

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.