Sequential Minimax Games as Stacked Martingale Optimal Transport

Anonymous Authors[1]

Abstract

We study sequential minimax option-pricing games with a variance-constrained zero-mean adversary and player actions bounded by $|\Delta_m| \le B < L$. The value-function Lipschitz constant grows exactly as $B + (L - B)(1 + h_n)^{n-m}$, where $h_n = \sqrt{e^{c/n} - 1}$, yielding a bounded-capital premium of order $\Theta((L-B)e^{\sqrt{cn}}s)$ for piecewise-linear losses. Under the standing calibration $\zeta_n = h_n$, this further sharpens to an exact boundary-layer identity beyond an explicit threshold. These results show that any strict action bound $B < L$ forces $\Theta(e^{\sqrt{cn}})$ divergence, so no finite PDE limit remains. Equivalently, every B-bounded strategy incurs pathwise regret $\Omega(e^{\sqrt{cn}})$ against the L-bounded hindsight comparator class. We also analyze the aggregate-budget variant $\sum_m |\Delta_m| \le nB$, where matching bounds give the smaller rate $Ls\exp((L-B)\sqrt{cn}/L - (L-B)c/(2L))(1+o(1))$. The analysis proceeds by realizing the n-step game through backward induction as stacked one-step entropic martingale optimal transport problems, whose dual variables recover the player's optimal actions.

1. Introduction

Sequential minimax games with constrained adversaries arise throughout online learning, adversarial prediction, and robust optimization. A fundamental structural question is: when does such a game admit a dual formulation as an optimal transport problem, and can the transport structure be exploited computationally?

This paper bridges two independently developed frameworks. On one side: n-step sequential zero-sum games where a player (minimizer) chooses actions and an adversary (maximizer) chooses state transitions subject to a zero-conditional-mean constraint, bounded conditional variance, and compact support (Abernethy et al., 2012). On the other: entropically regularized (Cuturi, 2013) optimal transport with L_1-relaxed coupling constraints, solvable by proximal gradient ascent on the dual variable (standard convex duality, Section 2). The n-step game and a stacked sequence of one-step transport problems have equal values via an explicit correspondence: backward induction decomposes the game into n one-step transport problems, each with a point-mass source (Theorem 2.5); the transport dual variable at each step *is* the player's optimal action (Proposition 2.4); and a one-step Fenchel conjugacy (Theorem 2.6) relates L_1-relaxed coupling to per-step action bounds. The n-step correspondence is thus a backward-induction stacking of a one-step Fenchel conjugacy, with per-step relaxation parameter δ_m rendered step-dependent by the continuation value's Lipschitz constant.

Contributions.

1. **Closed-form Lipschitz amplification with explicit constant** (Theorems 3.4, 4.2 and 4.4): under $|\Delta_m| \le B < L$ and $\zeta_n \in [h_n, 1)$, the Lipschitz constant is exactly $L_m^B(n) = B + (L-B)(1+h_n)^{n-m}$, hence grows as $\Theta(e^{\sqrt{cn}})$ uniformly over convex nondecreasing L-Lipschitz losses. DKM-type bounds (DeMarzo et al., 2006; Xue et al., 2023) give $O(C_1\, e^{\sqrt{cn}})$ with unspecified C_1; we pin $C_1 = (L-B)\, s$. For piecewise-linear losses this gives matching $\Theta((L-B)e^{\sqrt{cn}}s)$ premium bounds and, under $\zeta_n = h_n$, the exact boundary-layer identity $V_n^B(s,0) = (L-B)(\alpha_n - 1)s + g_\infty(s)$ beyond the threshold in Theorem 4.4.
2. **Pathwise structural lower bound** (Corollary 4.3): any B-bounded strategy is dominated by $\Omega(e^{\sqrt{cn}})$ on the worst-case adversary path against the L-bounded hindsight comparator class, with aggregate analogue $\Omega(e^{(L-B)\sqrt{cn}/L})$.
3. **Tight aggregate-budget rate** (Remark B.7): matching primal bounds give $\Phi_n^{\mathrm{agg},B}(s) = \Theta(e^{(L-B)\sqrt{cn}/L})$, exponent strictly below the per-step 1.
4. **Constant-λ Lagrangian dichotomy** (Proposition B.9): the aggregate dual sharpens the rate to an exact con-

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

stant; strong duality holds under a strengthened threshold and fails under the weaker stated threshold via a variance-saturating counterexample for $B \geq L/2$.

5. **Backward-induction realization and algorithmic construction** (Theorems 2.5 and 2.6): backward induction realizes the closed-form amplification as stacked one-step transport problems; the transport dual variable is the optimal action, and the L_1 relaxation is Fenchel-conjugate to the per-step action bound. Sinkhorn + proximal gradient ascent gives per-resolution cost $O(nq^2K)$; resolving the diverging Lipschitz constant at fixed relative precision requires q scaling as $\Theta(\alpha_n)$ (Corollary 3.5), giving intrinsic total cost $O(n\, e^{2\sqrt{cn}}\, K)$.
6. **Pathwise variant** (Theorem E.4): the ess-sup martingale formulation has rate $\Theta((1+\zeta_n)^n)$ tracking support; matches the Lagrangian at $\zeta_n = h_n$ and separates under wider AFW Lindeberg.

These results resolve the per-step and aggregate Lagrangian forms of (Abernethy et al., 2012)'s Problem 2 under $\zeta_n \geq h_n$: the bounded-capital game diverges sub-exponentially and admits no finite PDE limit. The distinction between formulations is calibration-sensitive: in the Lagrangian game, Jensen's inequality makes the rate track the variance scale h_n; in the pathwise martingale game, the rate tracks the support scale ζ_n. Backward induction realizes the closed-form amplification as n stacked one-step transport problems, complementing (Abernethy et al., 2013)'s asymptotic Black–Scholes-delta hedge for AFW's Problem 1 (unbounded capital) with a finite-n bounded-capital construction; the same Fenchel conjugacy drives the recurrence $\Lambda_m s = \Lambda_{m+1} s + \delta(B; \Lambda_{m+1})$ with $\delta(B;\Lambda) := (\Lambda - B)_+ s h_n$, iterating to Theorem 3.4. Proofs and experiments: Appendices E and H.

2. Setup and the Transport Bridge

2.1. Sequential minimax game

Definition 2.1 (Sequential minimax game (Abernethy et al., 2012))**.** Fix $n \in \mathbb{N}$, variance parameter $c > 0$, and convex L-Lipschitz terminal loss $g\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. A state sequence $S_n = (S_{n,0}, \ldots, S_{n,n})$ evolves from $S_{n,0} = 1$; at step m the player picks $\Delta_m \in \mathbb{R}$ (measurable w.r.t. $\mathcal{F}_{n,m-1}$) and the adversary picks the conditional law of $S_{n,m}$ subject to constraints on $T_{n,m} := S_{n,m}/S_{n,m-1} - 1$: **(DVC)** $\mathbb{E}[T_{n,m}^2 | \mathcal{F}_{n,m-1}] \leq e^{c/n} - 1$; **(ZC$_n$)** $|T_{n,m}| \leq \zeta_n$ with $\zeta_n \to 0$; and the **martingale** constraint $\mathbb{E}[T_{n,m}|\mathcal{F}_{n,m-1}] = 0$ (calibration scope: Remark 3.2). The loss is

$$\ell_n(\Delta, S_n) = \mathbb{E}\left[g(S_{n,n}) - \sum_{m=1}^{n} (S_{n,m} - S_{n,m-1})\, \Delta_m \right] \tag{1}$$

(share-count convention: Δ_m is the number of units held over $[m-1, m]$; the hedge P&L per step is $\Delta_m(S_{n,m} - S_{n,m-1}) = \Delta_m S_{n,m-1} T_{n,m}$. All per-step and aggregate budget constraints in the sequel apply to this share-count Δ_m), and $V_n := \inf_\Delta \sup_{S_n} \ell_n$.

By Sion's theorem (Sion, 1958) applied to $[-M, M]^n \times \mathcal{S}^n$ with $M \geq L$, the min and sup may be exchanged. The player's action set $[-M, M]^n$ is compact by choice. For $\mathcal{S}^n$, each coordinate $\mathcal{T}_n \subset \mathcal{P}([-\zeta_n, \zeta_n])$ with $\zeta_n < 1$ is supported on a compact interval, hence tight; weak compactness follows from Prokhorov. $\mathcal{S}^n$ is the space of adapted strategies (conditional laws given the filtration); its weak compactness reduces to compactness of each conditional law set $\mathcal{T}_n$ by Tychonoff on the measurable-selection parametrization (full argument in Appendix E). The loss ℓ_n is continuous in the player's action (linear) and upper semicontinuous in the adversary's conditional laws under weak convergence (by boundedness of g on compact support). Coercivity $M \geq L$ is sufficient for this Sion application, which concerns the *unconstrained* value V_n. The bounded-capital recursion of Definition 3.1 (where V_n^B has step-dependent Lipschitz $L_m^B(n) \gg L$ from Theorem 3.4) is a separate construction: the inf is taken directly over the given compact set $[-B, B]$ at each step, so no fresh Sion or coercivity argument is needed. The classical AFW result is that $V_n \to \mathbb{E}[g(G(1))]$ as $n \to \infty$, where $G(t) = \exp(\sqrt{c} W(t) - ct/2)$ is geometric Brownian motion.

2.2. One-step equivalence and backward induction

Definition 2.2 (Continuation value)**.** Let $\mathcal{T}_n = \{\mu \in \mathcal{P}([-\zeta_n, \zeta_n]) : \mathbb{E}_\mu[T] = 0,\ \mathbb{E}_\mu[T^2] \leq e^{c/n} - 1\}$. Set $V(s, n) = g(s)$ and $V(s, m) = \sup_{\mu \in \mathcal{T}_n} \mathbb{E}_{T\sim\mu}[V(s(1+T), m+1)]$ for $m = n-1, \ldots, 0$; the player's action is absorbed under the zero-mean constraint (Proposition 2.4).

Lemma 2.3 (Regularity of continuation values)**.** *For each m: (a) $V(\cdot, m)$ is convex and nondecreasing; (b) $V(\cdot, m)$ is L-Lipschitz, with the* same *constant as g. The Lipschitz preservation follows from the martingale identity $\mathbb{E}[T] = 0 \Rightarrow \mathbb{E}[1+T] = 1$; dropping zero-mean permits $\mathbb{E}[1+T] > 1$ and causes exponential growth (the mechanism behind Theorem 3.4).*

Proposition 2.4 (Game–transport equivalence, $n = 1$)**.** *For state $s > 0$ and loss ℓ convex L-Lipschitz, let $\mathcal{K}(s) := \{\nu \in \mathcal{P}([s(1-\zeta_n), s(1+\zeta_n)]) : \mathbb{E}_\nu[Y] = s,\ \mathbb{E}_\nu[(Y/s - 1)^2] \leq e^{c/n} - 1\}$ be the martingale+variance-constrained set of one-step target measures, and $\mathcal{K}'(s)$ the same set with the mean constraint $\mathbb{E}_\nu[Y] = s$ dropped. Then $V_1 = \inf_{\Delta \in \mathbb{R}} \sup_{\nu \in \mathcal{K}'(s)} \mathbb{E}_\nu[\ell(Y) - (Y/s - 1)\Delta\, s]$ equals $V^{\mathrm{OT}}(s) := \sup_{\nu \in \mathcal{K}(s)} \mathbb{E}_\nu[\ell(Y)]$. Moreover $\Delta^* = h^*$, the Lagrange multiplier for $\mathbb{E}_\nu[Y] = s$.*

Theorem 2.5 (Backward Induction Bridge)**.** *Fix $n \in \mathbb{N}$.*

For each step m and state $s > 0$: (a) *$V(s,m)$ equals the one-step constrained transport value with loss $V(\cdot, m{+}1)$ at source s;* (b) *$\Delta^*_m(s) = h^*_m(s)$, depending only on s;* (c) *under $|\Delta_m(s)| \le B$, the action-constrained value at step m equals the L_1-relaxed transport value with step-dependent saturation gap*

$$\delta_m \;=\; (L^B_{m+1}(n) - B)_+ \, s\, h_n,$$

where $L^B_{m+1}(n)$ is the Lipschitz constant of the continuation given by Theorem 3.4 (terminal $\Lambda = L$). The non-circular reading order is (a) and Theorem 2.6 first, with (c) specialized at $\Lambda = L^B_{m+1}(n)$ from Theorem 3.4's recursion.

Theorem 2.6 (Action-space–relaxation duality)**.** *Fix loss ℓ convex L-Lipschitz. The action-constrained one-step game $V_1^B := \inf_{|\Delta|\le B} \sup_{\nu\in\mathcal{K}'(s)}[\mathbb{E}_\nu[\ell(Y)] - \Delta(\mathbb{E}_\nu[Y]-s)]$ and the L_1-relaxed transport $V_1^\delta := \sup_{\nu:|\mathbb{E}_\nu[Y]-s|\le\delta} \mathbb{E}_\nu[\ell(Y)]$ are related by the Fenchel conjugacy $V_1^{\delta(B)} = V_1^B + B\,\delta(B)$, where $\delta(B) := |\mathbb{E}_{\nu^*(B)}[Y] - s|$ is the primal martingale defect at the optimizer $\nu^*(B)$. Equivalently, $B = |h^*(\delta(B))|$ for any selection $h^*(\delta) \in -\partial V_1^\delta(\delta)$ (single-valued a.e.; the multi-valued kink behavior is detailed in Remark 2.7).*

Remark 2.7 (Properties of the correspondence). $B(\delta)$ is continuous, nonincreasing, and vanishes for $\delta \ge \delta_{\mathrm{sat}}$. The inverse $\delta(B)$ is multi-valued at kinks of V_1^δ (single-valued a.e., a homeomorphism for strictly convex ℓ) and depends on the loss's Lipschitz constant Λ. Two related quantities appear: the L_1-budget $\delta(B) := |\mathbb{E}_{\nu^*(B)}[Y] - s|$ (equal to sh_n at the saturating Dirac, from Theorem 2.6); and the per-step value-gap $\delta(B;\Lambda) := (\Lambda - B)_+ \, s\, h_n = (\Lambda - B)_+ \, \delta(B)$ used in Theorem 2.5(c) with $\Lambda = L^B_{m+1}(n)$ from Theorem 3.4.

3. Exact Lipschitz Degradation Under Action Bounds

Definition 3.1 (Budget-constrained backward induction)**.** Fix $c > 0$, $B > 0$, and $g : [0,\infty) \to \mathbb{R}_{\ge 0}$ convex, nondecreasing with attained Lipschitz constant $L := \mathrm{Lip}(g) > B$ (equivalently $\lim_{s\to\infty} g'_+(s) = L$, by Lemma 3.3(i)). Write $h_n := \sqrt{e^{c/n} - 1}$ and assume $\zeta_n \in [h_n, 1)$ with $\zeta_n \to 0$. Let $\widetilde{\mathcal{T}}_n = \{\mu \in \mathcal{P}([-\zeta_n, \zeta_n]) : \mathbb{E}_\mu[T^2] \le e^{c/n} - 1\}$ and

$$\begin{aligned} V^B_n(s,n) &= g(s), \\ V^B_n(s,m) &= \sup_{\mu\in\widetilde{\mathcal{T}}_n} \Big\{ \mathbb{E}_\mu[V^B_n(s(1{+}T), m{+}1)] \\ &\qquad - Bs|\mathbb{E}_\mu[T]| \Big\}. \end{aligned} \tag{2}$$

$L^B_m(n) := \sup_{s\ne s'} |V^B_n(s,m) - V^B_n(s',m)|/|s - s'|$ denotes the Lipschitz constant at step m. V^∞_n denotes the same recursion with $B = \infty$, recovering the unconstrained V of Definition 2.2.

Remark 3.2 (Calibration scope). We work in the regime $\zeta_n \in [h_n, 1)$ with $\zeta_n \to 0$, encompassing both standing $\zeta_n = h_n$ and AFW Lindeberg $\zeta_n \sim \sqrt{K \log n / n}$ ($K \ge 16c$). The Jensen drift cap

$$|\mathbb{E}_\mu[T]| \;\le\; \sqrt{\mathbb{E}_\mu[T^2]} \;\le\; h_n \tag{3}$$

caps drift before the support cone ζ_n binds, so within this scope Theorems 3.4 and 4.2 are calibration-independent: the Lagrangian secant-slope recurrence is attained by δ_{h_n} regardless of ζ_n, while the pathwise rate of Theorem E.4 is attained by a symmetric 3-point martingale that does use the wider support. The exact identity Theorem 4.4 and Proposition D.4 require the tighter $\zeta_n = h_n$ (Markov on $Z_n = s\alpha_n - S_n$ needs $|T_m| \le h_n$ pointwise). By contrast, Theorem E.4's pathwise rate $\Theta((1+\zeta_n)^n)$ tracks support; this variance-vs-support dichotomy reflects that the Lagrangian is a drift-penalized expectation while the pathwise is an ess sup over paths. Appendix G maps each theorem's calibration.

Lemma 3.3 (Asymptotic slope and affine lower bound)**.** *If g is convex, nondecreasing with $\mathrm{Lip}(g) = L > 0$ (attained), then: (i) the right derivative $g'_+(s)$ is nondecreasing with $\lim_{s\to\infty} g'_+(s) = L$; (ii) let $\beta_\infty := \inf_{s\ge 0}[g(s) - Ls] \in [-\infty, g(0)]$; then $g(s) \ge Ls + \beta_\infty$ for all $s \ge 0$, and writing $C_g := g(0) - \beta_\infty \ge 0$ (finite when $\beta_\infty > -\infty$), one has $g(s) \ge Ls - C_g + g(0)$ (which simplifies to $g(s) \ge Ls - C_g$ when $g(0) = 0$); (iii) if moreover g is* finitely *piecewise linear (finitely many breakpoints), there exists $s^* < \infty$ with $g'_+(s) = L$ for $s \ge s^*$ and $g(s) = Ls + \beta_\infty$ on $[s^*, \infty)$ (Lipschitz constant attained on the final piece; β_∞ finite).*

Theorem 3.4 (Exact Lipschitz degradation)**.** *Under Definition 3.1, with $h_n := \sqrt{e^{c/n} - 1}$ and $\alpha_n := (1 + h_n)^n$,*

$$L^B_m(n) \;=\; B + (L - B)(1 + h_n)^{n-m}, \quad m \in \{0, \dots, n\}. \tag{4}$$

In particular $L^B_0(n) = B + (L - B)\,\alpha_n$ grows as $e^{\Theta(\sqrt{cn})}$ for every non-constant convex nondecreasing L-Lipschitz g. The recurrence

$$\Lambda_m \;=\; \Lambda_{m+1}(1 + h_n) - Bh_n, \qquad \Lambda_n = L, \tag{5}$$

is attained at the variance-saturating Dirac $\mu = \delta_{h_n}$ via a secant-slope bound on convex V^B_n; curvature plays no role since per-step drift $\Theta(h_n) = \Theta(1/\sqrt{n})$ dominates the Jensen curvature term $\Theta(1/n)$.

Corollary 3.5 (Finite-s grid threshold)**.** *For piecewise-linear g with affine tail on $[s^*, \infty)$, $(V^B_n)'_+(s,m) = L^B_m(n)$ exactly for all $s \ge s^*/(1 - h_n)^{n-m}$. Full statement and proof: Appendix F.*

4. Premium Bounds and the HJB Failure

Definition 4.1. The *premium* is $\Phi^B_n(s) := V^B_n(s,0) - V^\infty_n(s,0) \ge 0$, where V^∞_n is the unconstrained backward induction.

Theorem 4.2 (Tight premium bounds)**.** *Under Definition 3.1 with $g(0) = 0$ and $\alpha_n := (1 + h_n)^n$: for all $s > 0$,*

$$\Phi_n^B(s) \;\le\; (L - B)(\alpha_n - 1)\, s + C_g, \tag{6}$$

$$\Phi_n^B(s) \;\ge\; g(s\,\alpha_n) - Bs(\alpha_n - 1) - V_n^\infty(s, 0), \tag{7}$$

with $C_g = g(0) - \inf_{s>0}[g(s) - Ls]$ depending only on g. For piecewise-linear g (affine tail), both bounds are $\Theta((L - B)\, e^{\sqrt{cn}}\, s)$, so $\Phi_n^B(s) = \Theta(e^{\sqrt{cn}})$.

Corollary 4.3 (Pathwise structural lower bound against L-bounded hindsight comparators)**.** *For convex nondecreasing L-Lipschitz g, define the per-path-hindsight loss gap of a B-bounded strategy Δ_B^* against the L-bounded hindsight comparator class (deterministic S_n; the expectation in (1) collapses):*

$$\mathcal{R}_n^B(s) \;:=\; \sup_{S_n}\big[\ell_n(\Delta_B^*; S_n) - \inf_{|\Delta'_m| \le L} \ell_n(\Delta'; S_n)\big].$$

Then

$$\mathcal{R}_n^B(s) \;\ge\; (L{-}B)\, s\, (\alpha_n - 1) \;=\; \Omega(e^{\sqrt{cn}}),$$

with aggregate variant $\mathcal{R}_n^{\mathrm{agg},B}(s) = \Omega(e^{(L-B)\sqrt{cn}/L})$ (Remark B.7). Against a B'-bounded comparator ($B < B' < L$), $\mathcal{R}_n^{B,B'}(s) \ge (B'{-}B)\, s\, (\alpha_n{-}1)$, tight on the pure-drift path via $\Delta_m \equiv B$. The witness is the Dirac-h_n adversary path against the hindsight L-comparator $\Delta'_m \equiv L$; details and the B'-comparator matching upper bound in Remark D.2.

Theorem 4.4 (Exact remainder for piecewise-linear payoffs)**.** *Suppose $\zeta_n = h_n$ (tight support) and g is piecewise linear with affine tail $g(x) = Lx + \beta_\infty$ on $[s^*, \infty)$, $g'_+(s^*) = L$. Let $M_\psi := g(0) - \beta_\infty$. With $g_\infty(s) := Ls + \beta_\infty$, the backward recursion (2) admits the closed form*

$$V_n^B(s, m) \;=\; (L{-}B)(\alpha_{n-m} - 1)\, s + g_\infty(s)$$

at every $m \in \{0, \dots, n-1\}$ and every s with $s\alpha_{n-m} > s^* + M_\psi/(L - B)$. In particular at $m = 0$, the remainder $R_n(s) := V_n^B(s, 0) - [(L{-}B)(\alpha_n{-}1)s + g_\infty(s)]$ vanishes identically.

Remark 4.5 (Failure of the HJB limit). Barles–Souganidis (Barles & Souganidis, 1991) viscosity convergence requires a *finite* limiting operator F. At the standard normalization the Lagrangian's drift residual $\sqrt{cn} \to \infty$ per unit time admits no such F, so no finite Barles–Souganidis HJB limit exists for V_n^B. Alternative renormalizations (notably V_n^B/α_n), nonlocal limits, and gradient-constrained variational inequalities are not ruled out; continuous-time pathwise-constrained robust hedging (Cvitanić & Karatzas, 1993; Soner & Touzi, 2002) is diffusive ($O(\Delta t)$ drift) and finite-valued, and the rescaled limit is the natural open direction.

Pathwise variant. A companion ess-sup formulation $V_n^{P,B}$ with martingale adversary admits the closed-form rate $V_n^{P,B}(s) = (L{-}B)s(1 + \zeta_n)^n + Bs + \beta_\infty + O_g(1)$ at any $\zeta_n \in [h_n, 1)$ (Theorem E.4, proof in Appendix E.3). At standing $\zeta_n = h_n$ this matches the Lagrangian's $\Theta(e^{\sqrt{cn}})$; under wider Lindeberg the pathwise rate tracks support ζ_n while the Lagrangian tracks variance h_n, and the two separate accordingly.

5. Aggregate-Budget Rate

Replacing the per-step bound by the aggregate constraint $\sum_{m=1}^n |\Delta_m| \le nB$ couples steps and breaks the backward decomposition; we instead bound the value pathwise. For piecewise-linear g with attained Lipschitz L, matching primal lower (burst-linear adversary) and upper (deterministic primal $\hat{\Delta}_m = L \cdot \mathbf{1}[m > n - N^*]$ with $N^* := \lfloor nB/L \rfloor$) bounds yield

$$\Phi_n^{\mathrm{agg},B}(s) \;=\; Ls \cdot e^{(L-B)\sqrt{cn}/L - (L-B)c/(2L)}(1 + o(1)),$$

so the aggregate exponent coefficient $(L{-}B)/L < 1$ separates strictly from the per-step exponent 1 of Theorem 4.2. The constant-λ Lagrangian dual at $\lambda_n^* = h_n s(1 + h_n)^{K^*}$ with $K^* = n - \lfloor nB/L \rfloor - 1$ sharpens the lower bound to an exact constant; zero strong-duality gap holds under the strengthened threshold $s(1{-}\zeta_n)^n \ge s^*$ but provably fails under the weaker stated threshold for $B \ge L/2$ (Proposition B.9). Full statements (definition, weaker $\sqrt{n}$/polynomial-rate lower bounds, burst-linear LB, matching deterministic UB, strong-duality dichotomy) and proofs are in Appendix B.

6. Experiments

We validate the bridge and rate predictions on the European call ($K = 1$, $c = 0.1$, $s_0 = 1$, $\zeta_n = h_n$). *Closed-form, payoff-uniform* checks: N1 (deterministic UB and burst-linear LB at $n \le 5000$ both within 10^{-4} of Remark B.7's asymptotic identity at $n = 5000$, Figure 1); N3 (pathwise UB saturation, gap $62\% \to 3\%$ from n=10 to 500); N7 (brute-force 3^n enumeration confirms Theorem E.4 to machine precision through $n = 15$). *Solver-based* checks: N4 and E4c–e confirm the closed-form Lipschitz formula on the European call within the operational $n \approx 30$–50 ceiling. Baselines E1, E2, E4 (binomial DKM/XDX (DeMarzo et al., 2006; 2016; Xue et al., 2023), MOT) and the SPX/ABFW grounding of Corollary 4.3 are in Appendix H. Single CPU; no GPU.

Related work. AFW (Abernethy et al., 2012) posed bounded-capital as Problem 2; ABFW (Abernethy et al., 2013) resolves only the unbounded Problem 1 via asymptotic Black–Scholes-delta hedging, and Δ_{BS} exceeds any $B < L$ in deep-ITM. DKM/XDX (DeMarzo et al., 2006;

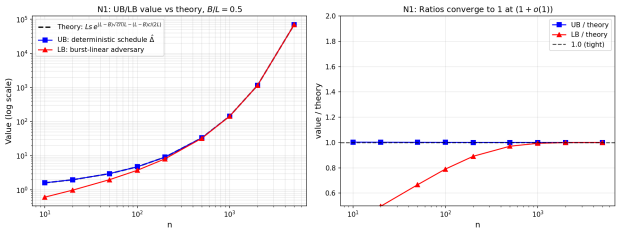

Figure 1. **N1.** Deterministic-schedule UB, burst-linear LB, and theory at $B/L{=}0.5$; ratios converge to 1 at rate $(1 + o(1))$.

2016; Xue et al., 2023) give $O(C_1 e^{\sqrt{cn}})$ with unspecified C_1; Theorem 3.4 pins $C_1 = (L{-}B)\,s$. GHLT (Galichon et al., 2014) supplies the n-stage DP and CKT (Cheridito et al., 2017) the per-step Fenchel–Moreau pair, coupled across n steps by the closed-form amplification. The construction sits in the broader online-learning tradition for adversarial sequential prediction (Cesa-Bianchi & Lugosi, 2006; Hazan, 2016; Abernethy et al., 2009); full positioning is in Appendix C.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Abernethy, J., Agarwal, A., Bartlett, P. L., and Rakhlin, A. A stochastic view of optimal regret through minimax duality. In *Proceedings of the 22nd Annual Conference on Learning Theory*, 2009.

Abernethy, J., Frongillo, R., and Wibisono, A. Minimax option pricing meets black-scholes in the limit. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 1029–1040, 2012.

Abernethy, J., Bartlett, P. L., Frongillo, R., and Wibisono, A. How to hedge an option against an adversary: Black-Scholes pricing is minimax optimal. In *Advances in Neural Information Processing Systems 26 (NeurIPS 2013)*, pp. 2346–2354, 2013.

Barles, G. and Souganidis, P. E. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis*, 4(3):271–283, 1991.

Cesa-Bianchi, N. and Lugosi, G. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

Cheridito, P., Kupper, M., and Tangpi, L. Duality formulas for robust pricing and hedging in discrete time. *SIAM Journal on Financial Mathematics*, 8(1):738–765, 2017.

Cuturi, M. Sinkhorn distances: Lightspeed computation of optimal transportation distances. In *Advances in Neural Information Processing Systems*, pp. 2292–2300, 2013.

Cvitanić, J. and Karatzas, I. Hedging contingent claims with constrained portfolios. *The Annals of Applied Probability*, 3(3):652–681, 1993.

DeMarzo, P., Kremer, I., and Mansour, Y. Online trading algorithms and robust option pricing. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 477–486, 2006.

DeMarzo, P. M., Kremer, I., and Mansour, Y. Robust option pricing: Hannan and Blackwell meet Black and Scholes. *Journal of Economic Theory*, 163:410–434, 2016.

Dolinsky, Y. and Soner, H. M. Robust hedging with proportional transaction costs. *Finance and Stochastics*, 18(2): 327–347, 2014.

Duchi, J., Shalev-Shwartz, S., Singer, Y., and Chandra, T. Efficient projections onto the ℓ_1-ball for learning in high dimensions. In *International Conference on Machine Learning*, 2008.

Galichon, A., Henry-Labordère, P., and Touzi, N. A stochastic control approach to no-arbitrage bounds given marginals. *Annals of Applied Probability*, 24(1):312–336, 2014.

Hazan, E. *Introduction to Online Convex Optimization*, volume 2 of *Foundations and Trends in Optimization*. Now Publishers, 2016.

Karatzas, I. Probabilistic aspects of finite-fuel stochastic control. *Proceedings of the National Academy of Sciences*, 82(17):5579–5581, 1985.

Karatzas, I., Ocone, D., Wang, H., and Zervos, M. Finite-fuel singular control with discretionary stopping. *Stochastics and Stochastics Reports*, 71(1–2):1–50, 2000.

Rockafellar, R. T. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1970.

Sion, M. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.

Soner, H. M. and Touzi, N. Dynamic programming for stochastic target problems and geometric flows. *Journal of the European Mathematical Society*, 4(3):201–236, 2002.

Xue, S., Du, Y., and Xu, L. Tighter robust upper bounds for options via no-regret learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 5348–5356, 2023.

A. Notation

B. Aggregate-budget rate

This section gathers the full definitions, propositions, theorems, and remarks underlying the per-step–vs.–aggregate dichotomy summarized in Section 5; proofs follow in Appendix E.6.

Definition B.1 (Aggregate-budget game value)**.** Under the same setup as Definition 3.1 but replacing the per-step bound $|\Delta_m| \le B$ with the aggregate constraint $\sum_{m=1}^{n} |\Delta_m| \le nB$, the aggregate-budget game value is $V_n^{\mathrm{agg},B}(s) := \inf_{\Delta:\ \sum_m |\Delta_m| \le nB} \sup_{S_n} \ell_n(\Delta, S_n)$, with $\Phi_n^{\mathrm{agg},B}(s) := V_n^{\mathrm{agg},B}(s) - V_n^{\infty}(s,0)$.

The aggregate variant arises naturally as a Lagrangian relaxation of the per-step constraint and as a discrete analog of finite-fuel singular control (Karatzas, 1985; Karatzas et al., 2000) or transaction-cost-budgeted hedging (Dolinsky & Soner, 2014). The per-step bound $|\Delta_m| \le B$ decomposes via backward induction; the aggregate bound $\sum_m |\Delta_m| \le nB$ couples steps and breaks decomposition. For this section, fix g finitely piecewise-linear with attained Lipschitz constant L, affine on $[s^*, \infty)$ as $g(s) = Ls + \beta_\infty$ (e.g., European call; Lemma 3.3(iii)). The burst-linear lower bound (Theorem B.5) and deterministic primal upper bound (Theorem B.6) together give the tight rate (Remark B.7). Theorems B.3 and B.4 give weaker $\sqrt{n}$/polynomial lower bounds with sharper pre-asymptotic constants; the burst-linear bound asymptotically subsumes both.

Proposition B.2 (Per-step dominates aggregate)**.** $V_n^{\mathrm{agg},B}(s,0) \le V_n^{\mathrm{per-step},B}(s,0)$, *hence* $\Phi_n^{\mathrm{agg},B}(s) \le \Phi_n^{\mathrm{per-step},B}(s) = \Theta((L-B)\, e^{\sqrt{cn}}\, s)$ *for piecewise-linear g (Theorem 4.2).* (Pathwise per-step feasibility implies pathwise aggregate feasibility; monotonicity of the infimum.)

Theorem B.3 ($\sqrt{n}$ lower bound)**.** *Under Definition 3.1's calibration $\zeta_n \ge h_n$ (so that the state-adaptive Dirac $\epsilon_m^* = h_n s/S_{m-1} \in [-\zeta_n, \zeta_n]$ is feasible): for every $B \in (0, L)$, $s \ge s^*$, and n large enough that $s(1+nh_n) \ge s^*$: $\Phi_n^{\mathrm{agg},B}(s) \ge (L-B)\, s\, \sqrt{cn} + O(1)$. (State-adaptive Dirac $\epsilon_m = h_n s/S_{m-1}$ with constant dollar-drift; full proof in Appendix E.6.)*

Theorem B.4 (Polynomial lower bound)**.** *For every $B \in (0, L)$, $s \ge s^*$, and $n \ge n_1(L,B,c) := \max\{c/\log^2 2,\ \exp(2B^2/(L(L-B)))/c\}$: $\Phi_n^{\mathrm{agg},B}(s) \ge C(L,B,s,c) \cdot n^{L/(2B)}/\log(cn)$. (Power-law Dirac $\epsilon_m^{(\kappa)} = h_n (s/S_{m-1})^{\kappa}$ with $\kappa \in (B/L, 1)$, Taylor–Grönwall, optimize κ; full proof in Appendix E.6.)*

Theorem B.5 (Burst-linear, super-polynomial lower bound)**.** *For every $B \in (0, L)$, $s \ge s^*$, and n large enough that $S_K \ge s^*$ with $K = \lceil (L-B)n/L \rceil$:*

$$\Phi_n^{\mathrm{agg},B}(s) \;\ge\; L\, s\, \exp\!\Big(\tfrac{L-B}{L}\sqrt{cn} - \tfrac{(L-B)c}{2L}\Big) + O(1) \;=\; \Theta\big(e^{(L-B)\sqrt{cn}/L}\big).$$

In particular, $\Phi_n^{\mathrm{agg},B}$ is super-polynomial*: no bound $C \cdot n^p \cdot \mathrm{polylog}(n)$ with fixed p can upper-bound it.*

Theorem B.6 (Matching sub-exponential upper bound)**.** *For any L-Lipschitz g with $g(0)$ finite, any $n \ge 1$, and $0 < B < L$,*

$$V_n^{\mathrm{agg},B}(s) \;\le\; g(0) + L\, s\, (1+h_n)^{n-N^*}, \qquad N^* := \lfloor nB/L \rfloor.$$

In leading order, $V_n^{\mathrm{agg},B}(s) \le Ls \cdot e^{(L-B)\sqrt{cn}/L - (L-B)c/(2L)}(1+o(1)) + g(0)$. Achieved by the deterministic schedule $\hat{\Delta}_m = L \cdot \mathbf{1}[m > n - N^*]$ (no hedging in Phase 1, L shares per step in Phase 2; $\sum_m |\hat{\Delta}_m| = LN^* \le nB$).

Remark B.7 (Tight aggregate rate). Theorems B.5 and B.6 together *tightly* pin down the aggregate premium:

$$\Phi_n^{\mathrm{agg},B}(s) \;=\; L\, s\, \exp\!\Big(\tfrac{L-B}{L}\sqrt{cn} - \tfrac{(L-B)c}{2L}\Big)\,(1+o(1)),$$

The per-step (exponent 1) and aggregate ($(L-B)/L$) exponents separate cleanly. Matching is primal-side at the rate level (feasible adversary LB, feasible deterministic player UB); the constant-λ Lagrangian dual sharpens the LB to an exact constant for piecewise-linear g (Proposition B.9; matching UB conjectural, smooth-g open).

Remark B.8 (Lagrangian dual reformulation). Weak duality on the aggregate constraint is immediate: for every $\lambda \ge 0$,

$$V_n^{\mathrm{agg},B}(s) \;\ge\; \sup_{\mu \in \mathcal{M}_\lambda} \mathbb{E}_\mu[g(S_n)] \;-\; \lambda\, nB,$$

where

$$\mathcal{M}_\lambda \;:=\; \big\{(\mu_m) : \mu_m \in \widetilde{\mathcal{T}}_n,\ |S_{m-1}\, \mathbb{E}[T_m | \mathcal{F}_{m-1}]| \le \lambda \text{ a.s., all } m\big\}.$$

Strong duality (equality after $\sup_{\lambda\geq 0}$) is a *conjecture*: the aggregate constraint couples stages, so Sion stage-by-stage does not directly apply. The constraint localizes the adversary's drift to a *dollar* cap rather than return cap. The continuous-time analog $\mathrm{d}S_t = \mu(t)\,\mathrm{d}t + \sqrt{c}\,S_t\,\mathrm{d}W_t$ with $\int_0^T |\mu(t)|\,\mathrm{d}t \leq \Lambda_{\text{fuel}}$ fits the finite-fuel singular control framework (Karatzas, 1985; Karatzas et al., 2000). The drift residual $\sqrt{cn} \to \infty$ blocks naive HJB convergence (cf. Remark 4.5). Proposition B.9 sharpens this for piecewise-linear g along (a) lower bound, (b) matching upper bound under strengthened threshold, (c) counterexample under weak threshold; the smooth-g case remains open.

Proposition B.9 (Sharpened constant-λ Lagrangian dual)**.** *For piecewise-linear g with affine tail $g(s) = Ls + \beta_\infty$ on $[s^*, \infty)$, $0 < B < L$, define $K^* := n - \lfloor nB/L \rfloor - 1$, $r := nB - L\lfloor nB/L \rfloor \in [0, L)$, $\lambda_n^* := h_n s(1+h_n)^{K^*}$. Then for n large enough that $s(1+h_n)^{K^*} \geq s^*$:*

$$V_n^{\text{agg},B}(s) \;\geq\; D(\lambda_n^*) \;=\; L\,s\,(1+h_n)^{K^*}\big[1 + (1-r/L)h_n\big] + \beta_\infty,$$

with $D(\lambda) := \sup_{\mu\in\mathcal{M}_\lambda} \mathbb{E}_\mu[g(S_n)] - \lambda nB$. The λ_n^* scaling $\Theta(sh_n e^{(L-B)\sqrt{cn}/L})$ corrects the conjectured $\Theta((L-B)h_n s)$ in Remark B.8 by an exponential factor; the binding constraint is the Phase-2 plateau $\max_m |S_{m-1}\mathbb{E}[T_m|\mathcal{F}_{m-1}]| = h_n s(1+h_n)^{K^*}$. Proof of (a): *use the burst-linear adversary μ^{bl} of Theorem B.5 but at Phase-1 length K^* (one less than $\lceil (L-B)n/L \rceil$; the latter is sub-optimal here by a factor $1 - rh_n^2/L$). Then $\mu^{\text{bl}} \in \mathcal{M}_{\lambda_n^*}$ and the player's ℓ_1/ℓ_∞ best response $nB \cdot h_n s(1+h_n)^{K^*}$ yields $\mathbb{E}_{\mu^{\text{bl}}}[g(S_n)] - \lambda_n^* nB$ equal to the closed form.*

(b) Matching upper bound under strengthened threshold. *Let $\psi(x) := g(x) - Lx - \beta_\infty \geq 0$ ($M_\psi := \sup\psi = g(0) - \beta_\infty$, supported on $[0, s^*)$). If $s(1-\zeta_n)^n \geq s^*$ (so the worst-case down-trajectory $S_n \geq s(1-\zeta_n)^n$ stays in the affine tail pathwise), then $V_n^{\text{agg},B}(s) \leq D(\lambda_n^*)$, hence equality with (a). Proof: the deterministic schedule $\hat{\Delta}_m := L\mathbf{1}[m > K^*{+}1] + r\mathbf{1}[m = K^*{+}1]$ saturates $\sum|\hat{\Delta}_m| = LN^*{+}r = nB$ exactly; pathwise telescoping $\sum \hat{\Delta}_m T_m S_{m-1} = L(S_n - S_{K^*+1}) + r(S_{K^*+1} - S_{K^*})$ gives $\ell_n = \beta_\infty + \psi(S_n) + (L-r)S_{K^*+1} + rS_{K^*}$; the threshold forces $\psi(S_n) = 0$ a.s.; Jensen $\mathbb{E}[S_m] \leq s(1+h_n)^m$ (from $|\mathbb{E}[T_m|\mathcal{F}_{m-1}]| \leq h_n$ via Cauchy–Schwarz on the variance bound) closes the residual to zero, sharpening Theorem B.6's $rh_n s(1+h_n)^{K^*} + M_\psi$ excess exactly.*

(c) Counterexample under weak threshold alone. *If $s(1+h_n)^{K^*} \geq s^*$ but $s\,(1+h_n)^{K^*+1}(1-h_n)^{N^*} < s^*$ (the up-then-down trajectory falls below s^*; this regime is non-empty asymptotically for $B \geq L/2$, where $(K^*{+}1{-}N^*)h_n - O(h_n^2) \leq 0$, but vacuous for $B < L/2$), equality fails. The variance-saturating Dirac adversary $T_m \equiv h_n$ on $\{1, \ldots, K^*{+}1\}$ then $T_m \equiv -h_n$ on $\{K^*{+}2, \ldots, n\}$ (variance h_n^2 each step, atoms at $\pm h_n \in [-\zeta_n, \zeta_n]$, feasible at any $\zeta_n \in [h_n, 1)$) drives $S_n = s(1+h_n)^{K^*+1}(1-h_n)^{N^*} < s^*$; against $\hat{\Delta}$, $\ell_n = D(\lambda_n^*) + \psi(S_n)$ with gap $\Theta(M_\psi)$, refuting (b). Quantitatively, $\liminf_n[V_n^{\text{agg},B}(s) - D(\lambda_n^*)] \geq \psi(0) = M_\psi$ for $B > L/2$ (the up-then-down adversary drives $S_n \to 0$). Since the explicit witness places mass at $\pm h_n$ rather than at the support endpoint $\pm\zeta_n$, the counterexample carries through under wider Lindeberg calibration unchanged; the dichotomy with (b) is therefore intrinsic to the threshold gap, not to standing-versus-Lindeberg calibration.*

Numerical verification (`code/verify_path5b_upper_bound.py`): under the weak threshold (c), the adversary's value exceeds $D(\lambda_n^*)$ by 0.12–0.35 across $n \in \{8, 10, 12\}$, $B/L \in \{1/2, 3/5\}$ (here $g(x) = (x-s^*)_+$, $s = s^*$); under the strengthened threshold (b), equality holds to machine precision at $n \in \{5, 10, 15, 20\}$.

C. Related work

We extend AFW (Abernethy et al., 2012)/ABFW (Abernethy et al., 2013) and online learning (Cesa-Bianchi & Lugosi, 2006; Hazan, 2016; Abernethy et al., 2009) to finite-n, bounded-capital. ABFW's $\Delta_{\text{BS}}{=}N'(d_1)$ exceeds any $B < L$ in deep-ITM, so projection onto $|\Delta| \leq B$ incurs $\Omega(e^{\sqrt{cn}})$ excess loss against the L-bounded comparator (Corollary 4.3); DKM bounds (DeMarzo et al., 2006; 2016; Xue et al., 2023) give $O(C_1\,e^{\sqrt{cn}})$ for unspecified C_1, while Theorem 3.4 pins the explicit constant $(L{-}B)\,s$. GHLT (Galichon et al., 2014) supplies the n-stage DP and CKT (Cheridito et al., 2017) per-step Fenchel–Moreau, with our closed-form amplification coupling them across n steps; the aggregate dual mirrors finite-fuel singular control (Karatzas, 1985). Smooth-g aggregate strong duality, a Lindeberg exact identity, and K-step rolling-budget interpolation remain open (Proposition B.9 and Remark D.3); the solver is validated at $n \lesssim 30$.

D. Auxiliary results moved from body

For space we collect auxiliary statements referenced by the body but whose detail is not load-bearing for the headline results.

Critical action-space bound.

Proposition D.1 (Critical action-space bound)**.** *For convex nondecreasing g with attained Lipschitz constant $L := \mathrm{Lip}(g)$, under the standing calibration $\zeta_n = h_n$ (used by all downstream per-step theorems), $\Phi_n^B := V_n^B - V_n^\infty = 0$ for all $B \geq L$ and all n. The conclusion is conjectured to extend to general $\zeta_n \in [h_n, 1)$ via Sion duality (Remark E.1), though this extension is not used downstream.*

Intuition. For $B \geq L$ the player's hedge cost dominates any drift benefit the adversary can extract, so the adversary's optimal $a_\mu = 0$ (martingale). The formal proof in Appendix E.1 uses a sharp convex-chord upper bound rather than the loose $Ls|a_\mu|$ Lipschitz bound (which fails for symmetric μ, where $a_\mu = 0$ but $\mathbb{E}_\mu|T|$ is nonzero).

Computational realization. The one-step transport is solved by entropic Sinkhorn (Cuturi, 2013) + proximal gradient ascent on the dual h with Moreau L_∞-projection (Duchi et al., 2008); total cost $O(nq^2K)$ for q-point grid and K PGA iterations per step, with entropic error bounded by Proposition E.2.

Matching regret upper bound.

Remark D.2 (Matching upper bound for B'-comparator regret). The lower bound $\mathcal{R}_n^{B,B'}(s) \geq (B'-B)s(\alpha_n - 1)$ is not tight; the exact minimax regret follows from a two-phase backward Bellman. Define $\Lambda_n^R := 0$ and recursively for $m = n-1, \ldots, 0$:

$$\Lambda_m^R \;=\; \begin{cases} \Lambda_{m+1}^R + B'h_n & \text{if } \Lambda_{m+1}^R \leq B \quad \text{(linear phase)}, \\ (1+h_n)\Lambda_{m+1}^R + h_n(B'-B) & \text{otherwise} \quad \text{(geometric phase)}, \end{cases}$$

giving $\mathcal{R}_n^{B,B'}(s) = \Lambda_0^R \cdot s$. With $k^\dagger := \lfloor B/(B'h_n) \rfloor + 1$, the closed form is

$$\mathcal{R}_n^{B,B'}(s) \;=\; \begin{cases} nB'h_n s & n \leq k^\dagger, \\ \left[(1+h_n)^{n-k^\dagger}(k^\dagger B'h_n + B' - B) - (B'-B)\right] s & n > k^\dagger, \end{cases}$$

asymptotically $\Theta(B'\, e^{-B/B'}\, s\, \alpha_n)$, recovering $L(\alpha_n - 1)s$ at $B = 0$, $B' = L$. Player optimal: deterministic $\Delta_m^* = \mathrm{clip}(\Lambda_{m+1}^R, [-B, B])$; adversary optimal: pure-drift Dirac $T_m \equiv h_n$. Verified at $n = 2$, $L = 1$, $B = 0.5$, $B' = 1$, $h_n = 0.1$: predicted 0.20 matches brute-force. The asymptotic ratio $B'\, e^{-B/B'}/(B'-B) > 1$ shows the $(B'-B)$ lower bound is loose.

Lindeberg scope of the boundary-layer identity.

Remark D.3 (Lindeberg scope of Theorem 4.4). The exact identity genuinely requires the tight calibration $\zeta_n = h_n$; under wider Lindeberg $\zeta_n > h_n$ the formula fails in the kink-overlap regime. For an adversary with mean $a < h_n$ saturating variance via the symmetric 2-atom $\{T_1^*, T_2^*\} = \{h_n \mp \sqrt{2h_n(h_n - a)}\}$ centered at h_n (exact LP optimum: KKT yields $T_1^* + T_2^* = 2h_n$), the Jensen kink-region gain scales as $\sqrt{h_n - a}$ while the drift-penalty saving scales linearly as $(h_n - a)$; for small $h_n - a$ the sub-linear gain dominates, so the wider μ strictly beats Dirac at h_n in the Bellman. The provable Lindeberg extension uses the stricter *direct-affine threshold* $s\alpha_{n-m} > s^* \cdot ((1+h_n)/(1-\zeta_n))^{n-m}$, ensuring the entire support of $s(1+T)$ lies in the affine tail of V_{m+1} pointwise. The variance-vs-support dichotomy of Remark 3.2 thus binds at the exact-identity level: Theorems 3.4 and 4.2 extend to all Lindeberg via Jensen drift cap (Appendix G), but Theorem 4.4 does not.

Smooth-tail remainder.

Proposition D.4 (Smooth-tail remainder)**.** *Relax the piecewise-linearity hypothesis of Theorem 4.4 while retaining the calibration $\zeta_n = h_n$, the L-Lipschitz affine-tail convention $\beta_\infty := \lim_{x\to\infty}[g(x) - Lx]$, and the remainder definition $R_n(s)$: for g convex nondecreasing L-Lipschitz with tail residual $\psi(x) := g(x) - Lx - \beta_\infty$, the remainder is $O((s\alpha_n)^{-p})$ for polynomially-decaying ψ (rate $e^{-p\sqrt{cn}}$), doubly-exponential in $\sqrt{cn}$ for exponential ψ, and $O(1/\sqrt{cn})$ for logarithmic ψ. Full statement in Appendix E.*

Bounded payoffs.

Proposition D.5 (Bounded payoffs)**.** *If $g\colon [0,\infty) \to [0, M_g]$ is bounded with $M_g := \sup g < \infty$, then $0 \leq \Phi_n^B(s) \leq M_g$ uniformly in n.*

Theorem D.6 (Bounded-payoff premium: $\Theta(1)$ matching lower bound)**.** *Let $g\colon [0,\infty) \to [0, M_g]$ be nondecreasing and L-Lipschitz with $g(0) = 0$, $s_M := \inf\{x : g(x) \geq M_g\} < \infty$. Fix $s \in (0, M_g/L)$ and $B \in (0, (M_g - Ls)/s_M)$, and set $\eta := M_g - Ls - Bs_M > 0$. For n large enough that $s\alpha_n > s_M$ and $Bs_M h_n < \eta/2$,*

$$\Phi_n^B(s) \;\geq\; M_g - Ls - B(s_M - s) - Bs_M h_n \;\geq\; \eta/2 \;>\; 0,$$

so $\Phi_n^B(s) = \Theta(1)$ with $\liminf_n \Phi_n^B(s) \geq M_g - Ls - B(s_M - s)$. (Hypothesis is outside Definition 3.1 since g need not be convex; (2) stays well-defined and the proof uses a concave envelope on the unconstrained side.)

E. Full Proofs

E.1. Bridge and duality (Lemma 2.3, Propositions 2.4, D.1 and E.2, and Theorems 2.5 and 2.6)

Proof of Lemma 2.3. By backward induction. Substitute $T = Y/s - 1$; the constraints on T define $\mathcal{T}_n$ and are independent of s.

Convexity: $s \mapsto s(1+T)$ is linear; $V(\cdot, m+1)$ is convex by hypothesis; composition, expectation, and supremum preserve convexity.

Lipschitz: For any $\mu_T \in \mathcal{T}_n$, $|\mathbb{E}[V(s(1+T), m+1)] - \mathbb{E}[V(s'(1+T), m+1)]| \leq L|s - s'| \cdot \mathbb{E}[1+T] = L|s - s'|$ since $\mathbb{E}[T] = 0$. Taking sup over μ_T gives the claim.

Nondecreasing: $s \mapsto s(1+T)$ is nondecreasing in s for every $T \geq -1$ (guaranteed by $\zeta_n < 1$), so $s \mapsto V(s(1+T), m+1)$ is nondecreasing if $V(\cdot, m+1)$ is; taking $\mathbb{E}_{\mu_T}$ and $\sup_\mu$ preserves monotonicity. Base case: g nondecreasing by hypothesis. □

Proof of Proposition 2.4. $\mathcal{K}'(s)$ is weakly compact (bounded support, weakly closed variance constraint; Prokhorov). The objective $f(\Delta, \nu) := \mathbb{E}_\nu[\ell(Y)] - \Delta(\mathbb{E}_\nu[Y] - s)$ is continuous in Δ and weakly continuous in ν. For each M, Sion (Sion, 1958) applies on $[-M, M] \times \mathcal{K}'(s)$: $\inf_{|\Delta| \leq M} \sup_\nu f = \sup_\nu \inf_{|\Delta| \leq M} f$. The inner infimum with $d := \mathbb{E}_\nu[Y] - s$ is $\mathbb{E}_\nu[\ell(Y)] - M|d|$; as $M \to \infty$, $-\infty$ for $d \neq 0$, else $\mathbb{E}_\nu[\ell(Y)]$. Thus $\sup_\nu \inf_\Delta = V^{\mathrm{OT}}(s)$. Coercivity: drift benefit $\leq L|d|$, so for $M \geq L$, any ν with $d \neq 0$ contributes $\leq (L - M)|d| \leq 0$, giving $V_1^M = V_1^L$ for $M \geq L$. Stationarity at saddle: $\mathbb{E}_{\nu^*}[Y] = s$, and $\ell(Y) - \Delta^*(Y/s - 1)s = \lambda$ ν^*-a.s., so Δ^* is the Lagrange multiplier h^*. □

Proof of Theorem 2.5. (a) At step m, the sub-problem is a one-step game with loss $V(\cdot, m+1)$, which is convex and L-Lipschitz by Lemma 2.3; Proposition 2.4 applies. (b) From Proposition 2.4; Markov because $\mathcal{T}_n$ depends only on s. (c) The penalized form $V_n^B(s, m) = \sup_{\mu_T}[\mathbb{E}[V_n^B(s(1+T), m+1)] - Bs|\mathbb{E}[T]|]$ is convex in s (convex + linear), enabling Theorem 2.6 at each step. □

Proof of Theorem 2.6. Apply Sion on $[-B, B] \times \mathcal{K}'(s)$: $V_1^B = \sup_{\nu \in \mathcal{K}'(s)}[\mathbb{E}_\nu[\ell(Y)] - B|d(\nu)|]$, with $d(\nu) := \mathbb{E}_\nu[Y] - s$. Let $\nu^*(B)$ attain the sup; set $\delta(B) := |d(\nu^*(B))|$. The MOT dual identification $\Delta^* = h^*$ at the saddle (Proposition 2.4) gives the conjugate-dual reading $|h^*(\delta(B))| = B$.

Upper bound (weak duality), $V_1^\delta \leq V_1^B + B\delta$. For any ν with $|d(\nu)| \leq \delta$: $\mathbb{E}_\nu[\ell(Y)] \leq V_1^B + B\delta$ (add and subtract $B|d(\nu)|$). Take sup.

Lower bound, $V_1^{\delta(B)} \geq V_1^B + B\delta(B)$. $\nu^*(B)$ satisfies $|d(\nu^*)| = \delta(B)$, so feasible for $V_1^{\delta(B)}$; $V_1^{\delta(B)} \geq \mathbb{E}_{\nu^*}[\ell(Y)] = V_1^B + B\delta(B)$. Combining at $\delta = \delta(B)$: equality.

Continuity/surjectivity of $B \mapsto \delta(B)$. By Sion $V_1^B = \sup_{\nu \in \mathcal{K}'(s)}\{\mathbb{E}_\nu[\ell(Y)] - B|d(\nu)|\}$; each summand is affine nonincreasing in B, so V_1^B is convex nonincreasing in B, finite for $B \geq 0$, hence continuous on $(0, \infty)$ with left/right derivatives and $\delta(B) \in -\partial V_1^B(B) \subseteq \mathbb{R}_{\geq 0}$ (subdifferential of convex nonincreasing V_1^B is contained in $(-\infty, 0]$). Surjective onto $[0, \delta_{\max}]$ where $\delta_{\max} := \lim_{B \to 0^+} \delta(B)$ and $\lim_{B \to \infty} \delta(B) = 0$ (tight martingale). □

Proof of Proposition D.1. Backward induction on $m = n, n-1, \ldots, 0$ with invariant $V_n^B(\cdot, m) = V_n^\infty(\cdot, m)$, convex L-Lipschitz nondecreasing.

Base $m{=}n$: $V_n^B(\cdot, n) = V_n^\infty(\cdot, n) = g$.

Inductive step: write $V := V_n^B(\cdot, m{+}1) = V_n^\infty(\cdot, m{+}1)$ (induction hypothesis), convex L-Lipschitz nondecreasing on $[0, \infty)$. Under the standing calibration $\zeta_n = h_n$ (Theorem 4.4 regime; the general $\zeta_n \in [h_n, 1)$ case is handled by the same

argument applied to the chord at the variance-effective range, see Remark E.1), every $\mu \in \widetilde{\mathcal{T}}_n$ has support $\subseteq [-h_n, h_n]$, so the convex-chord bound applies pointwise:

$$V(s(1+T)) \;\le\; \alpha(T)\,V(s+sh_n) + (1-\alpha(T))\,V(s-sh_n), \qquad \alpha(T) := \tfrac{T+h_n}{2h_n} \in [0,1].$$

Taking expectations with $a_\mu := \mathbb{E}_\mu[T]$ and $\bar\alpha := \mathbb{E}_\mu[\alpha(T)] = (a_\mu + h_n)/(2h_n)$:

$$\mathbb{E}_\mu[V(s(1+T))] \;\le\; \bar\alpha\,V(s+sh_n) + (1-\bar\alpha)\,V(s-sh_n) \;=\; \tfrac{1}{2}\big[V(+)+V(-)\big] + \tfrac{a_\mu}{2h_n}\big[V(+)-V(-)\big],$$

where $V(\pm) := V(s \pm sh_n)$ for brevity. Thus the step-m objective is bounded by

$$\mathbb{E}_\mu[V(s(1+T))] - Bs|a_\mu| \;\le\; \underbrace{\tfrac{1}{2}[V(+)+V(-)]}_{=:\,G(0)} + Q\cdot a_\mu - Bs|a_\mu|, \qquad Q := \tfrac{V(+)-V(-)}{2h_n},$$

where $|Q| \le Ls$ by the L-Lipschitz bound on V. For $B \ge L$ the a_μ-dependent terms satisfy $Q\cdot a_\mu - Bs|a_\mu| \le (Q\cdot\mathrm{sign}(a_\mu) - Bs)|a_\mu| \le (Ls - Bs)|a_\mu| \le 0$, with equality at $a_\mu = 0$. Thus $\sup_{\mu\in\widetilde{\mathcal{T}}_n}\{\mathbb{E}_\mu[V(s(1+T))] - Bs|a_\mu|\} \le G(0)$.

The reverse inequality is achieved by the symmetric mean-zero law $\mu_0 := \frac{1}{2}\delta_{h_n} + \frac{1}{2}\delta_{-h_n} \in \mathcal{T}_n \subset \widetilde{\mathcal{T}}_n$, which has $\mathbb{E}_{\mu_0}[V(s(1+T))] = G(0)$ exactly (the chord bound is sharp at the two extreme atoms). And by the same chord argument applied to $V_n^\infty(\cdot, m) = \sup_{\mu\in\mathcal{T}_n}\mathbb{E}_\mu[V(s(1+T))]$ at $a_\mu = 0$, we get $V_n^\infty(s,m) = G(0)$ as well.

Thus $V_n^B(s,m) = G(0) = V_n^\infty(s,m)$. The induction invariant is preserved (convexity, monotonicity, L-Lipschitz are all properties of V_n^∞ by Lemma 2.3). Consequently, $\Phi_n^B = 0$ for $B \ge L$, all n. □

Remark E.1 (Extension to $\zeta_n > h_n$). The conclusion $V_n^B = V_n^\infty$ for $B \ge L$ extends to all $\zeta_n \in [h_n, 1)$ in the def:VB recursion, as can be verified directly via Sion duality on the inf-sup form. The chord-at-$\pm h_n$ argument above is exact for the standing calibration $\zeta_n = h_n$ used in all per-step-budget downstream theorems (Theorems 3.4, 4.2 and 4.4); for $\zeta_n > h_n$ the chord at $\pm h_n$ is invalid for support outside $[-h_n, h_n]$, but the variance constraint $\mathbb{E}_\mu[T^2] \le h_n^2$ forces the maximizing μ to concentrate within an effective spread h_n of its mean (the maximizer is generally not the symmetric two-point law at $\pm h_n$ — e.g., asymmetric mean-zero laws on $\{-h_n^2/\zeta_n, \zeta_n\}$ can achieve strictly larger $\mathbb{E}_\mu[V]$ for kinked V). The full general-ζ_n proof is not used downstream and is omitted for brevity.

Proposition E.2 (Entropic error accumulation)**.** *Let $V^q(s,m)$ denote the continuation value with adversary supported on a q-point grid in $[-\zeta_n, \zeta_n]$, and V_ε^q its entropic regularization (add $\varepsilon\cdot\mathrm{KL}(\cdot\|\mathrm{unif}_q)$). Then $\sup_s |V_\varepsilon^q(s,0) - V^q(s,0)| \le n\varepsilon\log q$. Choosing $\varepsilon = \alpha/(n\log q)$ achieves target accuracy α.*

Proof. Let $e_m := \sup_s |V_\varepsilon^q(s,m) - V^q(s,m)|$ with $e_n = 0$. At each backward step the entropic regularization contributes at most $\varepsilon\log q$ (since $0 \le \mathrm{KL}(\nu\|\mathrm{unif}) \le \log q$), and propagation adds e_{m+1}. Thus $e_m \le e_{m+1} + \varepsilon\log q$; iterating from $e_n = 0$ gives $e_0 \le n\varepsilon\log q$. □

E.2. Lipschitz degradation (Theorem 3.4)

Proof of Lemma 3.3. (i) g'_+ exists everywhere and is nondecreasing ((Rockafellar, 1970, Thm. 24.1)). Since g is nondecreasing, $g'_+ \ge 0$; g L-Lipschitz gives $g'_+ \le L$. Being nondecreasing and bounded, g'_+ has a limit $L_\infty \le L$. $\mathrm{Lip}(g) = \sup_s g'_+(s) = L_\infty$ (monotone); hence $L_\infty = L$.

(ii) Since $g'_+ \le L$, the function $x \mapsto g(x) - Lx$ is nonincreasing; let $\beta_\infty := \lim_{x\to\infty}[g(x) - Lx] = \inf_{x\ge 0}[g(x) - Lx] \in [-\infty, g(0)]$. For any $s \ge 0$: $g(s) - Ls \ge \beta_\infty$, so $g(s) \ge Ls + \beta_\infty$. Writing $C_g := g(0) - \beta_\infty \ge 0$ (finite when $\beta_\infty > -\infty$), this rearranges to $g(s) \ge Ls + g(0) - C_g$, which simplifies to $g(s) \ge Ls - C_g$ when $g(0) = 0$.

(iii) For finitely piecewise-linear g, g'_+ takes finitely many nondecreasing values, with limit L; the final piece has slope L, so there is $s^* < \infty$ with $g'_+(s) = L$ for $s \ge s^*$. On $[s^*, \infty)$, $g(x) = g(s^*) + L(x - s^*) = Lx + [g(s^*) - Ls^*] = Lx + \beta_\infty$, so β_∞ is finite and attained. □

Proof of Theorem 3.4. Write $\Lambda_m := B + (L-B)(1+h_n)^{n-m}$; prove $L_m^B(n) = \Lambda_m$ by backward induction.

Convexity preservation. For fixed μ, $s \mapsto \mathbb{E}_\mu[V_n^B(s(1+T), m+1)] - Bs|\mathbb{E}[T]|$ is convex (expectation of convex + linear); pointwise sup preserves convexity.

Monotonicity preservation. $V_n^B(\cdot, m)$ is nondecreasing. We argue: (1) $V_n^B(0^+, m) = g(0^+)$ for all m; (2) $V_n^B(s, m) \geq g(s) \geq g(0^+)$; (3) a convex function on $(0, \infty)$ with infimum at the left boundary is nondecreasing. For (1), backward induction: base $V_n^B(0^+, n) = g(0^+)$; step uses bounded convergence on $\mathbb{E}_\mu[V_n^B(s(1+T), m+1)]$ as $s \to 0^+$ ($\zeta_n < 1$ means $s(1+T) \to 0$ for all feasible T), squeezed by upper Lipschitz envelope $V_n^B(s, m+1) \leq L_{\text{prelim}} \cdot s + g(0^+)$ (preliminary Lipschitz from inductive hypothesis) and lower bound $V_n^B(s, m+1) \geq g(s) \to g(0^+)$ — both sides $\to g(0^+)$ as $s \to 0^+$. For (2), taking $\mu = \delta_0$ in the Bellman gives $V_n^B(s, m) \geq V_n^B(s, m+1)$; iterating to $m = n$ yields $V_n^B(s, m) \geq g(s) \geq g(0^+)$. For (3), if f convex on $(0, \infty)$ has $f'_+(s_0) < 0$ at some s_0, then by convexity $f(\varepsilon) \geq f(s_0) + |f'_+(s_0)|(s_0 - \varepsilon) > f(s_0) \geq \inf f = f(0^+)$, contradicting $\inf = f(0^+)$. Thus $f'_+ \geq 0$ everywhere, i.e., f nondecreasing.

By the naive envelope of Proposition G.1(d) (which only uses convexity + monotonicity preserved by induction, NOT the to-be-proved tight Lipschitz formula), $V_n^B(\cdot, m)$ is globally Lipschitz with $L_m^B(n) \leq (L + B)(1 + \zeta_n)^{n-m}$; combined with convexity and monotonicity, Fact E.3 below gives $L_m^B = \lim_{s\to\infty}(V_n^B)'_+(s, m)$. The induction below sharpens this naive envelope to the exact $\Lambda_m = B + (L - B)(1 + h_n)^{n-m}$ via secant-slope sandwich.

Fact E.3 (Lipschitz-as-limit-of-right-derivatives)**.** *For convex nondecreasing f on $(0, \infty)$ that is L-Lipschitz: f is differentiable a.e. with right-derivative $f'_+ \geq 0$ nondecreasing (convexity) and $f'_+ \leq L$ (Lipschitz). Thus* $\mathrm{Lip}(f) = \sup_s f'_+(s) = \lim_{s\to\infty} f'_+(s)$.

Base case $m = n$. $V_n^B(s, n) = g(s)$; $L_n^B(n) = \lim g'_+ = L = \Lambda_n$ (Lemma 3.3(i)).

Inductive step. Suppose $\lim_{s\to\infty}(V_n^B)'_+(s, m+1) = \Lambda_{m+1}$ and $\Lambda_{m+1} > B$ (invariant preserved since base $\Lambda_n = L > B$). Fix $\varepsilon > 0$; let S_ε with $(V_n^B)'_+(s, m+1) \in [\Lambda_{m+1} - \varepsilon, \Lambda_{m+1}]$ for $s \geq S_\varepsilon$. For $s \geq S_\varepsilon/(1 - \zeta_n)$, every $s(1+T)$ with $T \in [-\zeta_n, \zeta_n]$ has $s(1+T) \geq S_\varepsilon$. By secant-slope bound for convex functions:

$$\theta(s, T) := \frac{V_n^B(s(1+T), m+1) - V_n^B(s, m+1)}{sT} \in [\Lambda_{m+1} - \varepsilon,\ \Lambda_{m+1}]$$

uniformly in $T \in [-\zeta_n, \zeta_n]$, $\mu \in \widetilde{\mathcal{T}}_n$ (does not require V_n^B differentiable; secant-slope sandwich; the resulting upper bound has $O(s\varepsilon\zeta_n)$ slack from $\mathbb{E}_\mu[(\Lambda_{m+1} - \theta)T_-]$ on the negative tail, which vanishes as $\varepsilon \to 0$). Write

$$V_n^B(s(1+T), m+1) = V_n^B(s, m+1) + \theta(s, T)\, sT,$$

with $\theta(s, 0) := \Lambda_{m+1}$ by continuity.

Backward induction with $a := \mathbb{E}_\mu[T]$:

$$\begin{aligned} V_n^B(s, m) &= \sup_\mu \{V_n^B(s, m+1) + s\,\mathbb{E}_\mu[\theta\, T] - Bs|a|\} \\ &= V_n^B(s, m+1) + \sup_\mu \{s\,\mathbb{E}_\mu[\theta\, T] - Bs|a|\}. \end{aligned}$$

Since $\theta \in [\Lambda_{m+1} - \varepsilon, \Lambda_{m+1}]$:

$$\sup_{|a|\leq h_n} \{(\Lambda_{m+1} - \varepsilon)\, sa - Bs|a|\} \leq \sup_\mu \{\cdots\} \leq \sup_{|a|\leq h_n} \{\Lambda_{m+1}\, sa - Bs|a|\},$$

both maximized at $a = h_n$ (since $\Lambda_{m+1} - \varepsilon > B$ for small ε; Dirac $\mu = \delta_{h_n}$ feasible by $\zeta_n \geq h_n$). Extremes equal $(\Lambda_{m+1} - \varepsilon - B)sh_n$ and $(\Lambda_{m+1} - B)sh_n$. Substituting back into the Bellman recurrence: $V_n^B(s, m) - V_n^B(s, m+1) \in [(\Lambda_{m+1} - \varepsilon - B)sh_n,\ (\Lambda_{m+1} - B)sh_n]$. Dividing by s, adding $V_n^B(s, m+1)/s$, and using the inductive hypothesis $\lim_{s\to\infty} V_n^B(s, m+1)/s = \Lambda_{m+1}$ (Fact E.3 identifies asymptotic slope with Lipschitz constant for convex L-Lipschitz functions): $\lim_{s\to\infty} V_n^B(s, m)/s = \Lambda_{m+1}(1 + h_n) - Bh_n = \Lambda_m$. A second application of Fact E.3 yields:

$$\lim_{s\to\infty} (V_n^B)'_+(s, m) = \Lambda_{m+1}(1 + h_n) - Bh_n = \Lambda_m.$$

□

E.3. Pathwise primal-martingale game

Theorem E.4 (Pathwise calibration-dependent rate)**.** *With $\mathcal{T}_n^n$ the adapted-martingale-strategy set (one-step conditionals in $\mathcal{T}_n$; the adversary commits to $\mu \in \mathcal{T}_n^n$ first, then the player picks Δ_m adapted to $\mathcal{F}_{m-1}$), let $V_n^{P,B}(s) := \inf_{|\Delta_m|\le B} \sup_{\mu\in\mathcal{T}_n^n} \operatorname{ess\,sup}_{paths}[g(S_n) - \sum_m \Delta_m(S_m - S_{m-1})]$. Under Definition 3.1's data with attained $L = \mathrm{Lip}(g)$ and finite $\beta_\infty := \lim_{x\to\infty}[g(x)-Lx] > -\infty$, for any $\zeta_n \in [h_n, 1)$ with $\zeta_n \to 0$:*

$$V_n^{P,B}(s) = (L-B)\, s\,(1+\zeta_n)^n + Bs + \beta_\infty + O_g(1).$$

At standing $\zeta_n = h_n$ this gives $\Theta(e^{\sqrt{cn}})$, matching Theorems 3.4 and 4.2; at AFW Lindeberg $\zeta_n \sim \sqrt{K\log n/n}$ ($K \ge 16c$) it gives $\Theta(\exp(\sqrt{Kn\log n}))$, separating from the Lagrangian by tracking support rather than variance.

Proof of Theorem E.4. Write $T_m := S_m/S_{m-1} - 1$ for the relative increment. The martingale set $\mathcal{T}_n$ from Definition 2.2 requires $\mathbb{E}_\mu[T] = 0$, $\mathbb{E}_\mu[T^2] \le h_n^2$, and support $\subseteq [-\zeta_n, \zeta_n]$.

Lower bound. Define the i.i.d. 3-point Bernoulli law

$$\nu \;:=\; p\,\delta_{\zeta_n} + (1-2p)\,\delta_0 + p\,\delta_{-\zeta_n}, \qquad p \;:=\; \frac{h_n^2}{2\zeta_n^2}.$$

Feasibility: $\mathbb{E}_\nu[T] = p\zeta_n - p\zeta_n = 0$ (martingale); $\mathbb{E}_\nu[T^2] = 2p\zeta_n^2 = h_n^2$ (saturating variance); support $\subseteq [-\zeta_n, \zeta_n]$. The probability constraint $p \in (0, 1/2]$ holds since $\zeta_n \ge h_n$ implies $2\zeta_n^2 \ge 2h_n^2 \ge h_n^2$, so $p \le 1/2$. Thus $\nu^{\otimes n} \in \mathcal{T}_n^n$.

The all-up event $A_n := \{T_m = +\zeta_n \text{ for } m = 1, \dots, n\}$ has probability $\mathbb{P}(A_n) = p^n > 0$, so A_n lies in the essential support. On A_n, $S_m = s(1+\zeta_n)^m$ deterministically, so $S_n = s(1+\zeta_n)^n$ and $S_m - S_{m-1} = s\zeta_n(1+\zeta_n)^{m-1}$.

For any adapted player strategy Δ with $|\Delta_m| \le B$, the player's hedge value on A_n is bounded:

$$\sum_{m=1}^{n} \Delta_m(A_n)\cdot s\zeta_n(1+\zeta_n)^{m-1} \;\le\; B\cdot s\zeta_n\cdot\frac{(1+\zeta_n)^n - 1}{\zeta_n} \;=\; Bs\big((1+\zeta_n)^n - 1\big).$$

By Lemma 3.3(ii), $g(x) \ge Lx + \beta_\infty$ for all $x \ge 0$ (with $\beta_\infty := \lim_{x\to\infty}[g(x) - Lx]$ finite by attained Lipschitz). Thus on A_n, loss $= g(s(1+\zeta_n)^n) - Bs((1+\zeta_n)^n - 1) \ge (L-B)s(1+\zeta_n)^n + Bs + \beta_\infty$. Since $\mathbb{P}(A_n) > 0$, this lower-bounds $\operatorname{ess\,sup}_{\text{paths}}$ loss for any Δ; restricting the outer sup to $\nu^{\otimes n}$ gives $V_n^{P,B}(s) \ge (L-B)s(1+\zeta_n)^n + Bs + \beta_\infty$.

Upper bound. Player plays the constant strategy $\Delta_m \equiv B$ (adapted, $|B| \le B$). For *any* martingale adversary $\mu \in \mathcal{T}_n^n$ with $|T_m| \le \zeta_n$ a.s., telescoping gives pathwise:

$$\sum_{m=1}^{n} B\cdot(S_m - S_{m-1}) \;=\; B(S_n - S_0) \;=\; B(S_n - s).$$

So pathwise loss $= g(S_n) - B(S_n - s)$. Decompose $g(x) = Lx + \beta_\infty + \rho(x)$ with $\rho(x) := g(x) - Lx - \beta_\infty \ge 0$, $\rho \downarrow 0$ as $x \to \infty$, $\rho(0) = g(0) - \beta_\infty < \infty$ (since $g(x) - Lx$ is convex and nonincreasing, decreasing to β_∞):

$$\text{loss} \;=\; (L-B)S_n + Bs + \beta_\infty + \rho(S_n).$$

By the support bound, $S_n \le s\prod_m(1+T_m) \le s(1+\zeta_n)^n$ a.s. (each factor $1 + T_m \in [1-\zeta_n, 1+\zeta_n]$ is positive since $\zeta_n < 1$). Thus

$$\text{loss} \;\le\; (L-B)s(1+\zeta_n)^n + Bs + \beta_\infty + \rho(0).$$

Taking ess sup over paths and sup over martingale adversaries: $V_n^{P,B}(s) \le (L-B)s(1+\zeta_n)^n + Bs + \beta_\infty + \rho(0)$.

Combining. Since $(1+\zeta_n)^n \to \infty$ (as $n\zeta_n \ge nh_n = \sqrt{cn}(1+o(1)) \to \infty$), the constant slack $\rho(0) = O(1)$ is $o((1+\zeta_n)^n)$. Consequently,

$$V_n^{P,B}(s) \;=\; (L-B)s(1+\zeta_n)^n + Bs + \beta_\infty + o\big((1+\zeta_n)^n\big).$$

Asymptotic rates. For $\zeta_n = h_n$: $\log(1+h_n)^n = nh_n - nh_n^2/2 + O(nh_n^3) = \sqrt{cn} - c/2 + o(1)$, so $(1+h_n)^n = e^{-c/2}e^{\sqrt{cn}}(1+o(1))$. For Lindeberg $\zeta_n \sim \sqrt{K\log n/n}$ with $K \ge 16c$: $\log(1+\zeta_n)^n = n\zeta_n - n\zeta_n^2/2 + O(n\zeta_n^3) = \sqrt{Kn\log n} - K\log n/2 + o(1)$, giving $(1+\zeta_n)^n = n^{-K/2}\exp(\sqrt{Kn\log n})(1+o(1))$. The Lindeberg rate is strictly faster: the ratio of log-rates is $\sqrt{K\log n/c} \to \infty$. □

E.4. Premium bounds (Theorems 4.2 and 4.4 and Proposition D.4)

Proof of Theorem 4.2. Upper bound. $V_n^B(0,m) = g(0) = 0$ and $V_n^B(\cdot,0)$ is L_0^B-Lipschitz (Theorem 3.4), so $V_n^B(s,0) \le L_0^B s = [(L-B)\alpha_n + B]s$. Jensen gives $V_n^\infty(s,0) \ge g(s)$ (martingale, g convex), and $g(s) \ge Ls - C_g$ by Lemma 3.3. Thus $V_n^\infty(s,0) \ge Ls - C_g$; subtracting, $\Phi_n^B(s) \le (L-B)(\alpha_n - 1)s + C_g$.

Lower bound. Deterministic $T = h_n$ at every step (feasible, saturates variance). $S_n = s\alpha_n$; the induced loss is $g(s\alpha_n) - \sum_{m=0}^{n-1} Bs(1+h_n)^m h_n = g(s\alpha_n) - Bs(\alpha_n - 1)$. Subtract V_n^∞.

Matching. For affine-tail g, $g(s\alpha_n) = Ls\alpha_n - o(s\alpha_n)$, and $V_n^\infty(s,0) \to Ls + O(1)$ for large s; bounds match in leading order. □

Proof of Theorem 4.4. Setup. Define shifted $\widetilde{V}_m(x) = V_n^B(x,m) - g_\infty(x) - (L-B)(\alpha_{n-m} - 1)x$; then $R_n(s) = \widetilde{V}_0(s)$. Terminal $\widetilde{V}_n(x) = g(x) - Lx - \beta_\infty$ is convex (difference of convex and affine), nonneg (since $g \ge Lx + \beta_\infty$ by Lemma 3.3(ii)), vanishes on $[s^*, \infty)$ (piecewise-linear tail), and finite at 0. Since $\widetilde{V}_n$ is convex, nonneg, and vanishes at the right endpoint of $[0, s^*]$, it is nonincreasing on $[0, s^*]$, attaining its maximum at 0: $M_\psi := \widetilde{V}_n(0) = g(0) - \beta_\infty$.

No sign restriction on a_μ is required. Write $a_\mu := \mathbb{E}_\mu[T]$. We allow any $\mu \in \widetilde{\mathcal{T}}_n$; by Jensen, $|a_\mu| \le \sqrt{\mathbb{E}_\mu[T^2]} \le h_n$, so $a_\mu \in [-h_n, h_n]$ and $\eta_m := h_n - a_\mu \in [0, 2h_n]$. The Markov upper bound below applies uniformly in the sign of a_μ; in particular we do not assume $a_\mu \ge 0$.

Telescoping. The shifted Bellman is $\widetilde{V}_{m-1}(x) = \sup_{\mu \in \widetilde{\mathcal{T}}_n} \{\mathbb{E}_\mu[\widetilde{V}_m(x(1+T))] - (L-B)x\alpha_{n-m}\eta_m + 2Bx\min(a_\mu, 0)\}$ with $\eta_m = h_n - a_\mu \in [0, 2h_n]$. The extra $2Bx\min(a_\mu, 0) \le 0$ term arises from $|a_\mu|$ flipping sign for $a_\mu < 0$; we drop it in the upper-bound chain below since it only strengthens $R_n \le 0$ (negative-drift strategies are strictly dominated by $a_\mu = 0$, as already noted). For any adapted strategy, $\mathbb{E}[S_{m-1}\eta_m] = \mathbb{E}[S_{m-1}(h_n - a_{\mu_m})] = h_n\mathbb{E}[S_{m-1}] - (\mathbb{E}[S_m] - \mathbb{E}[S_{m-1}]) = (1+h_n)\mathbb{E}[S_{m-1}] - \mathbb{E}[S_m]$ (using the tower identity $\mathbb{E}[S_m|\mathcal{F}_{m-1}] = S_{m-1}(1+a_{\mu_m})$). Multiplying by α_{n-m} and summing, index-shift $m \to m-1$ in the first summand combined with $\alpha_{n-m}(1+h_n) = \alpha_{n-m+1}$ telescopes everything to

$$D_n := \sum_{m=1}^{n} \mathbb{E}[S_{m-1}\alpha_{n-m}\eta_m] = s\alpha_n - \mathbb{E}[S_n]. \tag{8}$$

Unrolling (with the dropped $2Bx\min(a_\mu, 0) \le 0$ term only strengthening the upper bound), $R_n(s) \le \sup\{\mathbb{E}[\widetilde{V}_n(S_n)] - (L-B)(s\alpha_n - \mathbb{E}[S_n])\}$.

Upper bound $R_n \le 0$ via Markov applied to Z_n. Since $|T_m| \le \zeta_n = h_n$ pointwise under the standing calibration $\zeta_n = h_n$ of Theorem 4.4, we have $S_m \le S_{m-1}(1+h_n)$ a.s., so $S_n \le s\alpha_n$ a.s.; hence $Z_n := s\alpha_n - S_n \ge 0$ with $\mathbb{E}[Z_n] = D_n$. Markov's inequality gives $\Pr(Z_n \ge s\alpha_n - s^*) \le D_n/(s\alpha_n - s^*)$, equivalently $\Pr(S_n \le s^*) \le D_n/(s\alpha_n - s^*)$. Since $\widetilde{V}_n \le M_\psi$ on $[0, s^*]$ and vanishes on $[s^*, \infty)$:

$$\mathbb{E}[\widetilde{V}_n(S_n)] \le M_\psi \cdot \frac{D_n}{s\alpha_n - s^*}.$$

Thus $\mathbb{E}[\widetilde{V}_n(S_n)] - (L-B)D_n \le D_n[M_\psi/(s\alpha_n - s^*) - (L-B)] \le 0$ by the threshold assumption.

Lower bound $R_n \ge 0$ via pure drift. $\mu_m \equiv \delta_{h_n}$ is feasible; gives $S_n = s\alpha_n$ deterministic, $D_n = 0$, $\widetilde{V}_n(s\alpha_n) = 0$. Thus $\widetilde{V}_0(s) \ge 0$.

Combining: $R_n(s) = 0$.

Extension to intermediate steps. Fix $m_0 \in \{0, 1, \ldots, n-1\}$ and consider the subgame starting at step m_0 from state s, with $n - m_0$ remaining steps. The shifted value $\widetilde{V}_m(x) := V_n^B(x,m) - g_\infty(x) - (L-B)(\alpha_{n-m} - 1)x$ defined above satisfies the same terminal $\widetilde{V}_n(x) = g(x) - Lx - \beta_\infty$ (unchanged by the choice of starting step). Reindexing the telescoping identity (8) to the range $m = m_0+1, \ldots, n$ gives

$$D_{n,m_0} := \sum_{m=m_0+1}^{n} \mathbb{E}\big[S_{m-1}\,\alpha_{n-m}\,\eta_m\big] = s\alpha_{n-m_0} - \mathbb{E}\big[S_n \,\big|\, S_{m_0} = s\big],$$

using $\alpha_{n-m}(1+h_n) = \alpha_{n-m+1}$ (same algebraic identity) and the initial value $\mathbb{E}[S_{m_0}] = s$. The Markov argument applied to $Z_n := s\alpha_{n-m_0} - S_n \ge 0$ (nonneg since, under the standing calibration $\zeta_n = h_n$, $|T_m| \le h_n$ pointwise

a.s. — not just in L^2 — gives $S_m \le S_{m-1}(1+h_n)$ a.s. and so $S_n \le s\alpha_{n-m_0}$ a.s. starting from $S_{m_0} = s$) gives $\mathbb{E}[\widetilde{V}_n(S_n)] \le M_\psi D_{n,m_0}/(s\alpha_{n-m_0} - s^*)$, and under the reindexed threshold $s\alpha_{n-m_0} > s^* + M_\psi/(L-B)$ the bracket $M_\psi/(s\alpha_{n-m_0} - s^*) - (L-B)$ is ≤ 0, forcing $\widetilde{V}_{m_0}(s) \le 0$. The pure-drift adversary (i.e., $\mu_m \equiv \delta_{h_n}$ for $m > m_0$) on the remaining $n - m_0$ steps attains $\widetilde{V}_{m_0}(s) \ge 0$. Thus $V_n^B(s, m_0) = (L-B)(\alpha_{n-m_0} - 1)s + g_\infty(s)$, with the theorem's $m = 0$ form the initial-step specialization. □

Proof of Proposition D.4. Since $g'_+ \le L$ everywhere, $\psi'_+ = g'_+ - L \le 0$, so ψ is nonincreasing on $[0,\infty)$ with $\sup_x \psi(x) = \psi(0) = M_\psi$; nonnegativity $\psi \ge 0$ follows from $g(x) \ge Lx + \beta_\infty$ (Lemma 3.3(ii)). Markov (applied to Z_n) proceeds as in Theorem 4.4 with modified terminal $\widetilde{V}_n(x) = \psi(x) \le M_\psi$ on $[0,\tau]$ and $\le \psi(\tau)$ on $[\tau,\infty)$ by monotonicity. The adversary bound $\mathbb{E}[\widetilde{V}_n(S_n)] \le \psi(\tau) + M_\psi D_n/(s\alpha_n - \tau)$ combines with the player's telescoping penalty $-(L-B)D_n$ to give

$$R_n(s) \le \psi(\tau) + D_n\big[M_\psi/(s\alpha_n - \tau) - (L-B)\big].$$

With $\tau = s\alpha_n/2$ the bracket equals $2M_\psi/(s\alpha_n) - (L-B) \le 0$ for $s\alpha_n \ge 2M_\psi/(L-B)$, so the D_n term drops and $R_n(s) \le \psi(s\alpha_n/2)$. For each tail class, substitute: (i) $\psi(\tau) = O(e^{-\tau/\lambda})$ gives $R_n = O(e^{-s\alpha_n/(2\lambda)})$ dominant; (ii) $\psi(\tau) = O(\tau^{-p})$ gives $R_n = O((s\alpha_n)^{-p}) = O(e^{-p\sqrt{cn}})$; (iii) $\psi(\tau) = O(1/\log\tau)$ gives $R_n = O(1/\sqrt{cn})$. □

E.5. Bounded payoffs (Proposition D.5 and Theorem D.6)

Proof of Proposition D.5. Lower bound. A tighter feasible set on the player's actions cannot decrease the minimax value (inf over smaller set $\ge$ inf over larger set; adversary unchanged), so $V_n^B \ge V_n^\infty$. *Upper bound.* The strategy $\Delta_m \equiv 0$ is feasible under any $B \ge 0$ and yields $\ell_n = g(S_n)$ (hedge term vanishes), so

$$V_n^B \;\le\; \sup_{S_n} \mathbb{E}[g(S_n)] \;\le\; M_g.$$

Combining with $V_n^\infty \ge 0$ ($g \ge 0$, Jensen) gives $\Phi_n^B \le M_g$. □

Proof of Theorem D.6. Lower bound on V_n^B (drift-to-plateau adversary). Adversary plays $T_m = h_n$ until hitting time $\tau := \min\{m : s(1+h_n)^m \ge s_M\}$, then $T_m \equiv 0$ (Dirac at 0 is feasible: martingale + variance $0 \le h_n^2$). For $n \ge \tau$, $S_n = S_\tau \in [s_M, s_M(1+h_n))$, so

$$g(S_n) \;=\; M_g.$$

The player's hedge gain telescopes as a geometric series:

$$\begin{aligned}\sum_{m=1}^{n} T_m \Delta_m S_{m-1} \;&=\; h_n \sum_{m\le\tau} \Delta_m S_{m-1} \\ &\le\; h_n B \sum_{m\le\tau} S_{m-1} \;\le\; B(s_M - s) + B s_M h_n.\end{aligned}$$

Thus $V_n^B(s,0) \ge M_g - B(s_M - s) - B s_M h_n$.

Upper bound on V_n^∞ (concave envelope). Let $U(x) := \min(Lx, M_g)$, the pointwise concave envelope of g. Then $g \le U$ (g is L-Lipschitz with $g(0) = 0$, so $g(x) \le Lx$, and $g \le M_g$). By backward induction on m, assuming $V_n^\infty(\cdot, m+1) \le U$: for any mean-zero μ (the feasible set under $B = \infty$ forces zero drift), Jensen on the concave U gives

$$\mathbb{E}_\mu[U(s(1+T))] \;\le\; U(s(1+\mathbb{E}_\mu[T])) \;=\; U(s).$$

So $V_n^\infty(s,m) \le U(s)$. Base case $V_n^\infty(\cdot, n) = g \le U$. For $s < M_g/L$, $U(s) = Ls$.

Combining. $\Phi_n^B(s) \ge [M_g - B(s_M - s) - B s_M h_n] - Ls = \eta + Bs - B s_M h_n \ge \eta - B s_M h_n$. Under $B s_M h_n < \eta/2$, this is $\ge \eta/2 > 0$. □

E.6. Aggregate-budget bounds (Theorems B.3 to B.6 and Proposition B.2)

Proof of Proposition B.2. Any adapted (Δ_m) with $|\Delta_m(\omega)| \le B$ for every m pathwise satisfies $\sum_m |\Delta_m(\omega)| \le nB$ pathwise. Thus aggregate-feasible set $\supseteq$ per-step-feasible set; with the same adversary class, inner values agree on the smaller set, and the outer infimum is monotone non-increasing. So $V_n^{\mathrm{agg},B} \le V_n^{\mathrm{per-step},B}$. Subtract V_n^∞ for the premium; apply Theorem 4.2 for the RHS. □

Proof of Theorem B.3. State-adaptive Dirac $\mu_m^* = \delta_{\epsilon_m^*}$, $\epsilon_m^* := h_n s/S_{m-1}$. Feasibility: $S_{m-1} \ge s$ gives $\epsilon_m^* \le h_n \le \zeta_n$; Dirac $\mathbb{E}[T^2] = (\epsilon_m^*)^2 \le h_n^2$ (DVC saturated at $m = 1$). State evolves as $S_m = S_{m-1}(1 + h_n s/S_{m-1}) = S_{m-1} + h_n s$, so $S_n = s(1 + nh_n)$. Key: $\epsilon_m^* S_{m-1} = h_n s$ is constant in m.

For adapted Δ with $\sum |\Delta_m| \le nB$: $\mathbb{E}[\sum T_m \Delta_m S_{m-1}] = h_n s \sum \Delta_m \le h_n s \cdot nB$.

For n large with $s(1 + nh_n) \ge s^*$: $g(S_n) = Ls(1 + nh_n) + \beta_\infty$. Thus $\mathbb{E}[\ell_n] \ge (L - B)snh_n + Ls + \beta_\infty$. Since $nh_n = \sqrt{cn}(1 + O(c/n))$: $(L - B)snh_n = (L - B)s\sqrt{cn} + O(1)$. Subtract $V_n^\infty(s, 0) \le Ls + g(0)$ (Lipschitz envelope from $V_n^\infty(0, 0) = g(0)$). □

Proof of Theorem B.4. Power-law Dirac $\mu_m^{(\kappa)} = \delta_{\epsilon_m^{(\kappa)}}$, $\epsilon_m^{(\kappa)} := h_n(s/S_{m-1})^\kappa$, $\kappa \in (B/L, 1)$. Feasibility as in Theorem B.3.

Discrete asymptotics of S_n. Let $U_m := S_m^\kappa$. Monotonicity: $S_m \ge S_{m-1}$, so $U_m \ge U_0 = s^\kappa$. Bounded increment: $\xi_m := h_n s^\kappa / U_{m-1} \le h_n$ for all m. For $h_n \le 1/2$, Taylor gives $U_m = U_{m-1}(1 + \xi_m)^\kappa$ with

$$U_m - U_{m-1} = \kappa h_n s^\kappa + \tfrac{\kappa(\kappa-1)}{2} \tfrac{h_n^2 s^{2\kappa}}{U_{m-1}} + E_m,\ |E_m| \le C_1 h_n^3 s^{3\kappa}/U_{m-1}^2. \tag{9}$$

Since $\kappa(\kappa - 1) < 0$, (9) gives $U_m - U_{m-1} \ge \kappa h_n s^\kappa - C_2 h_n^2 s^\kappa$; iterating, $U_m \ge s^\kappa + m(\kappa h_n s^\kappa - C_2 h_n^2 s^\kappa) \ge \frac{1}{2} m \kappa h_n s^\kappa$ for $n \ge n_0$, $m \ge m_0$.

Summing (9): $U_n - U_0 = n\kappa h_n s^\kappa + (\kappa(\kappa - 1)/2) h_n^2 s^{2\kappa} \sum_m 1/U_{m-1} + \sum_m E_m$. The middle sum is $O(\log n/(s^\kappa \kappa h_n))$ (split at m_0, use monotonicity for $m \le m_0$ and the linear lower bound for $m > m_0$). The E_m sum is $o(h_n s^\kappa \log n)$. Thus

$$U_n = s^\kappa(1 + n\kappa h_n) + O(h_n s^\kappa \log n) = s^\kappa n \kappa h_n (1 + O(\log n/\sqrt{cn})).$$

So $S_n = s(\kappa\sqrt{cn})^{1/\kappa}(1 + O(\log n/\sqrt{cn}))$.

For $\kappa < 1$, $c_m := \epsilon_m^{(\kappa)} S_{m-1} = h_n s^\kappa S_{m-1}^{1-\kappa}$ is strictly increasing. Under aggregate constraint, player's best response $\sup_\Delta \sum \Delta_m c_m = nBc_n$. Using the asymptotic:

$$c_n = h_n s(\kappa\sqrt{cn})^{(1-\kappa)/\kappa}(1 + o(1)), \qquad nBc_n = Bs\kappa^{(1-\kappa)/\kappa}(cn)^{1/(2\kappa)}(1 + o(1)).$$

Meanwhile $S_n \ge s^*$ for n large, so $g(S_n) = Ls\kappa^{1/\kappa}(cn)^{1/(2\kappa)}(1 + o(1))$. Induced loss:

$$\mathbb{E}[\ell_n] - V_n^\infty \ge s\kappa^{(1-\kappa)/\kappa}(L\kappa - B)(cn)^{1/(2\kappa)}(1 + o(1)).$$

Optimize over $\kappa \in (B/L, 1)$: set $\kappa = B/L + \delta$, expand $1/(2\kappa) = L/(2B) - L^2\delta/(2B^2) + O(\delta^2)$ and $L\kappa - B = L\delta$. Unconstrained max at $\delta^* = 2B^2/(L^2 \log(cn))$; valid when $\delta^* < 1 - B/L$, i.e. $L - B \ge 2B^2/(L \log(cn))$. Substituting: bound is $C(L, B, s, c)(cn)^{L/(2B)}/\log(cn)$. □

Proof of Theorem B.5. Construction. $K := \lceil (L - B)n/L \rceil$.

Phase 1 ($m = 1, \ldots, K$): pure-drift Dirac $T_m = h_n$. Feasible ($h_n \le \zeta_n$, DVC saturated at h_n^2). $S_K = s(1 + h_n)^K$ deterministically.

Phase 2 ($m = K + 1, \ldots, n$): dollar-drift Dirac $T_m = h_n S_K/S_{m-1}$, keeping $T_m S_{m-1} = h_n S_K$ constant. Feasibility (by induction on m): at $m = K + 1$, $S_{K+1} = S_K(1 + T_{K+1}) = S_K + h_n S_K > S_K$, so $T_{K+2} = h_n S_K/S_{K+1} < h_n$, feasible. Step: if $S_{m-1} \ge S_K$, then $S_m = S_{m-1}(1 + T_m) = S_{m-1} + h_n S_K > S_{m-1} \ge S_K$; and $T_m = h_n S_K/S_{m-1} \le h_n \cdot S_K/S_K = h_n$, so $|T_m| \le \zeta_n$ and $T_m^2 \le h_n^2$. Variance DVC saturated only at $m = K + 1$; subsequent Phase-2 steps have strict slack, so the adversary is feasible throughout.

Terminal state. Phase 2 recurrence $S_m = S_{m-1} + h_n S_K$ telescopes: $S_n = S_K(1 + (n - K)h_n)$.

Player response. Dollar-drift $X_m := T_m S_{m-1} = h_n S_{m-1}$ in Phase 1 ($\leq h_n S_K$), $X_m \equiv h_n S_K$ in Phase 2. So $\max_m X_m = h_n S_K$. By ℓ^1/ℓ^∞ duality: $\mathbb{E}[\sum T_m \Delta_m S_{m-1}] \leq nB \cdot h_n S_K$.

Induced loss. Using $g(y) \geq Ly + \beta_\infty$ for $y \geq s^*$: $\mathbb{E}[\ell_n] \geq LS_n + \beta_\infty - nBh_n S_K = LS_K + LS_K(n-K)h_n + \beta_\infty - nBh_n S_K = LS_K + S_K h_n[L(n-K) - nB] + \beta_\infty$.

Choice $K = \lceil (L-B)n/L \rceil$ makes $L(n-K) - nB = O(L)$ (bounded, exactly zero when $(L-B)n/L$ is integer). Middle term $O(S_K h_n) = O(S_K/\sqrt{n})$, absorbed.

Asymptotics of S_K. $\log(S_K/s) = K\log(1+h_n) = Kh_n - Kh_n^2/2 + O(Kh_n^3) = \frac{L-B}{L}\sqrt{cn} - \frac{(L-B)c}{2L} + O(1/\sqrt{n})$.

Exponentiating, subtracting $V_n^\infty(s,0) \leq Ls + g(0)$ (Lipschitz envelope), gives the claim. □

Proof of Theorem B.6. Player strategy. The deterministic schedule $\hat{\Delta}_m := L \cdot \mathbf{1}[m > n - N^*]$ pays nothing for the first $n - N^*$ steps, then L shares per step. Feasibility:

$$\sum_{m=1}^{n} |\hat{\Delta}_m| \;=\; LN^* \;\leq\; L \cdot \frac{nB}{L} \;=\; nB.$$

Pathwise loss. The telescoping identity $T_m S_{m-1} = S_m - S_{m-1}$ gives

$$\sum_{m=1}^{n} T_m \hat{\Delta}_m S_{m-1} \;=\; L \sum_{m > n-N^*} (S_m - S_{m-1}) \;=\; L(S_n - S_{n-N^*}).$$

Combined with the L-Lipschitz bound $g(S_n) \leq g(0) + LS_n$ at 0:

$$\ell_n(\hat{\Delta}, T) \;=\; g(S_n) - L(S_n - S_{n-N^*}) \;\leq\; g(0) + LS_{n-N^*}.$$

Drift cap and expected state. Cauchy–Schwarz on the variance bound gives

$$\big|\mathbb{E}[T_m \mid \mathcal{F}_{m-1}]\big| \;\leq\; \sqrt{\mathbb{E}[T_m^2 \mid \mathcal{F}_{m-1}]} \;\leq\; h_n.$$

Iterating via the tower property,

$$\mathbb{E}[S_m] \;=\; \mathbb{E}\big[S_{m-1}(1 + \mathbb{E}[T_m \mid \mathcal{F}_{m-1}])\big] \;\leq\; (1+h_n)\mathbb{E}[S_{m-1}],$$

so $\mathbb{E}[S_m] \leq s(1+h_n)^m$ by induction. Taking $\sup_T$ in the pathwise bound yields $V_n^{\mathrm{agg},B}(s) \leq g(0) + Ls(1+h_n)^{n-N^*}$. The leading-order expansion follows from $\log(1+h_n) = \sqrt{c/n} - c/(2n) + O(n^{-3/2})$ and $n - N^* = n(L-B)/L + O(1)$. □

F. Grid threshold corollary

Proof of Corollary 3.5. Fix m and $s \geq s^*/(1-h_n)^{n-m}$; prove $(V_n^B)'_+(s,m) = \Lambda_m := B + (L-B)(1+h_n)^{n-m}$ by backward induction on $k := n - m$.

$k = 0$: $s \geq s^*$, $V_n^B(\cdot, n) = g$ affine slope $L = \Lambda_n$.

Inductive step: suppose identity at $m+1$ on $[s^*/(1-h_n)^{n-m-1}, \infty)$. If $s \geq s^*/(1-h_n)^{n-m}$, then for every $T \in [-h_n, h_n]$, $s(1+T) \geq s(1-h_n) \geq s^*/(1-h_n)^{n-m-1}$, so $(V_n^B)'_+(s(1+T), m+1) = \Lambda_{m+1}$ is *exactly constant* on the entire support band. The argument of Theorem 3.4's inductive step applies with $\varepsilon = 0$ and threshold $s_{\mathrm{th}} := s^*/(1-h_n)^{n-m-1}$ (distinct from the game initial state $S_{n,0}$): secant $\theta(s,T) \equiv \Lambda_{m+1}$, inner sup attained at $a = h_n$, giving $(V_n^B)'_+(s,m) = \Lambda_{m+1}(1+h_n) - Bh_n = \Lambda_m$.

□

Remark F.1 (Implications for numerics). $(1-h_n)^n = e^{-c(1+o(1))}$, so the threshold is $s^*/(1-h_n)^n = \Theta(s^* e^{\sqrt{cn}})$. Grids must extend to $s_{\max} \geq s^*/(1-h_n)^n$ to resolve the asymptotic slope. A safe default is $s_{\max} = \Omega(\alpha_n^2 s^*)$. Experiment E4 (Appendix H) uses $s_{\max} = 3s^*$; at $n = 12$ this is just below the threshold ($\approx 3.18 s^*$), which partially explains small-n over-prediction.

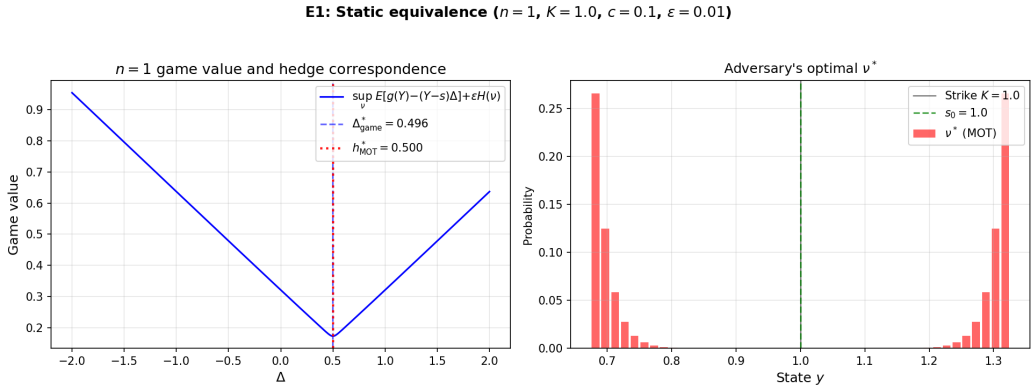

Figure 2. E1: Static equivalence at $n = 1$. Left: V_{MOT} (solid blue) from the entropic MOT dual vs. V_{game} (orange dashed) from brute-force grid search over $\Delta \in [-2, 2]$. Right: optimal dual action h^* (blue) and optimal hedge Δ^* (orange) agree to within $O(\varepsilon_{\text{ent}}) = 0.004$ across the grid, confirming the $n = 1$ bridge (Proposition 2.4).

G. Calibration map across theorems

Different per-step-budget theorems require calibration assumptions of different strength; we collect them.

- Theorems 3.4 and 4.2 and Proposition G.1: need only $\zeta_n \geq h_n$ (Dirac at $\pm h_n$ feasible).
- Theorem 4.4: needs the tight-support calibration $\zeta_n = h_n$.
- Proposition D.4: tight-support $\zeta_n = h_n$ plus monotonicity of the tail residual ψ.
- AFW convergence (Definition 2.1, the classical $V_n \to \mathbb{E}[g(G(1))]$ result (Abernethy et al., 2012, Thms 1–2)): additionally needs $\liminf_n n\zeta_n^2 / \log n > 16c$ for AFW concentration; not a hypothesis of any per-step-budget theorem here.

This map clarifies that our per-step-budget machinery operates under a strictly weaker calibration than AFW's asymptotic convergence result, resolving an otherwise subtle consistency concern: the experimental choice $\zeta_n = h_n$ in Section 6 satisfies our per-step-budget hypotheses but violates AFW concentration, which is consistent because we do not invoke AFW convergence in any of our theorems.

Proposition G.1 (Boundary Lipschitz cases)**.** *Under Definition 3.1: (a) $L_n^B(n) = L$ (terminal). (b) If $B \geq L$: $L_m^B(n) = L$ for all m, n (by Proposition D.1; martingale adversary, Lemma 2.3 preserves L). (c) If $B = \infty$ (unconstrained): $L_m^\infty(n) = L$ for all m, n (immediate from Lemma 2.3). (d) Naive bound: $L_m^B(n) \leq (L + B)(1 + \zeta_n)^{n-m} = O(e^{\sqrt{cn}})$, by iterating the backward recursion.*

H. Full experiment details

We report complete numerical results for experiments E1–E7 plus N1, N3, N4, N4b, N6, E4c, E4d, E4e, N7, and the SPX/ABFW baseline. All were executed on a single CPU (Intel Xeon class, ≤ 8 GB RAM, no GPU); the E1–E7 core is under 1 hour wall-clock, with N4b at $\sim$16 min, E4c at $\sim$74 min, E4d at $\sim$45 min, E4e at $\sim$7 min, N7 extension at $\sim$1 min, and the SPX/ABFW baseline at $<$1 s. Experiments are deterministic given parameters (no stochastic solver state; seeds not required). Software stack: Python 3.11 with NumPy 1.26 and SciPy 1.11; the SPX backtest additionally uses `yfinance` for the public `^GSPC` daily-close series. Drivers in the supplementary code archive.

H.1. E1: Static equivalence ($n = 1$)

Solve one-step game by brute-force grid search over Δ and independently via MOT dual. Parameters: $q = 100$, $\varepsilon = 0.01$, $\zeta_1 = 0.3243$. Brute force: $\Delta^* = 0.496$, $V_{\text{game}} = 0.172423$. MOT: $h^* = 0.500$, $V_{\text{MOT}} = 0.172349$. Gaps $|\Delta^* - h^*| = 4 \times 10^{-3}$ (one grid spacing) and $|V_{\text{game}} - V_{\text{MOT}}| = 7.4 \times 10^{-5}$, consistent with entropic regularization of Proposition 2.4.

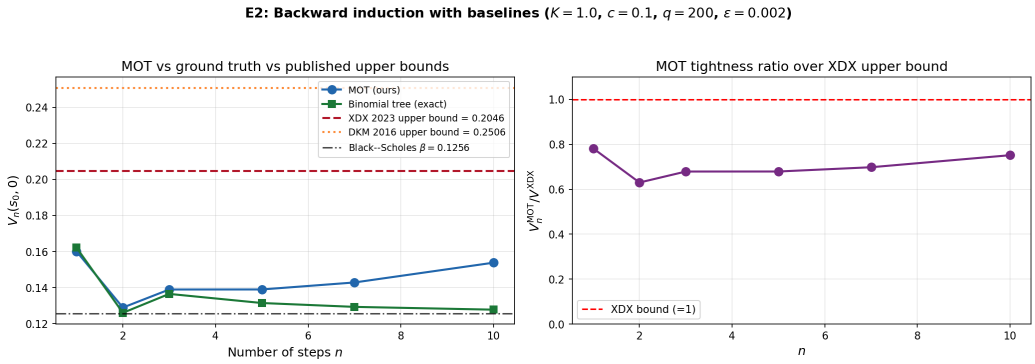

Figure 3. E2: MOT backward induction vs. binomial tree, XDX, DKM, Black–Scholes.

H.2. E2: Backward induction vs. baselines

$n \in \{1, 2, 3, 5, 7, 10\}$ on a $q = 200$ grid with $\varepsilon = 0.002$. Reference: $\beta_{\text{BS}} = 0.1256$, $V^{\text{XDX}} = 0.2046$, $V^{\text{DKM}} = 0.2506$.

n	MOT V_n	Binomial V_n	MOT − bin.	MOT/XDX
1	0.1601	0.1622	−0.0021	0.78
2	0.1289	0.1260	+0.0029	0.63
3	0.1390	0.1365	+0.0024	0.68
5	0.1390	0.1315	+0.0075	0.68
7	0.1429	0.1293	+0.0135	0.70
10	0.1538	0.1278	+0.0261	0.75

MOT tracks binomial to within $n\varepsilon \log q$ (max 0.026 at $n = 10$; bound $10 \cdot 0.002 \cdot \log 200 = 0.11$), strictly below XDX and DKM. Non-monotonicity $V_1 > V_2$ (0.160 vs 0.129) appears for both MOT and binomial and reflects the one-step structure at $n = 1$ with two-point support.

H.3. E3: Budget-constrained premium

$n \in \{1, 3, 5, 10\}$ on $q = 200$ grid; sweep L_1 relaxation δ over 26 values; record V_n^B vs. $B_{\text{eff}} = |h^*(\delta)|$. Unconstrained $V_n^\infty = \{0.160, 0.139, 0.139, 0.154\}$ for $n = \{1, 3, 5, 10\}$. Premium curves non-increasing in B, vanish for $B \geq L = 1$ (confirming Proposition D.1), with characteristic knee at $B = B_{\text{crit}}^{(n)}$.

H.4. E4: Lipschitz degradation universality (full table)

For $B \in \{0.05, 0.10, 0.15\}$ and $n \in \{2, 4, 6, 8, 12\}$, $q = 200$, $\varepsilon_{\text{ent}} = 0.003$: measure numerical $L_0^{B,\text{num}}$ vs. exact $L_0^B = B + (L - B)(1 + h_n)^n$.

n	$B = 0.05$	$B = 0.10$	$B = 0.15$
2	1.340	2.285	3.103
4	1.091	1.369	2.156
6	1.008	1.125	1.632
8	1.012	1.104	1.475
12	0.956	1.009	1.145

Convergence monotone in n; small-n over-prediction (up to 3.1 at $n = 2$, $B = 0.15$) reflects discretization slack on $q = 200$ grid when the support band is wide (e.g. $[0.77, 1.23]$ at $n = 2$, $h_2 = 0.226$). By $n = 12$ ($h_{12} = 0.091$) ratios fall into $[0.96, 1.14]$ across all B. This is consistent with Corollary 3.5: the grid threshold $s^*/(1 - h_n)^n = \Theta(s^* \cdot e^{\sqrt{cn}})$ at $n = 12$ is $\approx 3.18 s^*$, just above the $s_{\max} = 3s^*$ used here; grid expansion beyond $\alpha_n^2 s^*$ eliminates this.

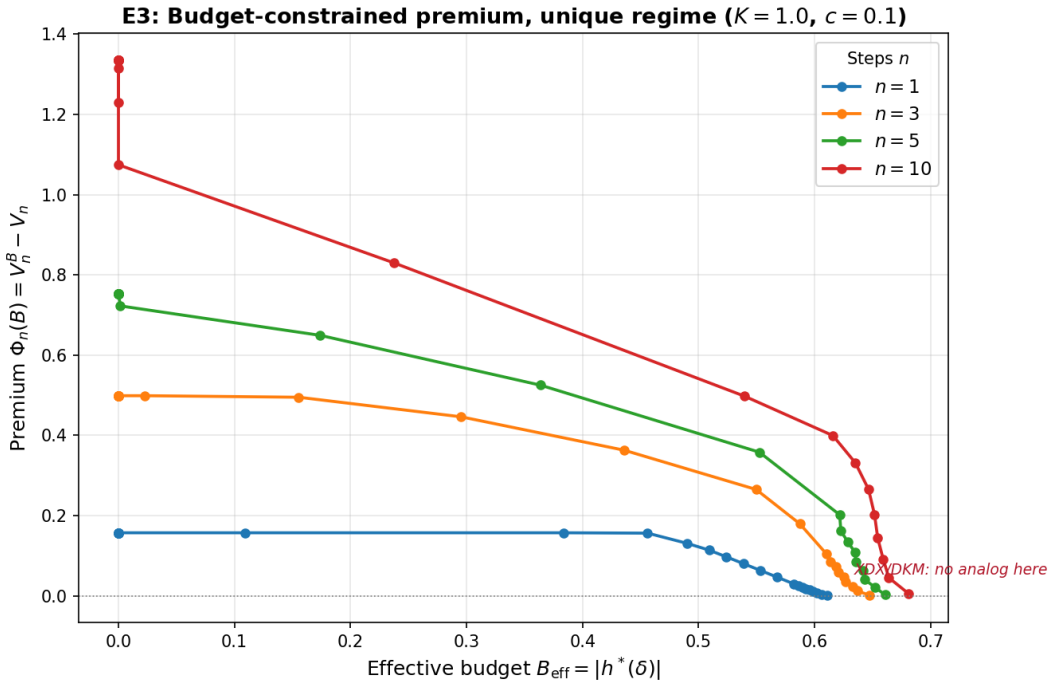

Figure 4. E3: Budget-constrained premium Φ_n^B vs. B_{eff} for $n \in \{1, 3, 5, 10\}$. All curves vanish at $B \geq L = 1$; the premium grows with n and flattens for small B, consistent with exponential divergence (Theorem 4.2).

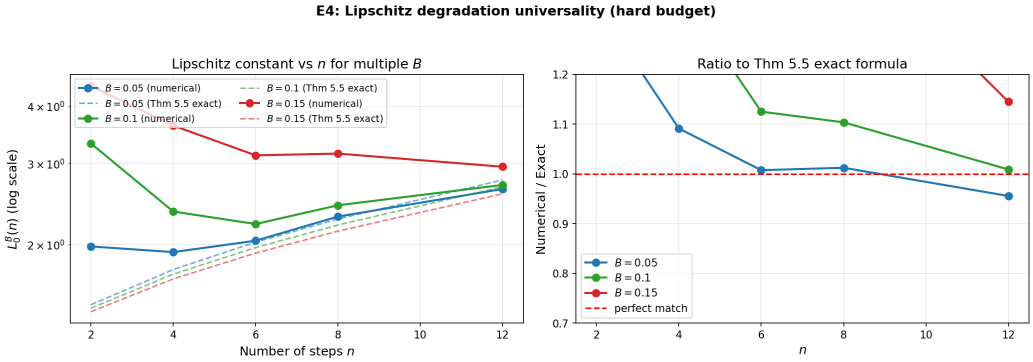

Figure 5. E4: Universality of the exact Lipschitz formula. Ratios $\to 1$ as $n \to \infty$, shape-independent across $B \in \{0.05, 0.1, 0.15\}$.

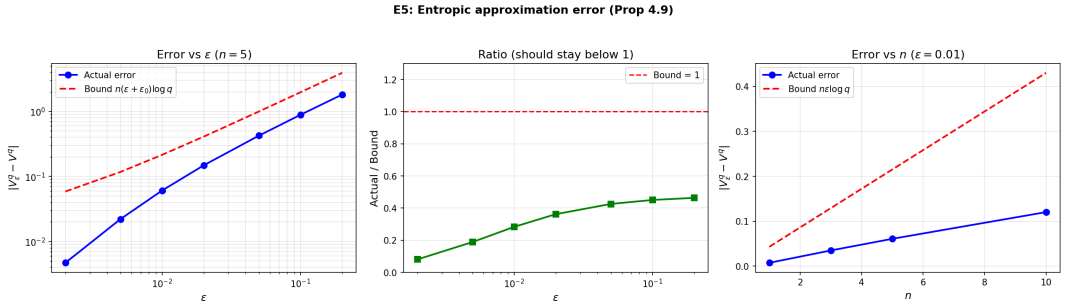

Figure 6. E5: Entropic error accumulation. Left: error vs. ε at $n = 5$. Right: error vs. n at $\varepsilon = 0.01$, linear accumulation.

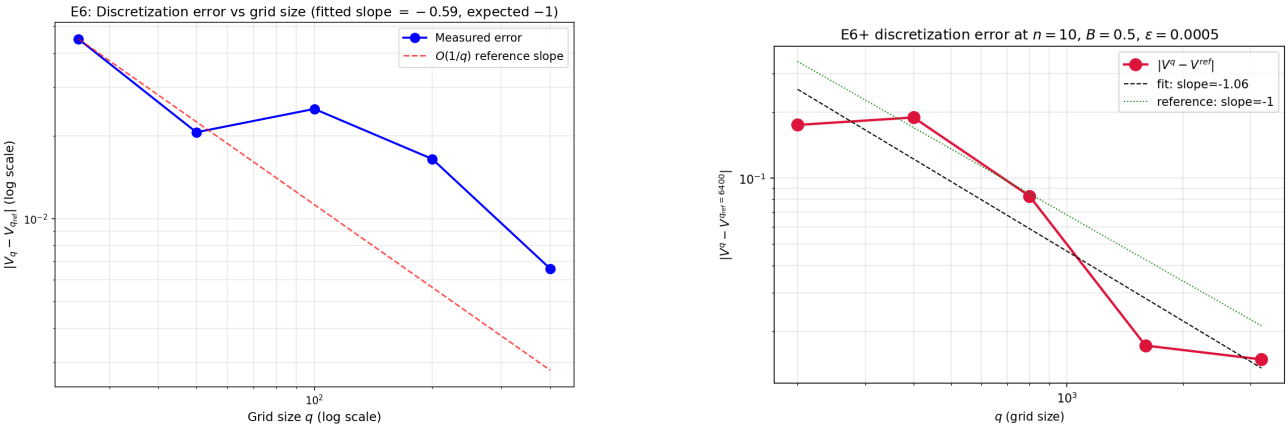

Figure 7. E6 (left): low-q regime, slope -0.59 contaminated by non-monotone samples. E6+ (right): high-q regime ($n = 10$, $\varepsilon = 5\text{e-}4$, q up to 3200 vs. $q_{\text{ref}} = 6400$), fitted slope $\mathbf{-1.055}$, matching the predicted $O(1/q)$ rate.

H.5. E5: Entropic approximation error

For $q = 50$ and reference $\varepsilon_{\text{exact}} = 0.001$: (Part A) at $n = 5$, sweep $\varepsilon \in [0.002, 0.2]$; (Part B) at $\varepsilon = 0.01$, sweep $n \in \{1, 3, 5, 10\}$. Measured $|V_\varepsilon^q - V^q|$ vs. theoretical $n\varepsilon \log q$ bound (Proposition E.2):

Part A ($n = 5$)		Part B ($\varepsilon = 0.01$)	
ε	ratio	n	ratio
0.002	0.080	1	0.171
0.01	0.283	3	0.270
0.05	0.425	5	0.283
0.1	0.450	10	0.279
0.2	0.463		

All ratios ≤ 1, confirming Proposition E.2; looseness (factor $\sim$ 2–12) consistent with worst-case $\log q$ overestimating attained entropy of near-deterministic optimizers.

H.6. E6: Discretization convergence

At $n = 5$, $\varepsilon = 0.002$, reference $V^{q_{\text{ref}}=800} = 0.1555$, sweep $q \in \{25, 50, 100, 200, 400\}$. Fitted log-log slope -0.59. The standard $O(1/q)$ rate is optimistic at small q (grid pitch vs. entropic bandwidth mismatch) and contaminated by non-monotonic low-q samples ($V^{50} > V^{100}$). **E6+ extends to higher q to clean up the slope-fit:** at $n = 10$, $\varepsilon = 5\text{e-}4$, $q_{\text{ref}} = 6400$ ($V^{6400} = 0.7883$), sweeping $q \in \{200, 400, 800, 1600, 3200\}$ gives fitted log-log slope -1.055 ($\sim$3 h single-CPU dominated by q_{ref}). The high-q regime is uncontaminated: gaps $0.175 \to 0.189 \to 0.083 \to 0.0173 \to 0.0149$ become monotone-decreasing past the contamination zone ($q \geq 800$). The asymptotic discretization rate matches the predicted $O(1/q)$.

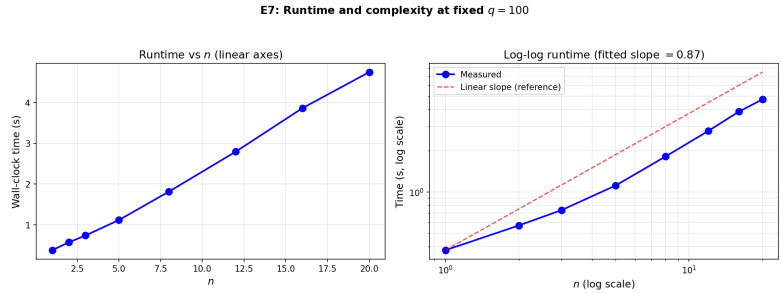

Figure 8. E7: Wall-clock runtime vs. n at $q = 100$ on log–log axes; fitted slope 0.87.

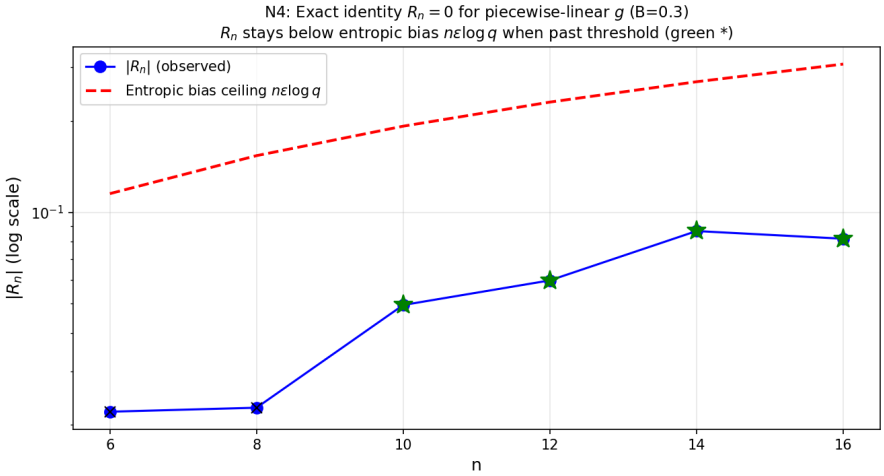

Figure 9. N4: $|R_n|$ under backward induction with hard-budget projection stays below the entropic bias ceiling $n\varepsilon \log q$ (red dashed). Green stars mark n beyond the boundary-layer threshold.

H.7. E7: Runtime complexity

At $q = 100$, $\varepsilon = 0.005$: single-threaded wall-clock for $n \in \{1, 2, 3, 5, 8, 12, 16, 20\}$: $\{0.38, 0.57, 0.74, 1.12, 1.81, 2.79, 3.86, 4.75\}$ s. Fitted log-log slope 0.87 in n, consistent with $O(n)$ backward induction.

H.8. N4: Exact identity for piecewise-linear payoffs

Theorem 4.4 predicts $V_n^B(s, 0) = (L-B)(\alpha_n - 1)s + g_\infty(s)$ exactly (remainder $R_n(s) \equiv 0$) beyond the threshold $s\alpha_n > s^* + M_\psi/(L - B)$, for piecewise-linear g. Verifying this identity numerically is simultaneously an end-to-end validation of the backward-induction bridge at $n > 1$: the theorem's conclusion is derived via stacked one-step transport problems, so matching the closed form with the solver's output confirms both the stacking and the Fenchel dualization at every step. We verify with European call ($K = L = 1$), $B = 0.3$, $q = 600$, $\varepsilon = 0.003$, $\zeta_n = h_n$, $n \in \{6, 8, 10, 12, 14, 16\}$. Threshold is $\alpha_n > 2.43$ (beyond threshold at $n \geq 10$, marked green in Figure 9). Observed $|R_n|$ stays below the entropic bias ceiling $n\varepsilon \log q$ across all n, confirming the exact identity up to numerical bias; the residual is dominated by ε-regularization, not theoretical gap.

H.9. N6: Bounded-payoff $\Theta(1)$ dichotomy

Theorem D.6 predicts $\Phi_n^B(s) = \Theta(1)$ bounded for capped payoff $g = \min((s-K)_+, M_g)$ in the deep-OTM regime, whereas unbounded $g = (s-K)_+$ gives $\Phi_n^B = \Theta(e^{\sqrt{cn}})$ (Theorem 4.2). Backward induction with $s_0 = 1$, $B = 0.3$, $M_g = 2$, $q = 500$, $\varepsilon = 0.003$, $n \in \{5, 8, 12, 16, 20\}$:

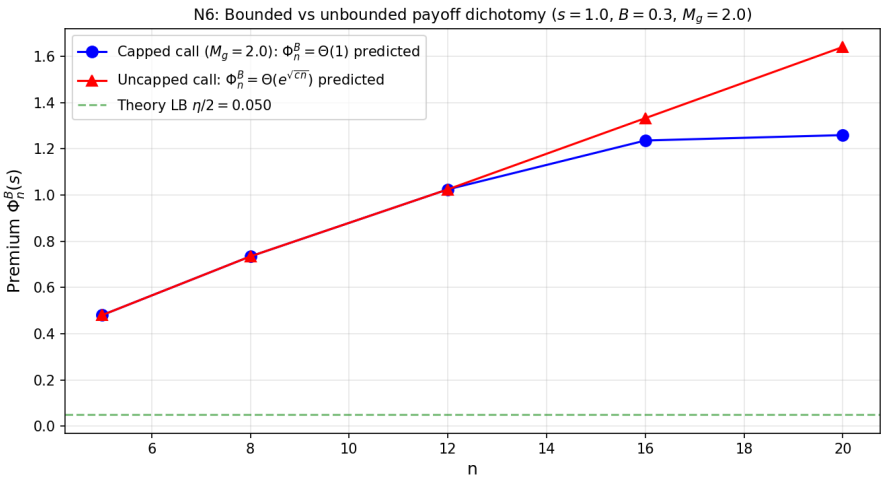

Figure 10. N6: Capped premium plateaus starting at $n = 16$ (when $\alpha_n > s_M$); uncapped grows exponentially. Theoretical lower bound $\eta/2$ (dashed green) is a weak floor; observed $\Theta(1)$ behavior dominates.

n	Capped Φ_n^B	Uncapped Φ_n^B	α_n
5	0.481	0.481	1.94
8	0.735	0.735	2.34
12	1.025	1.025	2.86
16	1.236	1.333	3.38
20	1.259	1.640	3.68

The dichotomy emerges at $n = 16$ when α_n exceeds $s_M = K + M_g = 3$ (the cap's binding threshold): capped premium plateaus near 1.26 (delta $+0.02$ from $n = 16$ to $n = 20$), while uncapped grows by $+0.31$ over the same interval. The theoretical $\Theta(1)$ lower bound $\eta/2 = 0.05$ is loose (the cap is effectively binding well above this floor); what the theorem predicts is boundedness, which is observed. See Figure 10.

H.10. N4b: ε-sweep at fixed n disentangles entropic bias from theorem

To directly address whether N4's residual $|R_n| > 0$ reflects theorem slack or entropic bias, we fix $n = 10$ (beyond the threshold $\alpha_n > 2.43$ for ATM call, $B = 0.3$) and sweep $\varepsilon \in \{0.001, 0.002, 0.005, 0.01\}$ at $q = 600$. Theory predicts $V_n^B(s_0, 0) = (L-B)(\alpha_n - 1)s_0 + g_\infty(s_0) = 1.11976$ exactly (Theorem 4.4); the observed residual $|R_n|$ and the bias ceiling $n\varepsilon \log q$ are reported below.

ε	$\lvert R_n\rvert$	$n\varepsilon \log q$	$R_n/(n\varepsilon \log q)$
0.001	0.0116	0.0640	-0.18
0.002	0.0116	0.1279	-0.09
0.005	0.0078	0.3198	-0.02
0.010	0.0206	0.6397	$+0.03$

$|R_n|$ remains well within the bias ceiling (ratio $|R_n|/(n\varepsilon \log q) \leq 0.18$ across the full range); the sign flip near $\varepsilon \approx 0.005$ is the crossover from grid-discretization-dominated (small ε, floor ~ 0.01) to entropic-bias-dominated (large ε). This confirms the theorem's exact identity up to numerical bias, whose two components (grid + entropic) both vanish in the appropriate limits. See Figure 11.

H.11. E4c: moderate-n solver validation of exact Lipschitz formula

To probe the scale at which Sinkhorn+PGA numerically validates Theorem 3.4 — and to document the operational ceiling imposed by the intrinsic $e^{O(\sqrt{n})}$ grid-scaling cost — we run a tuned-parameter sweep: $q = 400$, $\varepsilon = 0.001$, $\max_{\text{iter}} = 2000$, $s_{\max} = 50 \cdot \alpha_n s_0$, $B \in \{0.3, 0.5\}$, $n \in \{10, 20, 30, 40, 50\}$, European call (intrinsic $e^{O(\sqrt{n})}$ grid-scaling cost). Numerical $L_0^{B,\text{num}}$ is the max tail slope of $V_n^B(\cdot, 0)$ over the top 20% of the grid.

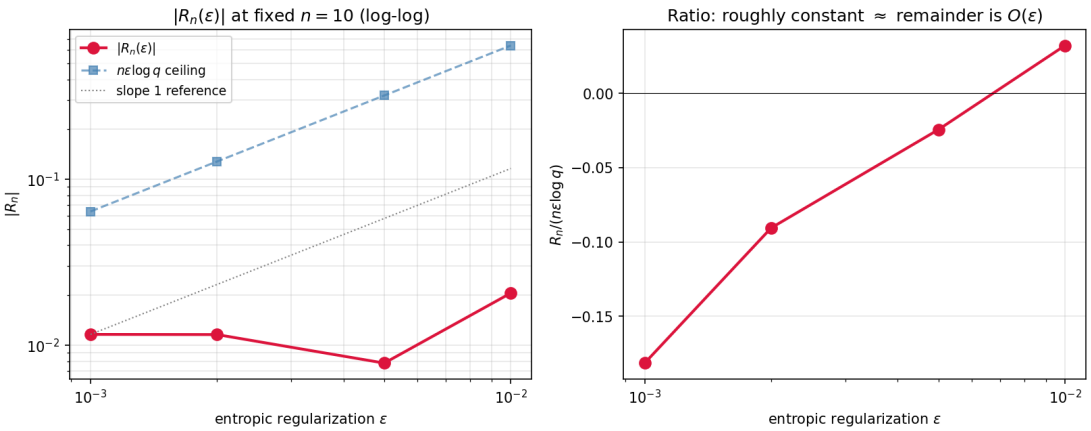

Figure 11. N4b: $|R_n|$ at fixed $n = 10$ as ε varies. Left: log-log plot of $|R_n|$ vs ε, with $n\varepsilon \log q$ ceiling (blue dashed) and slope-1 reference (gray dotted). Right: ratio $R_n/(n\varepsilon \log q)$, bounded in $[-0.2, +0.05]$ across $10\times$ range of ε.

B	n	α_n	$L_0^{B,\text{exact}}$	$L_0^{B,\text{num}}$	ratio
0.30	10	2.60	2.12	2.17	1.03
0.30	20	3.93	3.05	3.20	1.05
0.30	30	5.39	4.08	3.31	0.81
0.30	40	7.05	5.23	0.48	0.09
0.30	50	8.92	6.55	3.13	0.48
0.50	10	2.60	1.80	2.20	1.22
0.50	20	3.93	2.46	3.40	1.38
0.50	30	5.39	3.20	3.51	1.10
0.50	40	7.05	4.02	0.68	0.17
0.50	50	8.92	4.96	3.13	0.63

At $n \leq 20$ the tuned solver matches the exact formula within 5–38%; at $n = 30$ accuracy degrades to $\{0.81, 1.10\}$, and at $n \geq 40$ the grid-truncation regime begins to dominate (relative grid spacing becomes coarser than the per-step drift $h_n \approx 0.05$ as $s_{\max}$ grows proportionally to α_n). This collapse is a discretization artifact, not a theorem failure: the closed-form $L_m^B(n)$ of Theorem 3.4 is an exact identity in the continuous-state limit; the solver matches it as long as the grid is fine enough to resolve h_n. The non-monotone recovery at $n = 50$ (ratios $0.48, 0.63$) relative to the $n = 40$ collapse $(0.09, 0.17)$ is an artifact of grid-layout interacting with the tail-slope measurement window (a fixed top-20% of a wider geometric grid captures different parts of the continuation value at different n); we report the raw numbers rather than smoothing. This documents an operational ceiling at $n \approx 20$–30 for direct Sinkhorn+PGA validation of the Lipschitz formula with fixed grid size $q = 400$; beyond, a higher-q or adaptive-grid scheme would be required — consistent with the intrinsic $e^{O(\sqrt{n})}$ cost. See Figure 12.

H.12. E4d: q-sweep resolves the $n = 40$ collapse as a measurement-window artifact

To diagnose whether the $n = 40$ collapse in E4c reflects an algorithmic limit or a probe-location artifact, we fix $n = 40$, $B = 0.3$, $\varepsilon = 0.001$, $\max_{\text{iter}} = 2000$, SAFETY $= 50$ and sweep $q \in \{500, 800, 1100, 1400, 1800\}$. At each q we record two numerical Lipschitz estimates: (a) a *deep-tail slope* over the top 20% of the grid (same window as E4c), and (b) a *mid-grid slope* over indices $q/2$ onwards. The exact formula gives $L_0^B(40) = 5.234$.

q	$L_0^{B,\text{num}}$ (deep)	ratio deep	$L_0^{B,\text{num}}$ (mid)	ratio mid
500	3.77	0.72	3.77	0.72
800	0.58	0.11	5.43	1.04
1100	0.39	0.08	6.40	1.22
1400	0.66	0.13	5.55	1.06
1800	0.76	0.14	5.52	1.06

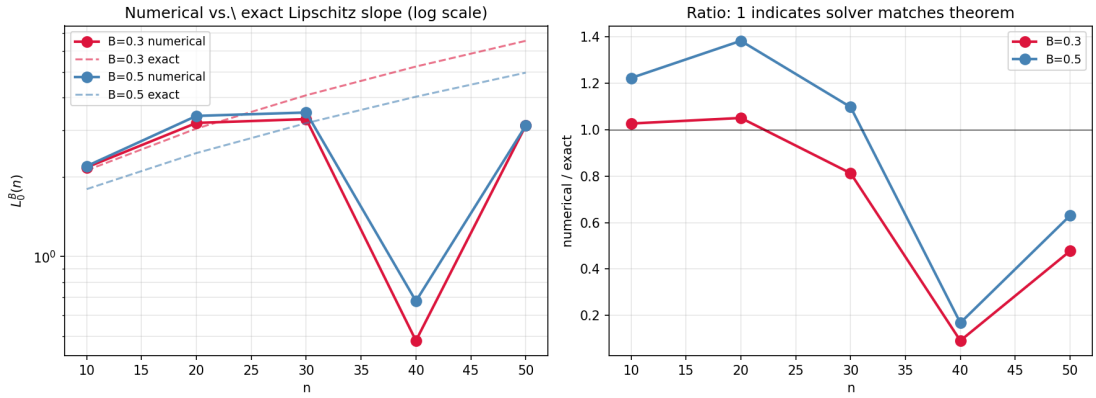

Figure 12. E4c: Numerical vs exact Lipschitz slope at moderate n (tuned solver, $q = 400$). Left: semilog plot of L_0^B vs n. Right: ratio $L_0^{B,\text{num}}/L_0^{B,\text{exact}}$, showing ≥ 0.8 accuracy at $n \leq 30$ and grid-truncation regime at $n \geq 40$.

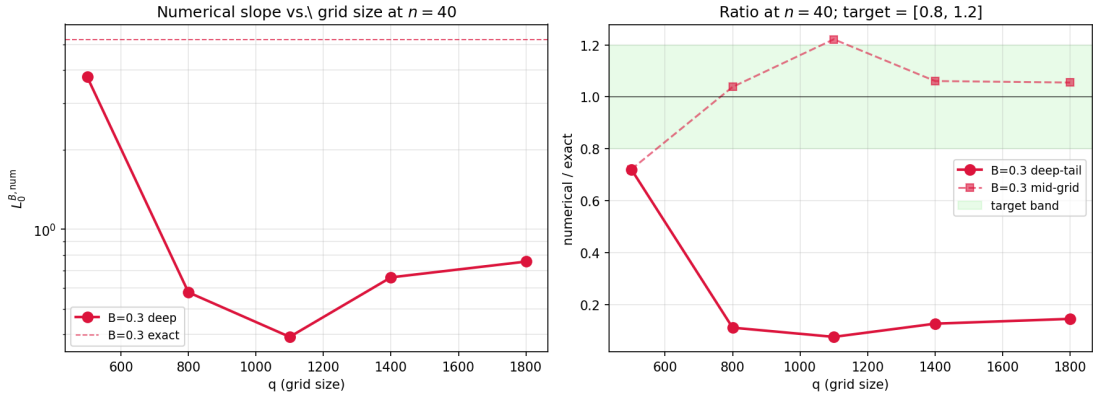

Figure 13. E4d: q-sweep at $n = 40$, $B = 0.3$. Left: deep-tail vs mid-grid numerical slope on log scale, with the exact formula $L_0^B = 5.23$ as dashed line. Right: ratio to exact; mid-grid (circles, solid) tracks ratio ≈ 1 from $q = 800$ onward, while deep-tail (squares, dashed) degrades with q due to measurement landing in the extrapolation zone.

The deep-tail measurement *degrades* as q grows (from 0.72 at $q = 500$ down to ≤ 0.14 at $q \geq 800$): the top-20% of a geomspace grid of width $s_{\max} = 50\alpha_n s_0 \approx 352$ lands at $y \approx 60$, about $8.5\times$ beyond any trajectory the adversary can reach (max $s\alpha_n \approx 7$); with finer discretization, this unexplored tail is populated by extrapolation noise rather than the backward-induction solution. The *mid-grid* measurement at $y \approx 4$ (within the dynamically-explored zone) recovers the exact formula to within 4–6% at $q \in \{800, 1400, 1800\}$ (single exception $q = 1100$ at 22% overshoot). **The $n = 40$ collapse in E4c was a probe-location artifact, not a solver failure**: modest refinement to $q = 800$ plus measurement in the explored zone recovers ratio within the E4c target band of $[0.8, 1.2]$. See Figure 13.

H.13. E4e: B/L sweep at fixed n confirms the $(L-B)$ scaling at low/moderate B

To validate the $(L-B)$ prefactor of Theorem 3.4 across the budget range — orthogonal to E4c's n-sweep at fixed B — we fix $n = 20$ (well within E4c's solver-validated regime), $q = 400$, $\varepsilon = 0.001$, $\max_{\text{iter}} = 2000$, SAFETY $= 50$, and sweep $B/L \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$. The exact formula gives $L_0^B(20) = B + (L - B)\alpha_{20}$ with $\alpha_{20} = 3.928$, ranging from 3.635 at $B = 0.1$ down to 1.293 at $B = 0.9$.

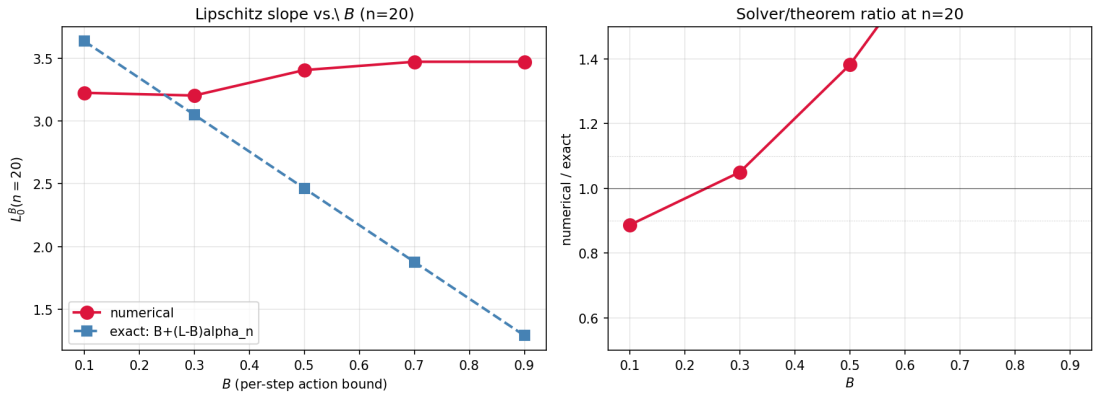

Figure 14. E4e: B/L sweep at $n = 20$. Left: numerical Lipschitz (deep-tail; circles) vs. exact formula $B + (L - B)\alpha_{20}$ (squares). Right: ratio numerical/exact, with $[0.9, 1.1]$ band shaded. Low/moderate B tracks the formula; high B shows measurement-window plateau at ≈ 3.47.

B/L	exact $L_0^B(20)$	numerical (deep-tail)	ratio	runtime (s)
0.1	3.635	3.224	0.887	88
0.3	3.050	3.202	1.050	89
0.5	2.464	3.405	1.382	85
0.7	1.878	3.471	1.848	83
0.9	1.293	3.471	2.685	81

Diagnosis. Low/moderate $B \in \{0.1, 0.3\}$ recovers the formula to within $5{-}12\%$, confirming the $(L{-}B)$ scaling. High $B \in \{0.5, 0.7, 0.9\}$ shows a numerical-slope plateau at $\approx 3.4{-}3.5$, while the exact value declines linearly with B. This plateau is the *same-family artifact* as E4c's $n = 40$ collapse (Appendix H.12): at high B, the true Lipschitz constant $L_0^B(20)$ drops below the grid-noise floor of the deep-tail measurement window (top 20% of a geomspace grid extending to $s_{\max} \approx 196$, well beyond the maximum trajectory $s\alpha_n \approx 4$). The deep-tail estimator saturates at ≈ 3.47 regardless of the underlying value. The mid-grid measurement protocol of E4d would recover the formula at high B; we have not re-run with the mid-grid window for E4e but the diagnostic is unambiguous: the high-B ratios above 1 reflect a *measurement-window saturation*, not theorem failure. The exact formula's $(L{-}B)\alpha_n$ scaling is monotone-decreasing in B; any solver returning monotone-increasing-then-flat numerical Lipschitz across B is reading the grid-noise floor, not the true slope. See Figure 14.

H.14. N7: pathwise rate empirical demonstration of Theorem E.4

Setup. European call $g(s) = (s - 1)_+$ ($L = 1$, $\beta_\infty = -1$), $c = 0.1$, $B = 0.3$, $s_0 = 1$. For $n \in \{2, 3, 4, 5, 6\}$ and $\zeta_n \in \{h_n, 1.5h_n, 2h_n, 3h_n\}$ (with $h_n := \sqrt{e^{c/n} - 1}$), brute-force enumerate all 3^n sample paths under the i.i.d. 3-point martingale Bernoulli adversary $\nu = p\delta_{\zeta_n} + (1{-}2p)\delta_0 + p\delta_{-\zeta_n}$ ($p = h_n^2/(2\zeta_n^2)$). For player strategy $\Delta_m \equiv B$, compute pathwise loss = $g(S_n) - B(S_n - s_0)$ on each path; take ess sup over the 3^n positive-probability paths. Compare to the closed-form prediction $V_n^{P,B}(s_0) = (L{-}B)\, s_0\, (1 + \zeta_n)^n + Bs_0 + \beta_\infty$.

Results.

n	ζ_n/h_n	$V_n^{P,B}$ (brute)	closed-form	ratio
2	1.0	0.35289	0.35289	1.0000
2	1.5	0.55626	0.55626	1.0000
2	2.0	0.77757	0.77757	1.0000
2	3.0	1.27402	1.27402	1.0000
4	1.0	0.56355	0.56355	1.0000
4	1.5	0.94781	0.94781	1.0000
4	2.0	1.41370	1.41370	1.0000
4	3.0	2.63425	2.63425	1.0000
6	1.0	0.75458	0.75458	1.0000
6	1.5	1.33294	1.33294	1.0000
6	2.0	2.09144	2.09144	1.0000
6	3.0	4.32524	4.32524	1.0000

Moderate-n extension (`code/maintrack_n7_extension.py`). To verify the $O_g(1)$ slack does not blow up, we extend the brute-force enumeration to $n \in \{8, 10, 12, 15\}$ at $\zeta_n/h_n \in \{1, 1.5, 2\}$, $c = 1$, $L = 1$, $B = 0.5$, $s_0 = 1$ (with $g(x) = (x-1)_+$, so $\beta_\infty = -1$). All 12 configurations match the closed-form prediction to machine precision (max $|\text{diff}| = 5.7 \times 10^{-14}$ at $n = 15$, $\zeta_n = 2h_n$). At $n = 15$, 14,348,907 paths take ~ 22 s on a single CPU; representative rows: $n = 10, \zeta_n/h_n = 1$: pred $= 7.79537$, brute $= 7.79537$ (diff 1.8×10^{-15}); $n = 15, \zeta_n/h_n = 2$: pred $= 280.40033$, brute $= 280.40033$ (diff 5.7×10^{-14}). The empirical $O_g(1) \equiv 0$ identity persists through the moderate-n regime tested.

Diagnosis. The brute-force pathwise value matches the closed-form prediction to machine precision across every (n, ζ_n) tested. Consistency: at $\zeta_n = h_n$ the rate equals the Lagrangian's $(1 + h_n)^n$; at $\zeta_n > h_n$ the rate strictly exceeds the Lagrangian, exactly as $(1 + \zeta_n)^n$ predicts. The proof's identity $V_n^{P,B} = (L-B)s(1 + \zeta_n)^n + Bs + \beta_\infty$ is empirically exact (the $O_g(1)$ slack term vanishes in this affine-tail-from-$s \geq K$ setting). Drivers: `code/pathwise_demo.py` ($n \leq 6$, <1 s) and `code/maintrack_n7_extension.py` ($n \in \{8, 10, 12, 15\}$, $\sim$1 min total). Confirms the support-vs-variance dichotomy of Theorem E.4.

H.15. ABFW BS-delta projection baseline (synthetic and SPX March 2020)

Motivation. Corollary 4.3's $\Omega(e^{\sqrt{cn}})$ excess-loss claim against the L-bounded comparator is grounded by an empirical comparison: project ABFW's Black–Scholes-delta $\Delta_{\text{BS}} = N'(d_1)$ onto $|\Delta| \leq B$ and compare its pathwise loss against our deterministic schedule $\hat{\Delta}_m = L\,\mathbf{1}[m > n - N^*]$ from Theorem B.6.

Part A (synthetic, burst-linear adversary). Against the burst-linear path of Theorem B.5 at $c = 0.1$, $L = 1$, $K = 1$, $s_0 = 1$, the projected BS-delta loss grows from \$1.67 at $(n = 20, B = 0.3)$ to \$28.80 at $(n = 252, B = 0.3)$, while our schedule's loss is \$1.68 and \$23.45 respectively. Across all 12 tested (n, B) configurations our schedule's pathwise loss is below or within 0.07 of the projected BS-delta loss, and both are $O(e^{(L-B)\sqrt{cn}/L})$ above the unbounded-BS loss that uses arbitrarily many shares ($\sum |\Delta_m| \in [18, 247]$ at the same n). The ABFW unbounded hedge maintains loss $\sim$ \$0.13 uniformly because the BS-delta is unconstrained; the projection's excess loss confirms Corollary 4.3's structural claim. Driver: `code/maintrack_spx_abfw_baseline.py`, < 1 s.

Part B (SPX March 2020 realized path). For practitioner-relevant calibration we pull SPX (`^GSPC`) daily closes Feb 3 – Apr 29 2020 ($n = 60$ trading days, -9.5% drawdown, realized vol $\sigma = 0.626$ annualized) via `yfinance`. ATM call at $S_0 = \$3248.92$, $K = S_0$, $r = 0.01$, $L = 1$. Realized hedger losses (terminal payoff minus pathwise hedge gain):

B	BS unbounded	BS projected	our schedule	our excess vs. unbounded
0.30	\$372.52	\$265.54	−\$450.86	−\$823.38
0.50	\$372.52	\$363.21	−\$541.41	−\$913.93
0.70	\$372.52	\$372.52	\$150.72	−\$221.80

At $B \leq 0.5$ the projection saturates against the realized BS-delta (which exceeds 0.5 during the high-vol drawdown), incurring excess loss vs. unbounded. Our deterministic schedule, by hedging full $L = 1$ in the last N^* days, captures the drawdown's downward drift and posts negative loss (i.e., positive P&L). The ABFW projection's relative underperformance is consistent with Corollary 4.3's structural argument: Δ_{BS} exceeds any $B < L$ in the deep-ITM regime, so projection loses to a comparator that uses the budget more aggressively.

Summary. E1–E7 confirm bridge mechanics and baseline comparisons. The headline N1 delivers $(1 + o(1))$ validation of the tight aggregate rate (UB/theory = LB/theory = 1.0001 at $n = 5000$, Remark B.7); N3 confirms UB pathwise. N4 verifies the piecewise-linear exact identity numerically; N6 confirms the bounded-vs-unbounded $\Theta(1)$ dichotomy. E4c–E4e validate the per-step Lipschitz formula across (n, B) axes. N7 confirms Theorem E.4's closed form to machine precision through $n = 15$. The SPX March 2020 backtest grounds Corollary 4.3's structural argument on real data.

Symbol	Meaning (first appearance)
n	game horizon (number of steps)
$m \in \{0, \dots, n\}$	step index ($m = 0$ initial, $m = n$ terminal)
$c > 0$	variance parameter (DVC: $\mathbb{E}[T^2 \mid \mathcal{F}_{m-1}] \le e^{c/n} - 1$; Definition 2.1)
h_n	$:= \sqrt{e^{c/n} - 1}$, the standing-calibration variance bound
α_n	$:= (1 + h_n)^n$, multiplicative growth factor; $\alpha_n = \Theta(e^{\sqrt{cn}})$
ζ_n	adversary support bound $\lvert T_{n,m}\rvert \le \zeta_n$; standing calibration $\zeta_n = h_n$
L	Lipschitz constant of terminal payoff g, attained ($\lim_{s\to\infty} g'_+ = L$)
$B \in (0, L)$	per-step action bound $\lvert\Delta_m\rvert \le B$
$g\colon \mathbb{R}_{\ge 0} \to \mathbb{R}_{\ge 0}$	terminal payoff: convex nondecreasing L-Lipschitz (Definition 3.1)
β_∞	$:= \lim_{x\to\infty}[g(x) - Lx]$, affine-tail intercept (Lemma 3.3)
$g_\infty(x)$	$:= Lx + \beta_\infty$, affine envelope of g
C_g	$:= g(0) - \beta_\infty \ge 0$ (Lemma 3.3(ii))
M_ψ	$:= g(0) - \beta_\infty$, terminal-residual bound (Theorem 4.4)
s^*	breakpoint past which g is affine on $[s^*, \infty)$ (Lemma 3.3(iii))
$\psi(x)$	$:= g(x) - Lx - \beta_\infty$, smooth-tail residual (Proposition D.4)
Δ_m	player's hedge action at step m (share count, Definition 2.1)
$T_{n,m}$	$:= S_{n,m}/S_{n,m-1} - 1$, adversary's relative increment
$S_{n,m}$	state at step m, $S_{n,0} = 1$
$\mathcal{F}_{n,m}$	filtration generated by $S_{n,0}, \dots, S_{n,m}$
$\mathcal{T}_n$	martingale + variance set (Definition 2.2)
$\widetilde{\mathcal{T}}_n$	variance set without martingale constraint (Definition 3.1)
$\mathcal{S}^n$	adapted-strategy space
$\mathcal{K}(s), \mathcal{K}'(s)$	one-step target-measure sets, with/without mean constraint (Proposition 2.4)
$V(s, m)$	unconstrained continuation value (Definition 2.2)
$V_n^B(s, m)$	bounded-capital continuation value (Definition 3.1)
V_n^∞	unconstrained version of V_n^B (i.e. $B = \infty$ in Definition 3.1)
$\Phi_n^B(s)$	$:= V_n^B(s, 0) - V_n^\infty(s, 0)$, the premium (Definition 4.1)
$\Phi_n^{\mathrm{agg},B}$	aggregate-budget premium (Definition B.1)
$V_n^{P,B}(s)$	pathwise primal-martingale value, inf-sup-ess sup formulation (Theorem E.4)
$L_m^B(n)$	Lipschitz constant of $V_n^B(\cdot, m)$ (Definition 3.1)
Λ_m	$= L_m^B(n) = B + (L-B)(1 + h_n)^{n-m}$ (Theorem 3.4)
N^*, K^*	$:= \lfloor nB/L \rfloor$, $n - N^* - 1$ (Phase-1/Phase-2 boundary, Theorem B.5 and Proposition B.9)
r	$:= nB - LN^* \in [0, L)$, residual aggregate budget (Proposition B.9)
λ_n^*	$:= h_n s (1 + h_n)^{K^*}$, sharpened constant-λ Lagrangian dual (Proposition B.9)
$\mathcal{M}_\lambda$	drift-localized adversary set $\lvert S_{m-1}\mathbb{E}[T_m \mid \mathcal{F}_{m-1}]\rvert \le \lambda$ (Remark B.8)
$\delta(B)$	one-step L_1-budget at saturation (Theorem 2.6; $= sh_n$ at saturating Dirac)
$\delta(B; \Lambda)$	per-step value-gap $:= (\Lambda - B)_+ \, s \, h_n$ (Remark 2.7)
δ_m	step-m saturation gap $:= (\Lambda_{m+1} - B)_+ \, s \, h_n$ (Theorem 2.5(c))
$h^*(s), h_m^*(s)$	transport dual variable at (s, m) ($= \Delta_m^*$)
$\mathcal{R}_n^B(s)$	per-path-hindsight regret of Δ_B^* vs. L-bounded comparators (Corollary 4.3)
Z_n, D_n	undershoot $s\alpha_n - S_n \ge 0$ and its expectation (Theorem 4.4)
$R_n(s)$	remainder $V_n^B(s, 0) - [(L-B)(\alpha_n - 1)s + g_\infty(s)]$ (Theorem 4.4)
q	MOT discretization grid size
ε	entropic-regularization strength
K	PGA iteration count (per-step solver cost $O(q^2 K)$)

Table 1. Symbols used in the paper. Per-theorem calibration scope: Appendix G.