

PART I

**PERSPECTIVE AND
PREPARATION**

COPYRIGHTED MATERIAL

CHAPTER 1

INTRODUCTION AND OVERVIEW

Our subject in this book is financial assets—how people choose them, how their prices are determined, and how their prices relate to each other and behave over time. To begin, it helps to have a clear and simple conception of what assets *are*, why people desire to hold and trade them, and how the allocation of resources to financial firms and markets can be justified.

1.1 AN ELEMENTAL VIEW OF ASSETS AND MARKETS

Economists usually think of assets as “bundles” of time–state–contingent claims. A metaphor helps to see what they mean by this. When events unfold through time it is as if we are moving along a sequence of time-stamped roulette wheels. At time t nature spins the appropriate wheel and we watch to see in which slot the ball settles. That slot defines the “state of the world” at t . When the state is realized, so is the cash value of each asset at time t , which is thus *contingent* on the state and the time. From our point of view the state itself is just a description of current reality in sufficient detail that we know what each asset is worth at the time.

1.1.1 Assets as Bundles of Claims

The simplest conceivable financial asset is one that entitles the holder to receive one unit of cash when the wheel for some particular date selects one particular state—and nothing otherwise. There are no exact counterparts in the real financial world, but the closest would be an insurance contract that pays a fixed amount under a narrowly defined condition. The next simpler conception is a “safe” asset that yields a specified cash payment at t regardless of where the wheel stops. A government-backed, default-free “discount” bond that matures at t would be the nearest example, since the issuer of the bond promises to pay a fixed number of units of cash regardless of the conditions at t . A default-free bond that matures at t_n and makes periodic payments of interest (“coupons”) at t_1, t_2, \dots, t_n is like a portfolio of these state-independent discount bonds. A forward contract to exchange a fixed number of units of cash at future date T for a fixed number of units of a commodity is a simple example of an asset whose value at T does depend on the state. One who is to pay the cash and receive the commodity has a state-independent liability (the cash that is owed) and a state-dependent receipt (the value of the commodity). At times before their maturities and expirations, values of marketable bonds and forward contracts alike are state dependent. Unlike either of these instruments, shares of stock have lifetimes without definite limit. A share of stock offers bundles of alternative state-contingent payments at alternative future dates out to some indefinite time at which a state is realized that corresponds to the company’s liquidation. Dividends are other time-stamped, state-contingent claims that might be paid along the way. A European-style *call option* on the stock offers claims that are tied to states defined explicitly in terms of the stock’s price at a fixed expiration date. One who holds such an option that expires at date T can pay a fixed sum (the “strike” price) and receive the stock on that date, but would choose to do so only in states in which the stock’s price exceeds the required cash payment. If the option is so “exercised” at T , the former option holder acquires the same state-contingent rights as the stockholder from that time.

Each day vast numbers of these and other time–state-contingent claims are created and passed back and forth among individuals, financial firms, and nonfinancial businesses. Some of the trades take place in central marketplaces like the New York Stock Exchange (NYSE) and affiliated European exchanges in Euronext, the Chicago Mercantile Exchange (CME), the Chicago Board Options Exchange (CBOE), and exchanges in other financial centers from London to Beijing. Other trades occur over computer-linked networks of dealers and traders such as the NASDAQ market and Instinet. Still other trades are made through agreements and contracts negotiated directly between seller and buyer with no middleman involved. In modern times political boundaries scarcely impede the flow of these transactions, so we now think of there being a “world” financial market. Worldwide, the process involves a staggering expenditure of valuable human labor and physical resources. Yet, when the day’s trading is done, not one single intrinsically valued physical commodity has been produced. Is this not remarkable?

1.1.2 Financial Markets as Transportation Agents

What justifies and explains this expenditure of resources? Since the transactions are made freely between consenting parties, each party to a trade must consider that what has been received compensates for what has been given up. Each party, if asked the reason for the trade it has made, would likely give an explanation that was highly circumstantial, depending on the transactor's particular situation and beliefs. Nevertheless, when we view assets through the economist's lens as time–state-contingent claims, a coherent image emerges: *Trading assets amounts to transferring resources across time and across states.* Thus, one who uses cash in a liquid, well managed money-market fund to buy a marketable, default-free, T -maturing discount bond gives up an indefinite time sequence of (almost) state-independent claims for a sequence of alternative state-*dependent* claims terminating with a state-independent receipt of principal value at T . The claims prior to T are those arising from potential sales of the bond before maturity, the amounts received depending on current conditions. Of course, the claims at all dates after any date $t \leq T$ are forfeited if the bond is sold at t . One who commits to hold the bond to T just makes a simple transfer across time. By contrast, one who trades the money-market shares for shares of common stock in XYZ company gives up the (almost) state-independent claims for an *indefinite* time sequence of claims that are highly state dependent. The exchange amounts to transferring or *transporting* claims from states that are unfavorable or merely neutral for XYZ to states that are favorable.

Once we recognize trading as such a transportation process, it is not so hard to understand why individuals would devote resources to the practice, any more than it is difficult to understand why we pay to have goods (and ourselves) moved from one place to another. We regard assets as being valued not for themselves but for the opportunities they afford for consumption of goods and services that do have intrinsic value. Just as goods and services are more valuable to us in one place than in another, opportunities for consumption are more valued at certain times and in certain states. Evidently, we are willing to pay enough to brokers, market makers, and financial firms to attract the resources they need to facilitate such trades. Indeed, we are sufficiently willing to allow governments at various levels to siphon off consumption opportunities that are generated by the transfers.

1.1.3 Why Is Transportation Desirable?

What is it that accounts for the subjective differences in value across times and states? Economists generally regard the different subjective valuations as arising from an inherent desire for “smoothness” in consumption, or, to turn it around, as a distaste for variation. We take out long-term loans to acquire durable goods that yield flows of benefits that last for many years; for example, we “issue” bonds in the form of mortgages to finance the purchases of our dwellings. This provides an alternative to postponing consumption at the desired level until enough is saved to finance it ourselves. We take the other side of the market, lending to banks through saving accounts and certificates of deposit (CDs) and buying bonds, to provide for consumption in

later years when other resources may be lacking. While the consumption opportunities that both activities open up are to some extent state dependent, the usual primary motivation is to transfer over time.

Transfers across states are made for two classes of reasons. One may begin to think that certain states are more likely to occur than considered previously, or one may begin to regard consumption in those states as more valuable if they do occur. In both cases it becomes more desirable to place “bets” on the roulette wheel’s stopping at those states. One places such bets by buying assets that offer higher cash values in the more valuable states—that is, by trading assets of lesser value in such states for those of higher value. Two individuals with different beliefs about the likelihood of future states, the value of consumption in those states, or a given asset’s entitlements to consumption in those states will want to trade the asset. They will do so if the consumption opportunities extracted by the various middlemen and governments are not too large. The “speculator” in assets is one who trades primarily to expand consumption opportunities in certain states. The “hedger” is one who trades mainly to preserve existing state-dependent opportunities. Claims for payoffs in the various states are continually being passed back and forth between and within these two classes of transactors.

1.1.4 What Vehicles Are Available?

The financial instruments that exist for making time–state transfers are almost too numerous to name. Governments at all levels issue bonds to finance current expenditures for public goods or transfers among citizens that are thought to promote social welfare. Some of these are explicitly or implicitly backed by the taxation authority of the issuer; others are tied to revenues generated by government-sponsored or government-aided entities. Corporate debt of medium to long maturity at initiation is traded on exchanges, and short-term corporate “paper” is shuffled around in the institutional “money” market. Such debt instruments of all sorts—short or long, corporate or government—are referred to as “fixed income” securities. Equity shares in corporations consist of “common” and “preferred” stocks, the latter offering prior claim to assets on liquidation and to revenues that fund payments of dividends. Most corporate equity is tradable and traded in markets, but private placements are sometimes made directly to institutions. There are exchange-traded funds that hold portfolios of bonds and of equities of various special classes (e.g., by industry, firm size, and risk class). Shares of these are traded on exchanges just as are listed stocks. Mutual funds offer stakes in other such portfolios of equities and bonds. These are managed by financial firms, with whom individuals must deal directly to purchase and redeem shares. There are physical commodities such as gold—and nowadays even petroleum—that do have intrinsic consumption value but are nevertheless held mainly or in part to facilitate time–state transfers. However, since the production side figures heavily in determining value, we do not consider these to be *financial* assets.

We refer to stocks, bonds, and investment commodities as *primary* assets, because their values in various states are not linked *contractually* to values of other assets. The classes of assets that are so contractually linked are referred to as *derivatives*, as

their values are derived from those of “underlying” primary financial assets or commodities. Thus, stock options—puts and calls—yield cash flows that are specified in terms of values of the underlying stocks during a stated period; values of commodity futures and forward contracts are specifically linked to prices of the underlying commodities; options and futures contracts on stock and bond indexes yield payoffs determined by the index levels, which in turn depend on prices of the component assets; values of interest-rate caps and swaps depend directly on the behavior of interest rates and ultimately on the values of debt instruments traded in fixed-income markets, lending terms set by financial firms, and actions of central banks. Terms of contracts for ordinary stock and index options and for commodity futures can be sufficiently standardized as to support the liquidity needed to trade in organized markets, such as the CBOE and CME. This permits one easily both to acquire the obligations and rights conferred by the instruments and to terminate them before the specified expiration dates. Thus, one buys an option either to get the right to exercise or to terminate the obligation arising from a previous net sale. Direct agreements between financial firms and individuals and nonfinancial businesses result in “structured” or “tailor-made” products that suit the individual circumstances. Typically, such specialized agreements must be maintained for the contractually specified terms or else terminated early by subsequent negotiation between the parties.

1.1.5 What Is There to Learn about Assets and Markets?

Viewing assets as time–state claims and markets as transporters of those claims does afford a useful conceptual perspective, but it does not give practical normative guidance to an investor, nor does it lead to specific predictions of how investors react to changing circumstances or of how their actions determine what we observe at market level. Without an objective way to define the various states of the world, their chances of occurring, and their implications for the values of specific assets, we can neither advise someone which assets to choose nor understand the choices they have made. We would like to do both these things. We would also like to have some understanding of how the collective actions of self-interested individuals and the functioning of markets wind up determining the prices of primary assets. We would like to know why there are, on average, systematic differences between the cash flows (per unit cost) that different classes of assets generate. We would like to know what drives the fluctuations in their prices over time. We would like to know whether there are in these fluctuations certain patterns that, if recognized, would enable one with some consistency to achieve higher cash flows; likewise, whether there is other publicly available information that would make this possible. Finally, we would like to see how prices of derivative assets prior to expiration relate to the prices of traded primary assets and current conditions generally. In the chapters that follow we will see some of the approaches that financial economists have taken over the years to address issues such as these. Although the time–state framework is not directly used, thinking in these terms can sometimes help us see the essential features of other approaches.

1.1.6 Why the Need for *Quantitative Finance*?

We want to know not just what typically happens but *why* things happen as they do, and attaining such understanding requires more than merely documenting empirical regularities. Although we concede up front that the full complexity of markets is beyond our comprehension, we still desire that the abstractions and simplifications on which we must rely yield useful predictions. We desire, *in addition*, that our abstract theories make us feel that they capture the essence of what is going on or else we would find them unsatisfying. The development of satisfying, predictive theories about quantifiable things requires building formal models, and the language in which we describe quantifiable things and express models is that of mathematics. Moreover, we need certain specific mathematical tools. If we regard the actors and transactors in financial markets as purposeful individuals, then we must think of them as having some way of ranking different outcomes and of striving to achieve the most preferred of these. Economists regard such endeavor as *optimizing* behavior and model it using the same tools of calculus and analysis that are used to find extrema of mathematical functions—that is, to find the peaks and troughs in the numerical landscape. But in financial markets nothing is certain; the financial landscape heaves and tosses through time in ways that we can by no means fully predict. Thus, the theories and predictions that we encounter in finance inevitably refer to *uncertain* quantities and future events. We must therefore supplement the tools of calculus and analysis with the developed mathematical framework for characterizing uncertainty—probability theory. Through the use of mathematical analysis and probability theory, *quantitative finance* enables us to attain more ambitious goals of understanding and predicting what goes on in financial markets.

1.2 WHERE WE GO FROM HERE

The two remaining chapters of this preliminary part of the book provide the necessary preparation in analysis and probability. For some, much of this will be a review of familiar concepts, and paging through it will refresh the memory. For others much of it will be new, and more thoughtful and deliberate reading will be required. However, no one who has not seen it before should expect to master the material on the first pass. The objective should be to get an overall sense of the concepts and remember where to look when they are needed. The treatment here is necessarily brief, so one will sometimes want to consult other sources.

Part II presents what most would consider the core of the “investment” side of financial theory. Starting with the basic arithmetic of bond prices and interest rates in Chapter 4, it progresses in the course of Chapters 5–10 through single-period portfolio theory and pricing models, theories and experimental evidence on choices under uncertainty, and empirical findings about marginal distributions of assets’ returns and about how prices vary over time. Chapter 11, “Stochastic Calculus,” is another “tools” chapter, placed here in proximity to the first exposure to models of prices that evolve in continuous time. Chapters 12 and 13 survey dynamic portfolio theory, which recognizes that people need to consider how current decisions affect constraints and

opportunities for the future. Chapter 14 looks at the implications of optimal dynamic choices and optimal information processing for the dynamic behavior of prices. Part II concludes with some empirical evidence of how well information is actually processed and how prices actually do vary over time.

The pricing models of Part II are based on a concept of market equilibrium in which prices attain values that make everyone content with their current holdings. Part III introduces an alternative paradigm of pricing by “arbitrage.” Within the time–state framework, pricing an asset by arbitrage amounts to assembling and valuing a collection of traded assets that offers (or can offer on subsequent reshuffling) the same time–state-dependent payoffs. If such a replicating package could be bought and sold for a price different from that of the reference asset, then buying the cheaper of the two and selling the other would yield an immediate, riskless profit. This is one type of arbitrage. Another would be a trade that confers for free some positive-valued time–state-contingent claim—that is, a free bet on some slot on the wheel. Presuming that markets of self-interested and reasonably perceptive individuals do not let such opportunities last for long, we infer that the prices of any asset and its replicating portfolio should quickly converge. Chapter 16 takes a first look at arbitrage pricing within a static setting where replication can be accomplished through buy-and-hold portfolios. Chapter 17 introduces the Black–Scholes–Merton theory for pricing by dynamic replication. We will see there that options and other derivatives can be replicated by acquiring and rebalancing portfolios over time, so long as prices of underlying assets are not too erratic. The implications of the model and its empirical relevance in today’s markets are considered in the chapter that follows. When underlying prices are discontinuous or are buffeted by extraneous influences that cannot be hedged away, not even dynamic replication will be possible. Nevertheless, through “martingale pricing” it is possible at least to set prices for derivatives and structured products in a way that affords no opportunity for arbitrage. The required techniques are explained and applied in the book’s concluding chapters.

