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.

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 *co*-variance, 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,

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

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 $\gamma \to 1$ is[1]

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

[1] To think about this limit precisely, add a constant to the utility function and write it as

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

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

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

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

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

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,

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

or

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

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; $E_t\left[\beta u'(c_{t+1}) x_{t+1}\right]$ 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}

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

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

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

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

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

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,

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

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	$R^e_{t+1} = R^a_{t+1} - R^b_{t+1}$
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 $x_{t+1}z_t$, 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^e_t, 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 and power utility,

$$r^f_t = \delta + \gamma E_t(\Delta \ln c_{t+1}) - \frac{\gamma^2}{2}\sigma^2_t(\Delta \ln c_{t+1}).$$

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

$$r^f_t = \delta + \gamma E_t(\Delta \ln c_{t+1}) - \frac{\gamma^2}{2}\sigma^2_t(\Delta \ln c_{t+1}), \tag{1.7}$$

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

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

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

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

(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

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

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

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

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

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

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

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

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,

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

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

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

or

$$E(R^i) - R^f = -\frac{\operatorname{cov}\left[u'(c_{t+1}), R^i_{t+1}\right]}{E\left[u'(c_{t+1})\right]}. \tag{1.13}$$

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

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

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,

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

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.