Game-theoretic versions of Kolmogorov's strong law of large numbers

V. G. Vovk

Technical Report

CLRC-TR-99-05

June 1995

Royal Holloway
University of London

Department of Computer Science
Egham, Surrey TW20 0EX, England

Abstract

We prove two variants of Kolmogorov's strong law of large numbers in a completely worst-case framework, eschewing any probabilistic assumptions. The first variant is an assertion about a game involving the Bookmaker predicting the values of unprobabilized random variables; in an intuitive sense it is much stronger than the usual strong law of large numbers for martingales. The second variant is an assertion about a security market.

The final version of this paper was prepared for the seminar on the foundations of probability (Aalborg University, June 1995) organized by Steffen Lauritzen.

1 Predictive strong law of large numbers

In this section we consider the following perfect-information game between 3 players, the Bookmaker, the Statistician, and the Nature. The game proceeds in trials. At each trial i, $i = 1, 2, \ldots$, the Bookmaker tries to predict the real number X_i the Nature is to produce at the end of the trial. His prediction consists of two numbers, E_i and $D_i \geq 0$; roughly, E_i is the Bookmaker's expectation of X_i, and D_i is his expectation of the accuracy $L_i := (X_i - E_i)^2$ (measured by the *Brier scoring rule*, see Dawid [1]) of the prediction E_i.

Along the lines of Chapter 3 of de Finetti [2] we give an operative interpretation to the numbers E_i and D_i as follows. Before X_i is disclosed, the Bookmaker lets the Statistician buy any amount, positive or negative, of *X_i-tickets* for \$$E_i$ each and *L_i-tickets* for \$$D_i$ each. An X_i-ticket (resp. L_i-ticket) is a contract which obliges the Bookmaker to pay the Statistician \$$X_i$ (resp. \$$(X_i - E_i)^2$) after X_i is disclosed. We will use the notation e_i and d_i for the number of X_i-tickets and L_i-tickets, respectively, bought by the Statistician at trial i. We describe the unfolding of the game, including the evolution of the Statistician's capital $\mathcal{K}_i$, as follows:

FOR $i = 1, 2, \ldots$:
  Bookmaker selects $E_i \in \mathbb{R}$ and $D_i \geq 0$
  Statistician selects $e_i, d_i \in \mathbb{R}$
  $$\mathcal{K}_i := \mathcal{K}_{i-1} - e_i E_i - d_i D_i \tag{1}$$
  Nature selects $X_i \in \mathbb{R}$
  $$\mathcal{K}_i := \mathcal{K}_i + e_i X_i + d_i (X_i - E_i)^2. \tag{2}$$

Initially, the Statistician's capital is $\mathcal{K}_0 := 1$. We call the pair (e_i, d_i) the *portfolio* held by the Statistician at trial i. Equation (1) shows how the Statistician's capital decreases when he buys the new portfolio and (2) describes the proceeds from holding the portfolio. We can summarize (1) and (2) as follows:

$$\mathcal{K}_i := \mathcal{K}_{i-1} + e_i(X_i - E_i) + d_i \left((X_i - E_i)^2 - D_i \right). \tag{3}$$

An important difference between our framework and de Finetti's is that de Finetti assumes that the Bookmaker can directly price any ticket of interest, and we try to minimize the quantity of tickets priced by the Bookmaker. In particular, for discussing Kolmogorov's strong law of large numbers it suffices to assume that the Bookmaker can price only the X_i-tickets and L_i-tickets.

To complete the description of our game, which will be denoted by $\mathcal{G}$, it remains to specify the rule for determining who won the game given its *path*

$$E_1 D_1 e_1 d_1 X_1 E_2 D_2 e_2 d_2 X_2 \ldots. \tag{4}$$

We say that this path *satisfies SLLN* if

$$\lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} (X_i - E_i) = 0. \tag{5}$$

Intuitively,

- the aim of the Statistician is to demonstrate that the realized path (4) satisfies SLLN; however, even if the path does not satisfy SLLN, the Statistician will be consoled if his capital increases infinitely:

$$\mathcal{K}_i \to \infty \quad (i \to \infty); \tag{6}$$

  the Statistician is prohibited to borrow: he is declared to lose the game as soon as $\mathcal{K}_i$ becomes negative;

- the aim of the Nature is to prevent the Statistician from attaining his goal;

- the Bookmaker does not have his own goals and can cooperate with either the Statistician or the Nature.

Formally, $\mathcal{G}$ is won by the Statistician if (5) or (6) holds and his capital $\mathcal{K}_i \geq 0$, $\forall i$; otherwise, $\mathcal{G}$ is won by the Nature. (Thus the Bookmaker never wins the game.)

Now we can state the "predictive version" of Kolmogorov's strong law of large numbers.

Theorem 1 *Let $\mathcal{G}^{<\infty}$ (resp. $\mathcal{G}^{=\infty}$) be the game $\mathcal{G}$ with the additional requirement that the Bookmaker should ensure*

$$\sum_{i=1}^{\infty} \frac{D_i}{i^2} < \infty \tag{7}$$

(resp.

$$\sum_{i=1}^{\infty} \frac{D_i}{i^2} = \infty). \tag{8}$$

The Statistician has a winning strategy in $\mathcal{G}^{<\infty}$, and the Nature has a winning strategy in $\mathcal{G}^{=\infty}$.

Let $\overline{\mathcal{G}}$ be the game whose only difference from $\mathcal{G}$ is that the Statistician is required to select $d_i \geq 0$. Intuitively, this means that D_i is the Bookmaker's upper estimate on the variance of X_i. We can state for $\overline{\mathcal{G}}$ an analogue of Theorem 1.

Theorem 2 *Let $\overline{\mathcal{G}}^{<\infty}$ (resp. $\overline{\mathcal{G}}^{=\infty}$) be the game $\overline{\mathcal{G}}$ with the additional requirement that the Bookmaker should ensure (7) (resp. (8)). The Statistician (resp. the Nature) has a winning strategy in $\overline{\mathcal{G}}^{<\infty}$ (resp. $\overline{\mathcal{G}}^{=\infty}$).*

2 Financial-theoretic strong law of large numbers

In this section we consider a security market where K securities are traded; its working is described by the following game $\mathcal{G}_K$:

FOR $i = 1, 2, \ldots$:
  Market selects $\mu_i \in]0,1]^K$ such that $\sum_{k=1}^K \mu_i(k) = 1$
  Gambler selects $\sigma_i \in \mathbb{R}^K$
  Statistician selects $p_i \in \mathbb{R}^K$
  Nature selects $\pi_i \in [-1, \infty[^K$ such that $\mu_i \cdot \pi_i = 0$
  $\mathcal{K}_i := \mathcal{K}_{i-1} + p_i \cdot \pi_i$.

The initial capital $\mathcal{K}_0$ is 1; "$\cdot$" stands for the dot product. Each $\mu_i = (\mu_i(1), \ldots, \mu_i(K))$ is interpreted as the "market portfolio" at trial i: say, $\mu_i(k)$ can be the total value (according to the current prices) of all the shares outstanding of security k; we normalize μ_i so that $\sum_k \mu_i(k) = 1$. Each $\pi_i = (\pi_i(1), \ldots, \pi_i(K))$ is interpreted as the return on the securities at trial i: say, $\pi_i(k) = -1$ means that security k became worthless; $\pi_i(k) = 0$ means that k's price did not change; $\pi_i(k) = 0.07$ means that k's price was 7% up. Each $\sigma_i = (\sigma_i(1), \ldots, \sigma_i(K))$ is the portfolio held by the Gambler at trial i: $\sigma_i(k)$ is the capital he invested into security k (it is possible that $\sigma_i(k) < 0$; in other words, unlimited short selling is allowed). Analogously, each p_i is the Statistician's portfolio at trial i.

The requirement $\mu_i \cdot \pi_i = 0$ means that the market portfolio does not change its value; in other words, we are using the value of the market portfolio as the numeraire. (So saying that p_i is the Statistician's portfolio at trial i, we mean that the Statistician invests $p_i(k)$ into security k, $k = 1, \ldots, K$, and invests the remaining $\mathcal{K}_{i-1} - \sum_k p_i(k)$ into the market portfolio[1].)

The Statistician's goal is to enforce his favorite strong law of large numbers, which says here that on the average the Gambler will neither profit nor lose:

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \sigma_i \cdot \pi_i = 0. \tag{9}$$

If he is unable to do it, the only thing that can comfort him is his becoming infinitely rich; as before, he does not tolerate any chance of getting into a position of debt. Formally, the Statistician wins the game if $\mathcal{K}_i \geq 0$, $\forall i$, and also (9) or (6) holds.

Since everything happens "modulo the market portfolio", it is natural to consider the "reduced portfolio"

$$\overline{\sigma}_i(k) := \sigma_i(k) - c_i \mu_i(k),$$

[1] To allow the investors to keep part of their capital in cash, we can formally add money to the list of securities traded in the market.

where c_i is the constant that makes $\sum_k \overline{\sigma}_i(k) = 0$ (namely, $c_i = \sum_k \sigma_i(k)$). We will say that $\overline{\sigma}_i$ is an *arbitrage portfolio.*

Theorem 3 *Let $\mathcal{G}_K^{<\infty}$ be the game $\mathcal{G}_K$ where the Gambler is required to ensure that*

$$\sum_{i=1}^{\infty} i^{-2} \sum_{k=1}^{K} \frac{\overline{\sigma}_i^2(k)}{\mu_i(k)} < \infty.$$

The Statistician has a winning strategy in $\mathcal{G}_K^{<\infty}$.

3 Proofs

3.1 Martingale-theoretic strong law of large numbers

First we introduce a generalization of the games $\mathcal{G}$, $\overline{\mathcal{G}}$, and $\mathcal{G}_K$ considered in Sections 1 and 2. Fix some sets $\mathcal{M}_{\mathrm{B}}$, $\mathcal{M}_{\mathrm{S}}$, and $\mathcal{M}_{\mathrm{N}}$ from which the Bookmaker, the Statistician, and the Nature, respectively, will select their moves. Let

$$\lambda\colon \mathcal{M}_{\mathrm{B}} \times \mathcal{M}_{\mathrm{S}} \times \mathcal{M}_{\mathrm{N}} \to \mathbb{R}$$

be the *reward function* for the Statistician. The game $\mathcal{G}_\lambda(\mathcal{K}_0)$, where $\mathcal{K}_0 \in \mathbb{R}$ is the Statistician's initial capital, is defined as follows:

FOR $i = 1, 2, \ldots$:
  Bookmaker selects $b_i \in \mathcal{M}_{\mathrm{B}}$
  Statistician selects $p_i \in \mathcal{M}_{\mathrm{S}}$
  Nature selects $x_i \in \mathcal{M}_{\mathrm{N}}$
  $\mathcal{K}_i := \mathcal{K}_{i-1} + \lambda(b_i, p_i, x_i)$.

A *situation* in our game is a finite sequence $b_1x_1 \ldots b_nx_n$ with all $b_i \in \mathcal{M}_{\mathrm{B}}$ and all $x_i \in \mathcal{M}_{\mathrm{N}}$; that is, it is a record of the Statistician's opponents' moves during the first n trials. To $n = 0$ corresponds the empty sequence $\square$. A *path* is an infinite sequence $b_1x_1b_2x_2\ldots$; the *sample space* of the game is the set of all paths. If ω is a path $b_1x_1b_2x_2\ldots$ and n is a nonnegative integer, ω^n stands for the situation $b_1x_1\ldots b_nx_n$. A *variable*[2] is a real-valued function defined on the sample space. A *process* is a real-valued (if not explicitly stated otherwise) function defined on the set of all situations. As usual in probability theory, we identify a process S with the sequence $S_0, S_1, \ldots$ of variables defined by

$$S_n(\omega) := S(\omega^n).$$

(Thus S_0 is constant.)

A process A is *predictable* if, for all n, $A_n(b_1x_1b_2x_2\ldots)$ does not depend on x_n. A *martingale* is a process S for which there exists the Statistician's strategy

[2] Following Shafer [5], we shorten "random variable" and "stochastic process" to "variable" and "process", respectively. (There is nothing random or stochastic in our framework.)

$p = p_n(\omega)$ (note that the Statistician's strategy is a predictable $\mathcal{M}_S$-valued process) ensuring

$$\mathcal{K}_n = S(b_1x_1 \ldots b_nx_n),\ \forall n,$$

in the game $\mathcal{G}_\lambda(S_0)$. We will say that p *gives rise* to S *in* $\mathcal{G}_\lambda(S_0)$. A process A is *increasing* if $A_0 = 0$ and $A_n \le A_{n+1}$, $\forall n$, on all paths. A *submartingale* is a process of the form $T = S + A$, where S is a martingale and A is an increasing predictable process; any such A will be called a *compensator* for T. A *supermartingale* is a process T for which there exists a martingale S such that everywhere

$$T_{n+1} - T_n \le S_{n+1} - S_n,\ \forall n; \tag{10}$$

in other words, a supermartingale is a process T of the form $T = S - A$ with S a martingale and A an increasing (but not necessarily predictable) process. (Therefore, the difference between the notions of supermartingale and martingale is that in the case of supermartingale the Statistician is allowed to throw away part of his money at every trial.) A *semimartingale* is a process of the form $T = S + A$, where S is a supermartingale and A is an increasing predictable process; any such A will be called a *supercompensator* for T. Any strategy p which gives rise to a martingale S satisfying (10) will be said to *bound* the supermartingale T.

Now we impose some restrictions on the reward function λ (it is instructive to compare them with the restrictions in [7]):

1. $\mathcal{M}_S$ is a closed convex cone in a Euclidean space with norm $\|\cdot\|$.
2. For every fixed $b \in \mathcal{M}_B$ and $x \in \mathcal{M}_N$, the function $p \mapsto \lambda(b, p, x)$ is linear on $\mathcal{M}_S$.

In the case of the games $\mathcal{G}$ and $\overline{\mathcal{G}}$, $\mathcal{M}_S$ is a cone in $\mathbb{R}^2$, and in the case of the games $\mathcal{G}_K$ (where the Market and the Gambler are aggregated into the Bookmaker), $\mathcal{M}_S$ is a cone in (in fact, the whole of) $\mathbb{R}^K$.

In this section we will consider supermartingales and semimartingales, but to prove the assertions in Section 1 it would be sufficient to consider only the supermartingales which are martingales and semimartingales which are submartingales.

Our assumptions immediately imply

Lemma 1 *The space of supermartingales is a convex cone: if $S^1, \ldots, S^m$ are supermartingales and $t_1 \ge 0, \ldots, t_m \ge 0$, then $t_1S^1 + \ldots + t_mS^m$ is a supermartingale as well.*

Proof If S^j is bounded by the Statistician's strategy p^j in $\mathcal{G}_\lambda(S_0^j)$, $j = 1, \ldots, m$, then $t_1S^1 + \ldots + t_mS^m$ is bounded by the strategy $t_1p^1 + \ldots + t_mp^m$ in $\mathcal{G}_\lambda(S_0^1 + \ldots + S_0^m)$. $\square$

With each supermartingale S we associate the predictable process $\|S\|$ defined by

$$\begin{aligned} \|S\|_n(\omega) := \inf\{\|p\| : \\ S(\omega^{n-1} * b * x) - S(\omega^{n-1}) \le \lambda(b, p, x), \ \forall\, b \in \mathcal{M}_{\mathrm{B}}, x \in \mathcal{M}_{\mathrm{N}}\} \end{aligned} \tag{11}$$

($s * a$ stands for finite sequence s extended by adding one more element a on the right). We say that the Statistician's strategy p in $\mathcal{G}_\lambda(S_0)$ is a *minimal bounding strategy for* S if p bounds S and $\|S\|_n = \|p_n\|$ everywhere for all n.

Lemma 2 *For each supermartingale S there exists a minimal bounding strategy.*

Proof We are only required to prove that the inf in (11) is attained. Fix n and ω and denote by $P(b, x)$ the set of all $p \in \mathcal{M}_{\mathrm{S}}$ which satisfy the inequality in (11). Each $P(b, x)$ is closed; therefore, their intersection P is also closed. We are required to prove that $\inf_{p \in P} \|p\|$ is attained; this follows from the compactness of every closed ball in a Euclidean space. □

Lemma 3 *Let S^1, S^2,... be a sequence of nonnegative supermartingales and t_1, t_2,... be a sequence of nonnegative numbers such that*

$$\sum_{j=1}^{\infty} t_j S_0^j < \infty$$

and, for any situation s,

$$\sum_{j=1}^{\infty} t_j \|S^j\|(s) < \infty.$$

Then the sum

$$S := \sum_{j=1}^{\infty} t_j S^j$$

is a nonnegative supermartingale and

$$\|S\| \le \sum_{j=1}^{\infty} t_j \|S^j\|.$$

Proof It suffices to fix for each S^j a minimal bounding strategy p^j and put

$$p := \sum_{j=1}^{\infty} t_j p^j;$$

it is easy to prove by simultaneous induction in n that $S_n < \infty$, $\forall n$, and p bounds S. □

Now our argument will follow [8], which, in its turn, follows Liptser's proof of Kolmogorov's strong law of large numbers. The following assertion is an analog of Doob's martingale convergence theorem for nonnegative supermartingales; by $\exists \lim_{n\to\infty} S_n$ we mean that the limit exists and is finite.

Lemma 4 *For any nonnegative supermartingale S there exists a nonnegative martingale S^* such that*

$$S_0^* = S_0, \ \|S^*\| \le \|S\|,$$

and always

$$\exists \lim_{n\to\infty} S_n \ \ or \ \ \lim_{n\to\infty} S_n^* = \infty. \tag{12}$$

Proof (Doob) Fix the Statistician's strategy p which bounds the supermartingale S in the game $\mathcal{G}_\lambda(S_0)$. Let a, b be positive rational numbers such that $a < b$. Define $\tau_0 := 0$ and, for $m = 1, 2, \ldots$,

$$\sigma_m := \min\{i > \tau_{m-1} : S_i > b\}, \ \tau_m = \min\{i > \sigma_m : S_i < a\}.$$

Let $p^{(a,b)}$ be the Statistician's strategy

$$p_i^{(a,b)} := \begin{cases} p_i, \text{ if } \exists m\colon \tau_{m-1} < i \le \sigma_m, \\ 0, \text{ otherwise,} \end{cases}$$

in the game $\mathcal{G}_\lambda(S_0)$, and $S^{a,b}$ be the martingale $p^{a,b}$ gives rise to. It is easy to see that

$$S_0^{(a,b)} = S_0, \ \|S^{(a,b)}\| \le \|S\|,$$

and always

$$\left.\begin{array}{l} \underline{\lim}_{n\to\infty} S_n < a \\ \overline{\lim}_{n\to\infty} S_n > b \end{array}\right\} \Longrightarrow \lim_{n\to\infty} S_n^{(a,b)} = \infty. \tag{13}$$

Arrange all such pairs (a, b) in a sequence $(a_1, b_1), (a_2, b_2), \ldots$ and put

$$S^* := \frac{1}{2} S + \sum_{m=1}^{\infty} 2^{-m-1} S^{(a_m, b_m)}.$$

By Lemma 3, S^* is a supermartingale with

$$\|S^*\| \le \frac{1}{2}\|S\| + \sum_{m=1}^{\infty} 2^{-m-1} \|S^{(a_m,b_m)}\| \le \|S\|.$$

To prove (12), it suffices to rewrite it as

$$\left(\neg\exists \lim_{n\to\infty} S_n\right) \Longrightarrow \lim_{n\to\infty} S_n^* = \infty,$$

notice that the antecedent is equivalent to

$$\exists (a, b)\colon \underline{\lim}_{n\to\infty} S_n < a \ \& \ \overline{\lim}_{n\to\infty} S_n > b,$$

and make use of (13). (Now S^* is guaranteed to be only a supermartingale, but it is easy to make it into a martingale.) ◻

Let us introduce some new terminology. We say that a property E of a path holds *almost surely* (a.s.) if there exists a nonnegative martingale S such that

$S_n \to \infty$ on the paths outside E; without loss of generality we can assume $S_0 = 1$. (We will say that S *witnesses* that E holds a.s.) In other words, E holds a.s. if the Statistician can, without risking to go bankrupt, either ensure E or become infinitely rich. We can see that Lemma 4 implies that each nonnegative supermartingale is convergent a.s. (this is a more familiar formulation of Doob's theorem).

Lemma 5 *If T is a nonnegative semimartingale and A is its supercompensator, then, almost surely,*

$$A_\infty < \infty \Rightarrow \exists \lim_{n\to\infty} T_n. \tag{14}$$

Proof Let $S := T - A$; S is a supermartingale. For $C = 1, 2, \ldots$, define the nonnegative supermartingales S^C by the requirements $S_0^C = C$ and

$$\Delta S_i^C = \begin{cases} \Delta S_i, \text{ if } A_i \le C, \\ 0, \text{ otherwise,} \end{cases} \qquad i = 1, 2, \ldots,$$

(Δa_i stands for $a_i - a_{i-1}$). Since $\|S^C\| \le \|S\|$, Lemma 3 implies that

$$R := \sum_{C=1}^{\infty} 2^{-C} (S^C)^*,$$

where * is the transformation from Lemma 4, witnesses that (14) holds a.s. □

We say that a process A is a *quadratic supervariation* of a martingale S if S^2 is a semimartingale and A is a supercompensator of S^2.

Lemma 6 *If S is a martingale and A is its quadratic supervariation, then, almost surely,*

$$A_\infty < \infty \Rightarrow \exists \lim_{n\to\infty} S_n.$$

Proof If A is a supercompensator of S^2, then A is a supercompensator of $(S+1)^2$ as well. By Lemma 5, S_n^2 and $(S_n+1)^2$ converge when $A_\infty < \infty$, almost surely. It remains to note that

$$S_n = \frac{1}{2}\Big((S_n+1)^2 - S_n^2 - 1\Big).$$

□

Now we can formulate the main result of this section.

Theorem 4 *If A is a quadratic supervariation of a martingale S, then, almost surely,*

$$\sum_{i=1}^{\infty} \frac{\Delta A_i}{i^2} < \infty \Longrightarrow \lim_{n\to\infty} \frac{S_n}{n} = 0.$$

Proof Noticing that

$$A'_n := \sum_{i=1}^{n} \frac{\Delta A_i}{i^2}$$

is a quadratic supervariation of the martingale

$$S'_n := \sum_{i=1}^{n} \frac{\Delta S_i}{i}$$

and applying Lemma 6, we deduce that, almost surely,

$$\sum_{i=1}^{\infty} \frac{\Delta A_i}{i^2} < \infty \Rightarrow \exists \sum_{i=1}^{\infty} \frac{\Delta S_i}{i}.$$

It remains to apply Kronecker's lemma (see, e.g., [6], Lemma 3.2.3), which implies that, for any sequence $t_1, t_2, \ldots$ of real numbers,

$$\exists \sum_{i=1}^{\infty} \frac{t_i}{i} \Rightarrow \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} t_i = 0.$$

$\square$

3.2 Statistician's strategy

It is easy to deduce from Theorem 4 the existence of the Statistician's strategies required in Theorems 1, 2, and 3. The case of Theorem 1 reduces to the case of Theorem 2, so we are only required to prove that the Statistician has winning strategies in $\overline{\mathcal{G}}^{<\infty}$ and $\mathcal{G}_K^{<\infty}$. The case of $\overline{\mathcal{G}}^{<\infty}$ is easy: it suffices to note that

$$S_n := \sum_{i=1}^{n} (X_i - E_i)$$

is a martingale and

$$A_n := \sum_{i=1}^{n} D_i$$

is its quadratic supervariation.

It remains to consider the game $\mathcal{G}_K^{<\infty}$. In this case, it suffices to apply Theorem 4 to the martingale

$$S_n := \sum_{i=1}^{n} \sigma_i \cdot \pi_i; \tag{15}$$

the next lemma ensures that

$$A_n := \sum_{i=1}^{n} \sum_{k=1}^{K} \frac{\overline{\sigma}_i^2(k)}{\mu_i(k)} \tag{16}$$

is a quadratic supervariation of S_n.

Lemma 7 *(16) is a quadratic supervariation of (15).*

Proof We are required to prove that there exists the Statistician's strategy p such that, for all $i \geq 1$,

$$(S_i^2 - A_i) - (S_{i-1}^2 - A_{i-1}) \leq p_i \cdot \pi_i,$$

i.e.,

$$(\sigma_i \cdot \pi_i)^2 + 2S_{i-1}(\sigma_i \cdot \pi_i) - \sum_k \frac{\overline{\sigma}_i^2(k)}{\mu_i(k)} \leq p_i \cdot \pi_i.$$

It is easy to see that the existence of such p is equivalent to the existence of p such that

$$(\sigma_i \cdot \pi_i)^2 - \sum_k \frac{\overline{\sigma}_i^2(k)}{\mu_i(k)} \leq p_i \cdot \pi_i,$$

i.e., dropping the indices i,

$$(\sigma \cdot \pi)^2 - p \cdot \pi \leq \sum_k \frac{\overline{\sigma}^2(k)}{\mu(k)}. \tag{17}$$

Let us see how the Statistician can choose such $p \in \mathbb{R}^K$. Recall that $\pi \in \mathbb{R}^K$ must satisfy

$$\pi(k) \geq -1,\ \forall k,\ \ \mu \cdot \pi = 0. \tag{18}$$

In the left-hand side of (17) we have a convex function of π, so it suffices for the Statistician to ensure that (17) holds for π the vertices of the polygon (18). These vertices are π^j, $j = 1, \ldots, K$, with

$$\pi^j(k) = \begin{cases} \frac{1}{\mu(j)} - 1, \text{ if } k = j, \\ -1, \text{ otherwise.} \end{cases}$$

Substituting this into (17), we obtain (assuming, without loss of generality, that σ is an arbitrage portfolio)

$$\left(\frac{\sigma(j)}{\mu(j)}\right)^2 - \frac{p(j)}{\mu(j)} + \sum_k p(k) \leq \sum_k \frac{\sigma^2(k)}{\mu(k)}.$$

It remains to put

$$p(j) := \frac{\sigma^2(j)}{\mu(j)} - \mu(j) \sum_k \frac{\sigma^2(k)}{\mu(k)}$$

(note that $\sum_k p(k) = 0$). ◻

3.3 Nature's strategy

The aim of this section is to prove the existence of the Nature's winning strategies under the conditions of Theorems 1 and 2. It suffices only to prove that the Nature has a winning strategy in $\mathcal{G}^{=\infty}$. By Martin's theorem about the determinacy of quasi-Borel games [4], either the Statistician helped by the Bookmaker or the Nature has a winning strategy in $\mathcal{G}^{=\infty}$. Therefore, it suffices to show that the Statistician helped by the Bookmaker cannot have a winning strategy. Assume that such a strategy exists; we will denote it by $\mathcal{S}$. We will attain our goal if we prove that the Nature has a randomized strategy beating $\mathcal{S}$ with probability 1. We will construct such a strategy following Kolmogorov [3].

The Nature's strategy is as follows: if $D_i \le i^2$,

$$X_i := \begin{pmatrix} E_i + i \\ E_i - i \\ E_i \end{pmatrix} \text{ with probability } \begin{pmatrix} D_i/(2i^2) \\ D_i/(2i^2) \\ 1 - D_i/i^2 \end{pmatrix}, \tag{19}$$

respectively; if $D_i > i^2$,

$$X_i := \begin{pmatrix} E_i + \sqrt{D_i} \\ E_i - \sqrt{D_i} \end{pmatrix} \text{ with probability } \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix}. \tag{20}$$

Let this strategy play against $\mathcal{S}$. Since always $\mathcal{K}_i \ge 0$, (3) implies that $\mathcal{K}_i$ is a nonnegative martingale (in the usual probabilistic sense). Therefore, $\mathcal{K}_i$ is bounded above almost surely (this follows from, e.g., Doob's inequality or Doob's convergence theorem). Besides, the Borel—Cantelli—Levy lemma (see, e.g., Shiryaev [6], Corollary VII.5.2) implies that, almost surely, $|X_i - E_i| \ge i$ for infinitely many i; therefore, almost surely, SLLN is not satisfied.

To make this argument rigorous, it remains to actually construct a probability space $(\Omega, \mathbf{P})$ with random variables X_i, $i = 1, 2, \ldots$, satisfying (19) and (20). It is easy: we put $\Omega := \{-1, 0, 1\}^\infty$,

$$X_i(\omega) := \begin{cases} E_i, \text{ if } \omega_i = 0, \\ E_i \pm i, \text{ if } \omega_i = \pm 1 \text{ and } D_i \le i^2, \\ E_i \pm \sqrt{D_i}, \text{ if } \omega_i = \pm 1 \text{ and } D_i > i^2, \end{cases}$$

where $\omega = \omega_1 \omega_2 \ldots \in \Omega$, and

$$\mathbf{P}(\omega_i = 1 \mid \mathcal{F}_{i-1}) := \mathbf{P}(\omega_i = -1 \mid \mathcal{F}_{i-1}) := \begin{cases} D_i/(2i^2), \text{ if } D_i \le i^2, \\ 1/2, \text{ otherwise,} \end{cases}$$

where $\mathcal{F}_k$ is the σ-algebra generated by $\omega_1, \ldots, \omega_k$. (Recall that the Statistician/Bookmaker's strategy is fixed, and so E_i and D_i are $\mathcal{F}_{i-1}$-measurable.)

Acknowledgments

The research described in this publication was made possible in part by Grant No. MRS000 from the International Science Foundation.

References

[1] A. P. Dawid. Probability forecasting. In: *Encyclopedia of Statistical Sciences* (ed. by S. Kotz and N. L. Johnson). Vol. 7, p. 210–218. New York: Wiley (1986)

[2] B. de Finetti. *Theory of probability.* Vol. 1. London: Wiley (1974)

[3] A. N. Kolmogorov. Sur la loi forte des grands nombres. *C. r. Acad. sci. Paris* **191**, 910–912 (1930)

[4] D. A. Martin. An extension of Borel determinacy. *Ann. Pure Appl. Logic* **49**, 279–293 (1990)

[5] G. Shafer. *The art of causal conjecture*, to appear

[6] A. N. Shiryayev. *Veroyatnost'.* 2nd ed. Moscow: Nauka (1989). English translation of the 1st edition: *Probability.* Berlin: Springer (1984).

[7] V. G. Vovk. A game of prediction with expert advice. *COLT'95*, to appear

[8] V. G. Vovk. A purely martingale version of Kolmogorov's strong law of large numbers. *Probab. Theory Appl.*, to appear