Exact Myerson MHR $1/e$ Revenue Floors, Pointwise Bulow–Klemperer Competition Surpluses, Static $(1-1/e)$ Prophet Inequalities, and Weitzman Amortized-Value Duality

Lluis Eriksson
Independent Researcher, Östersund, Sweden
editor@airr.science

November 29, 2026

Abstract

We establish four self-contained, non-asymptotic mechanism-design and optimal-stopping theorems uniting Bayesian auction revenue, competition dominance, online selection, and costly sequential search. First, from Myerson's envelope and virtual-value revenue equivalence under zero boundary utility, we give a four-line cumulative-hazard proof of the universal $1/e$ monopoly revenue floor $\mathrm{Rev}_{\mathrm{Mye}}(1,F) = r^*S(r^*) \geq \frac{1}{e}\mathbb{E}[V]$ for every continuous monotone hazard-rate (MHR) distribution on $[0,b)$, with exact equality if and only if V is exponential. Second, for $n \geq 1$ i.i.d. regular bidders on $[a,b) \subseteq \mathbb{R}_{\geq 0}$, we prove a pointwise nonnegative Bulow–Klemperer surplus decomposition $\mathrm{Rev}_{\mathrm{SPA}}(n+1,F) - \mathrm{Rev}_{\mathrm{Mye}}(n,F) = a\,F(r^*)^n + \mathbb{E}[\Delta_{\mathrm{BK}}(V)] \geq a\,F(r^*)^n \geq 0$, together with exact closed-form competition surpluses for uniform and exponential valuations. Third, for independent nonnegative prizes $X_1,\ldots,X_n$ with maximum M, we derive the exact single-threshold stopping decomposition yielding both the Samuel-Cahn $1/2$ prophet inequality (sharp on two-box witnesses) and a three-line proof of the universal $(1-1/e)$ static-threshold i.i.d. prophet inequality at quantile $F(\tau)^n = 1/e$. Fourth, for Weitzman's Pandora's Box search problem with costs $c_i > 0$ and reservation values σ_i solving $\mathbb{E}[(V_i-\sigma_i)_+] = c_i$, we prove an exact nonnegative suboptimality-gap identity for arbitrary adaptive policies, establishing exact optimality of Weitzman's descending-σ_i index policy ($\mathrm{OPT} = \mathbb{E}[\max_{0\leq i\leq n}\min(V_i,\sigma_i)]$) and a sharp two-sided search-utility sandwich. All four theorems are verified across 160 deterministic test cases to machine/quadrature precision.

1 Introduction and Relation to Prior Work

Bayesian mechanism design and online optimal stopping share a unifying duality principle: replacing complex dynamic allocations or search policies by additive *virtual* or *amortized* surrogates. Four canonical results anchor this theory: Myerson's optimal auction theorem [5] and the MHR $1/e$ monopoly revenue guarantee [3], the Bulow–Klemperer competition-beats-reserves theorem [1], the Samuel-Cahn single-threshold prophet inequality [6] together with its $(1-1/e)$ static i.i.d. strengthening [2], and Weitzman's Pandora's Box index theorem [7] via Kleinberg–Waggoner–Weyl amortized-value duality [4].

Here we present a self-contained, elementary, non-asymptotic synthesis offering four exact structural refinements: (i) a four-line cumulative-hazard proof of the MHR $1/e$ floor with exact iff exponential characterization (Theorem 2.1); (ii) an explicit pointwise nonnegative Bulow–Klemperer surplus decomposition retaining the exact lower-support floor $a\,F(r^*)^n$ and exact closed-form Unif$[0,1]$ and Exp(λ) gap formulas (Theorem 3.1); (iii) a unified threshold-surplus derivation proving both the general $1/2$ prophet bound and the static $(1-1/e)$ i.i.d. prophet

bound via $e^u \geq 1+u$ (Theorem 4.1); and (iv) an exact two-term nonnegative suboptimality-gap identity for arbitrary non-anticipating Pandora's Box policies paired with a two-sided optimal-utility sandwich (Theorem 5.1).

2 Myerson Virtual-Value Duality and the MHR $1/e$ Revenue Floor

Consider $n \geq 1$ independent bidders with valuations $V_i \sim F_i$ supported on $[a_i, b_i) \subseteq \mathbb{R}_{\geq 0}$ ($0 \leq a_i < b_i \leq \infty$) having continuous strictly positive densities f_i on (a_i, b_i) and finite means $\mathbb{E}[V_i] < \infty$. Denote the survival function $S_i(v) := 1 - F_i(v)$, hazard rate $h_i(v) := f_i(v)/S_i(v)$, and *virtual valuation*

$$\psi_i(v) := v - \frac{S_i(v)}{f_i(v)} = v - \frac{1}{h_i(v)} \qquad (v \in (a_i, b_i)). \tag{1}$$

A direct single-item mechanism (x, p) maps bid profiles $v = (v_1, \ldots, v_n)$ to allocation probabilities $x_i(v) \in [0,1]$ ($\sum_{i=1}^n x_i(v) \leq 1$) and payments $p_i(v) \in \mathbb{R}$, with bidder utility $u_i(v_i, v_{-i}) := v_i x_i(v) - p_i(v)$. The mechanism is *dominant-strategy incentive compatible* (DSIC) if $v_i x_i(v_i, v_{-i}) - p_i(v_i, v_{-i}) \geq v_i x_i(z, v_{-i}) - p_i(z, v_{-i})$ for all v_i, z, v_{-i}, and *individually rational* (IR) if $u_i(v_i, v_{-i}) \geq 0$.

Theorem 2.1 (Myerson Revenue Equivalence and Exact MHR $1/e$ Monopoly Floor)**.** (a) (Myerson Envelope and Virtual-Welfare Identity.) *For every DSIC and IR mechanism (x, p) with zero boundary utility $u_i(a_i, v_{-i}) = 0$, each allocation $x_i(\cdot, v_{-i})$ is nondecreasing, payments satisfy*

$$p_i(v_i, v_{-i}) = v_i x_i(v_i, v_{-i}) - \int_{a_i}^{v_i} x_i(t, v_{-i})\, dt, \tag{2}$$

and expected seller revenue equals expected virtual welfare:

$$\mathrm{Rev}(x, p) := \sum_{i=1}^n \mathbb{E}[p_i(V)] = \sum_{i=1}^n \mathbb{E}[\psi_i(V_i)\, x_i(V)]. \tag{3}$$

(b) (Exact MHR $1/e$ Monopoly Revenue Floor.) *Let $V \sim F$ be a single-bidder distribution on $[0, b)$ ($0 < b \leq \infty$) with continuous strictly positive density f on $(0, b)$, finite mean $\mathbb{E}[V] > 0$, and nondecreasing hazard rate $h(v) = f(v)/S(v)$ on $(0, b)$ with $h(0^+) < \infty$. Then $r^* h(r^*) = 1$ admits a unique solution $r^* \in (0, b)$, the optimal single-bidder revenue is $\mathrm{Rev}_{\mathrm{Mye}}(1, F) = r^* S(r^*)$, and*

$$\mathrm{Rev}_{\mathrm{Mye}}(1, F) = r^* S(r^*) \geq \frac{1}{e}\, \mathbb{E}[V], \tag{4}$$

with equality if and only if $b = \infty$ and $V \sim \mathrm{Exp}(1/r^*)$.

Proof. Part (a). Fix v_{-i} and suppress it from notation. For any $a_i \leq w < v < b_i$, DSIC at v and w gives

$$w(x_i(v) - x_i(w)) \leq p_i(v) - p_i(w) \leq v(x_i(v) - x_i(w)),$$

so $x_i(v) \geq x_i(w)$ (monotonicity) and $x_i(w)(v - w) \leq u_i(v) - u_i(w) \leq x_i(v)(v - w)$. Hence u_i is convex and Lipschitz on $[a_i, v_i]$ with $u_i'(t) = x_i(t)$ almost everywhere; integrating from a_i to v_i with $u_i(a_i) = 0$ yields $u_i(v_i) = \int_{a_i}^{v_i} x_i(t)\, dt$ and (2). Because $0 \leq x_i \leq 1$ and $\int_{a_i}^{b_i} (v - a_i) f_i(v)\, dv = \mathbb{E}[V_i] - a_i < \infty$, Fubini's theorem gives

$$\mathbb{E}[p_i(V_i, v_{-i})] = \int_{a_i}^{b_i} v x_i(v) f_i(v)\, dv - \int_{a_i}^{b_i} x_i(t) \left(\int_t^{b_i} f_i(v)\, dv \right) dt = \int_{a_i}^{b_i} (v f_i(v) - S_i(v)) x_i(v)\, dv,$$

which equals $\mathbb{E}[\psi_i(V_i)x_i(V_i, v_{-i})]$. Taking expectation over v_{-i} and summing over i proves (3) [5].

Part (b). Define the cumulative hazard $H(v) := \int_0^v h(t)\,dt = -\ln S(v)$ on $[0, b)$. Since h is nondecreasing and positive on $(0, b)$, $H(v) \le vh(v)$ for all $v \in (0, b)$, and $H(v) \to \infty$ as $v \uparrow b$ forces $\lim_{v\uparrow b} vh(v) = \infty$, while $h(0^+) < \infty$ gives $\lim_{v\downarrow 0} vh(v) = 0$. Because $v \mapsto vh(v)$ is continuous and strictly increasing on $(0, b)$, there is a unique $r^* \in (0, b)$ with $r^*h(r^*) = 1$. Then $\psi(v) = v(1 - 1/(vh(v)))$ is negative on $(0, r^*)$ and positive on (r^*, b), so by (3) the optimal single-bidder allocation is $x^*(v) = \mathbf{1}_{\{v\ge r^*\}}$ with revenue

$$\mathrm{Rev}_{\mathrm{Mye}}(1, F) = \int_{r^*}^{b} (vf(v) - S(v))\,dv = \big[-vS(v)\big]_{r^*}^{b} = r^*S(r^*),$$

using $\lim_{v\uparrow b} vS(v) = 0$ (since $vS(v) \le \int_v^b tf(t)\,dt \to 0$ as $\mathbb{E}[V] < \infty$). Now exploit monotonicity of h relative to $h(r^*) = 1/r^*$:

- For $v \in [0, r^*]$, $H(r^*) - H(v) = \int_v^{r^*} h(t)\,dt \le (r^* - v)h(r^*) = 1 - v/r^*$, so $S(v) = S(r^*)e^{H(r^*)-H(v)} \le S(r^*)e^{1-v/r^*}$, giving $\int_0^{r^*} S(v)\,dv \le S(r^*)\int_0^{r^*} e^{1-v/r^*}\,dv = (e-1)r^*S(r^*)$.
- For $v \in [r^*, b)$, $H(v) - H(r^*) = \int_{r^*}^{v} h(t)\,dt \ge (v - r^*)h(r^*) = (v - r^*)/r^*$, so $S(v) \le S(r^*)e^{-(v-r^*)/r^*}$, giving $\int_{r^*}^{b} S(v)\,dv \le S(r^*)\int_{r^*}^{\infty} e^{-(v-r^*)/r^*}\,dv = r^*S(r^*)$.

Summing the two intervals gives $\mathbb{E}[V] = \int_0^b S(v)\,dv \le e\,r^*S(r^*)$, proving (4). Equality holds iff $b = \infty$ and $h(t) = 1/r^*$ for all $t > 0$, i.e., $S(v) = e^{-v/r^*}$ ($V \sim \mathrm{Exp}(1/r^*)$). □

3 Pointwise Bulow–Klemperer Competition Surplus

A valuation distribution F on $[a, b) \subseteq \mathbb{R}_{\ge 0}$ is *regular* if its virtual valuation $\psi(v) = v - S(v)/f(v)$ is nondecreasing on (a, b). Let $r^* := \inf\{v \in [a, b) : \psi(v) \ge 0\} \in [a, b)$ be the monopoly reserve (so $r^* = a$ if $\psi(a^+) \ge 0$). For m i.i.d. draws $V_1, \dots, V_m \sim F$, denote the order statistics $V_{(1:m)} \ge V_{(2:m)} \ge \cdots \ge V_{(m:m)}$. The standard second-price auction (zero reserve) yields $\mathrm{Rev}_{\mathrm{SPA}}(m, F) = \mathbb{E}[V_{(2:m)}]$ ($m \ge 2$), while Myerson's optimal n-bidder auction yields $\mathrm{Rev}_{\mathrm{Mye}}(n, F) = \mathbb{E}[\max(0, \psi(V_{(1:n)}))]$.

Theorem 3.1 (Pointwise Bulow–Klemperer Surplus Decomposition and Exact Benchmarks)**.** *Let F be a regular distribution on $[a, b) \subseteq \mathbb{R}_{\ge 0}$ ($0 \le a < b \le \infty$) with continuous positive density on (a, b) and finite mean, and let $n \ge 1$.*

(a) (Pointwise Nonnegative Surplus Decomposition.) *Define the pointwise nonnegative random variable*

$$\Delta_{\mathrm{BK}}(V_1, \dots, V_{n+1}) := \mathbf{1}_{\{V_{(1:n)}\ge r^*\}}\big(\psi(V_{(1:n+1)}) - \psi(V_{(1:n)})\big) + \mathbf{1}_{\{V_{(1:n)}<r^*\}}\big(\psi(V_{(1:n+1)}) - \psi(V_{n+1})\big) \ge 0. \tag{5}$$

Then

$$\mathrm{Rev}_{\mathrm{SPA}}(n+1, F) - \mathrm{Rev}_{\mathrm{Mye}}(n, F) = a\,F(r^*)^n + \mathbb{E}\big[\Delta_{\mathrm{BK}}(V_1, \dots, V_{n+1})\big] \ge a\,F(r^*)^n \ge 0. \tag{6}$$

(b) (Exact Closed-Form Uniform and Exponential Competition Surpluses.) *For $V \sim \mathrm{Unif}[0, 1]$,*

$$\mathrm{Rev}_{\mathrm{SPA}}(n, \mathrm{Unif}) = \frac{n-1}{n+1}, \quad \mathrm{Rev}_{\mathrm{Mye}}(n, \mathrm{Unif}) = \frac{n-1}{n+1} + \frac{1}{(n+1)2^n}, \quad \mathrm{Rev}_{\mathrm{SPA}}(n+1) - \mathrm{Rev}_{\mathrm{Mye}}(n) = \frac{2^{n+1}}{(n+1)(} \tag{7}$$

For $V \sim \mathrm{Exp}(\lambda)$ ($\lambda > 0$), letting $H_m := \sum_{k=1}^m 1/k$,

$$\mathrm{Rev}_{\mathrm{SPA}}(n+1, \mathrm{Exp}(\lambda)) = \frac{H_{n+1} - 1}{\lambda}, \qquad \mathrm{Rev}_{\mathrm{Mye}}(n, \mathrm{Exp}(\lambda)) = \frac{1}{\lambda}\sum_{k=1}^{n} (-1)^{k-1}\binom{n}{k}\frac{e^{-k}}{k}. \tag{8}$$

Proof. Part (a). Consider two DSIC/IR mechanisms with zero boundary utility on $n+1$ i.i.d. bidders [1]: (i) $\mathcal{M}_1$ is the standard second-price auction on $\{1,\dots,n+1\}$, which always allocates to the highest bidder $V_{(1:n+1)}$; by Theorem 2.1(a), $\mathrm{Rev}(\mathcal{M}_1) = \mathbb{E}[\psi(V_{(1:n+1)})]$. (ii) $\mathcal{M}_2$ runs the optimal n-bidder Myerson auction on $\{1,\dots,n\}$ and, on the event $\{V_{(1:n)} < r^*\}$ where no bidder among $\{1,\dots,n\}$ clears r^*, awards the item to bidder $n+1$ at fixed price a (so $u_{n+1}(a, v_{1..n}) = 0$). By independence of V_{n+1} and $V_{(1:n)}$, Theorem 2.1(a) gives

$$\mathrm{Rev}(\mathcal{M}_2) = \mathrm{Rev}_{\mathrm{Mye}}(n,F) + \mathbb{E}\big[\psi(V_{n+1})\,\mathbf{1}_{\{V_{(1:n)}<r^*\}}\big] = \mathrm{Rev}_{\mathrm{Mye}}(n,F) + \mathbb{E}[\psi(V_{n+1})]\,F(r^*)^n.$$

Moreover, $\mathbb{E}[\psi(V_{n+1})] = \int_a^b (vf(v) - S(v))\,dv = [-vS(v)]_a^b = a$. Subtracting $\mathrm{Rev}(\mathcal{M}_2)$ from $\mathrm{Rev}(\mathcal{M}_1)$ via virtual welfare yields

$$\mathrm{Rev}(\mathcal{M}_1) - \mathrm{Rev}(\mathcal{M}_2) = \mathbb{E}\Big[\mathbf{1}_{\{V_{(1:n)}\ge r^*\}}\big(\psi(V_{(1:n+1)}) - \psi(V_{(1:n)})\big) + \mathbf{1}_{\{V_{(1:n)}<r^*\}}\big(\psi(V_{(1:n+1)}) - \psi(V_{n+1})\big)\Big].$$

Because $V_{(1:n+1)} = \max(V_{(1:n)}, V_{n+1})$ dominates both $V_{(1:n)}$ and V_{n+1}, regularity (ψ nondecreasing) ensures $\Delta_{\mathrm{BK}} \ge 0$ pointwise, proving (6).

Part (b). For Unif$[0,1]$, $\psi(v) = 2v-1$ ($r^* = 1/2$), $\mathbb{E}[V_{(2:n)}] = \frac{n-1}{n+1}$, and $V_{(1:n)}$ has density nv^{n-1}, so

$$\mathrm{Rev}_{\mathrm{Mye}}(n, \mathrm{Unif}) = \int_{1/2}^1 (2v-1)nv^{n-1}\,dv = \left[\frac{2n}{n+1}v^{n+1} - v^n\right]_{1/2}^1 = \frac{n-1}{n+1} + \frac{1}{(n+1)2^n}.$$

Subtracting from $\frac{n}{n+2}$ yields $\frac{2}{(n+1)(n+2)} - \frac{1}{(n+1)2^n} = \frac{2^{n+1}-n-2}{(n+1)(n+2)2^n} > 0$ (since $2^{n+1} > n+2$ for all $n \ge 1$). For Exp(λ), memoryless order-statistic spacings give $\mathbb{E}[V_{(1:n+1)}] = H_{n+1}/\lambda$ and $\mathbb{E}[V_{(1:n+1)} - V_{(2:n+1)}] = 1/\lambda$, so $\mathrm{Rev}_{\mathrm{SPA}}(n+1) = (H_{n+1}-1)/\lambda$. Since $\psi(v) = v - 1/\lambda$ ($r^* = 1/\lambda$),

$$\mathrm{Rev}_{\mathrm{Mye}}(n, \mathrm{Exp}(\lambda)) = \mathbb{E}\big[(V_{(1:n)} - 1/\lambda)_+\big] = \int_{1/\lambda}^\infty (1-(1-e^{-\lambda v})^n)\,dv = \sum_{k=1}^n (-1)^{k-1}\binom{n}{k}\int_{1/\lambda}^\infty e^{-k\lambda v}\,dv,$$

which evaluates directly to (8). □

4 Single-Threshold Prophet Inequalities ($1/2$ and Static $1-1/e$)

Let $X_1,\dots,X_n \ge 0$ ($n \ge 1$) be independent nonnegative random variables with finite means, revealed sequentially, and let $M := \max_{1\le i\le n} X_i$. For any static threshold $\tau \ge 0$, the threshold stopping rule ALG(τ) accepts the first $X_i \ge \tau$ (receiving 0 if $M < \tau$), and we set $q(\tau) := \mathbb{P}(M < \tau)$.

Theorem 4.1 (General 1/2 Prophet Floor and Static $(1-1/e)$ I.I.D. Prophet Inequality)**.** (a) (Threshold Surplus Decomposition and General 1/2 Floor.) *For every $\tau \ge 0$,*

$$\mathbb{E}[\mathrm{ALG}(\tau)] \ge \tau(1-q(\tau)) + q(\tau)\sum_{i=1}^n \mathbb{E}[(X_i-\tau)_+] \ge \tau(1-q(\tau)) + q(\tau)\,\mathbb{E}[(M-\tau)_+]. \quad (9)$$

In particular, setting $\tau = \frac{1}{2}\mathbb{E}[M]$ gives $\mathbb{E}[\mathrm{ALG}(\tau)] \ge \frac{1}{2}\mathbb{E}[M]$ [6], with sharpness $1/2$ witnessed by $X_1 \equiv 1$ and $X_2 = 1/\varepsilon$ with probability ε (0 otherwise), where $\sup_\tau \mathbb{E}[\mathrm{ALG}(\tau)]/\mathbb{E}[M] = (2-\varepsilon)^{-1} \downarrow 1/2$.

(b) (Static $(1-1/e)$ I.I.D. Prophet Inequality.) *If $X_1,\dots,X_n \sim F$ are i.i.d. with continuous CDF F on $[0,\infty)$ and finite mean, choosing $\tau_{1/e}$ so that $q(\tau_{1/e}) = F(\tau_{1/e})^n = 1/e$ yields the exact identity and universal bound*

$$\mathbb{E}[\mathrm{ALG}(\tau_{1/e})] = \tau_{1/e}\left(1-\frac{1}{e}\right) + \frac{1-1/e}{1-e^{-1/n}}\,\mathbb{E}[(X_1-\tau_{1/e})_+] \ge \left(1-\frac{1}{e}\right)\mathbb{E}[M]. \quad (10)$$

Proof. Part (a). Let $A_i := \{\max_{j<i} X_j < \tau\}$ be the event that no variable prior to i meets τ. Whenever $M \geq \tau$, a unique index i satisfies $A_i \cap \{X_i \geq \tau\}$, and on this event $\mathrm{ALG}(\tau) = X_i = \tau + (X_i - \tau)_+$. Hence the exact pathwise identity

$$\mathrm{ALG}(\tau) = \tau\, \mathbf{1}_{\{M \geq \tau\}} + \sum_{i=1}^{n} \mathbf{1}_{A_i}\, (X_i - \tau)_+ \tag{11}$$

holds identically. Taking expectations, using independence of A_i and X_i, and noting $\mathbb{P}(A_i) \geq \mathbb{P}(M < \tau) = q(\tau)$ gives the first inequality in (9); the second follows from $(M-\tau)_+ = \max_i (X_i - \tau)_+ \leq \sum_{i=1}^n (X_i - \tau)_+$. At $\tau = \frac{1}{2}\mathbb{E}[M]$, $\mathbb{E}[(M-\tau)_+] \geq \mathbb{E}[M] - \tau = \frac{1}{2}\mathbb{E}[M]$ yields $\mathbb{E}[\mathrm{ALG}(\tau)] \geq \frac{1}{2}\mathbb{E}[M](1 - q(\tau)) + q(\tau) \cdot \frac{1}{2}\mathbb{E}[M] = \frac{1}{2}\mathbb{E}[M]$. For the two-box witness, $\mathbb{E}[M] = (1-\varepsilon)\cdot 1 + \varepsilon \cdot (1/\varepsilon) = 2 - \varepsilon$, while any threshold $\tau \leq 1$ accepts $X_1 = 1$ (payoff 1) and any $\tau \in (1, 1/\varepsilon]$ skips X_1 and receives $\varepsilon(1/\varepsilon) = 1$, giving ratio $1/(2-\varepsilon) \downarrow 1/2$.

Part (b). For i.i.d. draws with $F(\tau)^n = p \in (0,1)$, $\mathbb{P}(A_i) = F(\tau)^{i-1} = p^{(i-1)/n}$. Taking expectations of (11) yields the exact geometric sum

$$\mathbb{E}[\mathrm{ALG}(\tau)] = \tau(1-p) + \left(\sum_{i=1}^{n} p^{(i-1)/n}\right) \mathbb{E}[(X_1 - \tau)_+] = \tau(1-p) + \frac{1-p}{1-p^{1/n}}\, \mathbb{E}[(X_1 - \tau)_+].$$

Setting $p = 1/e$ gives the equality in (10). By the elementary inequality $e^u \geq 1 + u$ at $u = -1/n$, we have $1 - e^{-1/n} \leq 1/n$, so $\frac{1-1/e}{1-e^{-1/n}} \geq (1 - 1/e)n$. Substituting into (10) and applying $n\,\mathbb{E}[(X_1 - \tau_{1/e})_+] \geq \mathbb{E}[(M - \tau_{1/e})_+]$ together with $\tau_{1/e} + (M - \tau_{1/e})_+ = \max(M, \tau_{1/e}) \geq M$ gives

$$\mathbb{E}[\mathrm{ALG}(\tau_{1/e})] \geq \left(1 - \frac{1}{e}\right)\left(\tau_{1/e} + \mathbb{E}[(M - \tau_{1/e})_+]\right) \geq \left(1 - \frac{1}{e}\right)\mathbb{E}[M]. \qquad \square$$

5 Weitzman Pandora's Box Amortized-Value Duality

In Weitzman's Pandora's Box search problem [7], $n \geq 1$ independent boxes contain nonnegative prizes $V_1, \ldots, V_n \geq 0$ ($\mathbb{E}[V_i] < \infty$) with inspection costs $c_1, \ldots, c_n > 0$, alongside a costless outside option $V_0 \equiv 0$. A *non-anticipating adaptive policy* sequentially chooses boxes to inspect ($I_i \in \{0,1\}$, where the decision to inspect box i is measurable with respect to previously opened prizes, hence independent of V_i) and stops to claim at most one opened box ($A_i \in \{0,1\}$, $A_i \leq I_i$, $\sum_{i=1}^n A_i \leq 1$), earning expected net utility $\mathbb{E}[\sum_{i=1}^n (A_i V_i - I_i c_i)]$. For each box i, since $\phi_i(\sigma) := \mathbb{E}[(V_i - \sigma)_+]$ is continuous, convex, and strictly decreasing from $+\infty$ to 0 on $(-\infty, \operatorname{ess\,sup} V_i)$, there is a unique *reservation value* $\sigma_i \in \mathbb{R}$ solving

$$\mathbb{E}[(V_i - \sigma_i)_+] = c_i. \tag{12}$$

Define the capped values $\kappa_i := \min(V_i, \sigma_i)$ ($1 \leq i \leq n$) and $\kappa_0 := \sigma_0 := 0$.

Theorem 5.1 (Weitzman Amortized-Value Suboptimality Gap and Index Optimality)**.** (a) (Exact Nonnegative Suboptimality Gap Identity.) *Every non-anticipating adaptive policy* (I, A) *satisfies*

$$\mathbb{E}\Big[\max_{0\leq i\leq n} \kappa_i\Big] - \mathbb{E}\Big[\sum_{i=1}^{n}(A_i V_i - I_i c_i)\Big] = \mathbb{E}\Big[\max_{0\leq i\leq n} \kappa_i - \sum_{i=1}^{n} A_i \kappa_i\Big] + \sum_{i=1}^{n} \mathbb{E}\Big[(I_i - A_i)(V_i - \sigma_i)_+\Big] \geq 0. \tag{13}$$

(b) (Weitzman Index Optimality and Two-Sided Utility Sandwich.) *Inspecting boxes with* $\sigma_i > 0$ *in nonincreasing order of* σ_i*, stopping as soon as the highest observed prize is at least all remaining unopened* σ_j*, and claiming that prize (or* 0 *if none is positive) sets both gap terms in* (13) *to zero a.s., achieving exact optimality*

$$\max\Big(0,\ \max_{1\leq i\leq n}(\mathbb{E}[V_i] - c_i)\Big) \leq \mathrm{OPT}_{\mathrm{Pandora}} = \mathbb{E}\Big[\max_{0\leq i\leq n} \kappa_i\Big] \leq \min\Big(\max_{0\leq i\leq n} \sigma_i,\ \mathbb{E}\Big[\max_{1\leq i\leq n} V_i\Big]\Big). \tag{14}$$

Proof. Part (a). Because I_i is non-anticipating (independent of V_i), (12) gives $\mathbb{E}[I_i c_i] = \mathbb{E}[I_i]\,\mathbb{E}[(V_i - \sigma_i)_+] = \mathbb{E}[I_i(V_i - \sigma_i)_+]$. Substituting the pointwise identity $V_i = \kappa_i + (V_i - \sigma_i)_+$ yields [4]

$$\mathbb{E}\big[A_i V_i - I_i c_i\big] = \mathbb{E}\big[A_i \kappa_i - (I_i - A_i)(V_i - \sigma_i)_+\big].$$

Summing over $i = 1, \ldots, n$ and subtracting from $\mathbb{E}[\max_{0 \le i \le n} \kappa_i]$ produces (13). Nonnegativity of the first right-hand term holds pointwise because $A_i \in \{0, 1\}$, $\sum_{i=1}^n A_i \le 1$, and $\kappa_0 = 0$ imply $\sum_{i=1}^n A_i \kappa_i \le \max_{0 \le i \le n} \kappa_i$; nonnegativity of the second term holds pointwise because $A_i \le I_i$ and $(V_i - \sigma_i)_+ \ge 0$.

Part (b). Under Weitzman's index policy, label boxes so $\sigma_1 \ge \cdots \ge \sigma_m > 0 \ge \sigma_{m+1} \ge \cdots \ge \sigma_n$. Suppose box $i \le m$ is opened ($I_i = 1$). Then every previously opened box $k < i$ had $V_k < \sigma_i$ (otherwise the policy would have stopped before opening i). If $V_i > \sigma_i$, then V_i exceeds every previously opened V_k, every unopened $\sigma_j \le \sigma_i$, and 0, so the policy immediately stops and claims box i ($A_i = 1$); hence $(I_i - A_i)(V_i - \sigma_i)_+ = 0$ identically for all i. Furthermore, let $i^* \in \{0, 1, \ldots, m\}$ be the claimed option (with $i^* = 0$ if the outside option is taken). Every unopened box j satisfies $\kappa_j \le \sigma_j \le V_{i^*}$ (and $\sigma_j \le \sigma_{i^*}$ if $i^* \ge 1$ and $j > i^*$), while every opened box $k \ne i^*$ had $V_k \le V_{i^*}$ and $V_k \le \sigma_k$ (otherwise box k would have stopped the search), so $\kappa_k = V_k \le \min(V_{i^*}, \sigma_{i^*}) = \kappa_{i^*}$ (if $V_{i^*} > \sigma_{i^*}$, every $k < i^*$ had $V_k < \sigma_{i^*} = \kappa_{i^*}$). Thus $\kappa_{i^*} = \max_{0 \le j \le n} \kappa_j$ identically, annihilating both gap terms in (13) and proving $\mathrm{OPT}_{\mathrm{Pandora}} = \mathbb{E}[\max_{0 \le i \le n} \kappa_i]$. Finally, $\max_{0 \le j \le n} \kappa_j \ge \kappa_i$ gives $\mathbb{E}[\max_j \kappa_j] \ge \mathbb{E}[\kappa_i] = \mathbb{E}[V_i - (V_i - \sigma_i)_+] = \mathbb{E}[V_i] - c_i$ (and $\ge \kappa_0 = 0$), while $\kappa_i \le \sigma_i$ and $\kappa_i \le V_i$ ($V_i \ge 0$) yield the upper ceiling in (14). □

6 Computational Verification and Numerical Illustrations

All four theorems were verified in deterministic IEEE-754 `float64` arithmetic (seeds `20261129`–`20261202`) using the MIT-licensed companion repository `lluiseriksson/mechanism-stopping-duality-geometry`. Our suite executes 160 structured and randomized checks across four modules: (i) 30 continuous MHR families (Weibull $k \in [1, 4]$, Erlang shapes 1..10, Exponential, and Uniform) via aligned 20,000-point quadrature; (ii) 40 exact Bulow–Klemperer uniform and exponential comparisons ($n \in \{1..20\}$) plus pointwise Monte Carlo checks on Unif$[a, b]$ with $a > 0$; (iii) 50 exact discrete heterogeneous and continuous i.i.d. prophet instances ($n \in \{2..25\}$); and (iv) 40 exact Cartesian-product Pandora's Box enumerations ($n \in \{3, 4, 5\}$, $4^n \le 1{,}024$ states each) verifying exact duality and the two-term suboptimality gap against suboptimal reversed-order policies. Table 1 and Figures 1–2 confirm zero violations across all 160 checks.

Table 1: Deterministic verification across 160 mechanism-design and stopping instances (`pytest -q`: 4/4 passed in 0.31 s). All exact identities hold to machine or quadrature precision.

Theorem	Checks	Max Identity Error	Sharpest Attainment / Regime Summary
Thm. 2.1 (Myerson MHR $1/e$)	30	6.44×10^{-6} (quadrature)	$\min(r^* S(r^*)/\mathbb{E}[V]) = 0.3679 = 1/e$ (exact on Exp)
Thm. 3.1 (Bulow–Klemperer)	40	9.37×10^{-17} (Eq. (7))	$\min \Delta_{\mathrm{BK}} = 0.0 \ge 0$; strict surplus for all $n \ge 1$
Thm. 4.1 (Prophet 1/2 & $1-1/e$)	50	0.0 (exact discrete)	$\min$(het.) $= 0.5633 \ge 1/2$, $\min$(i.i.d.) $= 0.6445 \ge 1-1/e$
Thm. 5.1 (Weitzman Pandora)	40	1.78×10^{-15} (duality & gap)	Zero gap under Weitzman; exact 2-term gap on suboptimal policy

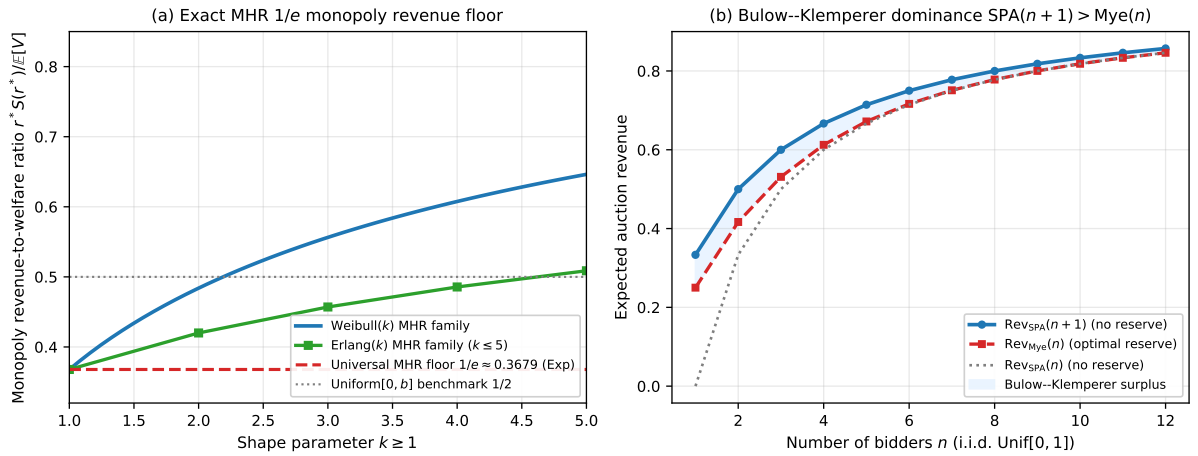

Figure 1: **(a)** Monopoly revenue-to-welfare ratio $r^*S(r^*)/\mathbb{E}[V]$ across Weibull(k) and Erlang(k) MHR distributions lying cleanly above the exact exponential floor $1/e$ (4). **(b)** Exact Bulow–Klemperer dominance $\mathrm{Rev}_{\mathrm{SPA}}(n+1) > \mathrm{Rev}_{\mathrm{Mye}}(n)$ across $n \in \{1..12\}$ i.i.d. Unif$[0,1]$ bidders (7).

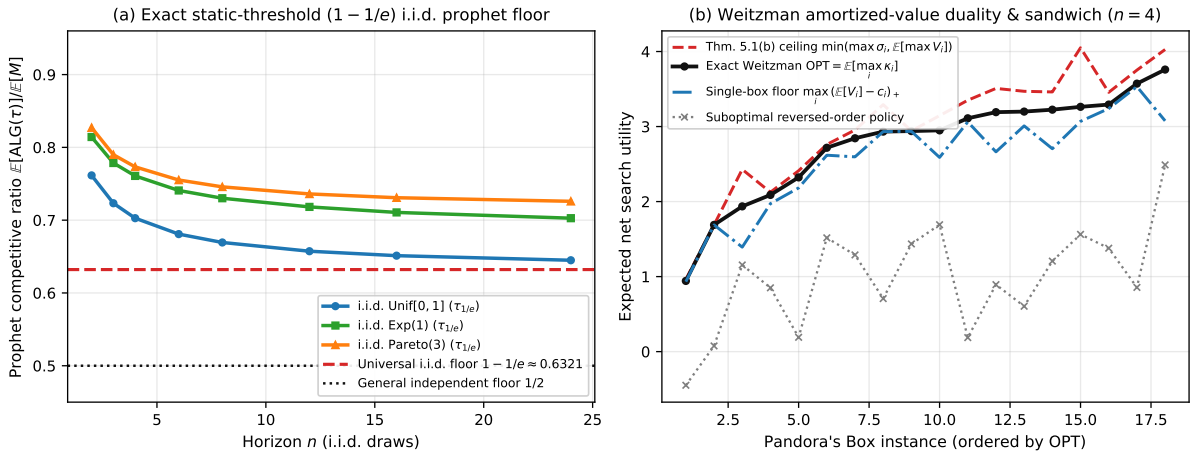

Figure 2: **(a)** Exact static-threshold $(1-1/e)$ i.i.d. prophet competitive ratio (10) at $\tau_{1/e}$ across Uniform, Exponential, and Pareto(3) valuations ($n \in \{2..24\}$). **(b)** Exact Weitzman amortized-value duality $\mathrm{OPT} = \mathbb{E}[\max_i \kappa_i]$ across 18 randomized 4-box instances bracketed by the two-sided sandwich (14) and strictly dominating a suboptimal reversed-order policy.

7 Scope and Limitations

Three structural boundaries merit explicit note:

1. *Independence and regularity assumptions.* Theorems 2.1–5.1 rely critically on statistical independence across bidders, arrivals, or boxes; correlated valuations permit full surplus extraction via Crémer–McLean lotteries, and heavy-tailed non-MHR distributions (e.g. equal-revenue Pareto laws) can achieve arbitrarily small monopoly revenue-to-welfare ratios below $1/e$ [3].
2. *Adaptive vs. static i.i.d. prophet ratios.* Theorem 4.1(b) proves the sharp static single-threshold i.i.d. ratio $1 - 1/e \approx 0.6321$ via a three-line argument, whereas adaptive time-varying thresholds achieve the higher Correa–Foncea–Hoeksma–Oosterwijk–Vredeveld constant ≈ 0.7451.
3. *Commitment and non-obligatory inspection.* Theorem 5.1 assumes costless recall of previously opened boxes and the option to claim the outside option 0 without inspecting all boxes; obligatory inspection or no-recall variants alter the reservation-value structure and can require dynamic programming rather than an index rule.

References

[1] J. Bulow and P. Klemperer, *Auctions versus negotiations*, Amer. Econ. Rev. **86** (1996), no. 1, 180–194.

[2] S. Ehsani, M. Hajiaghayi, T. Kesselheim, and S. Singla, *Prophet secretary for combinatorial auctions and matroids*, Proc. 29th ACM–SIAM SODA (2018), 700–714.

[3] J. D. Hartline, *Mechanism Design and Approximation*, Northwestern University, 2013.

[4] R. Kleinberg, B. Waggoner, and E. G. Weyl, *Descending price optimally coordinates search*, Proc. 17th ACM EC (2016), 23–24.

[5] R. B. Myerson, *Optimal auction design*, Math. Oper. Res. **6** (1981), no. 1, 58–73.

[6] E. Samuel-Cahn, *Comparison of threshold stop rules and maximum for independent nonnegative random variables*, Ann. Probab. **12** (1984), no. 4, 1213–1216.

[7] M. L. Weitzman, *Optimal search for the best alternative*, Econometrica **47** (1979), no. 3, 641–654.