1

Introduction: Probability and Finance as a Game[1]

We propose a framework for the theory and use of mathematical probability that rests more on game theory than on measure theory. This new framework merits attention on purely mathematical grounds, for it captures the basic intuitions of probability simply and effectively. It is also of philosophical and practical interest. It goes deeper into probability's conceptual roots than the established measure-theoretic framework, it is better adapted to many practical problems, and it clarifies the close relationship between probability theory and finance theory.

Jean Ville (1910–1988) as a student at the *École Normale Supérieure* in Paris. His study of martingales helped inspire our framework for probability.

From the viewpoint of game theory, our framework is very simple. Its most essential elements were already present in Jean Ville's 1939 book, *Étude critique de la notion de collectif*, which introduced martingales into probability theory. Following Ville, we consider only two players. They alternate moves, each is immediately informed of the other's moves, and one or the other wins.

[1]This is the first chapter of *Probability and Finance: It's Only a Game!*, by Glenn Shafer and Vladimir Vovk.

In such a game, one player has a winning strategy (§4.6), and so we do not need the subtle solution concepts now at the center of game theory in economics and the other social sciences.

Our framework is a straightforward but rigorous elaboration, with no extraneous mathematical or philosophical baggage, of two ideas that are fundamental to both probability and finance:

- **The Principle of Pricing by Dynamic Hedging.** When simple gambles can be combined over time to produce more complex gambles, prices for the simple gambles determine prices for the more complex gambles.
- **The Hypothesis of the Impossibility of a Gambling System.** Sometimes we hypothesize that no system for selecting gambles from those offered to us can both (1) be certain to avoid bankruptcy and (2) have a reasonable chance of making us rich.

The principle of pricing by dynamic hedging can be discerned in the letters of Blaise Pascal to Pierre de Fermat in 1654, at the very beginning of mathematical probability, and it re-emerged in the last third of the twentieth century as one of the central ideas of finance theory. The hypothesis of the impossibility of a gambling system also has a long history in probability theory, dating back at least to Cournot, and it is related to the efficient-markets hypothesis, which has been studied in finance theory since the 1970s. We show that in a rigorous game-theoretic framework, these two ideas provide an adequate mathematical and philosophical starting point for probability and its use in finance and many other fields. No additional apparatus such as measure theory is needed to get probability off the ground mathematically, and no additional assumptions or philosophical explanations are needed to put probability to use in the world around us.

Probability becomes game-theoretic as soon as we treat the expected values in a probability model as prices in a game. These prices may be offered to an imaginary player who stands outside the world and bets on what the world will do, or they may be offered to an investor whose participation in a market constitutes a bet on what the market will do. In both cases, we can learn a great deal by thinking in game-theoretic terms. Many of probability's theorems turn out to be theorems about the existence of winning strategies for the player who is betting on what the world or market will do. The theorems are simpler and clearer in this form, and when they are in this form, we are in a position to reduce the assumptions we make—the number of prices we assume are offered—down to the minimum needed for the theorems to hold. This parsimony is potentially very valuable in practical work, for it allows and encourages clarity about the assumptions we need and are willing to take seriously.

Defining a probability measure on a sample space means recommending a definite price for each uncertain payoff that can be defined on the sample space, a price at which one might buy or sell the payoff. Our framework requires much less than this. We may be given only a few prices, and some of them may be one-sided—certified only for selling, not for buying, or vice versa. From these given prices, using dynamic hedging, we may obtain two-sided prices for some additional payoffs, but only upper and lower prices for others.

The measure-theoretic framework for probability, definitively formulated by Andrei Kolmogorov in 1933, has been praised for its philosophical neutrality: it can guide our mathematical work with probabilities no matter what meaning we want to give to these probabilities. Any numbers that satisfy the axioms of measure may be called probabilities, and it is up to the user whether to interpret them as frequencies, degrees of belief, or something else. Our game-theoretic framework is equally open to diverse interpretations, and its greater conceptual depth enriches these interpretations. Interpretations and uses of probability differ not only in the source of prices but also in the role played by the hypothesis of the impossibility of a gambling system.

Our framework differs most strikingly from the measure-theoretic framework in its ability to model open processes—processes that are open to influences we cannot model even probabilistically. This openness can, we believe, enhance the usefulness of probability theory in domains where our ability to control and predict is substantial but very limited in comparison with the sweep of a deterministic model or a probability measure.

From a mathematical point of view, the first test of a framework for probability is how elegantly it allows us to formulate and prove the subject's principal theorems, especially the classical limit theorems: the law of large numbers, the law of the iterated logarithm, and the central limit theorem. In Part I, we show how our game-theoretic framework meets this test. We contend that it does so better than the measure-theoretic framework. Our game-theoretic proofs sometimes differ little from standard measure-theoretic proofs, but they are more transparent. Our game-theoretic limit theorems are more widely applicable than their measure-theoretic counterparts, because they allow reality's moves to be influenced by moves by other players, including experimenters, professionals, investors, and citizens. They are also mathematically more powerful; the measure-theoretic counterparts follow from them as easy corollaries. In the case of the central limit theorem, we also obtain an interesting one-sided generalization, applicable when we have only upper bounds on the variability of individual deviations.

In Part II, we explore the use of our framework in finance. We call Part II "Finance without Probability" for two reasons. First, the two ideas that we consider fundamental to probability—the principle of pricing by dynamic hedging and the hypothesis of the impossibility of a gambling system—are also native to finance theory, and the exploitation of them in their native form in finance theory does not require extrinsic stochastic modeling. Second, we contend that the extrinsic stochastic modeling that does sometimes seem to be needed in finance theory can often be advantageously replaced by the further use of markets to set prices. Extrinsic stochastic modeling can also be accommodated in our framework, however, and Part II includes a game-theoretic treatment of diffusion processes, the extrinsic stochastic models that are most often used in finance and are equally important in a variety of other fields.

In the remainder of this introduction, we elaborate our main ideas in a relatively informal way. We explain how dynamic hedging and the impossibility of a gambling system can be expressed in game-theoretic terms, and how this leads to game-theoretic formulations of the classical limit theorems. Then we discuss the diversity of ways in which game-theoretic probability can be used, and we summarize how our relentlessly game-theoretic point of view can strengthen the theory of finance.

1.1 A GAME WITH THE WORLD

At the center of our framework is a sequential game with two players. The game may have many—perhaps infinitely many—rounds of play. On each round, Player I bets on what will happen, and then Player II decides what will happen. Both players have perfect information; each knows about the other's moves as soon as they are made.

In order to make their roles easier to remember, we usually call our two players Skeptic and World. Skeptic is Player I; World is Player II. This terminology is inspired by the idea of testing a probabilistic theory. Skeptic, an imaginary scientist who does not interfere with what happens in the world, tests the theory by repeatedly gambling imaginary money at prices the theory offers. Each time, World decides what does happen and hence how Skeptic's imaginary capital changes. If this capital becomes too large, doubt is cast on the theory. Of course, not all uses of mathematical probability, even outside of finance, are scientific. Sometimes the prices tested by Skeptic express personal choices rather than a scientific theory, or even serve merely as a straw man. But the idea of testing a scientific theory serves us well as a guiding example.

In the case of finance, we sometimes substitute the names Investor and Market for Skeptic and World. Unlike Skeptic, Investor is a real player, risking real money. On each round of play, Investor decides what investments to hold, and Market decides how the prices of these investments change and hence how Investor's capital changes.

Dynamic Hedging

The principle of pricing by dynamic hedging applies to both probability and finance, but the word "hedging" comes from finance. An investor hedges a risk by buying and selling at market prices, possibly over a period of time, in a way that balances the

Table 1.1 Instead of the uninformative names Player I and Player II, we usually call our players Skeptic and World, because it is easy to remember that World decides while Skeptic only bets. In the case of finance, we often call the two players Investor and Market.

	PROBABILITY	**FINANCE**
Player I bets on what will happen.	**Skeptic** bets against the probabilistic predictions of a scientific theory.	**Investor** bets by choosing a portfolio of investments.
Player II decides what happens.	**World** decides how the predictions come out.	**Market** decides how the price of each investment changes.

risk. In some cases, the risk can be eliminated entirely. If, for example, Investor has a financial commitment that depends on the prices of certain securities at some future time, then he may be able to cover the commitment exactly by investing shrewdly in the securities during the rounds of play leading up to that future time. If the initial capital required is $\$\alpha$, then we may say that Investor has a strategy for turning $\$\alpha$ into the needed future payoff. Assuming, for simplicity, that the interest rate is zero, we may also say that $\$\alpha$ is the game's price for the payoff. This is the principle of pricing by dynamic hedging. (We assume throughout this chapter and in most of the rest of the book that the interest rate is zero. This makes our explanations and mathematics simpler, with no real loss in generality, because the resulting theory extends readily to the case where the interest rate is not zero: see §12.1.)

As it applies to probability, the principle of pricing by dynamic hedging says simply that the prices offered to Skeptic on each round of play can be compounded to obtain prices for payoffs that depend on more than one of World's moves. The prices for each round may include probabilities for what World will do on that round, and the global prices may include probabilities for World's whole sequence of play. We usually assume that the prices for each round are given either at the beginning of the game or as the game is played, and prices for longer-term gambles are derived. But when the idea of a probability game is used to study the world, prices may sometimes be derived in the opposite direction. The principle of pricing by dynamic hedging then becomes merely a principle of coherence, which tells us how prices at different times should fit together.

We impose no general rules about how many gambles are offered to Skeptic on different rounds of the game. On some rounds, Skeptic may be offered gambles on every aspect of World's next move, while on other rounds, he may be offered no gambles at all. Thus our framework always allows us to model what science models and to leave unmodeled what science leaves unmodeled.

The Fundamental Interpretative Hypothesis

In contrast to the principle of pricing by dynamic hedging, the hypothesis of the impossibility of a gambling system is optional in our framework. The hypothesis boils down, as we explain in §1.3, to the supposition that events with zero or low probability are unlikely to occur (or, more generally, that events with zero or low upper probability are unlikely to occur). This supposition is fundamental to many uses of probability, because it makes the game to which it is applied into a theory about the world. By adopting the hypothesis, we put ourselves in a position to test the prices in the game: if an event with zero or low probability does occur, then we can reject the game as a model of the world. But we do not always adopt the hypothesis. We do not always need it when the game is between Investor and Market, and we do not need it when we interpret probabilities subjectively, in the sense advocated by Bruno de Finetti. For de Finetti and his fellow neosubjectivists, a person's subjective prices are nothing more than that; they are merely prices that systematize the person's choices among risky options. See §1.4 and §2.6.

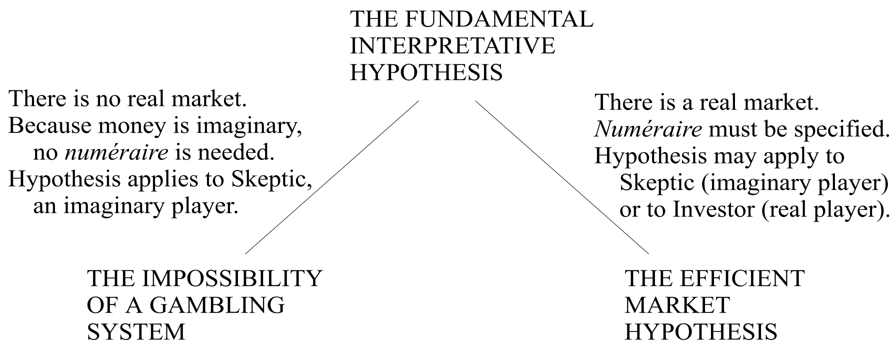

Fig. 1.1 The fundamental interpretative hypothesis in probability and finance.

We have a shorter name for the hypothesis of the impossibility of a gambling system: we call it the *fundamental interpretative hypothesis* of probability. It is interpretative because it tells us what the prices and probabilities in the game to which it is applied mean in the world. It is not part of our mathematics. It stands outside the mathematics, serving as a bridge between the mathematics and the world.

When we are working in finance, where our game describes a real market, we use yet another name for our fundamental hypothesis: we call it the *efficient-market hypothesis*. The efficient-market hypothesis, as applied to a particular financial market, in which particular securities are bought and sold over time, says that an investor (perhaps a real investor named Investor, or perhaps an imaginary investor named Skeptic) cannot become rich trading in this market without risking bankruptcy. In order to make such a hypothesis precise, we must specify not only whether we are talking about Investor or Skeptic, but also the *numéraire*—the unit of measurement in which this player's capital is measured. We might measure this capital in nominal terms (making a monetary unit, such as a dollar or a ruble, the *numéraire*), we might measure it relative to the total value of the market (making some convenient fraction of this total value the *numéraire*), or we might measure it relative to a risk-free bond (which is then the *numéraire*), and so on. Thus the efficient-market hypothesis can take many forms. Whatever form it takes, it is subject to test, and it determines upper and lower probabilities that have empirical meaning.

Since about 1970, economists have debated an efficient-markets hypothesis, with *markets* in the plural. This hypothesis says that financial markets are efficient in general, in the sense that they have already eliminated opportunities for easy gain. As we explain in Part II (§9.4 and Chapter 15), our efficient-market hypothesis has the same rough rationale as the efficient-markets hypothesis and can often be tested in similar ways. But it is much more specific. It requires that we specify the particular securities that are to be included in the market, the exact rule for accumulating capital, and the *numéraire* for measuring this capital.

Open Systems within the World

Our austere picture of a game between Skeptic and World can be filled out in a great variety of ways. One of the most important aspects of its potential lies in the possibility of dividing World into several players. For example, we might divide World into three players:

> Experimenter, who decides what each round of play will be about.
> Forecaster, who sets the prices.
> Reality, who decides the outcomes.

This division reveals the open character of our framework. The principle of pricing by dynamic hedging requires Forecaster to give coherent prices, and the fundamental interpretative hypothesis requires Reality to respect these prices, but otherwise all three players representing World may be open to external information and influence. Experimenter may have wide latitude in deciding what experiments to perform. Forecaster may use information from outside the game to set prices. Reality may also be influenced by unpredictable outside forces, as long as she acts within the constraints imposed by Forecaster.

Many scientific models provide testable probabilistic predictions only subsequent to the determination of many unmodeled auxiliary factors. The presence of Experimenter in our framework allows us to handle these models very naturally. For example, the standard mathematical formalization of quantum mechanics in terms of Hilbert spaces, due to John von Neumann, fits readily into our framework. The scientist who decides what observables to measure is Experimenter, and quantum theory is Forecaster (§8.4).

Weather forecasting provides an example where information external to a model is used for prediction. Here Forecaster may be a person or a very complex computer program that escapes precise mathematical definition because it is constantly under development. In either case, Forecaster will use extensive external information—weather maps, past experience, etc. If Forecaster is required to announce every evening a probability for rain on the following day, then there is no need for Experimenter; the game has only three players, who move in this order:

Forecaster, Skeptic, Reality.

Forecaster announces odds for rain the next day, Skeptic decides whether to bet for or against rain and how much, and Reality decides whether it rains. The fundamental interpretative hypothesis, which says that Skeptic cannot get rich, can be tested by any strategy for betting at Forecaster's odds.

It is more difficult to make sense of the weather forecasting problem in the measure-theoretic framework. The obvious approach is to regard the forecaster's probabilities as conditional probabilities given what has happened so far. But because the forecaster is expected to learn from his experience in giving probability forecasts, and because he uses very complex and unpredictable external information, it makes no sense to interpret his forecasts as conditional probabilities in a probability distribution formulated at the outset. And the forecaster does not construct a probability distribution along the way; this would involve constructing probabilities for what will happen on the next day not only conditional on what has happened so far but also conditional on what might have happened so far.

In the 1980s, A. Philip Dawid proposed that the forecasting success of a probability distribution for a sequence of events should be evaluated using only the actual outcomes and the sequence of forecasts (conditional probabilities) to which these outcomes give rise, without reference to other aspects of the probability distribution. This is Dawid's *prequential principle* [82]. In our game-theoretic framework, the prequential principle is satisfied automatically, because the probability forecasts provided by Forecaster and the outcomes provided by Reality are all we have. So long as Forecaster does not adopt a strategy, no probability distribution is even defined.

The explicit openness of our framework makes it well suited to modeling systems that are open to external influence and information, in the spirit of the nonparametric, semiparametric, and martingale models of modern statistics and the even looser predictive methods developed in the study of machine learning. It also fits the open spirit of modern science, as emphasized by Karl Popper [250]. In the nineteenth century, many scientists subscribed to a deterministic philosophy inspired by Newtonian physics: at every moment, every future aspect of the world should be predictable by a superior intelligence who knows initial conditions and the laws of nature. In the twentieth century, determinism was strongly called into question by further advances in physics, especially in quantum mechanics, which now insists that some fundamental phenomena can be predicted only probabilistically. Probabilists sometimes imagine that this defeat allows a retreat to a probabilistic generalization of determinism: science should give us probabilities for everything that might happen in the future. In fact, however, science now describes only islands of order in an unruly universe. Modern scientific theories make precise probabilistic predictions only about some aspects of the world, and often only after experiments have been designed and prepared. The game-theoretic framework asks for no more.

Skeptic and World Always Alternate Moves

Most of the mathematics in this book is developed for particular examples, and as we have just explained, many of these examples divide World into multiple players. It is important to notice that this division of World into multiple players does not invalidate the simple picture in which Skeptic and World alternate moves, with Skeptic betting on what World will do next, because we will continue to use this simple picture in our general discussions, in the next section and in later chapters.

One way of seeing that the simple picture is preserved is to imagine that Skeptic moves just before each of the players who constitute World, but that only the move just before Reality can result in a nonzero payoff for Skeptic. Another way, which we will find convenient when World is divided into Forecaster and Reality, is to add just one dummy move by Skeptic, at the beginning of the game, and then to group each of Forecaster's later moves with the preceding move by Reality, so that the order of play becomes

Skeptic, Forecaster, Skeptic, (Reality, Forecaster),
Skeptic, (Reality, Forecaster), ...

Either way, Skeptic alternates moves with World.

The Science of Finance

Other players sometimes intrude into the game between Investor and Market. Finance is not merely practice; there is a theory of finance, and our study of it will sometimes require that we bring Forecaster and Skeptic into the game. This happens in several different ways. In Chapter 14, where we give a game-theoretic reading of the usual stochastic treatment of option pricing, Forecaster represents a probabilistic theory about the behavior of the market, and Skeptic tests this theory. In our study of the efficient-market hypothesis (Chapter 15), in contrast, the role of Forecaster is played by Opening Market, who sets the prices at which Investor, and perhaps also Skeptic, can buy securities. The role of Reality is then played by Closing Market, who decides how these investments come out.

In much of Part II, however, especially in Chapters 10–13, we study games that involve Investor and Market alone. These may be the most important market games that we study, because they allow conclusions based solely on the structure of the market, without appeal to any theory about the efficiency of the market or the stochastic behavior of prices.

1.2 THE PROTOCOL FOR A PROBABILITY GAME

Specifying a game fully means specifying the moves available to the players—we call this the *protocol* for the game—and the rule for determining the winner. Both of these elements can be varied in our game between Skeptic and World, leading to many different games, all of which we call *probability games*. The protocol determines the sample space and the prices (in general, upper and lower prices) for variables. The rule for determining the winner can be adapted to the particular theorem we want to prove or the particular problem where we want to use the framework. In this section we consider only the protocol.

The general theory sketched in this section applies to most of the games studied in this book, including those where Investor is substituted for Skeptic and Market for World. (The main exceptions are the games we use in Chapter 13 to price American options.) We will develop this general theory in more detail in Chapters 7 and 8.

The Sample Space

The protocol for a probability game specifies the moves available to each player, Skeptic and World, on each round. This determines, in particular, the sequences of moves World may make. These sequences—the possible complete sequences of play by World—constitute the *sample space* for the game. We designate the sample space

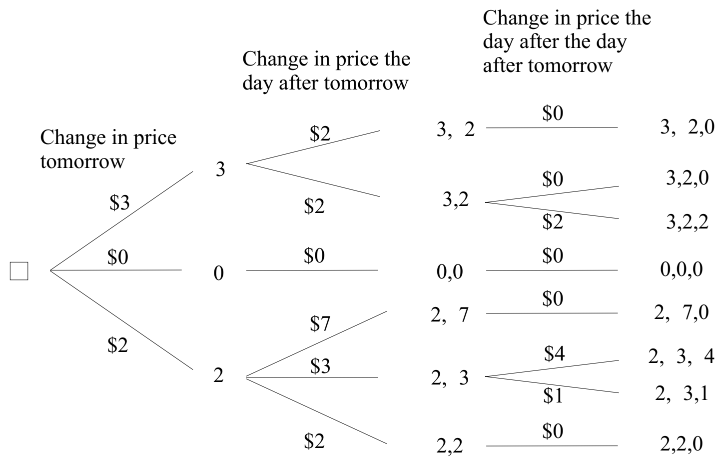

Fig. 1.2 An unrealistic sample space for changes in the price of a stock. The steps in the tree represent possible moves by World (in this case, the market). The nodes (situations) record the moves made by World so far. The initial situation is designated by □. The terminal nodes record complete sequences of play by World and hence can be identified with the paths that constitute the sample space. The example is unrealistic because in a real stock market there is a wide range of possible changes for a stock's price at each step, not just two or three.

by Ω, and we call its elements *paths*. The moves available to World may depend on moves he has previously made. But we assume that they do not depend on moves Skeptic has made. Skeptic's bets do not affect what is possible in the world, although World may consider them in deciding what to do next.

We can represent the dependence of World's possible moves on his previous moves in terms of a tree whose paths form the sample space, as in Figure 1.2. Each node in the tree represents a *situation*, and the branches immediately to the right of a nonterminal situation represent the moves World may make in that situation. The initial situation is designated by □.

Figure 1.2 is finite: there are only finitely many paths, and every path terminates after a finite number of moves. We do not assume finiteness in general, but we do pay particular attention to the case where every path terminates; in this case we say the game is *terminating*. If there is a bound on the length of the paths, then we say the game has a *finite horizon*. If none of the paths terminate, we say the game has an *infinite horizon*.

In general, we think of a situation (a node in the tree) as the sequence of moves made by World so far, as explained in the caption of Figure 1.2. So in a terminating game, we may identify the terminal situation on each path with that path; both are the same sequence of moves by World.

In measure-theoretic probability theory, a real-valued function on the sample space is called a *random variable*. Avoiding the implication that we have defined a

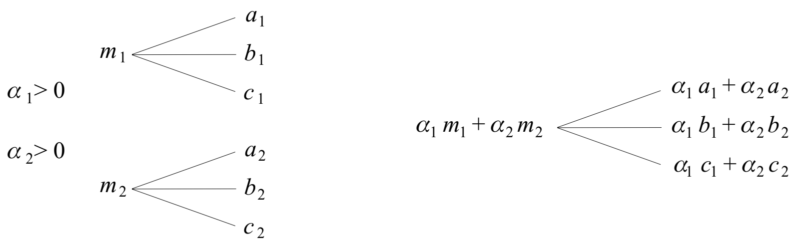

Fig. 1.3 Forming a nonnegative linear combination of two gambles. In the first gamble Skeptic pays m_1 in order to get a_1, b_1, or c_1 in return, depending on how things come out. In the second gamble, he pays m_2 in order to get a_2, b_2, or c_2 in return.

probability measure on the sample space, and also whatever other ideas the reader may associate with the word "random", we call such a function simply a *variable*. In the example of Figure 1.2, the variables include the prices for the stock for each of the next three days, the average of the three prices, the largest of the three prices, and so on. We also follow established terminology by calling a subset of the sample space an *event*.

Moves and Strategies for Skeptic

To complete the protocol for a probability game, we must also specify the moves Skeptic may make in each situation. Each move for Skeptic is a gamble, defined by a price to be paid immediately and a payoff that depends on World's following move. The gambles among which Skeptic may choose may depend on the situation, but we always allow him to combine available gambles and to take any fraction or multiple of any available gamble. We also allow him to borrow money freely without paying interest. So he can take any nonnegative linear combination of any two available gambles, as indicated in Figure 1.3.

We call the protocol *symmetric* if Skeptic is allowed to take either side of any available gamble. This means that whenever he can buy the payoff x at the price m, he can also sell x at the price m. Selling x for m is the same as buying $-x$ for $-m$ (Figure 1.4). So a symmetric protocol is one in which the gambles available to Skeptic in each situation form a linear space; he may take any linear combination of the available gambles, whether or not the coefficients in the linear combination are nonnegative. If we neglect bid-ask spreads and transaction costs, then protocols based on market prices are symmetric, because one may buy as well as sell a security at its market price. Protocols corresponding to complete probability measures are also symmetric. But many of the protocols we will study in this book are asymmetric.

A *strategy* for Skeptic is a plan for how to gamble in each nonterminal situation he might encounter. His strategy together with his initial capital determine his capital in every situation, including terminal situations. Given a strategy $\mathcal{P}$ and a situation t, we write $\mathcal{K}^{\mathcal{P}}(t)$ for Skeptic's capital in t if he starts with capital 0 and follows $\mathcal{P}$. In

Fig. 1.4 Taking the gamble on the left means paying m and receiving a, b, or c in return. Taking the other side means receiving m and paying a, b, or c in return—i.e., paying $-m$ and receiving $-a$, $-b$, or $-c$ in return. This is the same as taking the gamble on the right.

	Meaning	**Net payoff**	**Simulated satisfactorily by $\mathcal{P}$ if**
Buy x for α	Pay α, get x	$x - \alpha$	$\mathcal{K}^{\mathcal{P}} \geq x - \alpha$
Sell x for α	Get α, pay x	$\alpha - x$	$\mathcal{K}^{\mathcal{P}} \geq \alpha - x$

Table 1.2 How a strategy $\mathcal{P}$ in a probability game can simulate the purchase or sale of a variable x.

the terminating case, we may also speak of the capital a strategy produces at the end of the game. Because we identify each path with its terminal situation, we may write $\mathcal{K}^{\mathcal{P}}(\xi)$ for Skeptic's final capital when he follows $\mathcal{P}$ and World takes the path ξ.

Upper and Lower Prices

By adopting different strategies in a probability game, Skeptic can simulate the purchase and sale of variables. We can price variables by considering when this succeeds. In order to explain this idea as clearly as possible, we make the simplifying assumption that the game is terminating.

A strategy simulates a transaction satisfactorily for Skeptic if it produces at least as good a net payoff. Table 1.2 summarizes how this applies to buying and selling a variable x. As indicated there, $\mathcal{P}$ simulates buying x for α satisfactorily if $\mathcal{K}^{\mathcal{P}} \geq x - \alpha$. This means that

$$\mathcal{K}^{\mathcal{P}}(\xi) \geq x(\xi) - \alpha$$

for every path ξ in the sample space Ω. When Skeptic has a strategy $\mathcal{P}$ satisfying $\mathcal{K}^{\mathcal{P}} \geq x - \alpha$, we say he *can buy x for α*. Similarly, when he has a strategy $\mathcal{P}$ satisfying $\mathcal{K}^{\mathcal{P}} \geq \alpha - x$, we say he *can sell x for α*. These are two sides of the same coin: selling x for α is the same as buying $-x$ for $-\alpha$.

Given a variable x, we set

$$\overline{\mathbb{E}}\, x := \inf \left\{ \alpha \mid \text{ there is some strategy } \mathcal{P} \text{ such that } \mathcal{K}^{\mathcal{P}} \geq x - \alpha \right\}.^2 \qquad (1.1)$$

[2]We use := to mean "equal by definition"; the right-hand side of the equation is the definition of the left-hand side.

We call $\overline{\mathbb{E}}\,x$ the *upper price* of x or the *cost* of x; it is the lowest price at which Skeptic can buy x. (Because we have made no compactness assumptions about the protocol—and will make none in the sequel—the infimum in (1.1) may not be attained, and so strictly speaking we can only be sure that Skeptic can buy x for $\overline{\mathbb{E}}\,x + \epsilon$ for every $\epsilon > 0$. But it would be tedious to mention this constantly, and so we ask the reader to indulge the slight abuse of language involved in saying that Skeptic can buy x for $\overline{\mathbb{E}}\,x$.)

Similarly, we set

$$\underline{\mathbb{E}}\,x := \sup\left\{\alpha \mid \text{ there is some strategy } \mathcal{P} \text{ such that } \mathcal{K}^{\mathcal{P}} \geq \alpha - x\right\}. \tag{1.2}$$

We call $\underline{\mathbb{E}}\,x$ the *lower price* of x or the *scrap value* of x; it is the highest price at which Skeptic can sell x.

It follows from (1.1) and (1.2), and also directly from the fact that selling x for α is the same as buying $-x$ for $-\alpha$, that

$$\underline{\mathbb{E}}\,x = -\overline{\mathbb{E}}[-x]$$

for every variable x.

The idea of hedging provides another way of talking about upper and lower prices. If we have an obligation to pay something at the end of the game, then we hedge this obligation by trading in such a way as to cover the payment no matter what happens. So we say that the strategy $\mathcal{P}$ *hedges* the obligation y if

$$\mathcal{K}^{\mathcal{P}}(\xi) \geq y(\xi) \tag{1.3}$$

for every path ξ in the sample space Ω. Selling a variable x for α results in a net obligation of $x - \alpha$ at the end of the game. So $\mathcal{P}$ hedges selling x for α if $\mathcal{P}$ hedges $x - \alpha$, that is, if $\mathcal{P}$ simulates buying x for α. Similarly, $\mathcal{P}$ hedges buying x for α if $\mathcal{P}$ simulates selling x for α. So $\overline{\mathbb{E}}\,x$ is the lowest price at which selling x can be hedged, and $\underline{\mathbb{E}}\,x$ is the highest price at which buying it can be hedged, as indicated in Table 1.3.

These definitions implicitly place Skeptic at the beginning of the game, in the initial situation $\Box$. They can also be applied, however, to any other situation; we

Table 1.3 Upper and lower price described in terms of simulation and described in terms of hedging. Because hedging the sale of x is the same as simulating the purchase of x, and vice versa, the two descriptions are equivalent.

	Name	**Description in terms of the simulation of buying and selling**	**Description in terms of hedging**
$\overline{\mathbb{E}}\,x$	Upper price of x	Lowest price at which Skeptic can buy x	Lowest selling price for x Skeptic can hedge
$\underline{\mathbb{E}}\,x$	Lower price of x	Highest price at which Skeptic can sell x	Highest buying price for x Skeptic can hedge

simply consider Skeptic's strategies for play from that situation onward. We write $\overline{\mathbb{E}}_t\, x$ and $\underline{\mathbb{E}}_t\, x$ for the upper and lower price, respectively, of the variable x in the situation t.

Upper and lower prices are interesting only if the gambles Skeptic is offered do not give him an opportunity to make money for certain. If this condition is satisfied in situation t, we say that the protocol is *coherent* in t. In this case,

$$\underline{\mathbb{E}}_t\, x \le \overline{\mathbb{E}}_t\, x$$

for every variable x, and

$$\overline{\mathbb{E}}_t\, \mathbf{0} = \underline{\mathbb{E}}_t\, \mathbf{0} = 0,$$

where **0** denotes the variable whose value is 0 on every path in Ω.

When $\underline{\mathbb{E}}_t\, x = \overline{\mathbb{E}}_t\, x$, we call their common value the *exact price* or simply the *price* for x in t and designate it by $\mathbb{E}_t\, x$. Such prices have the properties of expected values in measure-theoretic probability theory, but we avoid the term "expected value" in order to avoid suggesting that we have defined a probability measure on our sample space. We do, however, use the word "variance"; when $\mathbb{E}_t\, x$ exists, we set

$$\overline{\mathbb{V}}_t\, x := \overline{\mathbb{E}}_t (x - \mathbb{E}_t\, x)^2 \quad \text{and} \quad \underline{\mathbb{V}}_t\, x := \underline{\mathbb{E}}_t (x - \mathbb{E}_t\, x)^2,$$

and we call them, respectively, the *upper variance* of x in t and the *lower variance* of x in t. If $\overline{\mathbb{V}}_t\, x$ and $\underline{\mathbb{V}}_t\, x$ are equal, we write $\mathbb{V}_t\, x$ for their common value; this is the (game-theoretic)*variance* of x in t.

When the game is not terminating, definitions (1.1), (1.2), and (1.3) do not work, because $\mathcal{P}$ may fail to determine a final capital for Skeptic when World takes an infinite path; if there is no terminal situation on the path ξ, then $\mathcal{K}^{\mathcal{P}}(t)$ may or may not converge to a definite value as t moves along ξ. Of the several ways to fix this, we prefer the simplest: we say that $\mathcal{P}$ hedges y if on every path ξ the capital $\mathcal{K}^{\mathcal{P}}(t)$ eventually reaches $y(\xi)$ and stays at or above it, and we similarly modify (1.1) and (1.2). We will study this definition in §8.3. On the whole, we make relatively little use of upper and lower price for nonterminating probability games, but as we explain in the next section, we do pay great attention to one special case, the case of probabilities exactly equal to zero or one.

1.3 THE FUNDAMENTAL INTERPRETATIVE HYPOTHESIS

The fundamental interpretative hypothesis asserts that no strategy for Skeptic can both (1) be certain to avoid bankruptcy and (2) have a reasonable chance of making Skeptic rich. Because it contains the undefined term "reasonable chance", this hypothesis is not a mathematical statement; it is neither an axiom nor a theorem. Rather it is an interpretative statement. It gives meaning in the world to the prices in the probability game. Once we have asserted that Skeptic does not have a reasonable chance of multiplying his initial capital substantially, we can identify other likely and unlikely events and calibrate just how likely or unlikely they are. An event is unlikely if its happening would give an opening for Skeptic to multiply his initial capital substantially, and it is the more unlikely the more substantial this multiplication is.

We use two distinct versions of the fundamental interpretative hypothesis, one *finitary* and one *infinitary*:

- **The Finitary Hypothesis.** No strategy for Skeptic can both (1) be certain to avoid bankruptcy and (2) have a reasonable chance of multiplying his initial capital by a large factor. (We usually use this version in terminating games.)

- **The Infinitary Hypothesis.** No strategy for Skeptic can both (1) be certain to avoid bankruptcy and (2) have a reasonable chance of making him infinitely rich. (We usually use this version in infinite-horizon games.)

Because our experience with the world is finite, the finitary hypothesis is of more practical use, but the infinitary hypothesis often permits clearer and more elegant mathematical statements. As we will show in Part I, the two forms lead to the two types of classical limit theorems. The finitary hypothesis leads to the weak limit theorems: the weak law of large numbers and the central limit theorem. The infinitary hypothesis leads to the strong limit theorems: the strong law of large numbers and the law of the iterated logarithm.

It is easy for World to satisfy the fundamental interpretative hypothesis in a probability game with a coherent protocol, for he can always move so that Skeptic does not make money. But becoming rich is not Skeptic's only goal in the games we study. In many of these games, Skeptic wins *either* if he becomes rich *or* if World's moves satisfy some other condition E. If Skeptic has a winning strategy in such a game, then the fundamental interpretative hypothesis authorizes us to conclude that E will happen. In order to keep Skeptic from becoming rich, World must move so as to satisfy E.

Low Probability and High Probability

In its finitary form, the fundamental interpretative hypothesis provides meaning to small upper probabilities and large lower probabilities.

We can define upper and lower probabilities formally as soon as we have the concepts of upper and lower price. As we mentioned earlier, an *event* is a subset of the sample space. Given an event E, we define its *indicator variable* $\mathbb{I}_E$ by

$$\mathbb{I}_E(\xi) := \begin{cases} 1 \text{ if } \xi \in E \\ 0 \text{ if } \xi \notin E. \end{cases}$$

Then we define its *upper probability* by

$$\overline{\mathbb{P}}\, E := \overline{\mathbb{E}}\, \mathbb{I}_E \tag{1.4}$$

and its *lower probability* by

$$\underline{\mathbb{P}}\, E := \underline{\mathbb{E}}\, \mathbb{I}_E\,. \tag{1.5}$$

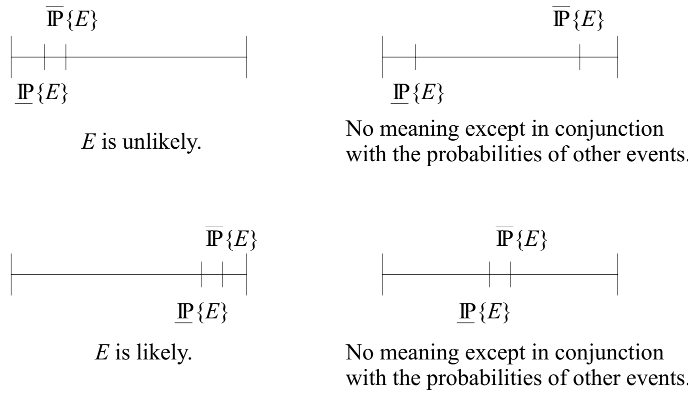

Fig. 1.5 Only extreme probabilities have meaning in isolation.

Assuming the protocol is coherent, upper and lower probability obey

$$0 \leq \underline{\mathbb{P}}\, E \leq \overline{\mathbb{P}}\, E \leq 1 \tag{1.6}$$

and

$$\underline{\mathbb{P}}\, E = 1 - \overline{\mathbb{P}}\, E^{\mathrm{c}}. \tag{1.7}$$

Here E^{c} is the *complement* of E in Ω—the set of paths for World that are not in E, or the event that E does not happen.

What meaning can be attached to $\overline{\mathbb{P}}\, E$ and $\underline{\mathbb{P}}\, E$? The fundamental interpretative hypothesis answers this question when the two numbers are very close to zero. Suppose, for example, that $\overline{\mathbb{P}}\, E = 0.001$. (In this case, $\underline{\mathbb{P}}\, E$ is also close to zero; by (1.6), it is between 0 and 0.001.) Then Skeptic can buy $\mathbb{I}_E$ for 0.001. Because $\mathbb{I}_E \geq 0$, the purchase does not open him to possible bankruptcy, and yet it results in a thousandfold increase in his investment if E happens. The fundamental interpretative hypothesis says that this increase is unlikely and hence implies that E is unlikely.

Similarly, we may say that E is very likely to happen if $\underline{\mathbb{P}}\, E$ and hence also $\overline{\mathbb{P}}\, E$ are very close to one. Indeed, if $\underline{\mathbb{P}}\, E$ is close to one, then by (1.7), $\overline{\mathbb{P}}\, E^{\mathrm{c}}$ is close to zero, and hence it is unlikely that E^{c} will happen—that is, it is likely that E will happen.

These interpretations are summarized in Figure 1.5. If $\overline{\mathbb{P}}\, E$ and $\underline{\mathbb{P}}\, E$ are neither both close to zero nor both close to one, as on the right in the figure, then they have little or no meaning in isolation. But if they are both close to zero, then we may say that E has "low probability" and is unlikely to happen. And if they are both close to one, then we may say that E has "high probability" and is likely to happen.

Strictly speaking, we should speak of the probability of E only if $\overline{\mathbb{P}}\, E$ and $\underline{\mathbb{P}}\, E$ are exactly equal, for then their common value may be called the (game-theoretic) probability of E. But as the figure indicates, it is much more meaningful for the two values to both be close to zero or both be close to one than for them to be exactly equal.

The most important examples of low and high probability in this book occur in the two weak laws that we study in Chapters 6 and 7: the weak law of large numbers and the central limit theorem. The weak law of large numbers, in its simplest form, says that when Skeptic is offered even odds on each of a long sequence of events, the probability is high that the fraction of the events that happen will fall within a small interval around $1/2$: an interval that may be narrowed as the number of the events increases. The central limit theorem gives numerical estimates of this high probability. According to our definition of high probability, these theorems say something about Skeptic's opportunities to make money. The law of large numbers says that Skeptic has a winning strategy in a game that he wins if World either stays close to $1/2$ or allows Skeptic to multiply his stake substantially, and the central limit theorem calibrates the tradeoff between how far World can stray from $1/2$ and how much he can constrain Skeptic's wealth.

Middling probabilities, although they do not have meaning in isolation, can acquire collective meaning from the limit theorems. The law of large numbers tells us, for example, that many probabilities for successive events all equal to $1/2$ produce a very high probability that the relative frequency of the events will approximate $1/2$.

Probability Zero and Probability One

As we have just emphasized, the finitary version of our fundamental hypothesis gives meaning to probabilities very close to zero or one. Skeptic is unlikely to become very rich, and therefore an event with a very low probability is unlikely to occur. The infinitary version sharpens this by giving meaning to probabilities exactly equal to zero or one. It is practically impossible for Skeptic to become infinitely rich, and therefore an event that makes this possible is practically certain not to occur.

Formally, we say that an event E is *practically impossible* if Skeptic, beginning with some finite positive capital, has a strategy that guarantees that

- his capital does not become negative (he does not go bankrupt), and
- if E happens, his capital increases without bound (he becomes infinitely rich).

We say that an event E is *practically certain*, or that it happens *almost surely*, if its complement E^{c} is practically impossible. It follows immediately from these definitions that a practically impossible event has upper probability (and hence also lower probability) zero, and that a practically certain event has lower probability (and hence also upper probability) one (§8.3).

The size of Skeptic's initial capital does not matter in the definitions of practical certainty and practical impossibility, provided it is positive. If the strategy $\mathcal{P}$ will do what is required when his initial capital is a, then the strategy $\frac{b}{a}\mathcal{P}$ will accomplish the same trick when his initial capital is b. Requiring that Skeptic's capital not become negative is equivalent to forbidding him to borrow money, because if he dared to gamble on borrowed money, World could force his capital to become negative. The real condition, however, is not that he never borrow but that his borrowing be bounded. Managing on initial capital a together with borrowing limited to b is the same as managing on initial capital $a + b$.

As we show in Chapters 3, 4, and 5, these definitions allow us to state and prove game-theoretic versions of the classical strong limit theorems—the strong law of large numbers and the law of the iterated logarithm. In its simplest form, the game-theoretic strong law of large numbers says that when Skeptic is offered even odds on each of an infinite sequence of events, the fraction of the events that happen will almost certainly converge to $1/2$. The law of the iterated logarithm gives the best possible bound on the rate of convergence.

Beyond Frequencies

As we explain in some detail in Chapter 2, the law of large numbers, together with the empiricist philosophy of the time, led in the nineteenth and early twentieth centuries to a widespread conviction that the theory of probability should be founded on the concept of relative frequency. If many independent trials are made of an event with probability p, then the law of large numbers says that the event will happen p of the time and fail $1 - p$ of the time. This is true whether we consider all the trials, or only every other trial, or only some other subsequence selected in advance. And this appears to be the principal empirical meaning of probability. So why not turn the theory around, as Richard von Mises proposed in the 1920s, and say that a probability is merely a relative frequency that is invariant under selection of subsequences?

As it turned out, von Mises was mistaken to emphasize frequency to the exclusion of other statistical regularities. The predictions about a sequence of events made by probability theory do not all follow from the invariant convergence of relative frequency. In the late 1930s, Jean Ville pointed out a very salient and decisive example: the predictions that the law of the iterated logarithm makes about the rate and oscillation of the convergence. Von Mises's theory has now been superseded by the theory of algorithmic complexity, which is concerned with the properties of sequences whose complexity makes them difficult to predict, and invariant relative frequency is only one of many such properties.

Frequency has also greatly receded in prominence within measure-theoretic probability. Where independent identically distributed random variables were once the central object of study, we now study stochastic processes in which the probabilities of events depend on preceding outcomes in complex ways. These models sometimes make predictions about frequencies, but instead of relating a frequency to a single probability, they may predict that a frequency will approximate the average of a sequence of probabilities. In general, emphasis has shifted from sums of independent random variables to martingales.

For some decades, it has been clear to mathematical probabilists that martingales are fundamental to their subject. Martingales remain, however, only an advanced topic in measure-theoretic probability theory. Our game-theoretic framework puts what is fundamental at the beginning. Martingales come at the beginning, because they are the capital processes for Skeptic. The fundamental interpretative hypothesis, applied to a particular nonnegative martingale, says that the world will behave in such a way that the martingale remains bounded. And the many predictions that follow include the convergence of relative frequencies.

1.4 THE MANY INTERPRETATIONS OF PROBABILITY

Contemporary philosophical discussions often divide probabilities into two broad classes:

- *objective probabilities*, which describe frequencies and other regularities in the world, and
- *subjective probabilities*, which describe a person's preferences, real or hypothetical, in risk taking.

Our game-theoretic framework accommodates both kinds of probabilities and enriches our understanding of them, while opening up other possibilities as well.

Three Major Interpretations

From our point of view, it makes sense to distinguish three major ways of using the idea of a probability game, which differ in how prices are established and in the role of the fundamental interpretative hypothesis, as indicated in Table 1.4.

Games of statistical regularity express the objective conception of probability within our framework. In a game of statistical regularity, the gambles offered to Skeptic may derive from a scientific theory, from frequencies observed in the past, or from some relatively poorly understood forecasting method. Whatever the source, we adopt the fundamental interpretative hypothesis, and this makes statistical regularity the ultimate authority: the prices and probabilities determined by the gambles offered to Skeptic must be validated by experience. We expect events assigned small upper probabilities not to happen, and we expect prices to be reflected in average values.

Table 1.4 Three classes of probability games.

	Authority for the Prices	**Role of the Fundamental Interpretative Hypothesis**
Games of Statistical Regularity	Statistical regularities	Adopted
Games of Belief	Personal choices among risks	Not adopted
Market Games	Market for financial securities	Optional

Games of belief bring the neosubjectivist conception of probability into our framework. A game of belief may draw on scientific theories or statistical regularities to determine the gambles offered on individual rounds. But the presence of these gambles in the game derives from some individual's commitment to use them to rank and choose among risks. The individual does not adopt the fundamental interpretative hypothesis, and so his prices cannot be falsified by what actually happens. The upper and lower prices and probabilities in the game are not the individual's hypotheses about what will happen; they merely indicate the risks he will take. A low probability does not mean the individual thinks an event will not happen; it merely means he is willing to bet heavily against it.

Market games are distinguished by the source of their prices: these prices are determined by supply and demand in some market. We may or may not adopt the hypothesis that the market is efficient. If we do adopt it, then we may test it or use it to draw various conclusions (see, e.g., the discussion of the Iowa Electronic Markets on p. 71). If we do not adopt it, even provisionally, then the game can still be useful as a framework for understanding the hedging of market risks.

Our understanding of objective and subjective probability in terms of probability games differs from the usual explanations of these concepts in its emphasis on sequential experience. Objective probability is often understood in terms of a population, whose members are not necessarily examined in sequence, and most expositions of subjective probability emphasize the coherence of one's belief about different events without regard to how those events might be arranged in time. But we do experience the world through time, and so the game-theoretic framework offers valuable insights for both the objective and the subjective conceptions. Objective probabilities can only be tested over time, and the idea of a probability game imposes itself whenever we want to understand the testing process. The experience anticipated by subjective probabilities must also be arrayed in time, and probability games are the natural framework in which to understand how subjective probabilities change as that experience unfolds.

Looking at Interpretations in Two Dimensions

The uses and interpretations of probability are actually very diverse—so much so that we expect most readers to be uncomfortable with the standard dichotomy between objective and subjective probability and with the equally restrictive categories of Table 1.4. A more flexible categorization of the diverse possibilities for using the mathematical idea of a probability game can be developed by distinguishing uses along two dimensions: (1) the source of the prices, and (2) the attitude taken towards the fundamental interpretative hypothesis. This is illustrated in Figure 1.6.

We use quantum mechanics as an example of a scientific theory for which the fundamental interpretative hypothesis is well supported. From a measure-theoretic point of view, quantum mechanics is sometimes seen as anomalous, because of the influence exercised on its probabilistic predictions by the selection of measurements by observers, and because its various potential predictions, before a measurement is selected, do not find simple expression in terms of a single probability measure.

SOURCE OF THE PRICES

	Irrelevant	Working Hypothesis	Believed	Well Supported
Scientific Theory				Quantum mechanics
Observed Regularities		Hypothesis testing	Statistical modeling and estimation	
Personal Choices	Neosubjective probability	Decision analysis Weather forecasting		
Market	Hedging	Testing the EMH	Inference based on the EMH	

STATUS OF THE FUNDAMENTAL INTERPRETATIVE HYPOTHESIS

Fig. 1.6 Some typical ways of using and interpreting a probability game, arrayed in two dimensions. (Here EMH is an acronym for the efficient-market hypothesis.)

From our game-theoretic point of view, however, these features are prototypical rather than anomalous. No scientific theory can support probabilistic predictions without protocols for the interface between the phenomenon being predicted and the various observers, controllers, and other external agents who work to bring and keep the phenomenon into relation with the theory.

Statistical modeling, testing, and estimation, as practiced across the natural and social sciences, is represented in Figure 1.6 in the row labeled "observed regularities". We speak of regularities rather than frequencies because the empirical information on which statistical models are based is usually too complex to be summarized by frequencies across identical or exchangeable circumstances.

As we have already noted, the fundamental interpretative hypothesis is irrelevant to the neosubjectivist conception of probability, because a person has no obligation to take any stance concerning whether his or her subjective probabilities and prices satisfy the hypothesis. On the other hand, an individual might conjecture that his or her probabilities and prices do satisfy the hypothesis, with confidence ranging from "working hypothesis" to "well supported". The probabilities used in decision analysis and weather forecasting can fall anywhere in this range. We must also consider another dimension, not indicated in the figure: With respect to whose knowledge is the fundamental interpretative hypothesis asserted? An individual might peers odds that he or she is not willing to offer to more knowledgeable observers.

Finally, the bottom row of Figure 1.6 lists some uses of probability games in finance, a topic to which we will turn shortly.

Folk Stochasticism

In our listing of different ways probability theory can be used, we have not talked about using it to study stochastic mechanisms that generate phenomena in the world. Although quite popular, this way of talking is not encouraged by our framework.

What is a stochastic mechanism? What does it mean to suppose that a phenomenon, say the weather at a particular time and place, is generated by chance according to a particular probability measure? Scientists and statisticians who use probability theory often answer this question with a self-consciously outlandish metaphor: A demigod tosses a coin or draws from a deck of cards to decide what the weather will be, and our job is to discover the bias of the coin or the proportions of different types of cards in the deck (see, e.g., [23], p. 5).

In *Realism and the Aim of Science*, Karl Popper argued that objective probabilities should be understood as the *propensities* of certain physical systems to produce certain results. Research workers who speak of stochastic mechanisms sometimes appeal to the philosophical literature on propensities, but more often they simply assume that the measure-theoretic framework authorizes their way of talking. It authorizes us to use probability measures to model the world, and what can a probability measure model other than a stochastic mechanism—something like a roulette wheel that produces random results?

The idea of a probability game encourages a somewhat different understanding. Because the player who determines the outcome in a probability game does not necessarily do so by tossing a coin or drawing a card, we can get started without a complete probability measure, such as might be defined by a biased coin or a deck of cards. So we can accommodate the idea that the phenomenon we are modeling might have only limited regularities, which permit the pricing of only some of its uncertainties.

The metaphor in which the flow of events is determined by chance drives statisticians to hypothesize full probability measures for the phenomena they study and to make these measures yet more extensive and complicated whenever their details are contradicted by empirical data. In contrast, our metaphor, in which outcomes are determined arbitrarily within constraints imposed by certain prices, encourages a minimalist philosophy. We may put forward only prices we consider well justified, and we may react to empirical refutation by withdrawing some of these prices rather than adding more.

We do, however, use some of the language associated with the folk stochasticism we otherwise avoid. For example, we sometimes say that a phenomenon is governed by a probability measure or by some more restrained set of prices. This works in our framework, because government only sets limits or general directions; it does not determine all details. In our games, Reality is governed in this sense by the prices announced by Forecaster: these prices set boundaries that Reality must respect in order to avoid allowing Skeptic to become rich. In Chapter 14 we explain what it means for Reality to be governed in this sense by a stochastic differential equation.

1.5 GAME-THEORETIC PROBABILITY IN FINANCE

Our study of finance theory in Part II is a case study of our minimalist philosophy of probability modeling. Finance is a particularly promising field for such a case study, because it starts with a copious supply of prices—market prices for stocks, bonds, futures, and other financial securities—with which we may be able to do something without hypothesizing additional prices based on observed regularities or theory.

We explore two distinct paths. The path along which we spend the most time takes us into the pricing of options. Along the other path, we investigate the hypothesis that market prices are efficient, in the sense that an investor cannot become very rich relative to the market without risking bankruptcy. This hypothesis is widely used in the existing literature, but always in combination with stochastic assumptions. We show that these assumptions are not always needed. For example, we show that market efficiency alone can justify the advice to hold the market portfolio.

We conclude this introductory chapter with a brief preview of our approach to option pricing and with some comments about how our framework handles continuous time. A more extensive introduction to Part II is provided by Chapter 9.

The Difficulty in Pricing Options

The worldwide market in derivative financial securities has grown explosively in recent years. The total nominal value of transactions in this market now exceeds the total value of the goods and services the world produces. Many of these transactions are in organized exchanges, where prices for standardized derivatives are determined by supply and demand. A larger volume of transactions, however, is in over-the-counter derivatives, purchased directly by individuals and corporations from investment banks and other financial intermediaries. These transactions typically involve hedging by both parties. The individual or corporation buys the derivative (a future payoff that depends, for example, on future stock or bond prices or on future interest or currency exchange rates) in order to hedge a risk arising in the course of their business. The issuer of the derivative, say an investment banker, buys and sells other financial instruments in order to hedge the risk acquired by selling the derivative. The cost of the banker's hedging determines the price of the derivative.

The bulk of the derivatives business is in futures, forwards, and swaps, whose payoffs depend linearly on the future market value of existing securities or currencies. These derivatives are usually hedged without considerations of probability [154]. But there is also a substantial market in options, whose payoffs depend nonlinearly on future prices. An option must be hedged dynamically, over the period leading up to its maturity, and according to established theory, the success of such hedging depends on stochastic assumptions. (See [128], p. xii, for some recent statistics on the total sales of different types of derivatives.)

For readers not yet familiar with options, the artificial examples in Figure 1.7 may be helpful. In both examples, we consider a stock that today sells for $\$8$ a share and tomorrow will either (1) go down in price to $\$5$, (2) go up in price to $\$10$, or (3)

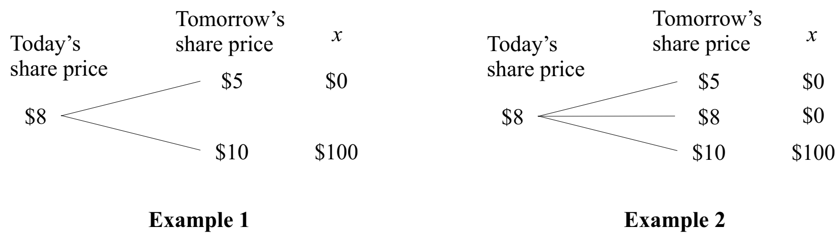

Fig. 1.7 The price of a share is now $\$8$. In Example 1, we assume that it will go up to $\$10$ or down to $\$5$ tomorrow. In Example 2, its price is also permitted to stay unchanged. In both cases, we are interested in the value today of the derivative x. In Example 1, x has a definite value: $\mathbb{E}\,x = \$60$. This price for x can be hedged exactly by buying 20 shares of the stock. If the stock down from $\$8$ to $\$5$, the loss of $\$3$ per share wipes out the $\$60$, but if it goes up to $\$10$, the gain of $\$2$ per share is just enough to provide the additional $\$40$ needed to provide x's $\$100$ payoff. In Example 2, no price for x can be hedged exactly. Instead we have $\overline{\mathbb{E}}\,x = \$60$ and $\underline{\mathbb{E}}\,x = \$0$. We should emphasize again that both examples are unrealistic. In a real financial market there is a whole range of possibilities—not just two or three possibilities—for how the price of a security can change over a single trading period.

(in Example 2) stay unchanged in price. Suppose you want to purchase an option to buy 50 shares tomorrow at today's price of $\$8$. If you buy this option and the price goes up, you will buy the stock at $\$8$ and resell it at $\$10$, netting $\$2$ per share, or $\$100$. What price should you pay today for the option? What is the value today of a payoff x that takes the value $\$100$ if the price of the stock goes up and the value $\$0$ otherwise? As explained in the caption to the figure, x is worth $\$60$ in Example 1, for this price can be hedged exactly. In Example 2, however, no price for x can be hedged exactly. The option in Example 1 can be priced because its payoff is actually a linear function of the stock price. When there are only two possible values for a stock price, any function of that price is linear and hence can be hedged. In Example 2, where the stock price has three possible values, the payoff of the option is nonlinear. In real stock markets, there is a whole range of possible values for the price of a stock at some future time, and hence there are many nonlinear derivatives that cannot be priced by hedging in the stock itself without additional assumptions.

A range of possible values can be obtained by a sequence of binary branchings. This fact can be combined with the idea of dynamic hedging, as in Figure 1.8, to provide a misleadingly simple solution to our problem. The solution is so simple that we might be tempted to believe in some imaginary shadow market, speedier and more liquid than the real market, where changes in stock prices really are binary but produce the less restricted changes seen in the slower moving real market. Unfortunately, there is no traction in this idea, for we can hedge only in real markets. In real stock markets, many different price changes are possible over the time periods during which we hold stock, and so we can never hedge exactly. The best we can do is hedge in a way that works on average, counting on the errors to average out. This is why probabilistic assumptions are needed.

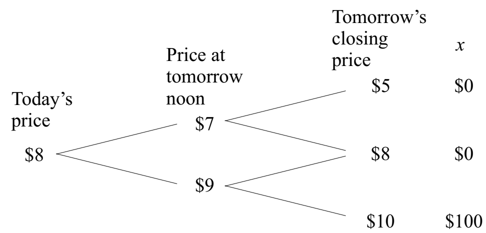

Fig. 1.8 In this example, $\mathbb{E}\,x = \$25$. To hedge this price, we first buy 25 shares of the stock today. We adjust this hedge at noon tomorrow, either by selling the 25 shares (if the price has gone down to \$7) or by buying another 25 shares (if the price has gone up to $\$8$).

The probabilistic models most widely used for option pricing are usually formulated, for mathematical tractability, in continuous time. These models include the celebrated Black-Scholes model, as well as models that permit jumps. As it turns out, binomial trees, although unrealistic as models of the market, can serve as computationally useful approximations to these widely used (although perhaps equally unrealistic, alas) continuous-time models. This point, first demonstrated in the late 1970s [66, 67, 256], has made binomial trees a standard topic in textbooks on option pricing.

Making More Use of the Market

The most common probability model for option pricing in continuous time, the Black-Scholes model, assumes that the underlying stock price follows a geometric Brownian motion. Under this assumption, options can be priced by a formula—the Black-Scholes formula—that contains a parameter representing the volatility of the stock price; the value of this parameter is usually estimated from past fluctuations. The assumption of geometric Brownian motion can be interpreted from our thoroughly game-theoretic point of view (Chapter 14). But if we are willing to make more use of the market, we can instead eliminate it (Chapters 10–13). The simplest options on some stocks now trade in sufficient volume that their prices are determined by supply and demand rather than by the Black-Scholes formula. We propose to rely on this trend, by having the market itself price one type of option, with a range of maturity dates. If this traded option pays a smooth and strictly convex function of the stock price at maturity, then other derivatives can be priced using the Black-Scholes formula, provided that we reinterpret the parameter in the formula and determine its value from the price of the traded option. Instead of assuming that the prices of the stock and the traded option are governed by some stochastic model, we assume only certain limits on the fluctuation of these prices. Our market approach also extends to the Poisson model for jumps (§12.3).

Probability Games in Continuous Time

Our discussion of option pricing in Part II involves an issue that is important both for our treatment of probability and for our treatment of finance: how can the game-theoretic framework accommodate continuous time? Measure theory's claim to serve as a foundation for probability has been based in part on its ability to deal with continuous time. In order to compete as a mathematical theory, our game-theoretic framework must also meet this challenge.

It is not immediately clear how to make sense of the idea of a game in which two players alternate moves continuously. A real number does not have an immediate predecessor or an immediate successor, and hence we cannot divide a continuum of time into points where Skeptic moves and immediately following points where World moves. Fortunately, we now have at our disposal a rigorous approach to continuous mathematics—nonstandard analysis—that does allow us to think of continuous time as being composed of discrete infinitesimal steps, each with an immediate predecessor and an immediate successor. First introduced by Abraham Robinson in the 1960s, long after the measure-theoretic framework for probability was established, nonstandard analysis is still unfamiliar and even intimidating for many applied mathematicians. But it provides a ready framework for putting our probability games into continuous time, with the great advantage that it allows a very clear understanding of how the infinite depends on the finite.

In Chapter 10, where we introduce our market approach to pricing options, we work in discrete time, just as real hedging does. Instead of obtaining an exact price for an option, we obtain upper and lower prices, both approximated by an expression similar to the familiar Black-Scholes formula. The accuracy of the approximation can be bounded in terms of the jaggedness of the market prices of the underlying security and the traded derivative. All this is very realistic but also unattractive and hard to follow because the approximations are crude, messy, and often arbitrary. In Chapter 11, we give a nonstandard version of the same theory. The nonstandard version, as it turns out, is simple and transparent. Moreover, the nonstandard version clearly says nothing that is not already in the discrete version, because it follows from the discrete version by the *transfer principle*, a general principle of nonstandard analysis that sometimes allows one to move between nonstandard and standard statements [136].

Some readers will see the need to appeal to nonstandard analysis as a shortcoming of our framework. There are unexpected benefits, however, in the clarity with which the transfer principle allows us to analyze the relation between discrete-time and continuous-time results. Although the discrete theory of Chapter 10 is very crude, its ability to calibrate the practical accuracy of our new purely game-theoretic Black-Scholes method goes well beyond what has been achieved by discrete-time analyses of the stochastic Black-Scholes method.

After introducing our approach to continuous time in Chapter 11, we use it to elaborate and extend our methods for option pricing (Chapters 12–13) and to give a general game-theoretic account of diffusion processes (Chapter 14), without working through corresponding discrete theory. This is appropriate, because the discrete theory will depend on the details of particular problems where the ideas are put to use. Discrete theory should be developed, however, in conjunction with any effort to put these ideas into practice. In our view, discrete theory should always be developed when continuous-time models are used, so that the accuracy of the continuous-time results can be studied quantitatively.

References

1. Patrick Billingsley. *Ergodic Theory and Information*. Wiley, New York, 1965.
2. John C. Cox and Stephen A. Ross. The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, 3:145–166, 1976.
3. John C. Cox, Stephen A. Ross, and Mark Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7:229–263, 1979.
4. A. Philip Dawid. Statistical theory: The prequential approach (with discussion). *Journal of the Royal Statistical Society. Series A*, 147:278–292, 1984.
5. Jack Clark Francis, William W. Toy, and J. Gregg Whittaker, editors. *The Handbook of Equity Derivatives*. Wiley, New York, revised edition, 2000.
6. Robert Goldblatt. *Lectures on the Hyperreals: An Introduction to Nonstandard Analysis*. Springer, New York, 1998. An excellent exposition based on ultrafilters. Contains a chapter on the Loeb measure.
7. John C. Hull. *Options, Futures, and Other Derivatives*. Prentice-Hall, Upper Saddle River, NJ, fourth edition, 2000. A popular introduction to the study of derivatives at the MBA level. It discusses both the theory and practice of derivatives and includes numerical procedures for implementing the most common pricing algorithms.
8. D. H. Mellor. *The Matter of Chance*. Cambridge University Press, London, 1971.
9. Karl R. Popper. *The Open Universe*. Hutchinson, London, 1982.
10. Karl R. Popper. *Realism and the Aim of Science*. Hutchinson, London, 1983.
11. Richard J. Rendleman and Brit J. Bartter. Two-state option pricing. *Journal of Finance*, 34:1093–1110, 1979.
12. Christian P. Robert. *The Bayesian Choice*. Springer, New York, 1994.

13. Abraham Robinson. *Non-Standard Analysis*. North-Holland, Amsterdam, second edition, 1974. First edition: 1966.

14. John von Neumann and Oscar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, third edition, 1953. Previous editions appeared in 1944 and 1947.