A Statistical Version of Prophet Inequalities[1]

David Assaf[2], Larry Goldstein[3] and Ester Samuel-Cahn[4]
April 30, 2001

Abstract

All classical "prophet inequalities" for independent random variables hold also in the case where only a noise corrupted version of those variables is observable. That is, if the pairs $(X_1, Z_1), \ldots, (X_n, Z_n)$ are independent with arbitrary, known joint distributions, and only the sequence $Z_1, \ldots, Z_n$ is observable, then all prophet inequalities which would hold if the X's were directly observable still hold, even though the expected X-values (i.e. the payoffs) for both the prophet and statistician, will be different. Our model includes, for example, the case when $Z_i = X_i + Y_i$, where the Y's are any sequence of independent random variables.

[1]This research was supported by Grant no. 94-00186 from the United States - Israel Binational Science Foundation (BSF) Jerusalem, Israel.
[2]Hebrew University, Department of Statistics
[3]University of Southern California, Department of Mathematics
[4]Hebrew University, Department of Statistics, and Center for Rationality and Interactive Decision Theory

1 Introduction and Summary

Let $\mathbf{X}_n = (X_1, \ldots, X_n)$ be a random vector having a *known* joint distribution, with $E|X_i| < \infty$ for all $i = 1, \ldots, n$. The setting of classical prophet inequalities is where the X sequence is observed sequentially, and the objective is to pick an X-value which is as large as possible. There, the quantity

$$V_p^0(\mathbf{X}_n) = E(\max_{1 \le i \le n} X_i),$$

denotes the value for a prophet who has foreknowledge of the entire X-sequence, and will thus select the largest value. A statistician, on the other hand, is limited to stopping rules $t \in T_X^n$, where $t \in T_X^n$ if and only if the event $\{t = i\}$ belongs to the σ-field generated by $X_1, \ldots, X_i$, and $P\{t \le n\} = 1$. Thus, there the statistician's value is

$$V_s^0(\mathbf{X}_n) = \sup_{t \in T_X^n} EX_t.$$

The restriction of the statistician to rules in T_X^n reflects the fact that the statistician's decision to stop at any given time i can depend only on the past and present observations, and not on ones in the future. This corresponds to a situation where no recall is allowed. Obviously $V_p^0(\mathbf{X}_n) \ge V_s^0(\mathbf{X}_n)$.

Ratio and difference prophet inequalities provide upper bounds on $V_p^0(\mathbf{X}_n)/V_s^0(\mathbf{X}_n)$ and $V_p^0(\mathbf{X}_n) - V_s^0(\mathbf{X}_n)$ over large classes of random variables. An excellent review of most of the earlier prophet inequalities is given in Hill and Kertz (1992). The first and probably the best known prophet inequality is the one in Krengel and Sucheston (1978) which states that for all $n \ge 2$ and all nonnegative independent $X_1, \ldots, X_n$ one has $V_p^0(\mathbf{X}_n)/V_s^0(\mathbf{X}_n) < 2$, and this bound cannot be improved. (See also Hill and Kertz, 1981a). For i.i.d. nonnegative X's the best bound a_n, on the above ratio, depends on n and it is generally believed that $\lim_{n\to\infty} a_n = 1.34\ldots$. (See Hill and Kertz, 1982).

For independent bounded random variables taking values in $[a, b]$ the difference prophet inequality $V_p^0(\mathbf{X}_n) - V_s^0(\mathbf{X}_n) \le (b-a)/4$ holds. (See Hill and Kertz, 1981b). Prophet inequalities with cost for observations (Jones, 1990) or with discounting (Boshuizen, 1991) have also been considered.

In the present paper an often more realistic model is considered. Though still both the statistician and the prophet are interested in as large an X-value as possible, the X's are not directly observable. The observed random variables are Z's, which may be thought of as the X's corrupted by "noise." That is, we take the general view that the pairs (X_i, Z_i) $i = 1, 2, \ldots, n$ are independent and have an arbitrary known joint distribution. An important special case of this model is the "additive noise" model where $Z_i = X_i + Y_i$, with the Y's mutually independent, identically distributed, and independent of the X's.

As a practical example, consider the following situation. Your firm wants to hire a typist. It receives n responses to an advertisement. The true score for typist i is X_i, perhaps some combined score depending on speed, accuracy and aesthetics of typed text. In this example it may be realistic to assume that the X's are i.i.d. The typists are interviewed sequentially, (e.g. one a day) and are given some sample task. The score of the i^{th} typist on this task is Z_i, probably X_i with some positive bias. Since the long-term performance for the i^{th} typist will be X_i, the X's are the variables of interest which should be maximized. Here the statistician must decide whether to hire the given candidate immediately after the interview, due possibly, to the non-availability of the candidate at a later stage, whereas the prophet may base his decision on an interview of all n candidates at no extra cost, and with no danger of non-availability.

In general, as only the sequence $\mathbf{Z}_n$ is observable, both the prophet and the statistician's choices must depend on this sequence only. In particular, the statistician is confined to the set of stopping rules T_Z^n defined similarly to T_X^n, but with rules based on the sequence $\mathbf{Z}_n$. Therefore, in this situation, the value for the statistician is defined as $V_s(\mathbf{X}_n, \mathbf{Z}_n) = \sup_{t \in T_Z^n} EX_t$. The value $V_p(\mathbf{X}_n, \mathbf{Z}_n)$ for the prophet in this framework will be defined in a similar manner in the next section. It should be noted that the classical prophet model is the one where $\mathbf{X}_n = \mathbf{Z}_n$, and is in particular a special case of the additive model with the noise Y identically zero. Thus the present model contains the previously described classical model.From this observation it follows that any bound for the noise corrupted case cannot be more stringent than those described earlier.

Our main result is that under very mild conditions all previously mentioned prophet inequalities remain valid in the most general "statistical setup". In particular, if $X_i \geq 0$, and the pairs $(X_1, Z_1), \ldots, (X_n, Z_n)$ are independent, with each (X_i, Z_i) having arbitrary joint distribution, the ratio $V_p(\mathbf{X}_n, \mathbf{Z}_n)/V_s(\mathbf{X}_n, \mathbf{Z}_n)$ is again bounded by 2, and if in addition these pairs are identically distributed the bound is the same a_n as above. Similarly if the pairs are independent and X_i takes values in $[a, b]$, the corresponding difference prophet inequality is again $(b - a)/4$. These results are obtained by showing that the noise corrupted case can be reduced to a "noise free" case for the independent random variables W_i, where W_i is the conditional expectation of X_i given Z_i.

In the noise corrupted case, one may consider other types of prophets. For example, a "perfect prophet" could be defined as one who observes the entire uncorrupted X-sequence and bases his decision on it, as in the classical setting. In Remark 2.8, we show that a perfect prophet may have an unbounded relative advantage over the statistician.

2 Main Results

Let $(X_1, Z_1), \ldots, (X_n, Z_n)$ be any sequence of pairs of random variables with known distribution and $E|X_i| < \infty$ for $i = 1, \ldots, n$. We begin with a precise definition of the value for the prophet in the "statistical setup" with only $\mathbf{Z}_n = (Z_1, \ldots, Z_n)$ observable, and the selection rule therefore based only on the Z's. The set of (nonrandomized) selection rules for the prophet is the set G_n of functions

$$G_n = \{g; g : R^n \to \{1, \ldots, n\}\},$$

with the interpretation that upon observing $Z_1, \ldots, Z_n$, the prophet selects the index $g(Z_1, \ldots, Z_n)$. The value for such a g is $EX_{g(\mathbf{Z}_n)}$. The value for the prophet is thus defined as

$$V_p(\mathbf{X}_n, \mathbf{Z}_n) = \sup_{g \in G_n} EX_{g(\mathbf{Z}_n)}.$$

Let $W_i^* = E(X_i|\mathbf{Z}_n)$ and $W_i = E(X_i|Z_i)$. The following proposition gives the full solution for the prophet.

Proposition 2.1 *Any rule* $g^*(\mathbf{Z}_n) = \text{argmax}\{W_i^* : 1 \le i \le n\}$ *is optimal in* G_n, *and*

$$V_p(\mathbf{X}_n, \mathbf{Z}_n) = EX_{g^*(\mathbf{Z}_n)} = E(\max_{1\le i\le n} W_i^*). \tag{1}$$

Proof. For any $g \in G_n$

$$E(X_{g(\mathbf{Z}_n)}|\mathbf{Z}_n) \le \max_{1\le i\le n} E(X_i|\mathbf{Z}_n) = \max_{1\le i\le n} W_i^* = E(X_{g^*(\mathbf{Z}_n)}|\mathbf{Z}_n).$$

Hence,

$$EX_{g(\mathbf{Z}_n)} \le E(\max_{1\le i\le n} W_i^*) = EX_{g^*(\mathbf{Z}_n)},$$

and (1) follows upon taking sup over all $g \in G_n$.□.

In the remainder of the present section we shall assume that the pairs $(X_1, Z_1), \ldots, (X_n, Z_n)$ are independent. In this case $W_i = W_i^*$, and hence Proposition 2.1 yields

Proposition 2.2 *Under independence,*

$$V_p(\mathbf{X}_n, \mathbf{Z}_n) = E(\max_{1\le i\le n} W_i) = V_p^0(\mathbf{W}_n).$$

We next consider the value for the statistician, who is limited to rules in T_Z^n.

Proposition 2.3 *For any* $t \in T_Z^n$

$$EW_t = EX_t.$$

Proof. Note that the random variable W_i is a function of Z_i, i.e., it is Z_i measurable. Also, by definition, the event $\{t = i\}$ is measurable with respect to $(Z_1, \ldots, Z_i)$. Therefore,

$$\begin{aligned} EX_t &= E\{E(X_t|\mathbf{Z}_n)\} \\ &= E\{E(\sum_{i=1}^n X_i \mathbf{1}(t=i)|\mathbf{Z}_n)\} \\ &= E\{\sum_{i=1}^n E(X_i\mathbf{1}(t=i)|\mathbf{Z}_n)\} \\ &= E\{\sum_{i=1}^n \mathbf{1}(t=i)E(X_i|\mathbf{Z}_n)\} \\ &= E\{\sum_{i=1}^n \mathbf{1}(t=i)E(X_i|Z_i)\} \\ &= E\{\sum_{i=1}^n \mathbf{1}(t=i)W_i)\} = EW_t. \end{aligned}$$

Independence is used in the fifth equality. □

(Clearly the above equality does *not* hold for $t \in T_X^n$.)

Remark 2.1 *Similarly, using* $X_{g(\mathbf{Z}_n)} = \sum X_i \mathbf{1}(g(\mathbf{Z}_n) = i)$ *and therefore* $E[X_{g(\mathbf{Z}_n)}|\mathbf{Z}_n] = \sum W_i^* \mathbf{1}(g(\mathbf{Z}_n) = i) = W_{g(\mathbf{Z}_n)}^*$, *it follows that for any* $g \in G_n$ $EX_{g(\mathbf{Z}_n)} = EW_{g(\mathbf{Z}_n)}^*$. *In the case of independence,* $W_{g(\mathbf{Z}_n)}^*$ *can be replaced by* $W_{g(\mathbf{Z}_n)}$.

Note that since W_i is a function of Z_i we have $T_W^n \subset T_Z^n$, hence

$$\sup_{t \in T_W^n} EW_t \leq \sup_{t \in T_Z^n} EW_t. \tag{2}$$

Proposition 2.5 shows that equality holds in (2). In order to prove this assertion we need to consider, as an intermediate step, randomized stopping rules for the W-sequence. (See Chow, Robbins and Siegmund, 1971, p. 111 for an exact definition).

Let $\mathbf{W}_i = (W_1, \ldots, W_i)$. Essentially a randomized stopping rule $\tilde{t}$ specifies the conditional probability $q_i(\mathbf{W}_i)$ of stopping at the i^{th} observation when $\mathbf{W}_i$ is observed, conditional on not having stopped earlier. Thus $q_i : \mathbf{W}_i \to [0,1]$ with $q_n(\mathbf{W}_n) \equiv 1$. Denote the unconditional probability of stopping at time i, when $\mathbf{W}_i$ is observed, by $p_i(\mathbf{W}_i)$. Then clearly

$$p_i(\mathbf{W}_i) = q_i(\mathbf{W}_i) \prod_{j=1}^{i-1} (1 - q_j(\mathbf{W}_j)), \qquad i = 1, \ldots, n.$$

Though it may be more natural to define a randomized stopping rule through the sequence of q-functions, there is clearly a one to one correspondence between the q and p sequences. The value $EW_{\tilde{t}}$ may be evaluated by $EW_{\tilde{t}} = E\{\sum_{i=1}^{n} W_i p_i(\mathbf{W}_i)\}$ (in the particular case where q_i (and p_i) are indicator functions, $\tilde{t}$ is nonrandomized).

Let $\tilde{T}_W^n$ denote the set of all randomized stopping rules for the W-sequence.

Proposition 2.4 *For every $t \in T_Z^n$ there exists a $\tilde{t} \in \tilde{T}_W^n$ such that $EW_t = EW_{\tilde{t}}$.*

Proof. Let $t = t_Z \in T_Z^n$ be given. Define $p_i(\mathbf{W}_i) = P(t_Z = i | \mathbf{W}_i)$. The sequence $p_i(\mathbf{W}_i)$ generates a randomized stopping rule $\tilde{t} \in \tilde{T}_W^n$. (Note that $P(\tilde{t} \leq n) = 1$, since t_Z is a stopping rule). Now

$$\begin{aligned} EW_{t_Z} &= E\{E(W_{t_Z} | \mathbf{W}_n)\} \\ &= E\{E(\sum_{i=1}^{n} W_i \mathbf{1}(t_Z = i) | \mathbf{W}_n)\} \\ &= E\{\sum_{i=1}^{n} E(W_i \mathbf{1}(t_Z = i) | \mathbf{W}_n)\} \\ &= E\{\sum_{i=1}^{n} W_i P(t_Z = i | \mathbf{W}_n)\} \\ &= E\{\sum_{i=1}^{n} W_i P(t_Z = i | \mathbf{W}_i)\} \\ &= E\{\sum_{i=1}^{n} W_i p_i(\mathbf{W}_i)\} = EW_{\tilde{t}}. \end{aligned}$$

where the fifth equality uses the fact that $\{t_Z = i\}$ is measurable with respect to $Z_1, \ldots, Z_i$. □

Proposition 2.5

$$V_s(\mathbf{X}_n, \mathbf{Z}_n) = V_s^0(\mathbf{W}_n)$$

Proof

$$\begin{aligned} V_s^0(\mathbf{W}_n) := \sup_{t \in T_W^n} EW_t = \sup_{\tilde{t} \in \tilde{T}_W^n} EW_{\tilde{t}} &= \sup_{t \in T_Z^n} EW_t \\ = \sup_{t \in T_Z^n} EX_t &:= V_s(\mathbf{X}_n, \mathbf{Z}_n). \end{aligned}$$

The first and last equalities are definitions, the second follows since the optimal rule can always be taken to be non-randomized (see Chow, Robbins and Siegmund, 1971), the third uses Propositions 2.4 and inequality (2), and the fourth equality follows from Proposition 2.3. $\square$.

Propositions 2.5 and 2.2 yield the following theorem

Theorem 2.1 *Let $(X_1, Z_1), \ldots, (X_n, Z_n)$ be independent pairs of random variables, with only $Z_1, \ldots, Z_n$ observable. Suppose $E|X_i| < \infty$, and set $W_i = E(X_i|Z_i)$. Then any prophet inequality (ratio or difference) which holds for the independent W's is valid also for $V_p(\mathbf{X}_n, \mathbf{Z}_n)$ and $V_s(\mathbf{X}_n, \mathbf{Z}_n)$. That is, if $A \subset \mathbf{R}^2$ is such that $(V_p^0(\mathbf{W}_n), V_s^0(\mathbf{W}_n)) \in A$, then $(V_p(\mathbf{X}_n, \mathbf{Z}_n), V_s(\mathbf{X}_n, \mathbf{Z}_n)) \in A$. In the case of an infinite sequence $(X_1, Z_1), (X_2, Z_2), \ldots$ corresponding statements hold, provided $E(\sup_{1 \leq i} X_i) < \infty$.*

Corollary 2.1 *1. If $W_i \geq 0$ then $V_p(\mathbf{X}_n, \mathbf{Z}_n)/V_s(\mathbf{X}_n, \mathbf{Z}_n) < 2$.*

2. If $W_i \geq 0$ are i.i.d. then $V_p(\mathbf{X}_n, \mathbf{Z}_n)/V_s(\mathbf{X}_n, \mathbf{Z}_n) \leq a_n$, where a_n is the bound of Hill and Kertz (1981b) for the i.i.d. case.

3. If $a \leq W_i \leq b$ for $i = 1, \ldots, n$ then $V_p(\mathbf{X}_n, \mathbf{Z}_n) - V_s(\mathbf{X}_n, \mathbf{Z}_n) \leq (b - a)/4$.
All three bounds are the best possible, and corresponding results are valid for the infinite case.

Remark 2.2 *Note that a sufficient condition for (1) in Corollary 2.1, i.e. for $W_i \geq 0$, is that $X_i \geq 0$, though this is clearly not necessary. Thus Z_i may take on negative values, as in the additive noise case where, say, $Y_i \sim \mathcal{N}(0, \sigma^2)$. A sufficient condition for (2) is that $X_i \geq 0$ and that (X_i, Z_i) be i.i.d. pairs. A sufficient condition for (3) is $a \leq X_i \leq b$.*

Remark 2.3 *Consider the case where there is a cost for sampling, i.e. $X_i = \hat{X}_i - c_i$ where the $\hat{X}_i$ are independent and c_i denotes the cost of sampling i units; usually $c_i = ci$ for some $c > 0$. Here $W_i = \hat{W}_i - c_i$ where $\hat{W}_i = E(\hat{X}_i|Z_i)$. Thus the "cost of sampling" structure carries over to the W_i and the corresponding prophet inequalities carry over as well. Similar results hold for the discounting structure where $X_i = \beta^{i-1}\hat{X}_i$, $\hat{X}_i \geq 0$ and $0 < \beta < 1$. Then $W_i = \beta^{i-1}\hat{W}_i$. i.e. the discounting structure carries over, as do the corresponding inequalities.*

Remark 2.4 *In Samuel-Cahn (1984) it is shown that the bound 2 for independent $X_i \geq 0$ corresponding to (1) in Corollary 2.1 holds even if the statistician is limited to "simple threshold rules" of the form $\inf\{i : X_i \geq b\} \wedge n$ for some constant b, instead of using an optimal rule. Thus in the present setting the above statement carries over for "simple threshold rules" applied to the W's.*

Remark 2.5 *Nowhere have we used the one-dimensional structure of the Z_i's. They could be multidimensional, or even just independent σ-fields, for all $i = 1, \ldots, n$.*

Remark 2.6 *All results carry over had we let the statistician resort to the randomized stopping rules $\tilde{T}_Z^n$ instead of T_Z^n. (A similar remark holds for the prophet).*

Remark 2.7 *It is not difficult to construct examples in which the noise does not vanish asymptotically but the corresponding bound is still sharp. A simple example of this type is an extension of the classical example for the independent case. For* $0 \le \mu \le 1$ *let* $X_1 = \ldots = X_{n-2} = 0, X_{n-1} = \mu$ *and* X_n *equal 1 with probability* μ *and 0 with probability* $1-\mu$. *Let* Y_i *be independent uniform on* $(-\frac{1}{2}, \frac{1}{2})$, *and consider the additive noise model. Then* $V_s(\mathbf{X}_n, \mathbf{Z}_n) = V_s^0(\mathbf{X}_n) = \mu$ *and* $V_p(\mathbf{X}_n, \mathbf{Z}_n) = V_p^0(\mathbf{X}_n) = \mu(2-\mu)$, *hence the ratio is* $2-\mu$, *and the bound of 2 is again sharp in the noise corrupted case as well.*

Remark 2.8 *Lastly, we consider the ratio prophet inequality in the case of a "perfect prophet" who is able to base his decision on the entire* X *sequence. In this case, the ratio of interest is* $V_p^0(\mathbf{X}_n)/V_s(\mathbf{X}_n, \mathbf{Z}_n)$. *The following example shows that the ratio may tend to infinity with* n.

Let the X*'s be independent Bernoulli* p *random variables. If any of these variables is equal to 1, the prophet will choose it; hence, his value equals* $1-(1-p)^n$. *For our example, we assume the statistician can only see the* Z *sequence where each of the ones of the original sequence have been flipped to zeros, and each of the zeros to ones, with probability* θ. *That is, let* $B_1, \ldots, B_n$ *be independent Bernoulli* $1-\theta$ *variables, and* $Z_i = B_i X_i + (1-B_i)(1-X_i)$. *Then, for* $\theta = 0$, *the* Z *sequence equals the* X *sequence, and the ratio* $V_p^0(\mathbf{X}_n)/V_s(\mathbf{X}_n, \mathbf{Z}_n)$ *is easily seen to be one. However, when* $\theta = 1/2$, *the* Z*'s and* X*'s are independent and so the statistician can only get the value* p; *thus the ratio tends to* n *as* $p \to 0$. *Intermediate cases may be obtained when* $0 < \theta < 1/2$.

Note that n *is the "best bound" on* $V_p^0(\mathbf{X}_n)/V_s(\mathbf{X}_n, \mathbf{Z}_n)$ *for all sequences of non-negative* X*'s and any* Z*'s, with any type of dependence, in the present "perfect prophet" setup. This follows since* $V_s(\mathbf{X}_n, \mathbf{Z}_n) \ge \max_{i=1,\ldots,n} EX_i$, *because the statistician can always obtain the value* $\max_{i=1,\ldots,n} EX_i$ *by choosing the (nonrandom) index which maximizes* EX_i, *whereas for the prophet*

$$V_p^0(\mathbf{X}) = E(\max_{i=1,\ldots,n} X_i) \le E(\sum_{i=1}^{n} X_i) \le n(\max_{i=1,\ldots,n} EX_i).$$

(Compare Hill and Kertz, 1981a, proposition 1).

References

Boshuizen, F. (1991). Prophet region for independent random variables with a discount factor. *J. Multivariate Anal.* **37**, pp. 76-84.

Chow, Y.S., Robbins, H. and Siegmund, D. (1971) *Great Expectations: The Theory of Optimal Stopping.* Houghton Mifflin, Boston.

Hill, T.P. and Kertz, R.P. (1981a) Ratio comparisons of supremum and stop rule expectations. *Z. Wahrschein verw. Gebiete.* **56**, pp. 283-285.

Hill, T.P. and Kertz, R.P. (1981b). Additive comparisons of stop rule and supremum expectations of uniformly bounded independent random variables. *Proc. Amer. Math. Soc.* **83**, pp. 582-585.

Hill, T.P. and Kertz, R.P. (1982) Comparisons of stop rules and supremum expectation for i.i.d random variables, *Ann. Probab.* **10**, pp. 336-345.

Hill, T.P. and Kertz, R.P, (1992). A survey of prophet inequalities in optimal stopping theory. *Contemporary Mathematics*, AMS, Vol. 125, pp. 191-207.

Jones, M., (1990). Prophet inequalities and cost of observation stopping problems. *J. Multivar. Anal.* **34**, pp. 238-253.

Krengel, U. and Sucheston, L. (1978). On semimarts, amarts and processes with finite value. *Prob. on Banach Spaces* (ed. J. Kuelbs), M. Dekker, NY, pp. 197-266.

Samuel-Cahn, E. (1984) Comparison of threshold stop rules and maximum for independent nonnegative random variables, *Ann. Probab.* **12**, pp. 1213-1216.