3

The Bounded Strong Law of Large Numbers

In this chapter, we formulate and prove the simplest forms of the game-theoretic strong law of large numbers.

In its very simplest form, the strong law of large numbers is concerned with a sequence of events, each with the same two possible outcomes, which we may call heads and tails. The law says that as one proceeds in the sequence, the proportion of heads converges to one-half almost surely. In symbols:

$$\lim_{n\to\infty} \frac{y_n}{n} = \frac{1}{2} \tag{3.1}$$

almost surely, where y_n is the number of heads among the first n events. A framework for mathematical probability must provide a precise mathematical context for this statement, including a mathematical definition of the term *almost surely*.

In the measure-theoretic framework, the mathematical context is provided by adopting a certain probability measure for the infinite sequence of events: heads has probability one-half each time, and the events are independent. The term *almost surely* means except on a set of measure zero. This makes our claim about convergence into the precise statement known as Borel's strong law: the sequences of outcomes for which the convergence to one-half fails have measure zero under the specified probability measure.

The game between Skeptic and Reality that we study in this chapter makes the claim about convergence precise in a different way. No probability measure is given, but before each event, Skeptic is allowed to bet as much as he wants on heads or on tails, at even odds. The meaning of *almost surely* is this: an event happens almost surely if Skeptic has a strategy for betting that does not risk bankruptcy and allows him to become infinitely rich if the event does not happen. This is all we need: the statement that the proportion of heads converges to one-half almost surely is now a precise statement in game theory, which we prove in this chapter.

The preceding paragraph does not mention our fundamental interpretative hypothesis, which says that Skeptic cannot become infinitely rich without risking bankruptcy. As we explained in Chapter 1, this hypothesis stands outside the mathematical theory. It has nothing to do with the mathematical meaning of the strong law of large numbers or with its proof. It comes into play only when we use the theorem. If we adopt the hypothesis for a particular sequence of events, then the statement "the proportion of these events that happen will converge to one-half almost surely" acquires a useful meaning: we think the convergence will happen. If we do not adopt the hypothesis for the particular sequence, then the statement does not have this meaning.

The game-theoretic formulation is more constructive than the measure-theoretic formulation. We construct a computable strategy for Skeptic that is sure to keep him from becoming bankrupt and allows him to become infinitely rich if the convergence of the proportion of heads to one-half fails. Moreover, our formulation is categorical—it is a statement about all sequences, not merely about sequences outside a set of measure zero. Every sequence either results in an infinite gain for Skeptic or else has a proportion of heads converging to one-half.

The game-theoretic formulation has a further unfamiliar and very interesting facet. In the folk picture of stochastic reality, outcomes are determined independently of how any observer bets. In the game between Skeptic and Reality, in contrast, Reality is allowed to take Skeptic's bets into account in deciding on outcomes. Yet this does not prevent Skeptic from constructing a winning strategy. No matter how diabolically Reality behaves, she cannot violate the required convergence without yielding an infinite gain to Skeptic.

We do not propose to replace stochastic reality with a rational, diabolical reality. We propose, rather, to eliminate altogether from the general theory of probability any particular assumption about how outcomes are determined. It is consistent with our framework to suppose that Reality is nothing more than the actual outcomes of the events—that Reality has no strategy, that there is no sense in the question of how she would have made the second event come out had the first event come out differently or had Skeptic bet differently. By lingering over this supposition, we underline the concreteness of the strong law of large numbers; it concerns only our beliefs about a single sequence of actual outcomes. But it is equally consistent with our framework to imagine an active, strategic Reality. This diversity of possible suppositions hints at the breadth of possible applications of probability, a breadth not yet, perhaps, fully explored.

We formalize our game of heads and tails in §3.1. Then, in §3.2, we generalize it to a game in which Reality decides on values for a bounded sequence of centered variables $x_1, x_2, \ldots$. The strong law for this game says that the average of the first N of the x_n will converge to zero as N increases. In order to explain this in the measure-theoretic framework, we postulate a complete probability distribution for all the variables (this amounts to a specification of prices for all measurable functions of the variables), and then we conclude that the convergence will occur except on a set of measure zero, provided the conditional expectation (given the information available before time n) of each x_n is zero. Our game-theoretic formulation dispenses not only with the use of measure zero but also with the complete probability distribution. We assume only that each x_n is offered to Skeptic at the price of zero just before it is announced by Reality. Skeptic may buy the x_n in any positive or negative amounts, but nonlinear functions of the x_n need not be priced. This formulation is more widely—or at least more honestly—applicable than the measure-theoretic formulation.

In §3.3, we generalize to the case where the successive variables have prices not necessarily equal to zero. In this case, the strong law says that the average difference between the variables and their prices converges to zero almost surely. The mathematical content of this generalization is slight (we continue to assume a uniform bound for the variables and their prices), but the generalization is philosophically interesting, because we must now discuss how the prices are set, a question that is very important for meaning and application. The diversity of interpretations of probability can be attributed to the variety of ways in which prices can be set.

In §3.4, we briefly discuss the generalization from two-sided prices—prices at which Skeptic is allowed both to buy and sell—to one-sided prices, at which he is allowed only to buy or only to sell. If Skeptic is only allowed, for example, to buy at given prices, and not to sell, then our belief that he cannot become infinitely rich implies only that the long-term average difference between the variables and their prices will almost surely not exceed zero.

In an appendix, §3.5, we comment on the computability of the strategies we construct and on the desirability of detailed investigation of their computational properties.

The main results of this chapter are special cases of more general results we establish in Chapter 4, where we allow Reality's moves, the variables x_n, to be unbounded. For most readers, however, this chapter will be a better introduction to the basic ideas of the game-theoretic framework than Chapter 4, because it presents these ideas without the additional complications that arise in the unbounded case.

3.1 THE FAIR-COIN GAME

Now we formalize our game of heads and tails. We call it the *fair-coin game*, but not too much meaning should be read into this name. The outcomes need not be determined by tossing a coin, and even if they are, there is no need for the coin to have any property that might be called fairness. All that is required is that Skeptic be allowed to bet at even odds on heads or on tails, as he pleases.

Skeptic begins with some initial capital, say $\$1$. He bets by buying some number, say M, of tickets in favor of heads; M may be any real number—positive, zero, or negative. Each ticket, which sells for $\$0$, pays the bearer $\$1$ if Reality chooses heads, and requires the bearer to forfeit $\$1$ if Reality chooses tails. So buying M tickets means agreeing to pay $\$M$ if Reality chooses tails in order to gain $\$M$ if Reality chooses heads; if M is negative, then this is really a bet in favor of tails. If we code tails as -1 and heads as 1, then the protocol for the game can be described as follows:

$\mathcal{K}_0 = 1$.
FOR $n = 1, 2, \ldots$:
  Skeptic announces $M_n \in \mathbb{R}$.
  Reality announces $x_n \in \{-1, 1\}$.
  $\mathcal{K}_n := \mathcal{K}_{n-1} + M_n x_n$.

The quantity $\mathcal{K}_n$ is Skeptic's capital just after the bet on the nth toss is settled.

Skeptic wins the game if (1) his capital $\mathcal{K}_n$ is never negative, and (2) either

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} x_i = 0 \tag{3.2}$$

or

$$\lim_{n\to\infty} \mathcal{K}_n = \infty \tag{3.3}$$

holds. Otherwise Reality wins.

Equation (3.2) says that the proportion of heads in the first n tosses converges to one-half. It is equivalent to (3.1), because $\sum_{i=1}^{n} x_i = 2y_n - n$. We use (3.2) instead of (3.1) only because it is suitable for the more general bounded forecasting game that we consider in the next section.

The rule for determining the winner, given by (3.2) and (3.2), completes our specification of the *fair-coin game*. It is a two-person, zero-sum, perfect-information game. Zero-sum because Skeptic wins if and only if Reality loses. (Since a win is conventionally scored as a 1 and a loss as -1, the two players' scores sum to zero.) Perfect-information because each player knows all the previous moves when he makes his own next move. Here is the strong law of large numbers for this game:

Proposition 3.1 *Skeptic has a winning strategy in the fair-coin game.*

This means the convergence (3.2) occurs almost surely, in the game-theoretic sense of this term explained on p. 17 and p. 61. Skeptic has a strategy that forces Reality to arrange the convergence if she is to keep him from becoming infinitely rich. To the extent that we believe that Skeptic cannot become infinitely rich, we should also believe that the convergence will happen. If we adopt the fundamental interpretative hypothesis, then we may simply assert that the convergence will occur.

Some readers might prefer to allow Skeptic to borrow money. Skeptic does not need any such concession, however; he has a winning strategy even if he is not allowed to borrow. Moreover, allowing Skeptic to borrow would not really change the picture so long as there were a limit to his borrowing; allowing him to borrow up to $\$\beta$ would have the same effect on our reasoning as changing his initial capital from $\$1$ to $\$(1+\beta)$, and so long as his initial capital is positive, its value makes no difference in our reasoning.

We will prove a generalization of Proposition 3.1 in §3.2. Our proof is constructive; we spell out Skeptic's strategy explicitly. The strategy can be described roughly as follows: If Skeptic establishes an account for betting on heads, and if at each step he bets a fixed proportion ϵ of the money then in the account on heads, then Reality can keep the account from getting indefinitely large only by eventually holding the average $\frac{1}{n}\sum_{i=1}^{n} x_i$ at or below ϵ. So Skeptic can force Reality to hold the average asymptotically at zero or less by splitting a portion of his capital into an infinite number of accounts for betting on heads, including accounts with ϵ that come arbitrarily close to zero. By also setting up similar accounts for betting on tails, Skeptic can force Reality to make the average converge exactly to zero.

The rule for determining the winner in our game will seem simpler if we break out the two players' *collateral duties* so as to emphasize the main question, whether (3.2) holds. The collateral duty of Skeptic is to make sure that his capital $\mathcal{K}_n$ is never negative. The collateral duty of Reality is to make sure that $\mathcal{K}_n$ does not tend to infinity. If a player fails to perform his or her collateral duties, he or she loses the game. (More precisely, the first player to fail loses. If Skeptic and Reality both fail, then Skeptic loses and Reality wins, because Skeptic's failure happens at some particular trial, while Reality's failure happens later, at the end of the infinite sequence of trials.) If both players perform their collateral duties, Skeptic wins if and only if (3.2) is satisfied.

Equation (3.2) is a particular event—a particular property of Reality's moves. We can define a whole gamut of analogous games with the same protocol and the same collateral duties but with other events in the place of (3.2). In this more general context, we say that a strategy *forces* an event E if it is a winning strategy for Skeptic in the game in which E replaces (3.2) as Skeptic's main goal. We say that Skeptic *can force* E if he has a strategy that forces E—that is, if E happens almost surely. As we will see in later chapters, Skeptic can force many events.

3.2 FORECASTING A BOUNDED VARIABLE

Suppose now that instead of being required to choose heads or tails (1 or -1) on each trial, Reality is allowed to choose any real number x between -1 and 1. This number becomes the payoff (positive, negative, or zero) in dollars for a ticket Skeptic can buy for $\$0$ before the trial. Skeptic is again allowed to buy any number M of such tickets; when he buys a positive number, he is betting Reality will choose x positive; when he buys a negative number, he is betting Reality will choose x negative.

Our new game generalizes the fair-coin game only in that Reality chooses from the closed interval $[-1, 1]$ rather than from the two-element set $\{-1, 1\}$:

BOUNDED FORECASTING GAME WITH FORECASTS SET TO ZERO
Players: Skeptic, Reality
Protocol:
  $\mathcal{K}_0 = 1$.
  FOR $n = 1, 2, \ldots$:
    Skeptic announces $M_n \in \mathbb{R}$.
    Reality announces $x_n \in [-1, 1]$.
    $\mathcal{K}_n := \mathcal{K}_{n-1} + M_n x_n$.
Winner: Skeptic wins if $\mathcal{K}_n$ is never negative and either (3.2) or (3.3) holds. Otherwise Reality wins.

Why do we call this a *forecasting* game? Who is forecasting what? The answer is that we have forecast Reality's move x_n, and the forecast is zero. This forecast has an economic meaning: Skeptic can buy x_n for zero. In §3.3, we generalize by allowing forecasts different from zero, made in the course of the game. These forecasts will also serve as prices for Skeptic.

Because the generalization from the fair-coin game to the bounded forecasting game with forecasts at zero involves only an enlargement of Reality's move space, the reformulation in terms of collateral duties and the concept of forcing developed in the preceding section apply here as well. We refrain from repeating the definitions.

The game-theoretic strong law of large numbers says that Skeptic can always win this game; if Reality is committed to avoiding making him infinitely rich, then he can force her to make the average of the $x_1, x_2, \ldots$ converge to zero.

Proposition 3.2 *Skeptic has a winning strategy in the bounded forecasting game with forecasts set to zero.*

Because Proposition 3.2 generalizes Proposition 3.1, our proof of it will establish Proposition 3.1 as well.

The proof of Proposition 3.2 will be facilitated by some additional terminology and notation.

As we explained in §1.2, a complete sequence of moves by World is called a *path*, and the set of all paths is called the *sample space* and designated by Ω. In the game at hand, the bounded forecasting game with forecasts set to zero, World consists of the single player Reality, Reality's moves always come from the interval $[-1, 1]$, and the game continues indefinitely. So Ω is the infinite Cartesian product $[-1, 1]^\infty$. Each path is an infinite sequence $x_1, x_2, \ldots$ of numbers in $[-1, 1]$.

As we said in §1.2, any function on the sample space is a *variable*. Here this means that any function of the $x_1, x_2, \ldots$ is a variable. In particular, the x_n themselves are variables.

A *situation* is a finite sequence of moves by Reality. For example, x_1x_2 is the situation after Reality has chosen x_1 as her first move and x_2 as her second move. We write $\Omega^\diamond$ for the set of all situations. In the game at hand, $\Omega^\diamond$ is the set of all finite sequences of numbers from $[-1, 1]$, including the sequence of length zero, the *initial situation*, which we designate by $\Box$.

We say that the situation s *precedes* the situation t if t, as a sequence, contains s as an initial segment—say $s = x_1x_2 \ldots x_m$ and $t = x_1x_2 \ldots x_m \ldots x_n$. We write $s \sqsubseteq t$ when s precedes t. If s is a situation and $x \in [-1, 1]$, we write sx for the situation obtained by concatenating s with x; thus if $s = x_1 \ldots x_n$, then $sx = x_1 \ldots x_n x$. If s and t are situations and neither precedes the other, then we say they are *divergent*. We write $|s|$ for the length of s; thus $|x_1x_2 \ldots x_n| = n$. If ξ is a path for Reality, say $\xi = x_1x_2 \ldots$, we write ξ^n for the situation $x_1x_2 \ldots x_n$. We say that s *begins* ξ whenever s is a situation, ξ is a path, and $s = \xi^n$ for some n.

We call a real-valued function on $\Omega^\diamond$ a *process*. Any process $\mathcal{P}$ can be interpreted as a *strategy* for Skeptic; for each situation s, we interpret $\mathcal{P}(s)$ as the number of tickets Skeptic is to buy in situation s. This definition of strategy puts no constraints on Skeptic. In particular, his initial capital does not constrain him; he is allowed to borrow money indefinitely. In the games in this chapter, however, a strategy that may require borrowing money cannot be a winning strategy for Skeptic. In these games, Skeptic loses if his capital becomes negative. If he adopts a strategy that would result in a negative capital in any situation, Reality can defeat him by choosing a path that goes through that situation.

In our game, Skeptic begins with the initial capital 1, but we can also consider the capital process that would result from his beginning with any given capital $\mathcal{K}_0$, positive, negative, or zero, and following a particular strategy $\mathcal{P}$. As in §1.2, we write $\mathcal{K}^{\mathcal{P}}$ for his capital process when he begins with zero: $\mathcal{K}^{\mathcal{P}}(\Box) = 0$ and

$$\mathcal{K}^{\mathcal{P}}(x_1x_2\dots x_n) := \mathcal{K}^{\mathcal{P}}(x_1x_2\dots x_{n-1}) + \mathcal{P}(x_1x_2\dots x_{n-1})x_n. \tag{3.4}$$

When he uses the $\mathcal{P}$ with any other initial capital α, his capital follows the process $\alpha + \mathcal{K}^{\mathcal{P}}$. We call a process a *martingale* if it is of the form $\alpha + \mathcal{K}^{\mathcal{P}}$—that is, if it is the capital process for some strategy and some initial capital.[1]

The capital processes that begin with zero form a linear space, for $\beta\mathcal{K}^{\mathcal{P}} = \mathcal{K}^{\beta\mathcal{P}}$ and $\mathcal{K}^{\mathcal{P}_1} + \mathcal{K}^{\mathcal{P}_2} = \mathcal{K}^{\mathcal{P}_1+\mathcal{P}_2}$. It follows that the set of all capital processes (the set of all martingales) is also a linear space.

If α_1 and α_2 are nonnegative numbers that add to one, and $\mathcal{P}_1$ and $\mathcal{P}_2$ are strategies, then the martingale that results from using the strategy $\alpha_1\mathcal{P}_1 + \alpha_2\mathcal{P}_2$ starting with capital 1 is given by the same convex combination of the martingales that result from using the respective strategies starting with capital 1:

$$1 + \mathcal{K}^{\alpha_1\mathcal{P}_1+\alpha_2\mathcal{P}_2} = \alpha_1(1 + \mathcal{K}^{\mathcal{P}_1}) + \alpha_2(1 + \mathcal{K}^{\mathcal{P}_2}). \tag{3.5}$$

We can implement the convex combination in (3.5) by dividing the initial capital 1 between two accounts, putting α_1 in one and α_2 in the other, and then applying the strategy $\alpha_k\mathcal{P}_k$ (which is simply $\mathcal{P}_k$ scaled down to the initial capital α_k) to the kth account.

We will also find occasion to form infinite convex combinations of strategies. If $\mathcal{P}_1, \mathcal{P}_2, \dots$ are strategies, $\alpha_1, \alpha_2, \dots$ are nonnegative real numbers adding to one, and the sum $\sum_{k=1}^{\infty} \alpha_k\mathcal{P}_k$ converges, then the sum $\sum_{k=1}^{\infty} \alpha_k\mathcal{K}^{\mathcal{P}_k}$ will also converge (by induction on (3.4)), and $1 + \sum_{k=1}^{\infty} \alpha_k\mathcal{K}^{\mathcal{P}_k}$ will be the martingale the strategy $\sum_{k=1}^{\infty} \alpha_k\mathcal{P}_k$ produces when it starts with initial capital 1. The strategy $\sum_{k=1}^{\infty} \alpha_k\mathcal{P}_k$ starting with 1 is implemented by dividing the initial capital of 1 among a countably infinite number of accounts, with α_k in the kth account, and applying $\alpha_k\mathcal{P}_k$ to the kth account.

Recall that an *event* is a subset of the sample space. We say that a strategy $\mathcal{P}$ for Skeptic *forces* an event E if

$$\mathcal{K}^{\mathcal{P}}(t) \geq -1 \tag{3.6}$$

for every t in $\Omega^\diamond$ and

$$\lim_{n\to\infty} \mathcal{K}^{\mathcal{P}}(\xi^n) = \infty \tag{3.7}$$

[1] As we explained in §2.4 (p. 53), we use the word "martingale" in this way only in symmetric probability protocols. The protocol we are studying now is symmetric: Skeptic can buy x_n-tickets in negative as well as in positive amounts.

for every path ξ not in E. This agrees with the definition given in §3.1; condition (3.6) says that Skeptic does not risk bankruptcy using the strategy starting with the capital 1, no matter what Reality does. We say that Skeptic *can force* E if he has a strategy that forces E; this is the same as saying that there exists a nonnegative martingale starting at 1 that becomes infinite on every path not in E.

We say that $\mathcal{P}$ *weakly forces* E if (3.6) holds and every path ξ not in E satisfies

$$\sup_n \mathcal{K}^{\mathcal{P}}(\xi^n) = \infty. \tag{3.8}$$

By these definitions, any strategy $\mathcal{P}$ for which (3.6) holds weakly forces $\sup_n \mathcal{K}^{\mathcal{P}}(\xi^n) < \infty$. We say that Skeptic *can weakly force* E if he has a strategy that weakly forces E; this is the same as saying that there exists a nonnegative martingale starting at 1 that is unbounded on every path not in E.

The following lemma shows that the concepts of forcing and weak forcing are nearly equivalent.

Lemma 3.1 *If Skeptic can weakly force E, then he can force E.*

Proof Suppose $\mathcal{P}$ is a strategy that weakly forces E. For any $C > 0$, define a new strategy $\mathcal{P}^{(C)}$ by

$$\mathcal{P}^{(C)}(s) := \begin{array}{l} \mathcal{P}(s) \text{ if } \mathcal{K}^{\mathcal{P}}(t) < C \text{ for all } t \sqsubseteq s \\ 0 \text{ otherwise.} \end{array}$$

This strategy mimics $\mathcal{P}$ except that it quits betting as soon as Skeptic's capital reaches C. Define a strategy $\mathcal{Q}$ by

$$\mathcal{Q} := \sum_{k=1}^{\infty} 2^{-k} \mathcal{P}^{(2^k)}.$$

Then $\lim_{n\to\infty} \mathcal{K}^{\mathcal{Q}}(\xi^n) = \infty$ for every ξ for which $\sup_n \mathcal{K}^{\mathcal{P}}(\xi^n) = \infty$. Since $\mathcal{K}^{\mathcal{P}} \geq -1$, $\mathcal{K}^{\mathcal{Q}} \geq -1$. Since $\sup_n \mathcal{K}^{\mathcal{P}}(\xi^n) = \infty$ for every ξ not in E, $\lim_{n\to\infty} \mathcal{K}^{\mathcal{Q}}(\xi^n) = \infty$ for every ξ not in E. So $\mathcal{Q}$ forces E. ∎

Proving Proposition 3.2 means showing Skeptic can force (3.2), and according to Lemma 3.1, it suffices to show he can weakly force (3.2). The next two lemmas will make this easy.

Lemma 3.2 *If Skeptic can weakly force each of a sequence $E_1, E_2, \ldots$ of events, then he can weakly force $\bigcap_{k=1}^{\infty} E_k$.*

Proof Let $\mathcal{P}_k$ be a strategy that weakly forces E_k. The capital process $1+\mathcal{K}^{\mathcal{P}_k}$ is nonnegative, and in our game this implies that it can at most double on each step:

$$1 + \mathcal{K}^{\mathcal{P}_k}(x_1 \ldots x_n) \leq 2^n.$$

Since $|\mathcal{P}_k| \leq 1 + \mathcal{K}^{\mathcal{P}_k}$ (see (3.4)), we can also say that

$$|\mathcal{P}_k(x_1 \ldots x_n)| \leq 2^n$$

for all k, which implies that a strategy $\mathcal{Q}$ can be defined by

$$\mathcal{Q} := \sum_{k=1}^{\infty} 2^{-k} \mathcal{P}_k.$$

Since $\mathcal{P}_k$ weakly forces E_k, $\mathcal{Q}$ also weakly forces E_k. So $\mathcal{Q}$ weakly forces $\bigcap_{k=1}^{\infty} E_k$. ∎

Lemma 3.3 *Suppose $\epsilon > 0$. Then Skeptic can weakly force*

$$\limsup_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} x_i \le \epsilon \tag{3.9}$$

and

$$\liminf_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} x_i \ge -\epsilon. \tag{3.10}$$

Proof We may suppose that $\epsilon < 1/2$. The game specifies that the initial capital is 1. Let $\mathcal{P}$ be the strategy that always buys $\epsilon\alpha$ tickets, where α is the current capital. Since Reality's move x is never less than -1, this strategy loses at most the fraction ϵ of the current capital, and hence the capital process $1 + \mathcal{K}^{\mathcal{P}}$ is nonnegative. It is given by $1 + \mathcal{K}^{\mathcal{P}}(\Box) = 1$ and

$$1 + \mathcal{K}^{\mathcal{P}}(x_1 \dots x_n) = (1 + \mathcal{K}^{\mathcal{P}}(x_1 \dots x_{n-1}))(1 + \epsilon x_n) = \prod_{i=1}^{n} (1 + \epsilon x_i) \,.$$

Let $\xi = x_1 x_2 \dots$ be a path such that $\sup_n \mathcal{K}^{\mathcal{P}}(x_1 \dots x_n) < \infty$. Then there exists a constant $C_\xi > 0$ such that

$$\prod_{i=1}^{n} (1 + \epsilon x_i) \le C_\xi$$

for all n. This implies that

$$\sum_{i=1}^{n} \ln (1 + \epsilon x_i) \le D_\xi$$

for all n for some D_ξ. Since $\ln(1+t) \ge t - t^2$ whenever $t \ge -\frac{1}{2}$, ξ also satisfies

$$\epsilon \sum_{i=1}^{n} x_i - \epsilon^2 \sum_{i=1}^{n} x_i^2 \le D_\xi,$$

$$\epsilon \sum_{i=1}^{n} x_i - \epsilon^2 n \le D_\xi,$$

$$\epsilon \sum_{i=1}^{n} x_i \le D_\xi + \epsilon^2 n,$$

or

$$\frac{1}{n} \sum_{i=1}^{n} x_i \le \frac{D_\xi}{\epsilon n} + \epsilon$$

for all n and hence satisfies (3.9). Thus $\mathcal{P}$ weakly forces (3.9). The same argument, with $-\epsilon$ in place of ϵ, establishes that Skeptic can weakly force (3.10). ∎

In order to complete the proof that Skeptic can weakly force (3.2), we now simply consider the events (3.9) and (3.10) for $\epsilon = 2^{-k}$, where k ranges over all natural numbers; this defines a countable number of events Skeptic can weakly force, and their intersection, which he can also weakly force (by Lemma 3.2), is (3.2).

3.3 WHO SETS THE PRICES?

We have formulated the bounded forecasting game in the simplest possible way. The variables $x_1, x_2, \ldots$ are all between -1 and 1, and they all have the same price: zero. But our proof applies equally well when each variable has a different price, say x_n has the price m_n, provided only that both x_n and m_n are uniformly bounded. (In fact, it is enough that the net payoffs of the tickets, the differences $x_n - m_n$, be uniformly bounded.) Of course, we must then replace (3.2) by

$$\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n}(x_i - m_i) = 0. \tag{3.11}$$

The price m_n can be chosen in whatever manner we please; we require only that it be announced before Skeptic places his bet M_n.

As we explained in §1.1, the idea that prices can be set freely can be expressed within our framework by introducing a third player, Forecaster, who sets them. The game then takes the following form:

BOUNDED FORECASTING GAME
Parameter: $C > 0$
Players: Forecaster, Skeptic, Reality
Protocol:
  $\mathcal{K}_0 := 1$.
  FOR $n = 1, 2, \ldots$:
    Forecaster announces $m_n \in [-C, C]$.
    Skeptic announces $M_n \in \mathbb{R}$.
    Reality announces $x_n \in [-C, C]$.
    $\mathcal{K}_n := \mathcal{K}_{n-1} + M_n(x_n - m_n)$.
Winner: Skeptic wins if $\mathcal{K}_n$ is never negative and either (3.11) or (3.3) holds. Otherwise Reality wins.

Since Forecaster can always choose the m_n to be zero, and since the bound C can be 1, this game generalizes the game of the preceding section. And Proposition 3.2 generalizes as well:

Proposition 3.3 *Skeptic has a winning strategy in the bounded forecasting game.*

The proof of Proposition 3.2 generalizes immediately to a proof of Proposition 3.3. Alternatively, any winning strategy in the game with zero prices can be adapted in an obvious way to produce a winning strategy in the game with arbitrary prices.

We can recover Proposition 3.2 from Proposition 3.3 by setting C equal to 1 and requiring Forecaster to set each m_n equal to zero. Imposing this constraint on Forecaster changes the game (because his move is entirely determined, he is no longer really in the game!), but this change obviously does not impair the validity of the proposition. A strategy for Skeptic that wins when his opponents, Forecaster and Reality, have complete freedom of action will obviously still win when they are constrained, partially or completely.

We introduce Forecaster, with his complete freedom of action, in order to emphasize that how the m_n are selected is quite immaterial to the reasoning by which we establish the strong law of large numbers. No matter how prices are set, Skeptic has a winning strategy. This is not to say that how prices are set is unimportant to users of the law of large numbers. On the contrary, the practical significance of the law depends both on how the prices m_n are determined and on how the outcomes x_n are determined. And they can be determined in a variety of ways. In physics, the m_n are furnished by theory, while the x_n are furnished by reality. In finance, both m_n and x_n are determined by a market; in many cases, m_n is the price of a stock at the beginning of day n, and x_n is its price at the end of day n.

It is possible for the price m_n to be set by a market even though x_n is determined outside the market. This happens, for example, in the Iowa Electronic Markets, which generate prices for events such as the outcomes of elections. Suppose the Iowa Electronic Markets continue to organize trading in contracts for the outcomes of U.S. presidential elections into the indefinite future: each November 1 before such an election, it determines a price for a contract that pays $\$1$ if the Democratic candidate wins. Then the game-theoretic strong law of large numbers tells us that either we can become infinitely rich without risking more than $\$1$ or else the market is calibrated, in the sense that the long-term average price of the contract approaches the long-term relative frequency with which the Democratic candidates win. Of course, this is a very idealized statement; we are assuming that the system that pits Democrats against Republicans will go on forever, that money is infinitely divisible, and that we can neglect transaction costs and bid-ask spreads. But in Chapter 6 we will prove a finitary game-theoretic law of large numbers, and the other idealizing assumptions can also be relaxed.

When we say that m_n is the price for x_n when Skeptic makes his move M_n, our manner of speaking is consistent with the general definition of price given in Chapter 1 (p. 14). Because it is determined by the path $x_1, x_2, \ldots$, we are entitled to call x_n a variable, and because Skeptic can buy it exactly for m_n, we have $\overline{\mathbb{E}}_t\, x_n = \underline{\mathbb{E}}_t\, x_n = m_n$ in the situation t where Forecaster has just announced m_n.

In measure-theoretic probability, the strong law for a sequence $x_1, x_2, \ldots$ of variables is formulated beginning with the assumption that the variables have a joint probability distribution, and the price m_n is the conditional expected value of x_n given $x_1, \ldots, x_{n-1}$; the conclusion is that (3.11) holds almost surely. We will leave for Chapter 8 the formal derivation of this measure-theoretic result from Proposition 3.3, but it is intuitively obvious that our game-theoretic formulation is more powerful, in the sense that it arrives at the same conclusion (Equation (3.11) holds almost surely) with fewer assumptions. Postulating a joint probability distribution for $x_1, x_2, \ldots$ amounts to assuming that for all n, every measurable function of x_n (and even of $x_n, x_{n+1}, \ldots$) is priced conditional on the outcomes $x_1, \ldots, x_{n-1}$. But the game-theoretic formulation assumes only that a price for x_n itself is given in light of $x_1, \ldots, x_{n-1}$. In the simplest case, where each x_n has only two possible values, heads or tails, there is no difference between pricing the x_n and pricing all measurable functions of it. But when each x_n can be chosen from a large range of values, the difference is immense, and consequently the game-theoretic result is much more powerful than the measure-theoretic result.

As we pointed out earlier, Proposition 3.3 continues to be true when the bounded forecasting game is modified by a restriction, partial or complete, on the freedom of action of Skeptic's opponents. When we constrain Forecaster by setting the m_n equal to a common value m at the beginning of the game, we obtain a game-theoretic generalization of the measure-theoretic strong law for the case where the $x_1, x_2, \ldots$ are independent random variables with a common mean m. If we then constrain Reality to choose x_n from the set $\{0, 1\}$, we obtain the game-theoretic result corresponding to the measure-theoretic strong law for a possibly biased coin; if 1 represents heads and 0 represents tails, then m corresponds to the probability of heads on each toss, and (3.11) says that $\frac{1}{n}\sum_{i=1}^n x_i$, the proportion of heads in the first n tosses, converges to m. If $m = 1/2$, then we are back to the fair-coin game with which we began the chapter, except that we are using 0 rather than -1 to represent heads.

3.4 ASYMMETRIC BOUNDED FORECASTING GAMES

Our bounded strong law of large numbers, Equation (3.11), can be decomposed into two parts:

$$\limsup_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} (x_i - m_i) \le 0, \tag{3.12}$$

and

$$\liminf_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} (x_i - m_i) \ge 0. \tag{3.13}$$

Moreover, the proof of Lemma 3.3 makes it clear that these two parts depend on different assumptions. If Skeptic can buy tickets at the prices m_n, then either (3.12) will hold or else he can become infinitely rich. If Skeptic can sell tickets at the prices m_n, then either (3.13) will hold or else he can become infinitely rich.

We can express this point more formally by adapting the game of the preceding section as follows:

BOUNDED UPPER FORECASTING GAME
Parameter: $C > 0$
Players: Forecaster, Skeptic, Reality
Protocol:
 $\mathcal{K}_0 := 1$.
 FOR $n = 1, 2, \ldots$:
  Forecaster announces $m_n \in [-C, C]$.
  Skeptic announces $M_n \ge 0$.
  Reality announces $x_n \in [-C, C]$.
  $\mathcal{K}_n := \mathcal{K}_{n-1} + M_n(x_n - m_n)$.

Winner: Skeptic wins if and only if $\mathcal{K}_n$ is never negative and either (3.12) or (3.3) holds. Otherwise Reality wins.

This is an asymmetric protocol, in the sense explained on p. 11. Using only the first half of Lemma 3.3, we obtain our usual result:

Proposition 3.4 *Skeptic has a winning strategy in the bounded upper forecasting game.*

The hypothesis of the impossibility of a gambling system, as applied to this game, says that the prices m_n are high enough that Skeptic cannot get infinitely rich by buying tickets. If we adopt this hypothesis, then we may conclude that (3.12) will hold for these prices and, a fortiori, for any higher prices.

We can similarly define a bounded lower forecasting game, in which (1) Skeptic selects a nonpositive rather than a nonnegative real number, and (2) Skeptic wins if his capital remains nonnegative and either (3.13) or (3.3) holds. Again Skeptic will have a winning strategy.

3.5 APPENDIX: THE COMPUTATION OF STRATEGIES

The scope of this book is limited to showing how the game-theoretic framework can handle traditional questions in probability theory (Part I) and finance theory (Part II). But our results raise many new questions, especially questions involving computation. All our theoretical results are based on the explicit construction of strategies, and it should be both interesting and useful to study the computational properties of these constructions.

All the strategies we construct are in fact computable. For example, the construction in §3.2 is obviously computable, and hence we can strengthen Proposition 3.3 to the following:

Proposition 3.5 *Skeptic has a computable winning strategy in the bounded forecasting game.*

This strengthening is relevant to points we have already made. For example, our argument on p. 71 concerning the Iowa Electronic Markets obviously requires that Skeptic's strategy be computable.

Proposition 3.5 is mathematically trivial, but it suggests many nontrivial questions. For example, fixing a computational model (such as the one-head and one-tape Turing machine), we can ask questions such as this:

> Does there exist a winning strategy for Skeptic in the fair-coin game such that the move at step n can be computed in time $O(n^c)$, for some c? If yes, what is the infimum of such c?

Similar questions, which may be of practical interest when one undertakes to implement the game-theoretic approach, can be asked about other computational resources, such as the required memory.

In a different direction, we can ask about the rate at which Skeptic can increase his capital if the sequence of outcomes produced by Reality in the bounded forecasting game does not satisfy (3.11). For example, the construction in §3.2 shows that the following is true.

Proposition 3.6 *Skeptic has a computable winning strategy in the bounded forecasting game with the condition (3.3) that his capital tends to infinity replaced by the condition*

$$\limsup_{n\to\infty} \frac{\log \mathcal{K}_n}{n} > 0$$

that his capital increases exponentially fast.

It might be interesting to study the trade-off (if any) between the computational efficiency of a strategy and the rate at which its capital tends to infinity.

Questions similar to that answered in Proposition 3.6 have been asked in algorithmic probability theory: see Schnorr ([267], [269]) in connection with the strong law of large numbers and Vovk (1987) in connection with the law of the iterated logarithm and the recurrence property.

References

86. A. Philip Dawid and Vladimir G. Vovk. Prequential probability: Principles and properties. *Bernoulli*, 5:125–162, 1999. This paper is the first to state formally a game-theoretic law of the iterated logarithm.

267. Claus Peter Schnorr. Klassifikation der Zufallsgesetze nach Komplexität und Ordnung. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 16:1–21, 1970.

269. Claus Peter Schnorr. *Zufälligkeit und Wahrscheinlichkeit: Eine algorithmische Bergründung der Wahrscheinlichkeitstheorie,* volume 218 of *Lecture Notes in Mathematics*. Springer, Berlin, 1971.

320. Vladimir G. Vovk. The law of the iterated logarithm for random Kolmogorov, or chaotic, sequences. *Theory of Probability and Its Applications*, 32:413–425, 1987. This paper gives an algorithmic proof of the law of the iterated logarithm which was adapted to a game-theoretic framework in Dawid and Vovk [86].