Hoeffding's inequality in game-theoretic probability

Vladimir Vovk

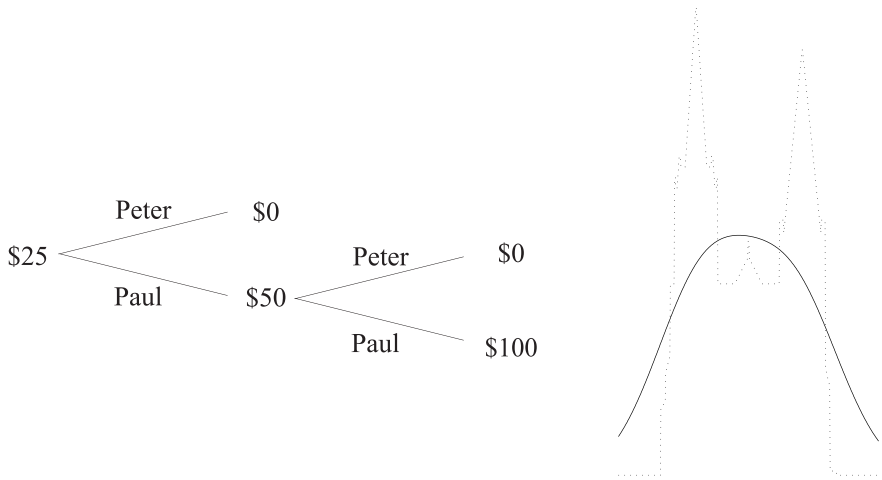

The Game-Theoretic Probability and Finance Project

Answer to Frequently Asked Question #3

August 19, 2007

Project web site:
http://www.probabilityandfinance.com

Abstract

This note makes the obvious observation that Hoeffding's original proof of his inequality remains valid in the game-theoretic framework. All details are spelled out for the convenience of future reference.

Contents

1 Introduction

The game-theoretic approach to probability was started by von Mises and greatly advanced by Ville [5]; however, it has been overshadowed by Kolmogorov's measure-theoretic approach [3]. The relatively recent book [4] contains game-theoretic versions of several results of probability theory, and it argues that the game-theoretic versions have important advantages over the conventional measure-theoretic versions. However, [4] does not contain any large-deviation inequalities. This note fills the gap by stating the game-theoretic version of Hoeffding's inequality ([2], Theorem 2).

2 Hoeffding's supermartingale

This section presents perhaps the most useful product of Hoeffding's method, a non-negative supermartingale starting from 1. This supermartingale will easily yield Hoeffding's inequality in the following section.

This is a version of the basic forecasting protocol from [4]:

GAME OF FORECASTING BOUNDED VARIABLES

Players: Sceptic, Forecaster, Reality

Protocol:

Sceptic announces $\mathcal{K}_0 \in \mathbb{R}$.
FOR $n = 1, 2, \dots$:
  Forecaster announces interval $[a_n, b_n] \subseteq \mathbb{R}$ and number $\mu_n \in (a, b)$.
  Sceptic announces $M_n \in \mathbb{R}$.
  Reality announces $x_n \in [a_n, b_n]$.
  Sceptic announces $\mathcal{K}_n \le \mathcal{K}_{n-1} + M_n(x_n - \mu_n)$.

On each round n of the game Forecaster outputs an interval $[a_n, b_n]$ which, in his opinion, will cover the actual observation x_n to be chosen by Reality, and also outputs his expectation μ_n for x_n. The forecasts are being tested by Sceptic, who is allowed to gamble against them. The expectation μ_n is interpreted as the price of a ticket which pays x_n after Reality's move becomes known; Sceptic is allowed to buy any number M_n, positive, zero, or negative, of such tickets. When x_n falls outside $[a_n, b_n]$, Sceptic becomes infinitely rich; without loss of generality we include the requirement $x_n \in [a_n, b_n]$ in the protocol; furthermore, we will always assume that $\mu_n \in (a_n, b_n)$. Sceptic is allowed to choose his initial capital $\mathcal{K}_0$ and is allowed to throw away part of his money at the end of each round.

It is important that the game of forecasting bounded variables is a perfect-information game: each player can see the other players' moves before making his or her (Forecaster and Sceptic are male and Reality is female) own move; there is no randomness in the protocol.

A *process* is a real-valued function defined on all finite sequences $(a_1, b_1, \mu_1, x_1, \dots, a_N, b_N, \mu_N, x_N)$, $N = 0, 1, \dots$, of Forecaster's and Reality's moves in the game of forecasting bounded variables. If we fix a strategy for Sceptic, Sceptic's capital $\mathcal{K}_N$, $N = 0, 1, \dots$, become a function of Forecaster's and Reality's previous moves; in other words, Sceptic's capital becomes a process. The processes that can be obtained this way are called (game-theoretic) *supermartingales*.

The following theorem is essentially inequality (4.16) in [2].

Theorem 1 *For any $h \in \mathbb{R}$, the process*

$$\prod_{n=1}^{N} \exp\left(h(x_n - \mu_n) - \frac{h^2}{8}(b_n - a_n)^2\right)$$

is a supermartingale.

Proof Assume, without loss of generality, that Forecaster is additionally required to always set $\mu_n := 0$. (Adding the same constant to a_n, b_n, and μ_n will not change anything for Sceptic.) Now we have $a_n < 0 < b_n$.

It suffices to prove that on round n Sceptic can make a capital of $\mathcal{K}$ into a capital of at least

$$\mathcal{K} \exp\left(hx_n - \frac{h^2}{8}(b_n - a_n)^2\right);$$

in other words, that he can obtain a payoff of at least

$$\exp\left(hx_n - \frac{h^2}{8}(b_n - a_n)^2\right) - 1$$

using the available tickets (paying x_n and costing 0). This will follow from the inequality

$$\exp\left(hx_n - \frac{h^2}{8}(b_n - a_n)^2\right) - 1 \le x_n \frac{e^{hb_n} - e^{ha_n}}{b_n - a_n} \exp\left(-\frac{h^2}{8}(b_n - a_n)^2\right),$$

which can be rewritten as

$$\exp(hx_n) \le \exp\left(\frac{h^2}{8}(b_n - a_n)^2\right) + x_n \frac{e^{hb_n} - e^{ha_n}}{b_n - a_n}. \tag{1}$$

Our goal is to prove (1). By the convexity of the function exp, it suffices to prove

$$\frac{x_n - a_n}{b_n - a_n} e^{hb_n} + \frac{b_n - x_n}{b_n - a_n} e^{ha_n} \le \exp\left(\frac{h^2}{8}(b_n - a_n)^2\right) + x_n \frac{e^{hb_n} - e^{ha_n}}{b_n - a_n},$$

i.e.,

$$\frac{b_n e^{ha_n} - a_n e^{hb_n}}{b_n - a_n} \le \exp\left(\frac{h^2}{8}(b_n - a_n)^2\right),$$

i.e.,

$$\ln\left(b_n e^{ha_n} - a_n e^{hb_n}\right) \le \frac{h^2}{8}(b_n - a_n)^2 + \ln(b_n - a_n). \tag{2}$$

The derivative of the left-hand side of (2) is

$$\frac{a_n b_n e^{ha_n} - a_n b_n e^{hb_n}}{b_n e^{ha_n} - a_n e^{hb_n}}$$

and the second derivative, after cancellations and regrouping, is

$$(b_n - a_n)^2 \frac{\left(b_n e^{ha_n}\right)\left(-a_n e^{hb_n}\right)}{\left(b_n e^{ha_n} - a_n e^{hb_n}\right)^2}.$$

The last ratio is of the form $u(1-u)$ where $0 < u < 1$. Hence it does not exceed $1/4$, and the second derivative itself does not exceed $(b_n - a_n)^2/4$. Inequality (2) now follows from the second-order Taylor expansion of the left-hand side around $h = 0$. ∎

The relation between the game-theoretic and measure-theoretic approaches to probability is described in [4], Chapter 8. Intuitively, the generality of the game-theoretic protocol stems from the fact that Forecaster is not asked to produce a full-blown probability forecast for x_n: only the elements (a_n, b_n, μ_n) that we really need for our mathematical result enter the game of forecasting bounded variables. Besides, the players are allowed to react to each other moves; in particular, Reality may react to Forecaster's moves and both Reality and Forecaster may react to Sceptic's moves (the latter is important in applications to defensive forecasting: see, e.g., [6]). It is remarkable that many measure-theoretic proofs carry over in a straightforward manner to game-theoretic probability.

3 Hoeffding's inequality

We start from the definition of upper probability, a game-theoretic counterpart (along with lower probability) of the standard measure-theoretic notion of probability. Suppose the game of forecasting bounded variables lasts a known number N of rounds. (See [4] for the general definition.) The *sample space* is the set of all sequences $(a_1, b_1, \mu_1, x_1, \ldots, a_N, b_N, \mu_N, x_N)$ of Forecaster's and Reality's moves in the game. An *event* is a subset of the sample space. The *upper probability* of an event E is the infimum of the initial value of non-negative supermartingales that take value at least 1 on E. (See [4], Chapter 8, for a demonstration that this definition agrees with measure-theoretic probability.)

Theorem 1 immediately gives Hoeffding's inequality (cf. [2], the proof of Theorem 2) when combined with the definition of game-theoretic probability:

Corollary 1 *Suppose the game of forecasting bounded variables lasts a fixed number N of rounds. If all a_n and b_n are given in advance and $t > 0$ is a* *known constant, the upper probability of the event*

$$\frac{1}{N}\sum_{n=1}^{N}(x_n - \mu_n) \geq t \tag{3}$$

does not exceed $e^{-2N^2t^2/C}$*, where* $C := \sum_{n=1}^{N}(b_n - a_n)^2$*.*

(The reader will see that it is sufficient for Sceptic to know only C at the start of the game, not the individual a_n and b_n.)

Proof The supermartingale of Theorem 1 starts from 1 and achieves

$$\prod_{n=1}^{N} \exp\left(h(x_n - \mu_n) - \frac{h^2}{8}(b_n - a_n)^2\right) \geq \exp\left(hNt - \frac{h^2}{8}C\right) \tag{4}$$

on the event (3). The right-hand side of (4) attains its maximum at $h := 4Nt/C$, which gives the statement of the corollary. ∎

Remark The measure-theoretic counterpart of Corollary 1 is sometimes referred to as the Hoeffding–Azuma inequality, in honour of Kazuoki Azuma (吾妻一興)[1]. The martingale version, however, is also stated in Hoeffding's paper ([2], the end of Section 2).

Acknowledgments

This work is inspired by a question asked by Yoav Freund. It has been partially supported by EPSRC (grant EP/F002998/1).

References

[1] Kazuoki Azuma. Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal*, 68:357–367, 1967.

[2] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.

[3] Andrei N. Kolmogorov. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Springer, Berlin, 1933. English translation: *Foundations of the Theory of Probability*. Chelsea, New York, 1950.

[4] Glenn Shafer and Vladimir Vovk. *Probability and Finance: It's Only a Game!* Wiley, New York, 2001.

[5] Jean Ville. *Etude critique de la notion de collectif*. Gauthier-Villars, Paris, 1939.

[6] VladimirVovk. Predictions as statements and decisions, The Game-Theoretic Probability and Finance project, `http://probabilityandfinance.com`, Working Paper #17, June 2006.