

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

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

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 $\{x_i\}$. Then the value of a project not already taken, $E \sum_j \beta^j [u(c_{t+j} + x_{t+j}) - u(c_{t+j})]$, 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.




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

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 (1985) 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 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 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




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.*

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

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

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

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

where Π = 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

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

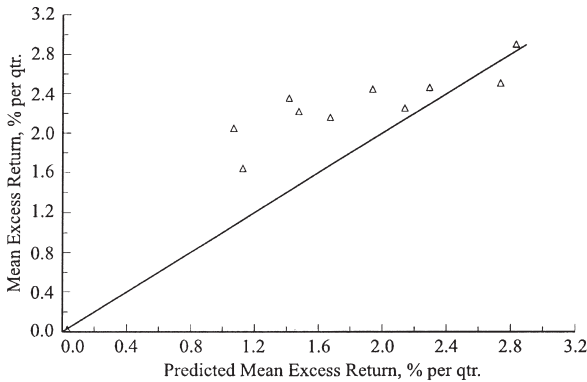


Figure 2.4. Mean excess returns of 10 CRSP size portfolios versus predictions of the power utility consumption-based model. The predictions are generated by $-R^f \text{cov}(m, R^i)$ with $m = \beta(c_{t+1}/c_t)^{-\gamma}$. $\beta = 0.98$ and $\gamma = 241$ are picked by first-stage GMM to minimize the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

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,

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