulations of this formula.

We have stopped short of a complete solution to the model, i.e., an expression with exogenous items on the right-hand side. We relate one endogenous variable, price, to two other endogenous variables, consumption and payoffs. One can continue to solve this model and derive the optimal consumption choice c_t, c_{t+1} in terms of more fundamental givens of the model. In the model I have sketched so far, those givens are the income sequence e_t, e_{t+1} and a specification of the full set of assets that the investor may buy and sell. We will in fact study such fuller solutions below. However, for many purposes one can stop short of specifying (possibly wrongly) all this extra structure, and obtain very useful predictions about asset prices from (1.2), even though consumption is an endogenous variable.

1.2 Marginal Rate of Substitution/Stochastic Discount Factor

We break up the basic consumption-based pricing equation into

$$p = E(mx),$$

$$m = \beta \frac{u'(c_{t+1})}{u'(c_t)},$$

where m_{t+1} is the *stochastic discount factor*.

A convenient way to break up the basic pricing equation (1.2) is to define the *stochastic discount factor* m_{t+1}

$$m_{t+1} \equiv \beta \frac{u'(c_{t+1})}{u'(c_t)}. \quad (1.3)$$

Then, the basic pricing formula (1.2) can simply be expressed as

$$p_t = E_t(m_{t+1}x_{t+1}). \quad (1.4)$$

When it is not necessary to be explicit about time subscripts or the difference between conditional and unconditional expectation, I will suppress the subscripts and just write $p = E(mx)$. The price always comes

at t , the payoff at $t + 1$, and the expectation is conditional on time- t information.

The term *stochastic discount factor* refers to the way m generalizes standard discount factor ideas. If there is no uncertainty, we can express prices via the standard present value formula

$$p_t = \frac{1}{R^f} x_{t+1}, \quad (1.5)$$

where R^f is the gross risk-free rate. $1/R^f$ is the *discount factor*. Since gross interest rates are typically greater than one, the payoff x_{t+1} sells “at a discount.” Riskier assets have lower prices than equivalent risk-free assets, so they are often valued by using risk-adjusted discount factors,

$$p_t^i = \frac{1}{R^i} E_t(x_{t+1}^i).$$

Here, I have added the i superscript to emphasize that each risky asset i must be discounted by an asset-specific risk-adjusted discount factor $1/R^i$.

In this context, equation (1.4) is obviously a generalization, and it says something deep: one can incorporate all risk corrections by defining a *single* stochastic discount factor—the same one for each asset—and putting it inside the expectation. m_{t+1} is *stochastic* or *random* because it is not known with certainty at time t . The correlation between the random components of the common discount factor m and the asset-specific payoff x^i generate asset-specific risk corrections.

m_{t+1} is also often called the *marginal rate of substitution* after (1.3). In that equation, m_{t+1} is the rate at which the investor is willing to substitute consumption at time $t + 1$ for consumption at time t . m_{t+1} is sometimes also called the *pricing kernel*. If you know what a kernel is and you express the expectation as an integral, you can see where the name comes from. It is sometimes called a *change of measure* or a *state-price density*.

For the moment, introducing the discount factor m and breaking the basic pricing equation (1.2) into (1.3) and (1.4) is just a notational convenience. However, it represents a much deeper and more useful separation. For example, notice that $p = E(mx)$ would still be valid if we changed the utility function, but we would have a different function connecting m to data. All asset pricing models amount to alternative ways of connecting the stochastic discount factor to data. At the same time, we will study lots of alternative expressions of $p = E(mx)$, and we can summarize many empirical approaches by applying them to $p = E(mx)$. By separating our models into these two components, we do not have to redo all that elaboration for each asset pricing model.

1.3 Prices, Payoffs, and Notation

The price p_t gives rights to a payoff x_{t+1} . In practice, this notation covers a variety of cases, including the following:

	Price p_t	Payoff x_{t+1}
Stock	p_t	$p_{t+1} + d_{t+1}$
Return	1	R_{t+1}
Price-dividend ratio	$\frac{p_t}{d_t}$	$\left(\frac{p_{t+1}}{d_{t+1}} + 1\right) \frac{d_{t+1}}{d_t}$
Excess return	0	$R_{t+1}^e = R_{t+1}^a - R_{t+1}^b$
Managed portfolio	z_t	$z_t R_{t+1}$
Moment condition	$E(p_t z_t)$	$x_{t+1} z_t$
One-period bond	p_t	1
Risk-free rate	1	R^f
Option	C	$\max(S_T - K, 0)$

The price p_t and payoff x_{t+1} seem like a very restrictive kind of security. In fact, this notation is quite general and allows us easily to accommodate many different asset pricing questions. In particular, we can cover stocks, bonds, and options and make clear that there is one theory for all asset pricing.

For stocks, the one-period payoff is of course the next price plus dividend, $x_{t+1} = p_{t+1} + d_{t+1}$. We frequently divide the payoff x_{t+1} by the price p_t to obtain a *gross return*

$$R_{t+1} \equiv \frac{x_{t+1}}{p_t}.$$

We can think of a return as a payoff with price one. If you pay one dollar today, the return is how many dollars or units of consumption you get tomorrow. Thus, returns obey

$$1 = E(mR),$$

which is by far the most important special case of the basic formula $p = E(mx)$. I use capital letters to denote *gross returns* R , which have

a numerical value like 1.05. I use lowercase letters to denote *net* returns $r = R - 1$ or log (continuously compounded) returns $r = \ln(R)$, both of which have numerical values like 0.05. One may also quote *percent* returns $100 \times r$.

Returns are often used in empirical work because they are typically stationary over time. (Stationary in the statistical sense; they do not have trends and you can meaningfully take an average. “Stationary” does not mean constant.) However, thinking in terms of returns takes us away from the central task of finding asset *prices*. Dividing by dividends and creating a payoff of the form

$$x_{t+1} = \left(1 + \frac{p_{t+1}}{d_{t+1}}\right) \frac{d_{t+1}}{d_t}$$

corresponding to a price p_t/d_t is a way to look at prices but still to examine stationary variables.

Not everything can be reduced to a return. If you borrow a dollar at the interest rate R^f and invest it in an asset with return R , you pay no money out-of-pocket today, and get the payoff $R - R^f$. This is a payoff with a *zero* price, so you obviously cannot divide payoff by price to get a return. Zero price does not imply zero payoff. It is a bet in which the value of the chance of losing exactly balances the value of the chance of winning, so that no money changes hands when the bet is made. It is common to study equity strategies in which one short-sells one stock or portfolio and invests the proceeds in another stock or portfolio, generating an excess return. I denote any such difference between returns as an *excess return*, R^e . It is also called a *zero-cost portfolio*.

In fact, much asset pricing focuses on excess returns. Our economic understanding of interest rate variation turns out to have little to do with our understanding of risk premia, so it is convenient to separate the two phenomena by looking at interest rates and excess returns separately.

We also want to think about the *managed portfolios*, in which one invests more or less in an asset according to some signal. The “price” of such a strategy is the amount invested at time t , say z_t , and the payoff is $z_t R_{t+1}$. For example, a market timing strategy might make an investment in stocks proportional to the price-dividend ratio, investing less when prices are higher. We could represent such a strategy as a payoff using $z_t = a - b(p_t/d_t)$.

When we think about conditioning information below, we will think of objects like z_t as *instruments*. Then we take an unconditional expectation of $p_t z_t = E_t(m_{t+1} x_{t+1}) z_t$, yielding $E(p_t z_t) = E(m_{t+1} x_{t+1} z_t)$. We can think of this operation as creating a “security” with payoff $x_{t+1} z_{t+1}$, and “price” $E(p_t z_t)$ represented with unconditional expectations.

A one-period bond is of course a claim to a unit payoff. Bonds, options, investment projects are all examples in which it is often more useful to think of prices and payoffs rather than returns.

Prices and returns can be real (denominated in goods) or nominal (denominated in dollars); $p = E(mx)$ can refer to either case. The only difference is whether we use a real or nominal discount factor. If prices, returns, and payoffs are nominal, we should use a nominal discount factor. For example, if p and x denote nominal values, then we can create real prices and payoffs to write

$$\frac{p_t}{\Pi_t} = E_t \left[\left(\beta \frac{u'(c_{t+1})}{u'(c_t)} \right) \frac{x_{t+1}}{\Pi_{t+1}} \right],$$

where Π denotes the price level (cpi). Obviously, this is the same as defining a nominal discount factor by

$$p_t = E_t \left[\left(\beta \frac{u'(c_{t+1})}{u'(c_t)} \frac{\Pi_t}{\Pi_{t+1}} \right) x_{t+1} \right].$$

To accommodate all these cases, I will simply use the notation price p_t and payoff x_{t+1} . These symbols can denote 0, 1, or z_t and R_t^e, r_{t+1} , or $z_t R_{t+1}$, respectively, according to the case. Lots of other definitions of p and x are useful as well.

1.4 Classic Issues in Finance

I use simple manipulations of the basic pricing equation to introduce classic issues in finance: the economics of interest rates, risk adjustments, systematic versus idiosyncratic risk, expected return-beta representations, the mean-variance frontier, the slope of the mean-variance frontier, time-varying expected returns, and present-value relations.

A few simple rearrangements and manipulations of the basic pricing equation $p = E(mx)$ give a lot of intuition and introduce some classic issues in finance, including determinants of the interest rate, risk corrections, idiosyncratic versus systematic risk, beta pricing models, and mean-variance frontiers.

Risk-Free Rate

The risk-free rate is related to the discount factor by

$$R^f = 1/E(m).$$

With lognormal consumption growth,

$$r_t^f = \delta + \gamma E_t(\Delta \ln c_{t+1}) - \frac{\gamma^2}{2} \sigma_t^2(\Delta \ln c_{t+1}).$$

Real interest rates are high when people are impatient (δ), when expected consumption growth is high (intertemporal substitution), or when risk is low (precautionary saving). A more curved utility function (γ) or a lower elasticity of intertemporal substitution ($1/\gamma$) means that interest rates are more sensitive to changes in expected consumption growth.

The risk-free rate is given by

$$R^f = 1/E(m). \tag{1.6}$$

The risk-free rate is known ahead of time, so $p = E(mx)$ becomes $1 = E(mR^f) = E(m)R^f$.

If a risk-free security is not traded, we can define $R^f = 1/E(m)$ as the “shadow” risk-free rate. In some models it is called the “zero-beta” rate. If one introduced a risk-free security with return $R^f = 1/E(m)$, investors would be just indifferent to buying or selling it. I use R^f to simplify formulas below with this understanding.

To think about the economics behind real interest rates in a simple setup, use power utility $u'(c) = c^{-\gamma}$. Start by turning off uncertainty, in which case

$$R^f = \frac{1}{\beta} \left(\frac{c_{t+1}}{c_t} \right)^\gamma.$$

We can see three effects right away:

1. Real interest rates are high when people are impatient, i.e. when β is low. If everyone wants to consume now, it takes a high interest rate to convince them to save.
2. Real interest rates are high when consumption growth is high. In times of high interest rates, it pays investors to consume less now, invest more, and consume more in the future. Thus, high interest rates lower the level of consumption today, while raising its growth rate from today to tomorrow.

3. Real interest rates are more sensitive to consumption growth if the power parameter γ is large. If utility is highly curved, the investor cares more about maintaining a consumption profile that is smooth over time, and is less willing to rearrange consumption over time in response to interest rate incentives. Thus it takes a larger interest rate change to induce him to a given consumption growth.

To understand how interest rates behave when there is some uncertainty, I specify that consumption growth is lognormally distributed. In this case, the real risk-free rate equation becomes

$$r_t^f = \delta + \gamma E_t(\Delta \ln c_{t+1}) - \frac{\gamma^2}{2} \sigma_t^2(\Delta \ln c_{t+1}), \quad (1.7)$$

where I have defined the log risk-free rate r_t^f and subjective discount rate δ by

$$r_t^f = \ln R_t^f; \quad \beta = e^{-\delta},$$

and Δ denotes the first difference operator,

$$\Delta \ln c_{t+1} = \ln c_{t+1} - \ln c_t.$$

To derive expression (1.7) for the risk-free rate, start with

$$R_t^f = 1/E_t \left[\beta \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma} \right].$$

Using the fact that normal z means

$$E(e^z) = e^{E(z) + (1/2)\sigma^2(z)}$$

(you can check this by writing out the integral that defines the expectation), we have

$$R_t^f = \left[e^{-\delta} e^{-\gamma E_t(\Delta \ln c_{t+1}) + (\gamma^2/2)\sigma_t^2(\Delta \ln c_{t+1})} \right]^{-1}.$$

Then take logarithms. The combination of lognormal distributions and power utility is one of the basic tricks to getting analytical solutions in this kind of model. Section 1.5 shows how to get the same result in continuous time.

Looking at (1.7), we see the same results as we had with the deterministic case. Real interest rates are high when impatience δ is high and when consumption growth is high; higher γ makes interest rates more sensitive to consumption growth. The new σ^2 term captures *precautionary savings*. When consumption is more volatile, people with this utility function are

more worried about the low consumption states than they are pleased by the high consumption states. Therefore, people want to save more, driving down interest rates.

We can also read the same terms backwards: consumption growth is high when real interest rates are high, since people save more now and spend it in the future, and consumption is less sensitive to interest rates as the desire for a smooth consumption stream, captured by γ , rises. Section 2.2 takes up the question of which way we should read this equation—as consumption determining interest rates, or as interest rates determining consumption.

For the power utility function, the curvature parameter γ simultaneously controls intertemporal substitution—aversion to a consumption stream that varies over time, risk aversion—aversion to a consumption stream that varies across states of nature, and precautionary savings, which turns out to depend on the third derivative of the utility function. This link is particular to the power utility function. More general utility functions loosen the links between these three quantities.

Risk Corrections

Payoffs that are positively correlated with consumption growth have lower prices, to compensate investors for risk.

$$p = \frac{E(x)}{R^f} + \text{cov}(m, x),$$

$$E(R^i) - R^f = -R^f \text{cov}(m, R^i).$$

Expected returns are proportional to the covariance of returns with discount factors.

Using the definition of covariance $\text{cov}(m, x) = E(mx) - E(m)E(x)$, we can write $p = E(mx)$ as

$$p = E(m)E(x) + \text{cov}(m, x). \quad (1.8)$$

Substituting the risk-free rate equation (1.6), we obtain

$$p = \frac{E(x)}{R^f} + \text{cov}(m, x). \quad (1.9)$$

The first term in (1.9) is the standard discounted present-value formula. This is the asset's price in a risk-neutral world—if consumption is constant or if utility is linear. The second term is a *risk adjustment*. An

asset whose payoff covaries positively with the discount factor has its price raised and vice versa.

To understand the risk adjustment, substitute back for m in terms of consumption, to obtain

$$p = \frac{E(x)}{R^f} + \frac{\text{cov}[\beta u'(c_{t+1}), x_{t+1}]}{u'(c_t)}. \quad (1.10)$$

Marginal utility $u'(c)$ declines as c rises. Thus, an asset's price is lowered if its payoff covaries positively with consumption. Conversely, an asset's price is raised if it covaries negatively with consumption.

Why? Investors do not like uncertainty about consumption. If you buy an asset whose payoff covaries positively with consumption, one that pays off well when you are already feeling wealthy, and pays off badly when you are already feeling poor, that asset will make your consumption stream more volatile. You will require a low price to induce you to buy such an asset. If you buy an asset whose payoff covaries negatively with consumption, it helps to smooth consumption and so is more valuable than its expected payoff might indicate. Insurance is an extreme example. Insurance pays off exactly when wealth and consumption would otherwise be low—you get a check when your house burns down. For this reason, you are happy to hold insurance, even though you expect to lose money—even though the price of insurance is greater than its expected payoff discounted at the risk-free rate.

To emphasize why the *covariance* of a payoff with the discount factor rather than its *variance* determines its riskiness, keep in mind that the investor cares about the volatility of consumption. He does *not* care about the volatility of his individual assets or of his portfolio, if he can keep a steady consumption. Consider then what happens to the volatility of consumption if the investor buys a little more ξ of payoff x . $\sigma^2(c)$ becomes

$$\sigma^2(c + \xi x) = \sigma^2(c) + 2\xi \text{cov}(c, x) + \xi^2 \sigma^2(x).$$

For small (marginal) portfolio changes, the *covariance* between consumption and payoff determines the effect of adding a bit more of each payoff on the volatility of consumption.

We use returns so often that it is worth restating the same intuition for the special case that the price is 1 and the payoff is a return. Start with the basic pricing equation for returns,

$$1 = E(mR^i).$$

I denote the return R^i to emphasize that the point of the theory is to distinguish the behavior of one asset R^i from another R^j .

The asset pricing model says that, although expected *returns* can vary across time and assets, expected *discounted* returns should always be the same, 1. Applying the covariance decomposition,

$$1 = E(m)E(R^i) + \text{cov}(m, R^i) \quad (1.11)$$

and, using $R^f = 1/E(m)$,

$$E(R^i) - R^f = -R^f \text{cov}(m, R^i) \quad (1.12)$$

or

$$E(R^i) - R^f = -\frac{\text{cov}[u'(c_{t+1}), R_{t+1}^i]}{E[u'(c_{t+1})]}. \quad (1.13)$$

All assets have an expected return equal to the risk-free rate, plus a risk adjustment. Assets whose returns covary positively with consumption make consumption more volatile, and so must promise higher expected returns to induce investors to hold them. Conversely, assets that covary negatively with consumption, such as insurance, can offer expected rates of return that are lower than the risk-free rate, or even negative (net) expected returns.

Much of finance focuses on expected returns. We think of expected returns increasing or decreasing to clear markets; we offer intuition that “riskier” securities must offer higher expected returns to get investors to hold them, rather than saying “riskier” securities trade for lower prices so that investors will hold them. Of course, a low initial price for a given payoff corresponds to a high expected return, so this is no more than a different language for the same phenomenon.

Idiosyncratic Risk Does Not Affect Prices

Only the component of a payoff perfectly correlated with the discount factor generates an extra return. *Idiosyncratic* risk, uncorrelated with the discount factor, generates no premium.

You might think that an asset with a volatile payoff is “risky” and thus should have a large risk correction. However, if the payoff is uncorrelated with the discount factor m , the asset receives *no* risk correction to its price,

and pays an expected return equal to the risk-free rate! In equations, if

$$\text{cov}(m, x) = 0,$$

then

$$p = \frac{E(x)}{R^f},$$

no matter how large $\sigma^2(x)$. This prediction holds even if the payoff x is highly volatile and investors are highly risk averse. The reason is simple: if you buy a little bit more of such an asset, it has no first-order effect on the variance of your consumption stream.

More generally, one gets no compensation or risk adjustment for holding *idiosyncratic* risk. Only *systematic* risk generates a risk correction. To give meaning to these words, we can decompose any payoff x into a part correlated with the discount factor and an idiosyncratic part uncorrelated with the discount factor by running a regression,

$$x = \text{proj}(x|m) + \varepsilon.$$

Then, the price of the residual or idiosyncratic risk ε is zero, and the price of x is the same as the price of its projection on m . The projection of x on m is of course that part of x which is perfectly correlated with m . The *idiosyncratic* component of any payoff is that part uncorrelated with m . Thus only the systematic *part* of a payoff accounts for its price.

Projection means linear regression without a constant,

$$\text{proj}(x|m) = \frac{E(mx)}{E(m^2)}m.$$

You can verify that regression residuals are orthogonal to right-hand variables $E(m\varepsilon) = 0$ from this definition. $E(m\varepsilon) = 0$ of course means that the price of ε is zero,

$$p(\text{proj}(x|m)) = p\left(\frac{E(mx)}{E(m^2)}m\right) = E\left(m^2 \frac{E(mx)}{E(m^2)}\right) = E(mx) = p(x).$$

The words “systematic” and “idiosyncratic” are defined differently in different contexts, which can lead to some confusion. In this decomposition, the residuals ε can be correlated with each other, though they are not correlated with the discount factor. The APT starts with a factor-analytic decomposition of the covariance of payoffs, and the word “idiosyncratic” there is reserved for the component of payoffs uncorrelated with all of the other payoffs.

Expected Return-Beta Representation

We can write $p = E(mx)$ as

$$E(R^i) = R^f + \beta_{i,m} \lambda_m.$$

We can express the expected return equation (1.12), for a return R^i , as

$$E(R^i) = R^f + \left(\frac{\text{cov}(R^i, m)}{\text{var}(m)} \right) \left(- \frac{\text{var}(m)}{E(m)} \right) \quad (1.14)$$

or

$$E(R^i) = R^f + \beta_{i,m} \lambda_m, \quad (1.15)$$

where β_{im} is the regression coefficient of the return R^i on m . This is a *beta pricing model*. It says that each expected return should be proportional to the regression coefficient, or beta, in a regression of that return on the discount factor m . Notice that the coefficient λ_m is the same for all assets i , while the $\beta_{i,m}$ varies from asset to asset. The λ_m is often interpreted as the *price of risk* and the β as the *quantity* of risk in each asset. As you can see, the price of risk λ_m depends on the volatility of the discount factor.

Obviously, there is nothing deep about saying that expected returns are proportional to betas rather than to covariances. There is a long historical tradition and some minor convenience in favor of betas. The betas refer to the projection of R on m that we studied above, so you see again a sense in which only the systematic component of risk matters.

With $m = \beta(c_{t+1}/c_t)^{-\gamma}$, we can take a Taylor approximation of equation (1.14) to express betas in terms of a more concrete variable, consumption growth, rather than marginal utility. The result, which I derive more explicitly and conveniently in the continuous-time limit below, is

$$E(R^i) = R^f + \beta_{i,\Delta c} \lambda_{\Delta c}, \quad (1.16)$$

$$\lambda_{\Delta c} = \gamma \text{ var}(\Delta c).$$

Expected returns should increase linearly with their betas on consumption growth itself. In addition, though it is treated as a free parameter in many applications, the factor risk premium $\lambda_{\Delta c}$ is determined by risk aversion and the volatility of consumption. The more risk averse people are, or the riskier their environment, the larger an expected return premium one must pay to get investors to hold risky (high beta) assets.

Mean-Variance Frontier

All asset returns lie inside a mean-variance frontier. Assets on the frontier are perfectly correlated with each other and with the discount factor. Returns on the frontier can be generated as portfolios of any two frontier returns. We can construct a discount factor from any frontier return (except R^f), and an expected return-beta representation holds using any frontier return (except R^f) as the factor.

Asset pricing theory has focused a lot on the means and variances of asset returns. Interestingly, the set of means and variances of returns is limited. All assets priced by the discount factor m must obey

$$|E(R^i) - R^f| \leq \frac{\sigma(m)}{E(m)} \sigma(R^i). \quad (1.17)$$

To derive (1.17) write for a given asset return R^i

$$1 = E(mR^i) = E(m)E(R^i) + \rho_{m, R^i} \sigma(R^i) \sigma(m)$$

and hence

$$E(R^i) = R^f - \rho_{m, R^i} \frac{\sigma(m)}{E(m)} \sigma(R^i). \quad (1.18)$$

Correlation coefficients cannot be greater than 1 in magnitude, leading to (1.17).

This simple calculation has many interesting and classic implications.

1. Means and variances of asset returns must lie in the wedge-shaped region illustrated in Figure 1.1. The boundary of the mean-variance region in which assets can lie is called the *mean-variance frontier*. It answers a naturally interesting question, “how much mean return can you get for a given level of variance?”
2. All returns on the frontier are perfectly correlated with the discount factor: the frontier is generated by $|\rho_{m, R^i}| = 1$. Returns on the upper part of the frontier are perfectly negatively correlated with the discount factor and hence positively correlated with consumption. They are “maximally risky” and thus get the highest expected returns. Returns on the lower part of the frontier are perfectly positively correlated with the discount factor and hence negatively correlated with consumption. They thus provide the best insurance against consumption fluctuations.

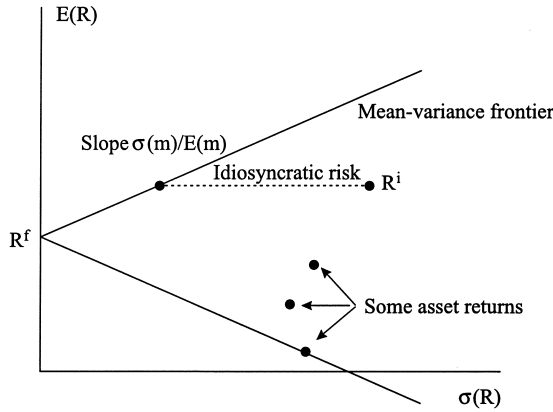


Figure 1.1. Mean-variance frontier. The mean and standard deviation of all assets priced by a discount factor m must lie in the wedge-shaped region.

3. All frontier returns are also perfectly correlated with each other, since they are all perfectly correlated with the discount factor. This fact implies that we can *span* or *synthesize* any frontier return from two such returns. For example, if you pick any single frontier return R^m , then all frontier returns R^{mv} must be expressible as

$$R^{mv} = R^f + a(R^m - R^f)$$

for some number a .

4. Since each point on the mean-variance frontier is perfectly correlated with the discount factor, we must be able to pick constants a, b, d, e such that

$$m = a + bR^{mv},$$

$$R^{mv} = d + em.$$

Thus, *any mean-variance efficient return carries all pricing information*. Given a mean-variance efficient return and the risk-free rate, we can find a discount factor that prices all assets and vice versa.

5. Given a discount factor, we can also construct a single-beta representation, so *expected returns can be described in a single-beta representation using any mean-variance efficient return* (except the risk-free rate),

$$E(R^i) = R^f + \beta_{i,mv}[E(R^{mv}) - R^f].$$

The essence of the beta pricing model is that, even though the means and standard deviations of returns fill out the space inside the mean-variance frontier, a graph of mean returns versus *betas* should yield a straight line. Since the beta model applies to every return including R^{mv} itself, and R^{mv} has a beta of 1 on itself, we can identify the factor risk premium as $\lambda = E(R^{mv} - R^f)$.

The last two points suggest an intimate relationship between discount factors, beta models, and mean-variance frontiers. I explore this relation in detail in Chapter 6. A problem at the end of this chapter guides you through the algebra to demonstrate points 4 and 5 explicitly.

6. We can plot the decomposition of a return into a “priced” or “systematic” component and a “residual,” or “idiosyncratic” component as shown in Figure 1.1. The priced part is perfectly correlated with the discount factor, and hence perfectly correlated with any frontier return. The residual or idiosyncratic part generates no expected return, so it lies flat as shown in the figure, and it is uncorrelated with the discount factor or any frontier return. Assets inside the frontier or even on the lower portion of the frontier are not “worse” than assets on the frontier. The frontier and its internal region characterize equilibrium asset returns, with rational investors happy to hold all assets. You would not want to put your whole portfolio in one “inefficient” asset, but you are happy to put some wealth in such assets.

*Slope of the Mean-Standard Deviation
Frontier and Equity Premium Puzzle*

The Sharpe ratio is limited by the volatility of the discount factor. The maximal risk-return trade-off is steeper if there is more risk or more risk aversion,

$$\left| \frac{E(R) - R^f}{\sigma(R)} \right| \leq \frac{\sigma(m)}{E(m)} \approx \gamma \sigma(\Delta \ln c).$$

This formula captures the equity premium puzzle, which suggests that either people are very risk averse, or the stock returns of the last 50 years were good luck which will not continue.

The ratio of mean excess return to standard deviation

$$\frac{E(R^i) - R^f}{\sigma(R^i)}$$

is known as the *Sharpe ratio*. It is a more interesting characterization of a security than the mean return alone. If you borrow and put more money

into a security, you can increase the mean return of your position, but you do not increase the Sharpe ratio, since the standard deviation increases at the same rate as the mean.

The slope of the mean-standard deviation frontier is the largest available Sharpe ratio, and thus is naturally interesting. It answers “how much more mean return can I get by shouldering a bit more volatility in my portfolio?”

Let R^{mv} denote the return of a portfolio on the frontier. From equation (1.17), the slope of the frontier is

$$\left| \frac{E(R^{mv}) - R^f}{\sigma(R^{mv})} \right| = \frac{\sigma(m)}{E(m)} = \sigma(m)R^f.$$

Thus, the slope of the frontier is governed by the volatility of the discount factor.

For an economic interpretation, again consider the power utility function, $u'(c) = c^{-\gamma}$,

$$\left| \frac{E(R^{mv}) - R^f}{\sigma(R^{mv})} \right| = \frac{\sigma[(c_{t+1}/c_t)^{-\gamma}]}{E[(c_{t+1}/c_t)^{-\gamma}]}. \quad (1.19)$$

The standard deviation on the right hand side is large if consumption is volatile or if γ is large. We can state this approximation precisely using the lognormal assumption. If consumption growth is lognormal,

$$\left| \frac{E(R^{mv}) - R^f}{\sigma(R^{mv})} \right| = \sqrt{e^{\gamma^2 \sigma^2(\Delta \ln c_{t+1})} - 1} \approx \gamma \sigma(\Delta \ln c). \quad (1.20)$$

(A problem at the end of the chapter guides you through the algebra of the first equality. The relation is exact in continuous time, and thus the approximation is easiest to derive by reference to the continuous-time result; see Section 1.5.)

Reading the equation, *the slope of the mean-standard deviation frontier is higher if the economy is riskier—if consumption is more volatile—or if investors are more risk averse.* Both situations naturally make investors more reluctant to take on the extra risk of holding risky assets. Both situations also raise the slope of the expected return-beta line of the consumption beta model, (1.16). (Or, conversely, in an economy with a high Sharpe ratio, low risk-averse investors should take on so much risk that their consumption becomes volatile.)

In postwar U.S. data, the slope of the historical mean-standard deviation frontier, or of average return-beta lines, is much higher than reasonable risk aversion and consumption volatility estimates suggest. This is the “equity premium puzzle.” Over the last 50 years in the United States,

real stock returns have averaged 9% with a standard deviation of about 16%, while the real return on treasury bills has been about 1%. Thus, the historical annual market Sharpe ratio has been about 0.5. Aggregate nondurable and services consumption growth had a mean and standard deviation of about 1%. We can only reconcile these facts with (1.20) if investors have a risk-aversion coefficient of 50!

Obvious ways of generalizing the calculation just make matters worse. Equation (1.20) relates consumption growth to the mean-variance frontier of all contingent claims. Market indices with 0.5 Sharpe ratios are if anything inside that frontier, so recognizing market incompleteness makes matters worse. Aggregate consumption has about 0.2 correlation with the market return, while the equality (1.20) takes the worst possible case that consumption growth and asset returns are perfectly correlated. If you add this fact, you need risk aversion of 250 to explain the market Sharpe ratio! Individuals have riskier consumption streams than aggregate, but as their risk goes up their correlation with any aggregate must decrease proportionally, so to first order recognizing individual risk will not help either.

Clearly, either (1) people are a *lot* more risk averse than we might have thought, (2) the stock returns of the last 50 years were largely good luck rather than an equilibrium compensation for risk, or (3) something is deeply wrong with the model, including the utility function and use of aggregate consumption data. This “equity premium puzzle” has attracted the attention of a lot of research in finance, especially on the last item. I return to the equity premium in more detail in Chapter 21.

Random Walks and Time-Varying Expected Returns

If investors are risk neutral, returns are unpredictable, and prices follow martingales. In general, prices scaled by marginal utility are martingales, and returns can be predictable if investors are risk averse and if the conditional second moments of returns and discount factors vary over time. This is more plausible at long horizons.

So far, we have concentrated on the behavior of prices or expected returns across assets. We should also consider the behavior of the price or return of a given asset over time. Going back to the basic first-order condition,

$$p_t u'(c_t) = E_t[\beta u'(c_{t+1})(p_{t+1} + d_{t+1})]. \quad (1.21)$$

If investors are risk neutral, i.e., if $u(c)$ is linear or there is no variation in consumption, if the security pays no dividends between t and $t + 1$, and

for short time horizons where β is close to 1, this equation reduces to

$$p_t = E_t(p_{t+1}).$$

Equivalently, prices follow a time-series process of the form

$$p_{t+1} = p_t + \varepsilon_{t+1}.$$

If the variance $\sigma_t^2(\varepsilon_{t+1})$ is constant, prices follow a *random walk*. More generally, prices follow a *martingale*. Intuitively, if the price today is a lot lower than investors' expectations of the price tomorrow, then investors will try to buy the security. But this action will drive up the price of the security until the price today does equal the expected price tomorrow. Another way of saying the same thing is that returns should not be predictable; dividing by p_t , expected returns $E_t(p_{t+1}/p_t) = 1$ should be constant; returns should be like coin flips.

The more general equation (1.21) says that prices should follow a martingale after adjusting for dividends and scaling by marginal utility. Since martingales have useful mathematical properties, and since risk neutrality is such a simple economic environment, many asset pricing results are easily derived by scaling prices and dividends by discounted marginal utility first, and then using "risk-neutral" formulas and risk-neutral economic arguments.

Since consumption and risk aversion do not change much day to day, we might expect the random walk view to hold pretty well on a day-to-day basis. This idea contradicts the still popular notion that there are "systems" or "technical analysis" by which one can predict where stock prices are going on any given day. The random walk view has been remarkably successful. Despite decades of dredging the data, and the popularity of media reports that purport to explain where markets are going, trading rules that reliably survive transactions costs and do not implicitly expose the investor to risk have not yet been reliably demonstrated.

However, more recently, evidence has accumulated that long-horizon excess returns are quite predictable, and to some this evidence indicates that the whole enterprise of economic explanation of asset returns is flawed. To think about this issue, write our basic equation for expected returns as

$$\begin{aligned} E_t(R_{t+1}) - R_t^f &= -\frac{\text{cov}_t(m_{t+1}, R_{t+1})}{E_t(m_{t+1})} \\ &= \frac{\sigma_t(m_{t+1})}{E_t(m_{t+1})} \sigma_t(R_{t+1}) \rho_t(m_{t+1}, R_{t+1}) \\ &\approx \gamma_t \sigma_t(\Delta c_{t+1}) \sigma_t(R_{t+1}) \rho_t(m_{t+1}, R_{t+1}). \end{aligned} \quad (1.22)$$

I include the t subscripts to emphasize that the relation applies to *conditional* moments. Sometimes, the *conditional* mean or other moment of a random variable is different from its *unconditional* moment. Conditional on tonight's weather forecast, you can better predict rain tomorrow than just knowing the average rain for that date. In the special case that random variables are i.i.d. (independent and identically distributed), like coin flips, the conditional and unconditional moments are the same, but that is a special case and not likely to be true of asset prices, returns, and macroeconomic variables. In the theory so far, we have thought of an investor, today, forming expectations of payoffs, consumption, and other variables tomorrow. Thus, the moments are really all *conditional*, and if we want to be precise we should include some notation to express this fact. I use subscripts $E_t(x_{t+1})$ to denote conditional expectation; the notation $E(x_{t+1}|I_t)$ where I_t is the information set at time t is more precise but a little more cumbersome.

Examining equation (1.22), we see that returns can be somewhat predictable—the expected return can vary over time. First, if the conditional variance of returns changes over time, we might expect the conditional mean return to vary as well—the return can just move in and out along a line of constant Sharpe ratio. This explanation does not seem to help much in the data; variables that forecast means do not seem to forecast variances and vice versa. Unless we want to probe the conditional correlation, predictable excess returns have to be explained by changing risk— $\sigma_t(\Delta c_{t+1})$ —or changing risk aversion γ . It is not plausible that risk or risk aversion change at daily frequencies, but fortunately returns are not predictable at daily frequencies. It is much more plausible that risk and risk aversion change over the business cycle, and this is exactly the horizon at which we see predictable excess returns. Models that make this connection precise are a very active area of current research.

Present-Value Statement

$$p_t = E_t \sum_{j=0}^{\infty} m_{t,t+j} d_{t+j}.$$

It is convenient to use only the two-period valuation, thinking of a price p_t and a payoff x_{t+1} . But there are times when we want to relate a price to the entire cash flow stream, rather than just to one dividend and next period's price.

The most straightforward way to do this is to write out a longer-term objective,

$$E_t \sum_{j=0}^{\infty} \beta^j u(c_{t+j}).$$

Now suppose an investor can purchase a stream $\{d_{t+j}\}$ at price p_t . As with the two-period model, his first-order condition gives us the pricing formula directly,

$$p_t = E_t \sum_{j=0}^{\infty} \beta^j \frac{u'(c_{t+j})}{u'(c_t)} d_{t+j} = E_t \sum_{j=0}^{\infty} m_{t,t+j} d_{t+j}. \quad (1.23)$$

You can see that if this equation holds at time t and time $t + 1$, then we can derive the two-period version

$$p_t = E_t [m_{t+1}(p_{t+1} + d_{t+1})]. \quad (1.24)$$

Thus, the infinite-period and two-period models are equivalent.

(Going in the other direction is a little tougher. If you chain together (1.24), you get (1.23) plus an extra term. To get (1.23) you also need the “transversality condition” $\lim_{t \rightarrow \infty} E_t [m_{t,t+j} p_{t+j}] = 0$. This is an extra first-order condition of the infinite-period investor, which is not present with overlapping generations of two-period investors. It rules out “bubbles” in which prices grow so fast that people will buy now just to resell at higher prices later, even if there are no dividends.)

From (1.23) we can write a risk adjustment to prices, as we did with one-period payoffs,

$$p_t = \sum_{j=1}^{\infty} \frac{E_t d_{t+j}}{R_{t,t+j}^f} + \sum_{j=1}^{\infty} \text{cov}_t(d_{t+j}, m_{t,t+j}),$$

where $R_{t,t+j}^f \equiv E_t(m_{t,t+j})^{-1}$ is the j period interest rate. Again, assets whose dividend streams covary negatively with marginal utility, and positively with consumption, have lower prices, since holding those assets gives the investor a more volatile consumption stream. (It is common instead to write prices as a discounted value using a risk-adjusted discount factor, e.g., $p_t^i = \sum_{j=1}^{\infty} E_t d_{t+j}^i / (R^i)^j$, but this approach is difficult to use correctly for multiperiod problems, especially when expected returns can vary over time.)

At a deeper level, the expectation in the two-period formula $p = E(mx)$ sums over states of nature. Equation (1.23) just sums over time as well and is mathematically identical.

1.5 Discount Factors in Continuous Time

Continuous-time versions of the basic pricing equations.	
Discrete	Continuous
$p_t = E_t \sum_{j=1}^{\infty} \beta^j \frac{u'(c_{t+j})}{u'(c_t)} D_{t+j}$	$p_t u'(c_t) = E_t \int_{s=0}^{\infty} e^{-\delta s} u'(c_{t+s}) D_{t+s} ds$
$m_{t+1} = \beta \frac{u'(c_{t+1})}{u'(c_t)}$	$\Lambda_t = e^{-\delta t} u'(c_t)$
$p = E(mx)$	$0 = \Lambda D dt + E_t[d(\Lambda p)]$
$E(R) = R^f - R^f \text{cov}(m, R)$	$E_t \left(\frac{dp}{p} \right) + \frac{D}{p} dt = r_t^f dt - E_t \left[\frac{d\Lambda}{\Lambda} \frac{dp}{p} \right]$

It is often convenient to express asset pricing ideas in the language of continuous-time stochastic differential equations rather than discrete-time stochastic difference equations as I have done so far. The appendix contains a brief introduction to continuous-time processes that covers what you need to know for this book. Even if you want to end up with a discrete-time representation, manipulations are often easier in continuous time. For example, relating interest rates and Sharpe ratios to consumption growth in the last section required a clumsy lognormal approximation; you will see the same sort of thing done much more cleanly in this section.

The choice of discrete versus continuous time is one of modeling convenience. The richness of the theory of continuous-time processes often allows you to obtain analytical results that would be unavailable in discrete time. On the other hand, in the complexity of most practical situations, you often end up resorting to numerical simulation of a discretized model anyway. In those cases, it might be clearer to start with a discrete model. But this is all a choice of language. One should become familiar enough with discrete- as well as continuous-time representations of the same ideas to pick the representation that is most convenient for a particular application.

First, we need to think about how to model securities, in place of price p_t and one-period payoff x_{t+1} . Let a generic security have price p_t at any moment in time, and let it pay dividends at the rate $D_t dt$. (I will continue to denote functions of time as p_t rather than $p(t)$ to maintain continuity with the discrete-time treatment, and I will drop the time subscripts where they are obvious, e.g., dp in place of dp_t . In an interval dt , the security

pays dividends $D_t dt$. I use capital D for dividends to distinguish them from the differential operator d .)

The instantaneous total return is

$$\frac{dp_t}{p_t} + \frac{D_t}{p_t} dt.$$

We model the price of risky assets as diffusions, for example,

$$\frac{dp_t}{p_t} = \mu(\cdot) dt + \sigma(\cdot) dz.$$

(I use the notation dz for increments to a standard Brownian motion, e.g., $z_{t+\Delta} - z_t \sim \mathcal{N}(0, \Delta)$. I use the notation (\cdot) to indicate that the drift and diffusions μ and σ can be functions of state variables. I limit the discussion to diffusion processes—no jumps.) What is nice about this diffusion model is that the increments dz are normal. However, the dependence of μ and σ on state variables means that the finite-time distribution of prices $f(p_{t+\Delta}|I_t)$ need not be normal.

We can think of a risk-free security as one that has a constant price equal to 1 and pays the risk-free rate as a dividend,

$$p = 1 \quad D_t = r_t^f, \quad (1.25)$$

or as a security that pays no dividend but whose price climbs deterministically at a rate

$$\frac{dp_t}{p_t} = r_t^f dt. \quad (1.26)$$

Next, we need to express the first-order conditions in continuous time. The utility function is

$$U(\{c_t\}) = E \int_{t=0}^{\infty} e^{-\delta t} u(c_t) dt.$$

Suppose the investor can buy a security whose price is p_t and that pays a dividend stream D_t . As we did in deriving the present-value price relation in discrete time, the first-order condition for this problem gives us the infinite-period version of the basic pricing equation right away,¹

$$p_t u'(c_t) = E_t \int_{s=0}^{\infty} e^{-\delta s} u'(c_{t+s}) D_{t+s} ds. \quad (1.27)$$

¹One unit of the security pays the dividend stream D_t , i.e., $D_t dt$ units of the numeraire consumption good in a time interval dt . The security costs p_t units of the consumption good. The investor can finance the purchase of ξ units of the security by reducing consumption from e_t to $c_t = e_t - \xi p_t/dt$ during time interval dt . The loss in utility from doing so is $u'(c_t)(e_t - c_t) dt = u'(c_t)\xi p_t$. The gain is the right-hand side of (1.27) multiplied by ξ .

This equation is an obvious continuous-time analogue to

$$p_t = E_t \sum_{j=0}^{\infty} \beta^j \frac{u'(c_{t+j})}{u'(c_t)} D_{t+j}.$$

It turns out that dividing by $u'(c_t)$ is not a good idea in continuous time, since the ratio $u'(c_{t+\Delta})/u'(c_t)$ is not well behaved for small time intervals. Instead, we can keep track of the *level* of marginal utility. Therefore, define the “discount factor” in continuous time as

$$\Lambda_t \equiv e^{-\delta t} u'(c_t).$$

Then we can write the pricing equation as

$$p_t \Lambda_t = E_t \int_{s=0}^{\infty} \Lambda_{t+s} D_{t+s} ds. \quad (1.28)$$

(Some people like to define $\Lambda_t = u'(c_t)$, in which case you keep the $e^{-\delta t}$ in the equation. Others like to scale Λ_t by the risk-free rate, in which case you get an extra $e^{-\int_{\tau=0}^s r_{t+\tau}^f d\tau}$ in the equation. The latter procedure makes it look like a risk-neutral or present-value formula valuation.)

The analogue to the one-period pricing equation $p = E(mx)$ is

$$0 = \Lambda D dt + E_t[d(\Lambda p)]. \quad (1.29)$$

To derive this fundamental equation, take the difference of equation (1.28) at t and $t + \Delta$ (or, start directly with the first-order condition for buying the security at t and selling it at $t + \Delta$),

$$p_t \Lambda_t = E_t \int_{s=0}^{\Delta} \Lambda_{t+s} D_{t+s} ds + E_t[\Lambda_{t+\Delta} p_{t+\Delta}].$$

For Δ small the term in the integral can be approximated as

$$p_t \Lambda_t \approx \Lambda_t D_t \Delta + E_t[\Lambda_{t+\Delta} p_{t+\Delta}]. \quad (1.30)$$

We want to get to d something, so introduce differences by writing

$$p_t \Lambda_t \approx \Lambda_t D_t \Delta + E_t[\Lambda_t p_t + (\Lambda_{t+\Delta} p_{t+\Delta} - \Lambda_t p_t)]. \quad (1.31)$$

Canceling $p_t \Lambda_t$,

$$0 \approx \Lambda_t D_t \Delta + E_t(\Lambda_{t+\Delta} p_{t+\Delta} - \Lambda_t p_t).$$

Taking the limit as $\Delta \rightarrow 0$,

$$0 = \Lambda_t D_t dt + E_t[d(\Lambda_t p_t)]$$

or, dropping time subscripts, equation (1.29).

Equation (1.29) looks different than $p = E(mx)$ because there is no price on the left-hand side; we are used to thinking of the one-period pricing equation as determining price at t given other things, including price at $t + 1$. But price at t is really here, of course, as you can see from equation (1.30) or (1.31). It is just easier to express the *difference* in price over time rather than price today on the left and payoff (including price tomorrow) on the right.

With no dividends and constant Λ , $0 = E_t(dp_t) = E_t(p_{t+\Delta} - p_t)$ says that price should follow a martingale. Thus, $E_t[d(\Lambda p)] = 0$ means that marginal utility-weighted price should follow a martingale, and (1.29) adjusts for dividends. Thus, it is the same as the equation (1.21), $p_t u'(c_t) = E_t(m_{t+1}(p_{t+1} + d_{t+1}))$ that we derived in discrete time.

Since we will write down price processes for dp and discount factor processes for $d\Lambda$, and to interpret (1.29) in terms of expected returns, it is often convenient to break up the $d(\Lambda_t p_t)$ term using Ito's lemma:

$$d(\Lambda p) = p d\Lambda + \Lambda dp + dp d\Lambda. \quad (1.32)$$

Using the expanded version (1.32) in the basic equation (1.29), and dividing by $p\Lambda$ to make it pretty, we obtain an equivalent, slightly less compact but slightly more intuitive version,

$$0 = \frac{D}{p} dt + E_t \left[\frac{d\Lambda}{\Lambda} + \frac{dp}{p} + \frac{d\Lambda dp}{\Lambda p} \right]. \quad (1.33)$$

(This formula only works when both Λ and p can never be zero. It is often enough the case that this formula is useful. If not, multiply through by Λ and p and keep them in numerators.)

Applying the basic pricing equations (1.29) or (1.33) to a risk-free rate, defined as (1.25) or (1.26), we obtain

$$r_t^f dt = -E_t \left(\frac{d\Lambda_t}{\Lambda_t} \right). \quad (1.34)$$

This equation is the obvious continuous-time equivalent to

$$R_t^f = \frac{1}{E_t(m_{t+1})}.$$

If a risk-free rate is not traded, we can use (1.34) to define a shadow risk-free rate or zero-beta rate.

With this interpretation, we can rearrange equation (1.33) as

$$E_t \left(\frac{dp_t}{p_t} \right) + \frac{D_t}{p_t} dt = r_t^f dt - E_t \left[\frac{d\Lambda_t dp_t}{\Lambda_t p_t} \right]. \quad (1.35)$$

This is the obvious continuous-time analogue to

$$E(R) = R^f - R^f \text{cov}(m, R). \quad (1.36)$$

The last term in (1.35) is the covariance of the return with the discount factor or marginal utility. Since means are order dt , there is no difference between covariance and second moment in the last term of (1.35). The interest rate component of the last term of (1.36) naturally vanishes as the time interval gets short.

Ito's lemma makes many transformations simple in continuous time. For example, the nonlinear transformation between consumption and the discount factor led us to some tricky approximations in discrete time. This transformation is easy in continuous time (diffusions are locally normal, so it is really the same trick). With $\Lambda_t = e^{-\delta t} u'(c_t)$ we have

$$\begin{aligned} d\Lambda_t &= -\delta e^{-\delta t} u'(c_t) dt + e^{-\delta t} u''(c_t) dc_t + \frac{1}{2} e^{-\delta t} u'''(c_t) dc_t^2, \\ \frac{d\Lambda_t}{\Lambda_t} &= -\delta dt + \frac{c_t u''(c_t)}{u'(c_t)} \frac{dc_t}{c_t} + \frac{1}{2} \frac{c_t^2 u'''(c_t)}{u'(c_t)} \frac{dc_t^2}{c_t^2}. \end{aligned} \quad (1.37)$$

Denote the local curvature and third derivative of the utility function as

$$\begin{aligned} \gamma_t &= -\frac{c_t u''(c_t)}{u'(c_t)}, \\ \eta_t &= \frac{c_t^2 u'''(c_t)}{u'(c_t)}. \end{aligned}$$

(For power utility, the former is the power coefficient γ and the latter is $\eta_t = \gamma(\gamma + 1)$.)

Using this formula we can quickly redo the relationship between interest rates and consumption growth, equation (1.7),

$$r_t^f = -\frac{1}{dt} E_t \left(\frac{d\Lambda_t}{\Lambda_t} \right) = \delta + \gamma_t \frac{1}{dt} E_t \left(\frac{dc_t}{c_t} \right) - \frac{1}{2} \eta_t \frac{1}{dt} E_t \left(\frac{dc_t^2}{c_t^2} \right).$$

We can also easily express asset prices in terms of consumption risk rather than discount factor risk, as in equation (1.16). Using (1.37) in (1.35),

$$E_t \left(\frac{dp_t}{p_t} \right) + \frac{D_t}{p_t} dt - r_t^f dt = \gamma E_t \left(\frac{dc_t}{c_t} \frac{dp_t}{p_t} \right). \quad (1.38)$$

Thus, assets whose returns covary more strongly with consumption get higher mean excess returns, and the constant relating covariance to mean return is the utility curvature coefficient γ .

Since correlations are less than 1, equation (1.38) implies that Sharpe ratios are related to utility curvature and consumption volatility directly; we do not need the ugly lognormal facts and an approximation that we needed in (1.20). Using $\mu_p \equiv E_t(dp_t/p_t)$; $\sigma_p^2 = E_t[(dp_t/p_t)^2]$; $\sigma_c^2 = E_t[(dc_t/c_t)^2]$,

$$\frac{\mu_p + \frac{D_t}{p_t} dt - r_t^f dt}{\sigma_p} \leq \gamma \sigma_c.$$

Problems—Chapter 1

1.

(a) The absolute risk-aversion coefficient is

$$\frac{u''(c)}{u'(c)}.$$

We scale by $u'(c)$ because expected utility is only defined up to linear transformations— $a + bu(c)$ gives the same predictions as $u(c)$ —and this measure of the second derivative is invariant to linear transformations. Show that the utility function with constant absolute risk aversion is

$$u(c) = -e^{-\alpha c}.$$

(b) The coefficient of relative risk aversion in a one-period model (i.e., when consumption equals wealth) is defined as

$$rra = \frac{cu''(c)}{u'(c)}.$$

For power utility $u'(c) = c^{-\gamma}$, show that the risk-aversion coefficient equals the power.

(c) The elasticity of intertemporal substitution is defined in a non-stochastic model with interest rate R as

$$\xi^I \equiv -\frac{c_2/c_1 d(c_1/c_2)}{dR/R}.$$

Show that with power utility $u'(c) = c^{-\gamma}$, the intertemporal substitution elasticity is equal to $1/\gamma$. (Hint: differentiate the first-order conditions)

2. Show that the “idiosyncratic risk” line in Figure 1.1 is horizontal.

3.

(a) Suppose you have a mean–variance efficient return R^{mv} and the risk-free rate. Using the fact that R^{mv} is perfectly correlated with the

discount factor, construct a discount factor m in terms of R^f and R^{mv} , with no free parameters. (You have to find the constants a and b in $m = a + bR^{mv}$. They will depend on things like $E(R^{mv})$.)

(b) Using this result, and the beta model in terms of m , show that expected returns can be described in a single-beta representation using any mean-variance efficient return (except the risk-free rate).

$$E(R^i) = R^f + \beta_{i,mv}[E(R^{mv}) - R^f].$$

4. Can the “Sharpe ratio” between two risky assets exceed the slope of the mean-variance frontier? That is, if R^{mv} is on the frontier, is it possible that

$$\frac{E(R^i) - E(R^j)}{\sigma(R^i - R^j)} > \frac{E(R^{mv}) - R^f}{\sigma(R^{mv})}?$$

5. Show that if consumption growth is lognormal, then

$$\left| \frac{E(R^{mv}) - R^f}{\sigma(R^{mv})} \right| = \frac{\sigma[(c_{t+1}/c_t)^{-\gamma}]}{E[(c_{t+1}/c_t)^{-\gamma}]} = \sqrt{e^{\gamma^2 \sigma^2(\Delta \ln c_{t+1})} - 1} \approx \gamma \sigma(\Delta \ln c).$$

(Start with $\sigma^2(x) = E(x^2) - E(x)^2$ and the lognormal property $E(e^z) = e^{Ez + (1/2)\sigma^2(z)}$.)

6. There are assets with mean return equal to the risk-free rate, but substantial standard deviation of returns. Long-term bonds are pretty close examples. Why would anyone hold such an asset? Wouldn't it be better to put your money in a mean-variance frontier asset?

7. The first-order conditions for an infinitely lived investor who can buy an asset with dividend stream $\{d_t\}$ are

$$p_t = E_t \sum_{j=1}^{\infty} \beta^j \frac{u'(c_{t+j})}{u'(c_t)} d_{t+j}. \quad (1.39)$$

The first-order conditions for buying a security with price p_t and payoff $x_{t+1} = d_{t+1} + p_{t+1}$ are

$$p_t = E_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} (p_{t+1} + d_{t+1}) \right]. \quad (1.40)$$

(a) Derive (1.40) from (1.39)

(b) Derive (1.39) from (1.40). You need an extra condition. Show that this extra condition is a first-order condition for maximization. To do this, think about what strategy the consumer could follow to improve utility if the condition did not hold, both if this is if the only security available and if the investor can trade all state- and date-contingent claims.

8. Suppose a consumer has a utility function that includes leisure. (This could also be a second good, or a good produced in another country.) Using the continuous-time setup, show that expected returns will now depend on *two* covariances, the covariance of returns with leisure and the covariance of returns with consumption, so long as leisure enters non-separably, i.e., $u(c, l)$ cannot be written $v(c) + w(l)$. (This is a three line problem, but you need to apply Ito's lemma to Λ .)

9. From

$$1 = E(mR)$$

show that the negative of the mean log discount factor must be larger than *any* mean return,

$$-E(\ln m) > E(\ln R).$$

How is it possible that $E(\ln R)$ is bounded—what about returns of the form $R = (1 - \alpha)R^f + \alpha R^m$ for arbitrarily large α ? (Hint: start by assuming m and R are lognormal. Then see if you can generalize the results using Jensen's inequality, $E(f(x)) > f(E(x))$ for f convex. The return that solves $\max_R E(\ln R)$ is known as the *growth optimal portfolio*.)

2

Applying the Basic Model

2.1 Assumptions and Applicability

Writing $p = E(mx)$, we do *not* assume

1. Markets are complete, or there is a representative investor
2. Asset returns or payoffs are normally distributed (no options), or independent over time
3. Two-period investors, quadratic utility, or separable utility
4. Investors have no human capital or labor income
5. The market has reached equilibrium, or individuals have bought all the securities they want to

All of these assumptions come later, in various special cases, but we *have not* made them yet. We *do* assume that the investor can consider a small marginal investment or disinvestment.

The theory of asset pricing contains lots of assumptions used to derive analytically convenient special cases and empirically useful representations. In writing $p = E(mx)$ or $pu'(c_t) = E_t[\beta u'(c_{t+1})x_{t+1}]$, we have *not* made most of these assumptions.

We have *not* assumed complete markets or a representative investor. These equations apply to each individual investor, for each asset to which he has access, independently of the presence or absence of other investors or other assets. Complete markets/representative agent assumptions are used if one wants to use aggregate consumption data in $u'(c_t)$, or other specializations and simplifications of the model.

We have *not* said anything about payoff or return distributions. In particular, we have not assumed that returns are normally distributed or that utility is quadratic. The basic pricing equation should hold for *any* asset, stock, bond, option, real investment opportunity, etc., and any monotone and concave utility function. In particular, it is often thought that mean-variance analysis and beta pricing models require these kinds of

limiting assumptions or quadratic utility, but that is not the case. A mean-variance efficient return carries all pricing information no matter what the distribution of payoffs, utility function, etc.

This is *not* a “two-period model.” The fundamental pricing equation holds for any two periods of a multiperiod model, as we have seen. Really, everything involves *conditional* moments, so we have not assumed i.i.d. returns over time.

I have written things down in terms of a time- and state-separable utility function and I have focused on the convenient power utility example. Nothing important lies in either choice. Just interpret $u'(c_t)$ as the partial derivative of a general utility function with respect to consumption at time t . State- or time-nonseparable utility (habit persistence, durability) complicates the relation between the discount factor and real variables, but does not change $p = E(mx)$ or any of the basic structure.

We do *not* assume that investors have no nonmarketable human capital, or no outside sources of income. The first-order conditions for purchase of an asset relative to consumption hold no matter what else is in the budget constraint. By contrast, the portfolio approach to asset pricing as in the CAPM and ICAPM relies heavily on the assumption that the investor has no nonasset income, and we will study these special cases below. For example, leisure in the utility function just means that marginal utility $u'(c, l)$ may depend on l as well as c .

We do not even really need the assumption (yet) that the market is “in equilibrium,” that the investor has bought all of the asset that he wants to, or even that he can buy the asset at all. We can interpret $p = E(mx)$ as giving us the value, or willingness to pay for, a small amount of a payoff x_{t+1} that the investor does not yet have. Here is why: If the investor had a little ξ more of the payoff x_{t+1} at time $t + 1$, his utility $u(c_t) + \beta E_t u(c_{t+1})$ would increase by

$$\begin{aligned} & \beta E_t [u(c_{t+1} + \xi x_{t+1}) - u(c_{t+1})] \\ &= \beta E_t \left[u'(c_{t+1}) x_{t+1} \xi + \frac{1}{2} u''(c_{t+1}) (x_{t+1} \xi)^2 + \dots \right]. \end{aligned}$$

If ξ is small, only the first term on the right matters. If the investor has to give up a small amount of money $v_t \xi$ at time t , that loss lowers his utility by

$$u(c_t - v_t \xi) = u'(c_t) v_t \xi + \frac{1}{2} u''(c_t) (v_t \xi)^2 + \dots$$

Again, for small ξ , only the first term matters. Therefore, in order to receive the small extra payoff ξx_{t+1} , the investor is willing to pay the small

amount $v_t \xi$, where

$$v_t = E_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} x_{t+1} \right].$$

If this private valuation is higher than the market value p_t , and if the investor can buy some more of the asset, he will. As he buys more, his consumption will change; it will be higher in states where x_{t+1} is higher, driving down $u'(c_{t+1})$ in those states, until the value to the investor has declined to equal the market value. Thus, *after* an investor has reached his optimal portfolio, the *market value* should obey the basic pricing equation as well, using post-trade or equilibrium consumption. But the formula can also be applied to generate the marginal *private* valuation, using pre-trade consumption, or to value a *potential*, not yet traded security.

We *have* calculated the value of a “small” or marginal portfolio change for the investor. For some investment projects, an investor cannot take a small (“diversified”) position. For example, a venture capitalist or entrepreneur must usually take all or nothing of a project with payoff stream $\{x_t\}$. Then the value of a project not already taken, $E \sum_j \beta^j u(c_{t+j} + x_{t+j})$, might be substantially different from its marginal counterpart, $E \sum_j \beta^j u'(c_{t+j}) x_{t+j}$. Once the project is taken, of course, $c_{t+j} + x_{t+j}$ becomes c_{t+j} , so the marginal valuation still applies to the ex post consumption stream. Analysts often forget this point and apply marginal (diversified) valuation models such as the CAPM to projects that must be bought in discrete chunks. Also, we have abstracted from short sales and bid/ask spreads; this modification changes $p = E(mx)$ from an equality to a set of inequalities.

2.2 General Equilibrium

Asset returns and consumption: which is the chicken and which is the egg? I present the exogenous return model, the endowment economy model, and the argument that it does not matter for studying $p = E(mx)$.

So far, we have not said where the joint statistical properties of the payoff x_{t+1} and marginal utility m_{t+1} or consumption c_{t+1} come from. We have also not said anything about the fundamental exogenous shocks that drive the economy. The basic pricing equation $p = E(mx)$ tells us only what the price should be, *given* the joint distribution of consumption (marginal utility, discount factor) and the asset payoff.

There is nothing that stops us from writing the basic pricing equation as

$$u'(c_t) = E_t[\beta u'(c_{t+1})x_{t+1}/p_t].$$

We can think of this equation as determining today's *consumption* given asset prices and payoffs, rather than determining today's *asset price* in terms of consumption and payoffs. Thinking about the basic first-order condition in this way gives the permanent income model of consumption.

Which is the chicken and which is the egg? Which variable is exogenous and which is endogenous? The answer is, neither, and for many purposes, it does not matter. The first-order conditions characterize any equilibrium; if you happen to know $E(mx)$, you can use them to determine p ; if you happen to know p , you can use them to determine consumption and savings decisions.

For most asset pricing applications we are interested in understanding a wide cross section of assets. Thus, it is interesting to contrast the cross-sectional variation in asset prices (expected returns) with cross-sectional variation in their second moments (betas) with a single discount factor. In most applications, the discount factor is a function of aggregate variables (market return, aggregate consumption), so it is plausible to hold the properties of the discount factor constant as we compare one individual asset to another. Permanent income studies typically dramatically restrict the number of assets under consideration, often to just an interest rate, and study the time-series evolution of aggregate or individual consumption.

Nonetheless, it is an obvious next step to complete the solution of our model economy; to find c and p in terms of truly exogenous forces. The results will of course depend on what the rest of the economy looks like, in particular the production or intertemporal transformation technology and the set of markets.

Figure 2.1 shows one possibility for a general equilibrium. Suppose that the production technologies are linear: the real, physical rate of return (the rate of intertemporal *transformation*) is not affected by how much is invested. Now consumption must adjust to these technologically given rates of return. If the rates of return on the intertemporal technologies were to change, the consumption process would have to change as well. This is, implicitly, how the permanent income model works. This is how many finance theories such as the CAPM and ICAPM and the Cox, Ingersoll, and Ross (1986) model of the term structure work as well. These models specify the return process, and then solve the consumer's portfolio and consumption rules.

Figure 2.2 shows another extreme possibility for the production technology. This is an "endowment economy." Nondurable consumption




Figure 2.1. Consumption adjusts when the rate of return is determined by a linear technology.

appears (or is produced by labor) every period. There is nothing anyone can do to save, store, invest, or otherwise transform consumption goods this period to consumption goods next period. Hence, asset prices must adjust until people are just happy consuming the endowment process. In this case consumption is exogenous and asset prices adjust. Lucas (1978) and Mehra and Prescott (1985) are two very famous applications of this sort of “endowment economy.”

Which of these possibilities is correct? Well, neither, of course. The real economy and all serious general equilibrium models look something like Figure 2.3: one can save or transform consumption from one date to the next, but at a decreasing rate. As investment increases, rates of return decline.

Does this observation invalidate the modeling we do with the linear technology (CAPM, CIR, permanent income) model, or the endowment




Figure 2.2. Asset prices adjust to consumption in an endowment economy.




Figure 2.3. General equilibrium. The solid lines represent the indifference curve and production possibility set. The dashed straight line represents the equilibrium rate of return. The dashed box represents an endowment economy that predicts the same consumption and asset return process.

economy model? No. Start at the equilibrium in Figure 2.3. Suppose we model this economy as a linear technology, but we happen to choose for the rate of return on the linear technologies exactly the same stochastic process for returns that emerges from the general equilibrium. The resulting joint consumption-asset return process is exactly the same as in the original general equilibrium! Similarly, suppose we model this economy as an endowment economy, but we happen to choose for the endowment process exactly the stochastic process for consumption that emerges from the equilibrium with a concave technology. Again, the joint consumption-asset return process is exactly the same.

Therefore, there is nothing wrong in adopting one of the following strategies for empirical work:

1. Form a statistical model of bond and stock returns, solve the optimal consumption-portfolio decision. Use the equilibrium consumption values in $p = E(mx)$.
2. Form a statistical model of the consumption process, calculate asset prices and returns directly from the basic pricing equation $p = E(mx)$.
3. Form a completely correct general equilibrium model, including the production technology, utility function, and specification of the market structure. Derive the equilibrium consumption and asset price process, including $p = E(mx)$ as one of the equilibrium conditions.

If the statistical models for consumption and/or asset returns are right, i.e., if they coincide with the equilibrium consumption or return process generated by the true economy, either of the first two approaches

will give correct predictions for the joint consumption-asset return process.

Most finance models, developed from the 1950s through the early 1970s, take the return process as given, implicitly assuming linear technologies. The endowment economy approach, introduced by Lucas (1978), is a breakthrough because it turns out to be much easier. It is much easier to evaluate $p = E(mx)$ for fixed m than it is to solve joint consumption-portfolio problems for given asset returns, all to derive the equilibrium consumption process. To solve a consumption-portfolio problem we have to model the investor's entire environment: we have to specify *all* the assets to which he has access, what his labor income process looks like (or wage rate process, and include a labor supply decision). Once we model the consumption stream directly, we can look at each asset in isolation, and the actual computation is almost trivial. This breakthrough accounts for the unusual structure of the presentation in this book. It is traditional to start with an extensive study of consumption-portfolio problems. But by modeling consumption directly, we have been able to study pricing directly, and portfolio problems are an interesting side trip which we can defer.

Most uses of $p = E(mx)$ do not require us to take any stand on exogeneity or endogeneity, or general equilibrium. This is a condition that must hold for any asset, for any production technology. Having a taste of the extra assumptions required for a general equilibrium model, you can now appreciate why people stop short of full solutions when they can address an application using only the first-order conditions, using knowledge of $E(mx)$ to make a prediction about p .

It is enormously tempting to slide into an interpretation that $E(mx)$ *determines* p . We routinely think of betas and factor risk prices—components of $E(mx)$ —as *determining* expected returns. For example, we routinely say things like “the expected return of a stock increased *because* the firm took on riskier projects, thereby increasing its beta.” But the whole consumption process, discount factor, and factor risk premia change when the production technology changes. Similarly, we are on thin ice if we say anything about the effects of policy interventions, new markets and so on. The equilibrium consumption or asset return process one has modeled statistically may change in response to such changes in structure. For such questions one really needs to start thinking in general equilibrium terms. It may help to remember that there is an army of permanent-income macroeconomists who make precisely the opposite assumption, taking our asset return processes as exogenous and studying (endogenous) consumption and savings decisions.

2.3 Consumption-Based Model in Practice

The consumption-based model is, in principle, a complete answer to all asset pricing questions, but works poorly in practice. This observation motivates other asset pricing models.

The model I have sketched so far can, in principle, give a complete answer to all the questions of the theory of valuation. It can be applied to *any* security—bonds, stocks, options, futures, etc.—or to any uncertain cash flow. All we need is a functional form for utility, numerical values for the parameters, and a statistical model for the conditional distribution of consumption and payoffs.

To be specific, consider the standard power utility function

$$u'(c) = c^{-\gamma}. \quad (2.1)$$

Then, excess returns should obey

$$0 = E_t \left[\beta \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma} R_{t+1}^e \right]. \quad (2.2)$$

Taking unconditional expectations and applying the covariance decomposition, expected excess returns should follow

$$E(R_{t+1}^e) = -R^f \text{cov} \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\gamma}, R_{t+1}^e \right]. \quad (2.3)$$

Given a value for γ , and data on consumption and returns, you can easily estimate the mean and covariance on the right-hand side, and check whether actual expected returns are, in fact, in accordance with the formula.

Similarly, the present-value formula is

$$p_t = E_t \sum_{j=1}^{\infty} \beta^j \left(\frac{c_{t+j}}{c_t} \right)^{-\gamma} d_{t+j}. \quad (2.4)$$

Given data on consumption and dividends or another stream of payoffs, you can estimate the right-hand side and check it against prices on the left.

Bonds and options do not require separate valuation theories. For example, an N -period default-free nominal discount bond (a U.S. Treasury strip) is a claim to one dollar at time $t + N$. Its price should be

$$p_t = E_t \left(\beta^N \left(\frac{c_{t+N}}{c_t} \right)^{-\gamma} \frac{\Pi_t}{\Pi_{t+N}} 1 \right),$$




Figure 2.4. Mean excess returns of 10 CRSP size portfolios versus predictions of the power utility consumption-based model. The predictions are generated by $-R^f \text{cov}(m, R^i)$ with $m = \beta(c_{t+1}/c_t)^{-\gamma}$. $\beta = 0.98$ and $\gamma = 241$ are picked by first-stage GMM to minimize the sum of squared pricing errors (deviation from 45° line). Source: Cochrane (1996).

where $\Pi =$ price level (\$/good). A European option is a claim to the payoff $\max(S_{t+T} - K, 0)$, where $S_{t+T} =$ stock price at time $t + T$, $K =$ strike price. The option price should be

$$p_t = E_t \left[\beta^T \left(\frac{c_{t+T}}{c_t} \right)^{-\gamma} \max(S_{t+T} - K, 0) \right].$$

Again, we can use data on consumption, prices, and payoffs to check these predictions.

Unfortunately, this specification of the consumption-based model does not work very well. To give a flavor of some of the problems, Figure 2.4 presents the mean excess returns on the ten size-ranked portfolios of NYSE stocks versus the predictions—the right-hand side of (2.3)—of the consumption-based model. I picked the utility curvature parameter $\gamma = 241$ to make the picture look as good as possible. (The section on GMM estimation below goes into detail on how to do this. The figure presents the first-stage GMM estimate.) As you can see, the model is not hopeless—there is some correlation between sample average returns and the consumption-based model predictions. But the model does not do very well. The pricing error (actual expected return – predicted expected return) for each portfolio is of the same order of magnitude as the spread in expected returns across the portfolios.

2.4 Alternative Asset Pricing Models: Overview

I motivate exploration of different utility functions, general equilibrium models, and linear factor models such as the CAPM, APT, and ICAPM as ways to circumvent the empirical difficulties of the consumption-based model.

The poor empirical performance of the consumption-based model motivates a search for alternative asset pricing models—alternative functions $m = f(\text{data})$. All asset pricing models amount to different functions for m . I give here a bare sketch of some of the different approaches; we study each in detail in later chapters.

- 1) *Different utility functions.* Perhaps the problem with the consumption-based model is simply the functional form we chose for utility. The natural response is to try different utility functions. Which variables determine marginal utility is a far more important question than the functional form. Perhaps the stock of durable goods influences the marginal utility of nondurable goods; perhaps leisure or yesterday's consumption affect today's marginal utility. These possibilities are all instances of *nonseparabilities*. One can also try to use micro data on individual consumption of stockholders rather than aggregate consumption. Aggregation of heterogeneous investors can make variables such as the cross-sectional variance of income appear in aggregate marginal utility.
- 2) *General equilibrium models.* Perhaps the problem is simply with the consumption *data*. General equilibrium models deliver equilibrium decision rules linking consumption to other variables, such as income, investment, etc. Substituting the decision rules $c_t = f(y_t, i_t, \dots)$ in the consumption-based model, we can link asset prices to other, hopefully better-measured macroeconomic aggregates.

In addition, true general equilibrium models completely describe the economy, including the stochastic process followed by all variables. They can answer questions such as *why* is the covariance (beta) of an asset payoff x with the discount factor m the value that it is, rather than take this covariance as a primitive. They can in principle answer structural questions, such as how asset prices might be affected by different government policies or the introduction of new securities. Neither kind of question can be answered by just manipulating investor first-order conditions.

- 3) *Factor pricing models.* Another sensible response to bad consumption data is to model marginal utility in terms of other variables directly. *Factor pricing* models follow this approach. They just specify that the

discount factor is a linear function of a set of proxies,

$$m_{t+1} = a + b_A f_{t+1}^A + b_B f_{t+1}^B + \dots, \quad (2.5)$$

where f^i are *factors* and a, b_i are parameters. (This is a different sense of the use of the word “factor” than “discount factor” or “factor analysis.” I did not invent the confusing terminology.) By and large, the factors are just selected as plausible proxies for marginal utility: events that describe whether typical investors are happy or unhappy. Among others, the Capital Asset Pricing Model (CAPM) is the model

$$m_{t+1} = a + bR_{t+1}^W,$$

where R^W is the rate of return on a claim to total wealth, often proxied by a broad-based portfolio such as the value-weighted NYSE portfolio. The Arbitrage Pricing Theory (APT) uses returns on broad-based portfolios derived from a factor analysis of the return covariance matrix. The Intertemporal Capital Asset Pricing Model (ICAPM) suggests macroeconomic variables such as GNP and inflation and variables that forecast macroeconomic variables or asset returns as factors. Term structure models such as the Cox–Ingersoll–Ross model specify that the discount factor is a function of a few term structure variables, for example the short rate of interest and a few interest rate spreads.

Many factor pricing models are derived as general equilibrium models with linear technologies and no labor income; thus they also fall into the general idea of using general equilibrium relations (from, admittedly, very stylized general equilibrium models) to substitute out for consumption.

- 4) *Arbitrage or near-arbitrage pricing.* The mere existence of a representation $p = E(mx)$ and the fact that marginal utility is positive $m \geq 0$ (these facts are discussed in the next chapter) can often be used to deduce prices of one payoff in terms of the prices of other payoffs. The Black–Scholes option pricing model is the paradigm of this approach: Since the option payoff can be replicated by a portfolio of stock and bond, any discount factor m that prices the stock and bond gives the price for the option. Recently, there have been several suggestions on how to use this idea in more general circumstances by using very weak further restrictions on m , and we will study these suggestions in Chapter 17.

We return to a more detailed derivation and discussion of these alternative models of the discount factor m below. First, and with this brief overview in mind, we look at $p = E(mx)$ and what the discount factor m represents in a little more detail.

Problems—Chapter 2

1. The representative consumer maximizes a CRRA utility function,

$$E_t \sum \beta^j c_{t+j}^{1-\gamma}.$$

Consumption is given by an endowment stream.

(a) Show that with log utility, the price/consumption ratio of the consumption stream is constant, no matter what the distribution of consumption growth.

(b) Suppose there is news at time t that future consumption will be higher. For $\gamma < 1$, $\gamma = 1$, and $\gamma > 1$, evaluate the effect of this news on the price. Make sense of your results. (There is a real-world interpretation here. It is often regarded as a puzzle that the market declines on good economic news. This is attributed to an expectation by the market that the Fed will respond to such news by raising interest rates. Note that $\gamma > 0$ in this problem gives a completely real and frictionless interpretation to this phenomenon! I thank Pete Hecht for this nice problem.)

2. The linear quadratic permanent income model is a very useful general equilibrium model that we can solve in closed form. It specifies a production technology rather than fixed endowments, and it easily allows aggregation of disparate consumers. (Hansen [1987] is a wonderful exposition of what one can do with this setup.)

The consumer maximizes

$$E \sum_{t=0}^{\infty} \beta^t \left(-\frac{1}{2} \right) (c_t - c^*)^2$$

subject to a linear technology

$$k_{t+1} = (1+r)k_t + i_t,$$

$$i_t = e_t - c_t.$$

e_t is an exogenous endowment or labor income stream. Assume $\beta = 1/(1+r)$; the discount rate equals the interest rate or marginal productivity of capital.

- (a) Show that optimal consumption follows

$$c_t = rk_t + r\beta \sum_{j=0}^{\infty} \beta^j E_t e_{t+j}, \quad (2.6)$$

$$c_t = c_{t-1} + (E_t - E_{t-1})r\beta \sum_{j=0}^{\infty} \beta^j e_{t+j}, \quad (2.7)$$

i.e., consumption equals permanent income, precisely defined, and consumption follows a random walk whose innovations are equal to innovations in permanent income.

(b) Assume that the endowment e_t follows an AR(1)

$$e_t = \rho e_{t-1} + \varepsilon_t$$

and specialize (2.6) and (2.7). Calculate and interpret the result for $\rho = 1$ and $\rho = 0$. (The result looks like a “consumption function” relating consumption to capital and current income, except that the slope of that function depends on the persistence of income shocks. Transitory shocks will have little effect on consumption, and permanent shocks a larger effect.)

(c) Calculate the one-period interest rate (it should come out to r , of course) and the price of a claim to the consumption stream. e and k are the only state variables, so the price should be a function of e and k . Interpret the time variation in the price of the consumption stream. (This consumer gets more risk averse as consumption rises to c^* . c^* is the bliss point, so at the bliss point there is no average return that can compensate the consumer for greater risk.)

3. Consider again CRRA utility,

$$E_t \sum \beta^j c_{t+j}^{1-\gamma}.$$

Consumption growth follows a two-state Markov process. The states are $\Delta c_t = c_t/c_{t-1} = h, l$, and a 2×2 matrix π governs the set of transition probabilities, i.e., $\text{pr}(\Delta c_{t+1} = h | \Delta c_t = l) = \pi_{l \rightarrow h}$. (This is the Mehra–Prescott [1986] model, but it will be faster to do it than to look it up. It is a useful and simple endowment economy.)

(a) Find the risk-free rate (price of a certain real payoff of 1) in this economy. This price is generated by

$$p_t^b = E_t(m_{t,t+1}1).$$

You are looking for two values, the price in the l state and the price in the h state.

(b) Find the price of the consumption stream (the price at t of $\{c_{t+1}, c_{t+2}, \dots\}$). To do this, guess that the price/consumption ratio must be a function of state (h, l) , and find that function. From

$$p_t^c = E_t(m_{t,t+1}(p_{t+1}^c + c_{t+1}))$$

find a recursive relation for p_t^c/c_t , and hence find the two values of p_t^c/c_t , one for the h state and one for the l state.

(c) Pick $\beta = 0.99$ and try $\gamma = 0.5, 5$. (Try more values if you feel like it.) Calibrate the consumption process to have a 1% mean and 1% standard deviation, and consumption growth uncorrelated over time. Calculate prices and returns in each state.

(d) Now introduce serial correlation in consumption growth with $\gamma = 5$. (You can do this by adding weight to the diagonal entries of the transition matrix π .) What effect does this change have on the model?