Bounded Best Buyers: Prophet Inequalities Revisited

Farbod Ekbatani, Rad Niazadeh, Pranav Nuti, Jan Vondrák

Abstract

We consider a seller with a single item to allocate, and a sequence of buyers arriving in an online fashion, where buyer i has independently drawn valuation v_i. The seller seeks to maximize the valuation of the buyer who receives the item, and the expected valuation is compared against the prophet benchmark $\mathbb{E}[\max_i v_i]$. Suppose the realized maximum valuation is always contained in a bounded multiplicative window: for some $\alpha > 0$ and $f > 0$,

$$\alpha \leq \max_i v_i \leq (1+f)\alpha.$$

One natural example where this might occur is if there is an initial deterministic buyer with valuation α, and all other buyers have valuations at most $(1+f)\alpha$. The posted price

$$T = \max\left\{\alpha, \frac{1+f}{1+2f}\,\mathbb{E}\left[\max_i v_i\right]\right\}$$

and allocation to the first buyer with $v_i \geq T$ obtain expected valuation at least

$$\frac{1+f}{1+2f}\,\mathbb{E}\left[\max_i v_i\right].$$

The same guarantee is also obtained by two other prices: a price T with $\mathbb{P}[\max_i v_i < T] = f^2/((1+f)(1+2f))$, and the price $T = \max\{T^*, \alpha\}$ where $\mathbb{E}[(\max_i v_i - T^*)_+] = \frac{f}{1+f}T^*$. The factor is tight. As $f \to \infty$, the different prices and the guarantee converge to results known for the classical 1/2 prophet inequality; as $f \to 0$, the assumption forces the maximum valuation to be known exactly and the guarantee converges to 1.

1 Introduction

Consider a seller tasked with selling a single item to a sequence of arriving buyers. Buyer i has valuation v_i, drawn independently from a known distribution F_i. Upon seeing v_i, the seller may choose to allocate the item to buyer i, or continue to the next buyer. We evaluate the seller by the valuation of the buyer who receives the item. The 'prophet' benchmark is the highest realized valuation

$$v_{\max} := \max_i v_i.$$

The classical prophet inequality says that the seller has an algorithm that obtains at least half of the expectation of the prophet:

$$\mathbb{E}[\text{ALG}] \geq \frac{1}{2}\mathbb{E}[v_{\max}].$$

This 1/2 guarantee originates in the work of [5], and is known to be tight. One particularly simple class of algorithms involves the seller posting a take-it-or-leave-it price T and selling to the first buyer whose valuation is at least T. Samuel-Cahn showed that such a policy can also obtain the 1/2 guarantee [7].

In this note, we identify a simple refinement when the seller has information about the scale of the best buyer. Assume that the realized maximum valuation is always in a known multiplicative range:

$$\alpha \le v_{\max} \le (1+f)\alpha.$$

Importantly, note that this is not the assumption that every buyer's valuation lies in $[\alpha, (1+f)\alpha]$. A natural setting where our assumption holds is if there is an initial deterministic buyer with valuation α, and all other buyers have valuations at most $(1+f)\alpha$.

Our main observation is that with this information about the maximum valuation, the seller can post a price that obtains at least

$$\frac{1+f}{1+2f}\,\mathbb{E}[v_{\max}],$$

in expectation, and no online allocation rule can guarantee a better factor under the same bounded-maximum assumption. This guarantee interpolates smoothly between the classical 1/2 result when the upper and lower bounds are very far apart, and 1 when the maximum valuation is known exactly.

The impossibility of beating this factor is already shown by a simple, classic example with two buyers. Let $v_1 = \alpha$ deterministically, and let

$$v_2 = \begin{cases} (1+f)\alpha, & \text{with probability } 1/(1+f), \\ 0, & \text{otherwise.} \end{cases}$$

Then $\alpha \le v_{\max} \le (1+f)\alpha$ always. A seller who sees the first buyer can get expected valuation at most α: accepting gives α, while rejecting also gives in expectation $(1+f)\alpha \cdot 1/(1+f) = \alpha$. On the other hand, the prophet obtains

$$\mathbb{E}[v_{\max}] = (1+f)\alpha \cdot \frac{1}{1+f} + \alpha \cdot \frac{f}{1+f} = \frac{1+2f}{1+f}\alpha,$$

so the best possible ratio is exactly $(1+f)/(1+2f)$.

2 Related work

The bounded-maximum assumption studied here is different than the assumption that every buyer valuation is supported on a common interval. For independent valuations supported on a common interval $[a, b]$, [4] proved the sharp additive comparison

$$\mathbb{E}[v_{\max}] - \mathbb{E}[\mathrm{ALG}] < \frac{1}{4}(b-a),$$

with the constant 1/4 best possible. For valuations normalized to $0 \le v_i \le 1$, [8] proved the comparison

$$\mathbb{E}[v_{\max}] \le 2\mathbb{E}[\mathrm{ALG}] - \mathbb{E}[\mathrm{ALG}]^2.$$

After affine normalization, this gives a guarantee that is stronger than ours when every valuation lies in a common range. However, we require only $\alpha \le v_{\max} \le (1+f)\alpha$, which is strictly weaker because individual valuations may be below α. The two-buyer example above shows that the factor $(1+f)/(1+2f)$ is the best possible under this weaker information.

This note is an exposition of the argument of [3], which studies prophet inequalities with cancellation costs. That paper studies a model in which the seller may revoke an earlier allocation by paying a cancellation cost proportional to f and the value of the earlier allocation. Note that in the bounded-maximum case, however, once the seller has accepted a valuation of at least α, no later valuation can improve the net payoff enough to justify a revocation when the maximum valuation lies within a factor $1+f$.

Finally, the lower bound α can be viewed as a conservative prediction of the maximum valuation. Recent work [1] studies prophet inequalities with conservative predictions of the maximum realized valuation. They obtain the same guarantee as the current note but under the weaker assumption that α is not necessarily deterministic. However, their proof is substantially more complicated.

See [6] and [2] for a more thorough survey of results.

3 The bounded-maximum theorem

Theorem 1. *Let $v_1, \ldots, v_n$ be independent nonnegative buyer valuations, revealed sequentially, and suppose that for known $\alpha > 0$ and $f > 0$,*

$$\alpha \le v_{\max} \le (1+f)\alpha.$$

Setting any of the prices

$$T = \max\left\{\frac{1+f}{1+2f}\mathbb{E}[v_{\max}], \alpha\right\}, \ \textit{or} \ \Pr[v_{\max} < T] = \frac{f^2}{(1+f)(1+2f)}, \ \textit{or}$$
$$T = \max\{T^*, \alpha\} \ \textit{where } T^* \textit{ is the solution of } \mathbb{E}[(v_{\max} - T^*)_+] = \frac{f}{1+f}T^*$$

obtains expected valuation at least

$$\frac{1+f}{1+2f}\mathbb{E}[v_{\max}].$$

Proof. Let ALG_T denote the valuation obtained by posting price T and allocating the item to the first buyer with valuation at least T. The standard lower bound on the performance of this algorithm is:

$$\mathbb{E}[\mathrm{ALG}_T] \ge T \cdot \Pr[v_{\max} \ge T] + \Pr[v_{\max} < T] \cdot \mathbb{E}[(v_{\max} - T)_+].$$

This is obtained by noting that for any price T, the obtained valuation decomposes as the revenue collected when the item sells plus the utility of the buyer.

The first term is simply the revenue. The second term is the utility that would be obtained if, after the real sequence failed to clear price T, an independent copy of the best buyer were to appear last. The actual best buyer in the original sequence gets at least this much expected utility, because the surplus of the best buyer is counted whenever all earlier buyers failed to clear the price.

Our plan is to strengthen this inequality and to show for our prices that:

$$\mathbb{E}[\mathrm{ALG}_T] \ge T \cdot \Pr[v_{\max} \ge T] + \Pr[v_{\max} < T] \cdot \frac{\mathbb{E}[(v_{\max} - T)_+]}{1 - \mathbb{E}[(v_{\max} - T)_+]/(fT)}. \tag{1}$$

To prove this inequality, start by noting that for any fixed price $T \ge \alpha$, the obtained valuation decomposes as follows:

$$\mathbb{E}[\mathrm{ALG}_T] = T \cdot \Pr[v_{\max} \ge T] + \sum_{i=1}^{n} \mathbb{E}[(v_i - T)_+] \cdot \Pr[v_1 < T, \ldots, v_{i-1} < T], \tag{2}$$

where we have decomposed the utility according to which buyer the item is allocated to. Next, note by independence that

$$\Pr[v_1 < T, \ldots, v_{i-1} < T] = \frac{\Pr[v_{\max} < T]}{\Pr[\max_{j \geq i} v_j < T]}.$$

Using (2), we get

$$\mathbb{E}[\mathrm{ALG}_T] = T \cdot \Pr[v_{\max} \geq T] + \Pr[v_{\max} < T] \cdot \sum_{i=1}^{n} \frac{\mathbb{E}[(v_i - T)_+]}{\Pr[\max_{j \geq i} v_j < T]}. \tag{3}$$

Since $T \geq \alpha$ and $\max_{j \geq i} v_j \leq (1+f)\alpha \leq (1+f)T$,

$$(\max_{j \geq i} v_j - T)_+ \leq fT \cdot \mathbf{1}\{\max_{j \geq i} v_j \geq T\}.$$

Consequently,

$$\Pr[\max_{j \geq i} v_j < T] \leq 1 - \frac{\mathbb{E}[(\max_{j \geq i} v_j - T)_+]}{fT}.$$

Substituting this into (3) gives

$$\mathbb{E}[\mathrm{ALG}_T] \geq T \cdot \Pr[v_{\max} \geq T] + \Pr[v_{\max} < T] \cdot \sum_{i=1}^{n} \frac{\mathbb{E}[(v_i - T)_+]}{1 - \frac{\mathbb{E}[(\max_{j \geq i} v_j - T)_+]}{fT}}.$$

In Lemma 1, proved below, we will show that

$$\sum_{i=1}^{n} \frac{\mathbb{E}[(v_i - T)_+]}{1 - \mathbb{E}[(\max_{j \geq i} v_j - T)_+]/(fT)} \geq \frac{\mathbb{E}[(v_{\max} - T)_+]}{1 - \mathbb{E}[(v_{\max} - T)_+]/(fT)}.$$

This finishes the proof of (1). We now verify the advertised prices. First consider

$$T = \max\left\{\frac{1+f}{1+2f}\mathbb{E}[v_{\max}], \alpha\right\}.$$

If $T = \alpha$, then since $v_{\max} \geq \alpha$, the item sells and

$$\mathbb{E}[\mathrm{ALG}_T] \geq \alpha \geq \frac{1+f}{1+2f}\mathbb{E}[v_{\max}].$$

Otherwise $T = \frac{1+f}{1+2f}\mathbb{E}[v_{\max}] > \alpha$. Since $v_{\max} \leq T + (v_{\max} - T)_+$ pointwise,

$$\mathbb{E}[(v_{\max} - T)_+] \geq \mathbb{E}[v_{\max}] - T.$$

The function $z \mapsto z/(1 - z/(fT))$ is increasing on the relevant interval, and a direct calculation gives

$$\frac{\mathbb{E}[v_{\max}] - T}{1 - (\mathbb{E}[v_{\max}] - T)/(fT)} = T.$$

Substituting into (1) gives

$$\mathbb{E}[\mathrm{ALG}_T] \geq T \cdot \Pr[v_{\max} \geq T] + \Pr[v_{\max} < T] \cdot T = T = \frac{1+f}{1+2f}\mathbb{E}[v_{\max}].$$

Next consider a price T satisfying

$$x := \Pr[v_{\max} < T] = \frac{f^2}{(1+f)(1+2f)}.$$

Note that we may assume such a price exists by applying a small random perturbation to the valuations if necessary. By scaling, assume $\mathbb{E}[v_{\max}] = 1$. Using $\mathbb{E}[(v_{\max} - T)_+] \geq 1 - T$ in (1), the guarantee for such a price is at least

$$(1-x)T + x \cdot \frac{fT(1-T)}{(1+f)T - 1}.$$

Since $v_{\max} \leq (1+f)\alpha$, our normalization implies $\alpha \geq 1/(1+f)$, and so $T \geq 1/(1+f)$. On the interval $T > 1/(1+f)$, the function

$$\frac{fT(1-T)}{(1+f)T-1}$$

is convex, and its tangent at $(1+f)/(1+2f)$ gives

$$\frac{fT(1-T)}{(1+f)T-1} \geq \frac{1+f}{1+2f} - \frac{f^2+3f+1}{f^2}\left(T - \frac{1+f}{1+2f}\right).$$

For $x = f^2/((1+f)(1+2f))$, the coefficient of $T - (1+f)/(1+2f)$ cancels after multiplying by x and adding $(1-x)T$. Thus

$$(1-x)T + x \cdot \frac{fT(1-T)}{(1+f)T-1} \geq \frac{1+f}{1+2f}.$$

Undoing the normalization proves

$$\mathbb{E}[\mathrm{ALG}_T] \geq \frac{1+f}{1+2f}\mathbb{E}[v_{\max}].$$

Finally consider the price $T = \max\{T^*, \alpha\}$, where T^* solves

$$\mathbb{E}[(v_{\max} - T^*)_+] = \frac{f}{1+f}T^*.$$

Such a solution exists because the left-hand side is continuous and decreasing in T^*, while the right-hand side is continuous and increasing in T^*. If $T^* \leq \alpha$, then

$$\mathbb{E}[(v_{\max} - \alpha)_+] \leq \mathbb{E}[(v_{\max} - T^*)_+] = \frac{f}{1+f}T^* \leq \frac{f}{1+f}\alpha.$$

Since $v_{\max} \geq \alpha$,

$$\mathbb{E}[v_{\max}] = \alpha + \mathbb{E}[(v_{\max} - \alpha)_+] \leq \frac{1+2f}{1+f}\alpha.$$

Thus price α sells and obtains at least $\alpha \geq \frac{1+f}{1+2f}\mathbb{E}[v_{\max}]$. If $T^* > \alpha$, then $T = T^*$ and

$$\frac{\mathbb{E}[(v_{\max} - T)_+]}{1 - \mathbb{E}[(v_{\max} - T)_+]/(fT)} = T.$$

The strengthened inequality (1) gives $\mathbb{E}[\mathrm{ALG}_T] \geq T$. Moreover,

$$\mathbb{E}[v_{\max}] \leq T + \mathbb{E}[(v_{\max} - T)_+] = \frac{1+2f}{1+f}T,$$

and hence $T \geq \frac{1+f}{1+2f}\mathbb{E}[v_{\max}]$.

□

4 The averaging lemma

In the previous proof, we omitted the proof of the following lemma:

Lemma 1. *Let $v_1, \ldots, v_n$ be independent nonnegative buyer valuations satisfying $v_i \le (1+f)T$ for every i. Then,*

$$\sum_{i=1}^{n} \frac{\mathbb{E}[(v_i - T)_+]}{1 - \mathbb{E}[(\max_{j \ge i} v_j - T)_+]/(fT)} \ge \frac{\mathbb{E}[(v_{\max} - T)_+]}{1 - \mathbb{E}[(v_{\max} - T)_+]/(fT)}.$$

Proof. It suffices to prove the statement for two independent random variables, since the general statement follows by induction over suffixes. For two variables, write $A = v_1$ and $B = v_2$, and define

$$Y_A := 1 - \frac{(A-T)_+}{fT}, \qquad Y_B := 1 - \frac{(B-T)_+}{fT}.$$

The assumption $A, B \le (1+f)T$ implies $0 \le Y_A, Y_B \le 1$. Moreover,

$$1 - \frac{\mathbb{E}[(A-T)_+]}{fT} = \mathbb{E}[Y_A], \qquad 1 - \frac{\mathbb{E}[(B-T)_+]}{fT} = \mathbb{E}[Y_B],$$

and

$$1 - \frac{\mathbb{E}[(\max\{A,B\} - T)_+]}{fT} = \mathbb{E}[\min\{Y_A, Y_B\}].$$

After canceling the common factor fT, the desired inequality is equivalent to

$$\frac{1 - \mathbb{E}[Y_A]}{\mathbb{E}[\min\{Y_A, Y_B\}]} + \frac{1 - \mathbb{E}[Y_B]}{\mathbb{E}[Y_B]} \ge \frac{1 - \mathbb{E}[\min\{Y_A, Y_B\}]}{\mathbb{E}[\min\{Y_A, Y_B\}]}.$$

Rearranging, this becomes

$$\mathbb{E}[\min\{Y_A, Y_B\}] \ge \mathbb{E}[Y_A]\mathbb{E}[Y_B].$$

Since $0 \le Y_A, Y_B \le 1$, we have the pointwise inequality

$$\min\{Y_A, Y_B\} \ge Y_A Y_B.$$

Taking expectations and using independence gives the required result. □

References

[1] Johannes Brüstle, Ilan Reuven Cohen, and Stefano Leonardi. Prophet inequality with conservative prediction. arXiv preprint arXiv:2602.17358, 2026.

[2] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Recent developments in prophet inequalities. *ACM SIGecom Exchanges*, 17(1):61–70, 2019.

[3] Farbod Ekbatani, Rad Niazadeh, Pranav Nuti, and Jan Vondrák. Prophet inequalities with cancellation costs. arXiv preprint arXiv:2404.00527v1, 2024.

[4] Theodore P. Hill and Robert P. Kertz. Additive comparisons of stop rule and supremum expectations of uniformly bounded independent random variables. *Proceedings of the American Mathematical Society*, 83(3):582–585, 1981.

[5] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. In *Probability on Banach Spaces*, Advances in Probability and Related Topics, vol. 4, pages 197–266. Marcel Dekker, 1978.

[6] Brendan Lucier. An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):26–49, 2017.

[7] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *The Annals of Probability*, 12(4):1213–1216, 1984.

[8] Ester Samuel-Cahn. Prophet inequalities for threshold rules for independent bounded random variables. In *Statistical Decision Theory and Related Topics IV*, vol. 2, pages 177–182. Springer, New York, 1988.