Pseudo-prophet inequalities in average-optimal stopping

Pieter Allaart[*], Michael Monticino

Mathematics Department

University of North Texas

April 18, 2001

Abstract

This note considers the average-optimal expected return of two players observing independent random variables $X_1, \ldots, X_n$, whose distributions are generated at random. One player, the pseudo prophet, knows the distributions prior to observing the random variables. The other player, the gambler, has no such foresight. Sharp difference and ratio comparisons of the two players' returns are given.

MSC: 60G40

Keywords: Optimal stopping; Random distributions; Prophet inequality

[*]Corresponding author: Pieter Allaart, Mathematics Department, Denton, TX 76203-1430, USA

1 Introduction

A classical problem in optimal stopping theory is to compare the optimal expected return of an ordinary gambler to that of a player with complete foresight. Such comparisons are called *prophet inequalities*, and a vast literature exists about them. See Hill and Kertz (1992) for a survey.

This note considers a related problem in which both players have less than complete foresight into the future. Specifically, let $X_1, \ldots, X_n$ be independent random variables, observed sequentially by two players. Assume that the distributions of the variables are known in advance to the first player, but unknown to the second. The first player will be called a "pseudo prophet", since he has partial, but not complete, foresight into the future. The second player will be referred to as the "gambler." Assume that the gambler adopts a stop rule that will maximize his expected return EX_τ on the average (in a sense to be made precise), over all possible distributions. How much larger, on the average, can the pseudo prophet's expected return be compared to that of the gambler?

There are two natural scenarios:

Scenario 1. Nature picks a distribution P at random according to some mechanism, and $X_1, \ldots, X_n$ are sampled sequentially from P.

Scenario 2. At each stage $1 \leq i \leq n$, nature picks a distribution P_i at random, independent of previous distributions, and X_i is sampled from P_i.

In Scenario 1, the gambler is able to gather information about P while playing. Thus, for large n, the pseudo prophet's advantage can be expected to be smaller than in Scenario 2, where no information gathering is possible. Samuels (1981) gives average-optimal stop rules in Scenario 1 when P is uniform with a two-sided Pareto prior on the endpoints, but for the general case no sharp upper bounds for the pseudo prophet's advantage seem to be known.

Scenario 2 is analytically more tractable, since here the gambler essentially observes an i.i.d. sequence of random variables from the average distribution $\bar{P}$. The purpose of this note is to show that the extremal case is when P is Dirac with probability one, thereby reducing the problem to an ordinary prophet/gambler comparison for i.i.d. random variables. A theorem by Hill and Kertz (1982) can then be invoked to obtain sharp ratio and difference inequalities.

2 The main result

Let $\mathcal{P}$ be the space of probability distributions on $\mathbb{R}^+ = [0, \infty)$, endowed with the weak* topology. Let $\mathcal{P}_0$ be the subset of $\mathcal{P}$ of distributions having a finite first moment. Note that $\mathcal{P}_0$ is Borel measurable since its complement in $\mathcal{P}$ can be written as

$$\bigcap_{k=1}^{\infty} \bigcup_{m=1}^{\infty} \left\{ P : \int_{[0,m]} x dP(x) > k \right\}.$$

A *prior* is a Borel probability measure on the space $\mathcal{P}$. Call a prior Q *integrable* if $Q(\mathcal{P}_0) = 1$. For an integrable prior Q, let $\mathbf{P}_1, \ldots, \mathbf{P}_n$ be independent random distributions sampled from Q, and let $X_1, \ldots, X_n$ be random variables satisfying

$$Prob(X_1 \in A_1, \ldots, X_n \in A_n | \mathbf{P}_1 = P_1, \ldots, \mathbf{P}_n = P_n) = P_1(A_1) \cdots P_n(A_n) \tag{1}$$

for all sets $A_1, \ldots, A_n$ and distributions $P_1, \ldots, P_n$. Thus, given $\mathbf{P}_1, \ldots, \mathbf{P}_n$, the variables $X_1, \ldots, X_n$ are independent with respective distributions $\mathbf{P}_1, \ldots, \mathbf{P}_n$. By considering product spaces in the usual way, it can be assumed that $X_1, \ldots, X_n$ and $\mathbf{P}_1, \ldots, \mathbf{P}_n$ are all defined on the same underlying probability space. Let $\mathcal{T}_n^G$ denote the set of stop rules τ satisfying $\tau \leq n$ and

$$\{\tau = k\} \in \sigma(\{X_1, \ldots, X_k\}), \qquad k = 1, \ldots, n,$$

and let $\mathcal{T}_n^P$ be the set of stop rules τ satisfying $\tau \leq n$ and

$$\{\tau = k\} \in \sigma(\{X_1, \ldots, X_k, \mathbf{P}_1, \ldots, \mathbf{P}_n\}), \qquad k = 1, \ldots, n.$$

Since the pseudo prophet knows $\mathbf{P}_1, \ldots, \mathbf{P}_n$ ahead of time, he can employ stop rules in $\mathcal{T}_n^P$. So his optimal expected return is

$$V_n^P(Q) := \sup_{\tau \in \mathcal{T}_n^P} EX_\tau.$$

The gambler, however, is restricted to using stop rules in $\mathcal{T}_n^G$. So his value is

$$V_n^G(Q) := \sup_{\tau \in \mathcal{T}_n^G} EX_\tau.$$

Let $D_n(Q) = V_n^P(Q) - V_n^G(Q)$, and $R_n(Q) = V_n^P(Q)/V_n^G(Q)$.

Theorem 2.1 *Let Q be an integrable prior. Then*

(i) $R_n(Q) \leq a_n$; and

(ii) if Q assigns measure one to distributions supported on $[a, b]$, then $D_n(Q) \leq b_n(b - a)$.

Here a_n and b_n are the same implicitly defined constants given in Theorems A and B of Hill and Kertz (1982). Both inequalities are sharp, and (ii) is attained.

The constants a_n and b_n satisfy $1.1 < a_n < 1.6$ and $0 < b_n < 1/4$, and can be easily approximated (see Hill and Kertz (1982) for the details). For example, $a_2 \approx 1.171$, $a_3 \approx 1.221$, $a_4 \approx 1.248$, $a_{10} \approx 1.301$, $a_{10,000} \approx 1.341$, and $b_2 = 1/16$, $b_3 \approx .077$, $b_4 \approx .085$, $b_{10} \approx .100$, $b_{10,000} \approx .111$.

The proof of the theorem uses the following lemma, whose proof is routine. Let $\bar{P}$ denote the average distribution of Q. That is,

$$\bar{P}(A) = \int_{\mathcal{P}} P(A) dQ(P).$$

Lemma 2.2 *For every Borel-measurable function $f : \mathbb{R}^+ \to \mathbb{R}$,*

$$\int_{\mathbb{R}^+} f(x) d\bar{P}(x) = \int_{\mathcal{P}} \int_{\mathbb{R}^+} f(x) dP(x) dQ(P).$$

Proof of Theorem 2.1. Let δ_x denote Dirac measure at x. Define the Borel mapping $\eta : x \to \delta_x$ from $\mathbb{R}^+$ to $\mathcal{P}$, and let Q^* be the prior on $\mathcal{P}$ defined by $Q^*(B) = \bar{P}(\eta^{-1}(B))$, for each Borel set B of $\mathcal{P}$. Let $Y_1, \ldots, Y_n$ be a sequence of i.i.d. random variables with common distribution $\bar{P}$, let $\mathcal{S}_n$ be the set of all stop rules for $Y_1, \ldots, Y_n$, and define $V(Y_1, \ldots, Y_n) = \sup_{\tau \in \mathcal{S}_n} EY_\tau$. From (1) it follows immediately that $X_1, \ldots, X_n$ are unconditionally i.i.d. with common distribution $\bar{P}$, and since the gambler observes only the values $X_1, \ldots, X_n$, it follows that

$$V_n^G(Q) = V(Y_1, \ldots, Y_n).$$

Since $\int P(A) dQ^*(P) = \int \delta_y(A) d\bar{P}(y) = \bar{P}(A)$, it follows likewise that

$$V_n^G(Q^*) = V(Y_1, \ldots, Y_n) = V_n^G(Q).$$

It will now be shown that

$$V_n^P(Q^*) \geq V_n^P(Q).$$

Define

$$V_n^P(P_1, \ldots, P_n) = \sup_{\tau \in \mathcal{T}_n^P} E[X_\tau | \mathbf{P}_1 = P_1, \ldots, \mathbf{P}_n = P_n].$$

Note that for fixed $P_1, \ldots, P_n$,

$$\begin{aligned} V_n^P(P_1, \ldots, P_n) &\leq E(X_1 \vee \cdots \vee X_n | \mathbf{P}_1 = P_1, \ldots, \mathbf{P}_n = P_n) \\ &= \int_{\mathbb{R}^+} \cdots \int_{\mathbb{R}^+} x_1 \vee \cdots \vee x_n \; dP_1(x_1) \ldots dP_n(x_n). \end{aligned}$$

Thus

$$\begin{aligned} V_n^P(Q) &= \int_{\mathcal{P}} \cdots \int_{\mathcal{P}} V_n^P(P_1, \ldots, P_n) \; dQ(P_1) \ldots dQ(P_n) \\ &\leq \int_{\mathcal{P}} \cdots \int_{\mathcal{P}} \left(\int_{\mathbb{R}^+} \cdots \int_{\mathbb{R}^+} x_1 \vee \cdots \vee x_n \; dP_1(x_1) \ldots dP_n(x_n) \right) \; dQ(P_1) \ldots dQ(P_n) \\ &= \int_{\mathbb{R}^+} \cdots \int_{\mathbb{R}^+} x_1 \vee \cdots \vee x_n \; d\bar{P}(x_1) \ldots d\bar{P}(x_n) \\ &= E(Y_1 \vee \cdots \vee Y_n) \\ &= V_n^P(Q^*). \end{aligned}$$

The second equality follows by a repeated application of Lemma 2.2. The last equality follows since under Q^*, the pseudo-prophet has complete foresight, and the unconditional distribution of each X_i is $\bar{P}$. It follows that

$$D_n(Q) \le V_n^P(Q^*) - V_n^G(Q^*) = E(Y_1 \vee \cdots \vee Y_n) - V(Y_1, \ldots, Y_n),$$

and

$$R_n(Q) \le \frac{V_n^P(Q^*)}{V_n^G(Q^*)} = \frac{E(Y_1 \vee \cdots \vee Y_n)}{V(Y_1, \ldots, Y_n)}.$$

Applying Theorems A and B of Hill and Kertz (1982) completes the proof. □

Remark 2.3 Although the worst-case prior Q^* in the proof of the theorem is supported on the set of Dirac measures, the bounds (i) and (ii) remain sharp if the support of Q is required to be all of $\mathcal{P}_0$. To see this for (i), fix $\varepsilon > 0$, let Q be a prior with $R_n(Q) \ge a_n - \varepsilon/2$, and let $\hat{Q}$ be any prior with full support on $\mathcal{P}_0$. (Such priors were described, for instance, by Dubins and Freedman (1967), Ferguson (1973), Mauldin, Sudderth and Williams (1992), and Hill and Monticino (1998)). For $0 \le t \le 1$, define

$$Q_t := t\hat{Q} + (1-t)Q.$$

Then Q_t has full support on $\mathcal{P}_0$ for each $t > 0$. It is straightforward to prove that $V_n^G(Q_t)$ and $V_n^P(Q_t)$ are continuous as functions of t. Hence, for $t > 0$ sufficiently small, $R_n(Q_t) \ge R_n(Q) - \varepsilon/2 \ge a_n - \varepsilon$.

Similarly, the inequalities (i) and (ii) remain sharp if Q gives measure one to distributions with full support in $\mathbb{R}^+$, or to absolutely continuous distributions.

Remark 2.4 If $n = 2$ and $\bar{P}$ is the uniform distribution on $[0, 1]$, a simple expression for $D_2(Q)$ in terms of the variance of the mean can be derived. For a distribution P, let $\mu_P = \int x dP$ denote the mean of P. Note that for $0 \le c \le 1$

$$\int_{[0,1]} x \vee c \, d\bar{P}(x) = \int_0^1 x \vee c \, dx = (1 + c^2)/2.$$

Hence $V_2^G(Q) = \int_{[0,1]} x \vee (1/2) \, d\bar{P}(x) = 5/8$. Similarly,

$$\begin{aligned}
V_2^P(Q) &= \int_{\mathcal{P}} \int_{\mathcal{P}} V_2^P(P_1, P_2) \, dQ(P_1) dQ(P_2) \\
&= \int_{\mathcal{P}} \int_{\mathcal{P}} \left(\int_{[0,1]} (x \vee \mu_{P_2}) \, dP_1(x) \right) \, dQ(P_1) dQ(P_2) \\
&= \int_{\mathcal{P}} \left(\int_{[0,1]} (x \vee \mu_{P_2}) \, d\bar{P}(x) \right) dQ(P_2) \\
&= \int_{\mathcal{P}} \frac{1 + \mu_P^2}{2} \, dQ(P) = \frac{1 + E\mu_{\mathbf{P}}^2}{2}.
\end{aligned}$$

Thus

$$D_2(Q) = \frac{Var(\mu_{\mathbf{P}})}{2}. \tag{2}$$

Using (2), $D_2(Q)$ may be computed explicitly for several well-known priors whose average distribution is uniform on $[0,1]$. For instance, if Q is the random-rescaling prior introduced by Dubins and Freedman (1967) with base measure equal to the uniform distribution on the vertical line segment $x = 1/2$, $0 \leq y \leq 1$, then $Var(\mu_{\mathbf{P}}) = 1/40$ (cf. Mauldin and Williams, 1990), so $D_2(Q) = 1/80$. Alternatively, if Q is a Dirichlet process prior with base measure $\alpha([0,x]) = cx$ on $[0,1]$ (see Ferguson, 1973), then it can be shown that $Var(\mu_{\mathbf{P}}) = [12(c+1)]^{-1}$, so $D_2(Q) = [24(c+1)]^{-1}$.

Remark 2.5 The authors do not know sharp pseudo-prophet inequalities for priors Q restricted to distributions with a given (fixed) mean μ, or priors satisfying other kinds of partial information.

References

Dubins, L.E., Freedman, D.A., 1967. Random distribution functions. Proc. Fifth Berkeley Symposium Math. Statist. Probl. 2, 183-214.

Ferguson, T., 1973. A Bayesian analysis of some nonparametric problems. Ann. Statist. 1, 209-230.

Hill, T.P., Kertz, R.P., 1982. Comparisons of stop rule and supremum expectations of i.i.d. random variables. Ann. Probab. 10, no.2, 336-345.

Hill, T.P., Kertz, R.P., 1992. A survey of prophet inequalities in optimal stopping theory. Contemp. Math. 125, 191-207.

Hill, T.P., Monticino, M.G., 1998. Constructions of random distributions via sequential barycenter arrays. Ann. Statist. 26, no. 4, 1242-1253.

Mauldin, R.D., Sudderth, W.D., Williams, S.C., 1992. Polya trees and random distributions. Ann. Statist. 20, no. 3, 1203-1221.

Mauldin, R.D., Williams, S.C., 1990. Reinforced random walks and random distributions. Proc. Amer. Math. Soc. 110, no. 1, 251-258.

Samuels, S., 1981. Minimax stopping rules when the underlying distribution is uniform. J. Amer. Stat. Assoc. 76, 188-197.