Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Optimal Harvest Timing under Nonlinear Biomass Growth and Environmental Memory: Stochastic Control, Annual Catch Quotas, and Ecological Risk

Jules SADEFO KAMDEM

jules.sadefo-kamdem@umontpellier.fr

Université de Montpellier

Research Article

Keywords:

Posted Date: November 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11763282/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Optimal Harvest Timing under Nonlinear Biomass Growth and Environmental Memory: Stochastic Control, Annual Catch Quotas, and Ecological Risk

Jules SADEFO KAMDEM
Université de Montpellier, Montpellier Recherche en Économie (MRE, UR 209)
Montpellier, France

Abstract

We investigate optimal harvest timing when renewable-resource biomass follows nonlinear Gilpin–Ayala growth driven by independent Brownian and fractional Brownian disturbances. For $H \in (1/2, 1)$, the latter is interpreted in the Young sense and introduces predictable environmental persistence. An explicit positive biomass representation and an exact finite-grid conditional dynamic program clarify why current stock need not be a sufficient decision state. We establish sufficient conditions for a unique history-conditioned intervention threshold on a bounded stock interval and show that favorable fractional innovations raise the threshold when a single intervention remains and one subsequent exercise opportunity is available. We align all admissible harvest dates with the numerical monitoring chronology and distinguish model-based horizon-crossing risk from fisheries-management time-frequency targets. New, independently implemented simulations use the corrected decision grid, paired common shocks, both annual-intensity and constant-step-variance normalizations, and independent ecological certification after policy selection. A nested comparison of constant, time-varying stock-based, and fractional-predictor-based rules quantifies the incremental predictive-information channel without equating it with the value of the full fractional history. A compact extension imposes exogenous annual catch quotas through a remaining-allowance state and establishes monotonicity of the constrained value in the available quota. A separate short simulation quantifies quota effects within a restricted policy family without claiming structurally calibrated quota levels. Public Pacific halibut stock-assessment summaries provide scale and motivation, not structural parameter identification. All numerical claims are limited to the implemented policy families and reported Monte Carlo uncertainty.

Keywords: optimal harvesting; nonlinear growth; fractional Gaussian noise; impulse control; non-Markovian information; ecological chance constraints; catch quotas; policy evaluation.
JEL classification: C61, C63, Q22, Q23, Q57.

1 Introduction

Renewable-resource harvesting involves a fundamental intertemporal trade-off: extraction yields revenue today but diminishes the biological capital needed for future growth. Even with a fixed economic environment, intervention costs make harvesting a timing and sizing problem. Stochastic shocks create an option value of waiting. When favorable environmental conditions persist, a manager may delay harvest to benefit from future stock growth; annual catch limits further affect whether waiting leaves enough scope to harvest later. This simple biological-economic tension motivates the informational and regulatory components below.

The classical escapement framework [3], stochastic fishery adjustment models [10–12], and fixed-cost real-options models [4, 6] mostly employ Markovian biological states. Gilpin–Ayala density dependence [5] offers a flexible alternative to logistic growth; recent work has analyzed such growth in regime-switching environments and with clustered Hawkes disturbances [8, 9].

Fractional Brownian motion, introduced by Mandelbrot and van Ness [16], provides a different type of dependence distributed across many temporal lags. Persistent environmental noise is empirically plausible [13, 14], but persistence alone does not identify a fractional model because structural change and regime switching can mimic long memory [15].

This paper isolates the theoretical consequence of fractional information for costly partial interventions. Under the mixed Itô–Young specification, we first establish positivity and an explicit solution between impulses. We then formulate a conditional finite-grid Bellman recursion and prove a sufficient stock-space single-crossing result conditional on realized environmental history. This is a local, finite-horizon theorem, not a global smooth-fit or free-boundary claim. A constructive example demonstrates that its conditions can hold simultaneously; numerical positivity of the local derivative is not treated as a computer-assisted certificate. Related non-Markovian impulse-control existence work includes Djehiche et al. [18], Perninge [19], Jönsson and Perninge [20], Perninge [21], while Liu and Zervos [22] illustrates the rich qualitative regimes that can arise even under Markovian impulse harvesting.

The empirical and computational portions are deliberately narrower than structural fisheries management. We first distinguish female relative spawning biomass used in the International Pacific Halibut Commission (IPHC) Harvest Strategy Policy from normalized Gilpin–Ayala biomass. We subsequently use a corrected monitoring chronology and independent reference code to evaluate simple candidate policies. Two fractional-noise normalizations help distinguish changes in dependence from changes in stepwise marginal variance. Finally, nested policy classes separate the value of time flexibility from the incremental value of a one-step fractional predictor, and an independent sample separates policy search from ecological certification. A short quota extension uses an exogenous annual catch allowance, rather than adding a separate regulator optimization problem, to clarify which results would change under catch limits [23].

2 Economic and stochastic environment

2.1 Filtered probability space and admissible information

Fix $T < \infty$ and a complete probability space $(\Omega, \mathcal{A}, \mathbb{P})$ carrying independent processes $W = (W_t)_{0\le t\le T}$ and $B^H = (B^H_t)_{0\le t\le T}$, with W a standard Brownian motion and B^H a centered fractional Brownian motion of Hurst index $H \in (1/2, 1)$. Let $\mathcal{N}$ denote the $\mathbb{P}$-null sets and define the right-continuous completed filtration

$$\mathcal{F}_t = \bigcap_{u\in(t,T]} (\sigma(W_s, B^H_s : 0 \le s \le u) \vee \mathcal{N}), \quad t < T, \qquad \mathcal{F}_T = \sigma(W_s, B^H_s : 0 \le s \le T) \vee \mathcal{N}. \tag{1}$$

These are the usual conditions. Independence ensures that W remains Brownian relative to $(\mathcal{F}_t)$; B^H is adapted and continuous, but is generally neither Markovian nor a martingale. We use the Itô integral with respect to W and the pathwise Young integral with respect to B^H. Since $H > 1/2$, one can choose a pathwise stock regularity exponent $\alpha \in (1 - H, 1/2)$ between harvest times, giving $\alpha + H > 1$ for Young integration. The resulting stock is adapted and continuous between intervention times, and is càdlàg after its multiplicative harvest jumps.

Admissible intervention dates τ_i are $(\mathcal{F}_t)$-stopping times valued in the predetermined monitoring grid, and effort or catch choices are $\mathcal{F}_{\tau_i}$-measurable, subject to the stated feasibility constraints. The observation model used for numerical prediction assumes the shocks are separately observable: at t_j the manager knows $m_j = (\xi_0, \dots, \xi_{j-1})$, as well as the stock and past actions. Formally, $\mathcal{G}_j$ is the completion of the sigma-field generated by $X^+_{t_0}$, $X^-_{t_k}$ for $1 \le k \le j$, and $X^+_{t_k}$, ΔW_k, ξ_k and intervention choices for $k < j$ (with the natural indexing of between-date innovations). Thus $\mathcal{G}_j \subseteq \mathcal{F}_{t_j}$ and, importantly, current post-decision stock $X^+_{t_j}$ is not used to choose its own harvest. All sampled-grid Bellman conditional expectations

below are taken with respect to $\mathcal{G}_j$; their Gaussian conditional means use exactly the observed fractional increments. Observing these latent shocks is an explicit information assumption, not an inference from stock-assessment data.

2.2 Nonlinear biomass dynamics

Let $X_t > 0$ denote standing biomass, $r > 0$ intrinsic growth, $K > 0$ carrying capacity, and $\theta > 0$ the density-shape exponent. The Gilpin–Ayala drift is

$$G_\theta(x) = rx\Big[1 - (x/K)^\theta\Big]. \tag{2}$$

The logistic law is recovered at $\theta = 1$. The biological maximum-sustainable-yield reference is

$$X_{\mathrm{MSY}}(\theta) = K(1+\theta)^{-1/\theta}, \qquad G_{\mathrm{MSY}}(\theta) = rK\theta(1+\theta)^{-(1+1/\theta)}. \tag{3}$$

Neither quantity is generally an economically optimal impulse trigger.

Let W be standard Brownian motion and B^H an independent fractional Brownian motion, with $H \in (1/2, 1)$ and

$$\mathbb{E}[B_t^H B_s^H] = \tfrac{1}{2}(t^{2H} + s^{2H} - |t-s|^{2H}). \tag{4}$$

Between harvests, biomass follows

$$dX_t = G_\theta(X_t)\,dt + \sigma_W X_t\,dW_t + \sigma_H X_t\,dB_t^H, \qquad X_0 = x > 0. \tag{5}$$

The Brownian integral is an Itô integral and the fractional integral is a Young integral. Thus σ_W has units time$^{-1/2}$ and σ_H units time^{-H}. No Brownian drift correction is assigned to the fractional component. Even when the independent mixed noise has a semimartingale representation for some values of H [17], biomass alone is not thereby made Markovian.

2.3 Harvest technology and payoff

An intervention with effort E at a monitoring date harvests $h = qEX^-$ and leaves $X^+ = (1-qE)X^-$. Fix $\eta \in (0,1)$ and a compact effort set

$$0 < \underline{e} \le E \le \overline{e} \le (1-\eta)/q. \tag{6}$$

With inverse demand $p(h) = ah^b$, $a > 0$, $-1 < b \le 0$, define $\beta = 1 + b \in (0, 1]$. The harvest payoff is

$$\pi_E(x) = a(qEx)^\beta - cE - f, \qquad a, c, f > 0. \tag{7}$$

In harvest-quantity coordinates, $R(x,h) = ah^\beta - ch/(qx) - f$. There is no fixed cost when no intervention occurs. Let $S \ge 0$ denote terminal salvage value; the numerical illustrations below use $S \equiv 0$.

2.4 Monitoring and admissible dates: corrected chronology

Fix $0 = t_0 < t_1 < \cdots < t_M = T$, with equal step $\Delta = T/M$ in the numerical experiments. Biomass $X_0^+ = x_0$ is observed at t_0 *after* the initial biological/management state has been set. To align biological propagation and management decisions, our admissible intervention set is

$$\mathcal{T}_{\mathrm{harvest}} = \{t_1, \ldots, t_{M-1}\}. \tag{8}$$

Thus neither t_0 nor terminal date $T = t_M$ is an intervention opportunity. The chronology is

$$X_{j-1}^+ \xrightarrow{\text{no-harvest evolution}} X_j^- \xrightarrow{\text{decision at } t_j} X_j^+, \quad 1 \le j \le M-1, \tag{9}$$

followed by a final no-harvest evolution to X_M. At most one intervention is allowed per date. The numerical policy class imposes no additional cap, so $N_{\max} = M - 1$; the theoretical recursion also permits a smaller stated cap.

For $\pi = \{(\tau_i, E_i)\}$ with $\tau_i \in \mathcal{T}_{\text{harvest}}$, the finite-grid objective is

$$J_T(x_0; \pi) = \mathbb{E}\left[\sum_{\tau_i \in \mathcal{T}_{\text{harvest}}} e^{-\rho\tau_i} \pi_{E_i}(X_{\tau_i}^-) + e^{-\rho T} S(X_T)\right]. \tag{10}$$

This convention deliberately differs from the previous numerical formulation that allowed intervention at t_M; none of those previous policy returns is carried over as a revised result.

3 Biomass representation and history-dependent control

3.1 Closed-form representation and positivity

Proposition 3.1 (Explicit positive biomass solution)**.** *For $r, K, \theta > 0$, $\sigma_W, \sigma_H \geq 0$ and $H \in (1/2, 1)$, the mixed Itô–Young equation* (5) *has, between impulses, the unique positive solution*

$$F_t = \exp\left[\left(-\theta r + \frac{\theta\sigma_W^2}{2}\right) t - \theta\sigma_W W_t - \theta\sigma_H B_t^H\right], \tag{11}$$

$$X_t = \left[F_t\left(x^{-\theta} + \frac{\theta r}{K^\theta}\int_0^t F_s^{-1}\, ds\right)\right]^{-1/\theta}. \tag{12}$$

Any finite sequence of admissible partial interventions preserves strict positivity on $[0, T]$.

Proof. Apply Itô's formula to $Y = X^{-\theta}$ for the Brownian part and the Young chain rule for the fractional part. Because B^H has zero quadratic variation when $H > 1/2$,

$$dY_t = \left[\left(-\theta r + \frac{\theta(\theta+1)}{2}\sigma_W^2\right) Y_t + \frac{\theta r}{K^\theta}\right] dt - \theta\sigma_W Y_t\, dW_t - \theta\sigma_H Y_t\, dB_t^H. \tag{13}$$

The integrating factor of the homogeneous linear equation is (11), with Brownian Itô correction $-\frac{1}{2}\theta^2\sigma_W^2 t$ in the logarithm. Variation of constants gives (12). Both F_t and the bracketed term are strictly positive and finite on a finite interval. At a partial harvest the stock is multiplied by at least $\eta > 0$; iterate over the finitely many dates. □

Corollary 3.2 (Positive stock versus ecological safety)**.** *The model cannot hit zero at finite time under admissible harvesting, but it can cross any positive critical biomass. Ecological shortfall should therefore be based on $\min_j X_j^+$ or, if studied separately, $\inf_{t \leq T} X_t$, rather than literal extinction at zero.*

3.2 Finite-dimensional Gaussian prediction

Let $\xi_j = B_{t_{j+1}}^H - B_{t_j}^H$ denote raw, not standardized, fractional increments and $m_j = (\xi_0, \ldots, \xi_{j-1})$. On an equispaced grid, the fGn covariance is

$$\operatorname{Cov}(\xi_j, \xi_{j+\ell}) = \Delta^{2H}\gamma_H(\ell), \quad \gamma_H(\ell) = \tfrac{1}{2}\left(|\ell+1|^{2H} - 2|\ell|^{2H} + |\ell-1|^{2H}\right). \tag{14}$$

For $j \geq 1$, write $\Gamma_{pp} = \operatorname{Cov}(m_j)$, $\Gamma_{fp} = \operatorname{Cov}(\xi_j, m_j)$ and $\gamma_0 = \Delta^{2H}$. Then

$$\xi_j \mid m_j \sim N(\mu_j(m_j), v_j), \quad \mu_j = \Gamma_{fp}\Gamma_{pp}^{-1} m_j, \quad v_j = \gamma_0 - \Gamma_{fp}\Gamma_{pp}^{-1}\Gamma_{pf} > 0. \tag{15}$$

At $j = 0$, $\mu_0 = 0$ and $v_0 = \Delta^{2H}$. With only the previous increment observed,

$$\mu_j = \kappa_H \xi_{j-1}, \qquad \kappa_H = 2^{2H-1} - 1 > 0. \tag{16}$$

A predictor of a single next increment is generally not a sufficient state for the entire future fractional path. The distinction matters for the later policy comparison.

3.3 Conditional dynamic programming

Let $\mathcal{H}_j$ denote the realized variables generating $\mathcal{G}_j$ from Section 2.1. Conditional on $\mathcal{G}_j$, the sampled mixed-noise scheme below is well defined. Here $\mathbb{E}[\cdot \mid \mathcal{H}_j]$ abbreviates conditioning on $\mathcal{G}_j$. Let $V_{j,n}$ denote optimal *time-t_j* value before deciding, with n harvest opportunities remaining. Set

$$V_{M,n}(x, \mathcal{H}_M) = S(x), \qquad n = 0, \dots, N_{\max}. \tag{17}$$

For $j = M - 1, \dots, 1$, define

$$C_{j,n}(x, \mathcal{H}_j) = e^{-\rho\Delta} \mathbb{E}\left[V_{j+1,n}(X_{j+1}^x, \mathcal{H}_{j+1}) \mid \mathcal{H}_j\right], \tag{18}$$

$$I_{j,n}(x, \mathcal{H}_j) = \max_{E \in [\underline{e}, \overline{e}]} \left\{\pi_E(x) + e^{-\rho\Delta} \mathbb{E}\left[V_{j+1,n-1}(X_{j+1}^{(1-qE)x}, \mathcal{H}_{j+1}) \mid \mathcal{H}_j\right]\right\}, \quad n \geq 1, \tag{19}$$

$$V_{j,n}(x, \mathcal{H}_j) = \max\{C_{j,n}(x, \mathcal{H}_j), I_{j,n}(x, \mathcal{H}_j)\}, \qquad I_{j,0} := -\infty. \tag{20}$$

Here X_{j+1}^z denotes between-date biomass propagated from the post-decision stock z using the same conditional innovations. Since there is no decision at t_0, the initial value is

$$V_{0,n}^{\text{initial}}(x_0, \mathcal{H}_0) = e^{-\rho\Delta} \mathbb{E}[V_{1,n}(X_1^x, \mathcal{H}_1) \mid \mathcal{H}_0]. \tag{21}$$

Proposition 3.3 (Finite-grid Bellman principle)**.** *Suppose terminal salvage is continuous with polynomial growth, the effort set is compact, and the stock transition is continuous in its initial stock with integrable polynomial moments. Equations* (17)–(21) *give the exact finite-grid value among policies adapting to the specified observation filtration and using dates* (8)*.*

Proof. At t_M only salvage is available. At each earlier decision date, either wait or harvest once and reduce remaining opportunities by one. Conditional integration, compactness, and backward induction give the recursion and existence of maximizing effort. No Markov property of biomass alone is invoked. At t_0, waiting is forced, giving (21). □

Remark 3.4 (Exact process versus monitored approximation). For the exact equation (5), an exact Bellman recursion conditional on the full environmental path is conceptually available at the same monitoring dates. However, the finite predictor (15) describes the *sampled-increment* information model and is exact for the split-step transition, not a claim that sampled increments summarize a continuously observed fractional bridge.

4 Conditional harvest boundary

4.1 Exact deterministic flow and split-step propagation

The deterministic Gilpin–Ayala flow is

$$\phi_u(x) = \frac{x e^{ru}}{\left[1 + (x/K)^\theta (e^{r\theta u} - 1)\right]^{1/\theta}}. \tag{22}$$

The positivity-preserving split-step transition is

$$\Psi_j(x,z) = \phi_{\Delta/2}(e^z \phi_{\Delta/2}(x)), \quad z = \sigma_W \Delta W_j - \tfrac{1}{2}\sigma_W^2 \Delta + \sigma_H \xi_j. \tag{23}$$

The spatial derivative satisfies

$$\phi_u'(x) = e^{ru}\left[1 + (x/K)^\theta (e^{r\theta u} - 1)\right]^{-1-1/\theta} > 0, \quad 0 < \partial_x \Psi_j(x,z) \le e^{r\Delta + z}. \tag{24}$$

Hence the discounted conditional slope factor is

$$e^{-\rho\Delta}\mathbb{E}_{m_j}\left[e^{r\Delta+z}\right] = \exp\left[(r-\rho)\Delta + \sigma_H \mu_j(m_j) + \tfrac{1}{2}\sigma_H^2 v_j\right]. \tag{25}$$

The Brownian Itô correction cancels the Brownian lognormal moment; it is not duplicated for the fractional component.

4.2 Lipschitz bounds and threshold existence

Assume S is nonnegative, increasing, globally Lipschitz and $S(0) = 0$. Let $C_{j,n}$ and $I_{j,n}$ now refer to the split-step model. Write

$$x_{\text{low}} = \frac{(f/a)^{1/\beta}}{q\overline{e}}, \qquad L_\pi = a\beta(q\overline{e})^\beta x_{\text{low}}^{\beta-1}, \qquad L_{M,n} = \text{Lip}(S). \tag{26}$$

For $j = M-1, \dots, 1$, define the conditional quantities

$$\Lambda_{j,n}(m_j) = e^{-\rho\Delta}\mathbb{E}_{m_j}\left[e^{r\Delta+z} L_{j+1,n}(m_{j+1})\right], \tag{27}$$

$$L_{j,0} = \Lambda_{j,0}, \qquad L_{j,n} = \max\{\Lambda_{j,n}, L_\pi + \Lambda_{j,n-1}\}, \quad n \ge 1. \tag{28}$$

Lemma 4.1 (Monotonicity and Lipschitz control)**.** *Under these assumptions, each $V_{j,n}(\cdot, m_j)$ is nonnegative, increasing, continuous, and globally Lipschitz, and $\text{Lip}_x(C_{j,n}) \le \Lambda_{j,n}(m_j)$ and $\text{Lip}_x(V_{j,n}) \le L_{j,n}(m_j)$.*

Proof. The transition is increasing in initial biomass, as are the gross-revenue and terminal-value functions; an additional opportunity cannot reduce value. For $x \le x_{\text{low}}$, one has $a(qEx)^\beta \le a(q\overline{e}x)^\beta \le f$ for any admissible E, so $\pi_E(x) \le -cE \le 0$. Actions with nonpositive current profit are dominated by waiting since they reduce stock and consume an opportunity. Therefore current profit can be replaced by $\pi_E^+ = \max\{0, \pi_E\}$ inside the outer maximum without changing V. On the support of positive profit, $x > x_{\text{low}}$ and $0 \le \pi_E'(x) = a\beta(qE)^\beta x^{\beta-1} \le L_\pi$. The positive-part payoff is thus uniformly Lipschitz. The transition slope bound (24), conditional integration, and backward induction prove (27)–(28). The required Gaussian exponential moments are finite for the finite horizon and fixed sampled histories. □

Define the intervention advantage $D_{j,n} = I_{j,n} - C_{j,n}$. Let $[0, B]$ be a relevant finite stock interval, $B > 0$.

Theorem 4.2 (Unique conditional threshold on a bounded interval)**.** *Fix $1 \le j < M$, $n \ge 1$, and a realized history m_j. Suppose there is $\ell \in (0, B)$ such that*

$$a(q\overline{e}\ell)^\beta < c\underline{e} + f, \tag{29}$$

$$a\beta(q\underline{e})^\beta B^{\beta-1} > \Lambda_{j,n}(m_j), \tag{30}$$

$$\pi_{E_0}(B) > C_{j,n}(B, m_j) \quad \textit{for some } E_0 \in [\underline{e}, \overline{e}]. \tag{31}$$

These are, respectively, the low-stock non-intervention, marginal-dominance, and high-stock intervention conditions. Then $D_{j,n}$ has a unique zero $x^*_{j,n}(m_j) \in (\ell, B)$, with waiting below the zero and harvesting above it. Moreover for $\ell \le x < y \le B$,

$$D_{j,n}(y, m_j) - D_{j,n}(x, m_j) \ge \delta_{j,n}(m_j)(y - x), \quad \delta_{j,n} := a\beta(q\underline{e})^\beta B^{\beta-1} - \Lambda_{j,n}(m_j) > 0. \tag{32}$$

When the three conditions hold uniformly with positive margins over a compact set of histories, the unique threshold varies continuously with history. Measurable maximizing efforts exist.

Proof. For a fixed effort E and $\ell \le x < y \le B$, concavity with $\beta \le 1$ gives

$$\pi_E(y) - \pi_E(x) = a(qE)^\beta(y^\beta - x^\beta) \ge a\beta(q\underline{e})^\beta B^{\beta-1}(y - x).$$

The post-harvest continuation value $C_{j,n-1}((1 - qE)x, m_j)$ is increasing in x, and the waiting value is $\Lambda_{j,n}$-Lipschitz by Lemma 4.1. Consequently each fixed-effort advantage satisfies (32). Taking its maximum over effort preserves the same lower increment bound. Condition (29) ensures all impulse payoffs are negative for $x \le \ell$; waiting dominates immediate intervention, including its reduced continuation stock and remaining opportunity count. Condition (31) implies positive advantage at B since subsequent continuation is nonnegative. Continuity and strict increase yield exactly one zero. Continuity in finite Gaussian history and the measurable maximum theorem give the final assertions. □

Remark 4.3 (Why the two effort endpoints differ). The *upper* effort $\overline{e}$ appears in L_π because it bounds the derivative of positive immediate profits *from above*. The *lower* effort $\underline{e}$ appears in (30) because the proof bounds the derivative of current revenue *from below uniformly over effort*. Confusing these bounds changes the certified single-crossing margin. The constant Λ certifies uniqueness but is not inserted into the threshold equation $D(x^*, m) = 0$.

Remark 4.4 (Grid kink versus continuous-time smooth fit). The strict crossing in (32) yields a kink at a regular discrete threshold; continuous-time smooth fit must not be imposed on this finite-grid value function. The theorem asserts neither a global half-line trigger nor continuous-time smooth-fit regularity.

4.3 Last remaining harvest: predictive-news effect

Let $S = 0$ and $g(x) = \max\{0, \max_{E \in [\underline{e}, \overline{e}]} \pi_E(x)\}$. Because t_{M-1} is the final intervention date,

$$V_{M-1,1}(x, m) = g(x), \qquad C_{M-2,1}(x, \mu) = e^{-\rho\Delta}\mathbb{E}\left[g(\Psi_{M-2}(x, Z_\mu))\right], \tag{33}$$

where $Z_\mu = \sigma_W\sqrt{\Delta}Z_W - \frac{1}{2}\sigma_W^2\Delta + \sigma_H(\mu + \sqrt{v_{M-2}}Z_H)$ and Z_W, Z_H are independent standard Gaussian variables.

Corollary 4.5 (Effect of favorable predictive innovations)**.** *Suppose Theorem 4.2 applies at $j = M - 2, n = 1$ for two predictive means $\mu_- < \mu_+$, with $\sigma_H > 0$, zero salvage, and g strictly increasing on a reachable set of positive probability. Then the two interior thresholds obey*

$$x^*(\mu_+) > x^*(\mu_-). \tag{34}$$

At differentiable regular crossings, $\partial_\mu x^* = \partial_\mu C/(g'(x^*) - \partial_x C) > 0$.

Proof. Under common-normal coupling, $Z_{\mu_+} > Z_{\mu_-}$ almost surely. The split-step stock transition is strictly increasing in its shock argument, and g is nondecreasing with strict increase on a reachable set. Therefore $C(x, \mu_+) > C(x, \mu_-)$ for relevant $x > 0$, whereas immediate last-harvest value $g(x)$ does not depend on predictive mean. By the unique positive-slope crossing, the root shifts strictly upward. Implicit differentiation yields the derivative formula. □

4.4 A sharper local crossing criterion

The stock-uniform margin in Theorem 4.2 is sufficient, not necessary. Its failure at economically relevant decision frequencies does not imply that the last-harvest advantage has multiple roots. For the final two admissible occasions, let $S = 0$, $D(x, \mu) = g(x) - C(x, \mu)$, with g and C as defined above. Where g is differentiable and strictly positive, its unique profit-maximizing effort and slope are

$$e^*(x) = \mathrm{clip}_{[\underline{e},\overline{e}]}\left[\left(\frac{a\beta(qx)^\beta}{c}\right)^{1/(1-\beta)}\right], \qquad g'(x) = a\beta(qe^*(x))^\beta x^{\beta-1}. \tag{35}$$

The split-step derivative and the actual conditional waiting-value derivative are

$$\partial_x \Psi_\Delta(x, z) = \phi'_{\Delta/2}\Big(e^z \phi_{\Delta/2}(x)\Big)\, e^z \phi'_{\Delta/2}(x), \tag{36}$$

$$\partial_x C(x, \mu) = e^{-\rho\Delta}\mathbb{E}\left[g'(\Psi_\Delta(x, Z_\mu))\partial_x \Psi_\Delta(x, Z_\mu)\right]. \tag{37}$$

The differentiation in (37) follows from the global Lipschitz bound on g, the stock-transition derivative bound, and integrability of Gaussian exponential moments; the derivative of g is interpreted almost everywhere.

Proposition 4.6 (Local single crossing with a positive-profit region)**.** *Fix a predictive mean μ and a finite $B > 0$. Suppose $g(x) = 0$ for $0 < x \le b$ and $g(x) > 0$ for $b < x \le B$, where $b \in (0, B)$. Assume $C(x, \mu) > 0$ for $0 < x \le b$, $g(B) > C(B, \mu)$, and, for some $\varepsilon_\mu > 0$,*

$$g'(x) - \partial_x C(x, \mu) \ge \varepsilon_\mu \quad \text{for almost every } x \in (b, B). \tag{38}$$

Then $D(\cdot, \mu)$ has exactly one zero $x^*(\mu) \in (b, B)$, and waiting is optimal below this zero while exercising the remaining harvest is optimal above it.

Proof. For $0 < x \le b$, one has $D(x, \mu) = -C(x, \mu) < 0$. On (b, B), D is absolutely continuous on compact subintervals. Integrating (38) shows $D(y, \mu) - D(x, \mu) \ge \varepsilon_\mu(y - x) > 0$ whenever $b < x < y < B$. Continuity, a negative value at b, and the positive endpoint value $D(B, \mu)$ imply a unique crossing, with the stated sign pattern. □

A positive derivative *only at a candidate root* establishes local transversality, not uniqueness on (b, B). Numerically scanning (38) is a useful diagnostic, but a computer-assisted proof requires validated bounds over every interval and control of quadrature error.

4.5 Exact between-date stock dynamics

The threshold argument is not tied to numerical splitting. If

$$\Xi_u = (r - \tfrac{1}{2}\sigma_W^2)u + \sigma_W \Delta W_u + \sigma_H \Delta B_u^H, \quad A_\Delta = \frac{\theta r}{K^\theta}\int_0^\Delta e^{\theta \Xi_u}\, du, \tag{39}$$

the exact no-harvest transition over a monitoring interval is

$$\mathcal{X}_\Delta(x) = \frac{x e^{\Xi_\Delta}}{(1 + A_\Delta x^\theta)^{1/\theta}}, \qquad 0 < \partial_x \mathcal{X}_\Delta(x) \le e^{\Xi_\Delta}. \tag{40}$$

The proof of Theorem 4.2 carries over on a fixed grid when conditional slope bounds are computed using the *full observed environmental history*. This does not extend the explicit two-normal comparison in Corollary 4.5 to arbitrary conditional fractional bridges without an additional order-preserving coupling assumption.

5 Annual catch quotas: a parsimonious extension

We introduce *exogenous* annual catch limits without a second-level regulator optimization. Let $Q_y \geq 0$ be an announced quota, measured in normalized stock units, and let $\mathcal{J}_y = \{8y+1, \ldots, 8y+8\} \cap \{1, \ldots, M-1\}$ be the eligible dates of year $y = 0, \ldots, 11$. A compliant harvest plan obeys, almost surely,

$$\sum_{j \in \mathcal{J}_y} h_j \leq Q_y, \qquad X_j^+ = X_j^- - h_j, \qquad h_j \geq 0. \tag{41}$$

Let R_j be the allowance just before the decision at t_j, initially $R_1 = Q_0$. For $x > 0$ define the positive-harvest action set and net payoff

$$\mathcal{A}^Q(x, R) = \big(0, \min\{R, q\bar{e}x, (1-\eta)x\}\big], \qquad \Pi(x, h) = ah^\beta - \frac{ch}{qx} - f. \tag{42}$$

Waiting ($h = 0$) is a separate action with zero current payoff and no fixed cost. The quota extension permits catch quantities corresponding to effort below the original $\underline{e}$. In the illustrative policies below this happens only through truncation of a target catch by a binding residual quota, avoiding artificial infeasibility; the quota-free theorem and its original effort constraints remain unchanged. At a subsequent eligible date, the remaining allowance updates as $\mathcal{U}_j(R, h) = R - h$ within a fishing year and as $Q_{y(j+1)}$ on entry into the next year (no banking). The conditional Bellman expressions are

$$C_{j,n}^Q(x, m, R) = e^{-\rho\Delta}\mathbb{E}_m\Big[V_{j+1,n}^Q(\Psi_j(x, Z), m', \mathcal{U}_j(R, 0))\Big], \tag{43}$$

$$I_{j,n}^Q(x, m, R) = \sup_{h \in \mathcal{A}^Q(x,R)} \Big\{\Pi(x, h) + e^{-\rho\Delta}\mathbb{E}_m\Big[V_{j+1,n-1}^Q(\Psi_j(x-h, Z), m', \mathcal{U}_j(R, h))\Big]\Big\}, \tag{44}$$

$$V_{j,n}^Q(x, m, R) = \max\{C_{j,n}^Q, I_{j,n}^Q\}, \qquad V_{M,n}^Q = S(x), \tag{45}$$

where $I^Q = -\infty$ if $n = 0$ or the action set is empty. The initial date t_0 still admits no intervention. Conditional expectation is over the next observed shocks, with the history updated before the next decision.

Proposition 5.1 (Value of quota flexibility)**.** *For fixed (x, m, j, n) and a fixed future quota schedule, $R_1 \leq R_2$ implies $V_{j,n}^Q(x, m, R_1) \leq V_{j,n}^Q(x, m, R_2)$, whenever the value functions are finite. Moreover, imposing finite annual quotas cannot raise the global optimum relative to the otherwise identical unrestricted* quantity-action *problem.*

Proof. Couple strategies to the same innovation path. Every catch sequence feasible from R_1 is feasible from R_2, with identical resulting stock and payoffs; in future years both sequences receive the same prescribed allowance. Thus the feasible strategy sets are nested and taking suprema proves the claim. The corresponding unrestricted quantity-action feasible set contains every quota-compliant strategy. This comparison uses the same quantity-action domain on both sides, not the narrower original minimum-effort domain. □

The quota restriction may change the geometry of the exercise region: Theorem 4.2 and Corollary 4.5 are not automatically transferred to the augmented state (x, m, R). The related bi-level model of van Dijk et al. [23] studies limits on *annual quota adjustments* and fleet responses; those distinct choices are intentionally absent here.

Minimal quota experiment. We reoptimize the constant trigger–target-effort rule separately in three training-sample scenarios: no quota, an annual quota at 100%, and an annual quota at 80% of the unrestricted rule's *training-sample mean annual catch*. The common benchmark is $\widehat{Q}_{\mathrm{ref}} = 0.097495$ in normalized units. The allowance is frozen before holdout evaluation. At each admissible date the candidate policy proposes qex if $x \geq s$, clips its proposed catch to the residual annual allowance, computes effort $h/(qx)$ and *does not harvest* if the resulting net gain is nonpositive. Thus the simulated rules are explicitly restricted heuristics in the action set (42), not solutions of the full Bellman recursion (45). The training set contains 320 paths and the independent economic holdout contains 1,500; all rows use paired Gaussian primitives and the same $H = 0.75$, $\sigma_H = 0.10$ reference dynamics. Relative to the unrestricted reference,

Table 1: Illustrative annual catch quotas: independently optimized restricted rules (1,500 holdout paths).

Scenario	Annual Q	Holdout NPV	Annual catch	$\mathbb{P}(\min X^+ < 0.30)$	Quota use
No quota	–	0.02369	0.0938	0.036	–
100% of ref.	0.0975	0.00917	0.0642	0.021	65.8%
80% of ref.	0.0780	0.00185	0.0355	0.015	45.5%

Quotas in normalized biomass units. $Q_{\mathrm{ref}} = 0.097495$ is determined solely from training-path average annual harvest. Quota use is average annual catch divided by the annual quota. All figures are pilot Monte Carlo estimates, not structural fisheries predictions or ecological certifications.

paired holdout NPV differences (reference minus quota policy) are 0.01452 (95% Monte Carlo interval $[0.01414, 0.01490]$) for the 100% quota and 0.02183 ($[0.02132, 0.02235]$) for the 80% quota. These are conditional comparisons of the selected rules on common validation paths, not claims about globally optimized welfare. The annual-average benchmark is restrictive for intermittent fixed-cost harvests: individual catches in active years can exceed the unconditional annual mean. Consequently the table measures a stylized trade-off, not welfare effects of any actual IPHC catch limit. Quota compliance is exact by construction, whereas the reported biomass-shortfall frequencies are descriptive holdout estimates rather than certified ecological guarantees. The selected-policy comparisons need not obey the monotonicity of *global* optimal values in Proposition 5.1, because optimization is restricted and sample based.

6 Pacific halibut anchors and identification limits

The International Pacific Halibut Commission (IPHC) publishes annual stock-assessment summaries and series of commercial landings, fishing mortality, catch-per-unit-effort, and female spawning biomass [25, 26]. Its 2025 assessment reported an estimated beginning-of-2026 female spawning biomass around 166 million pounds and a relative spawning biomass of approximately 38%; the assessment uses an ensemble of four models rather than their simple mean [25].

No identification of fractional memory from these anchors. Female spawning biomass in the assessment is an inferred, model-dependent annual series, not direct observations of X_t in equation (5). Annual persistence can reflect smoothing, structural breaks, recruitment shifts, and nonstationarity as well as long memory. Commercial mortality includes components that are not direct removals from the mature female spawning component. Accordingly, θ, H, σ_H and the separation between short-memory and fractional shocks are *scenario parameters* here. Benchmark revenue and cost coefficients are retained from the normalized experiment of Sarkar [6], not fitted to transaction-level halibut prices or fishing costs.

Table 2: Public Pacific-halibut context, not estimates of the normalized control model.

Quantity	Published value
Female spawning biomass, early 2026, ensemble median	166 million lb
Relative spawning biomass, early 2026	38%
2025 total fishing mortality (TSD–009)	28.80 million lb
2025 commercial landings (TSD–007)	16.35 million lb
Coastwide commercial WPUE (TSD–010)	110 lb/standardized skate

Source: publicly reported figures identified in the earlier working-paper data compilation and IPHC assessment records [25, 26]; no fresh reconstruction of the underlying workbooks is claimed here. The quantities have different biological and measurement definitions.

IPHC policy targets are not model thresholds. The 2025 IPHC Harvest Strategy Policy distinguishes (i) long-run time-frequency performance relative to RSB20% and (ii) a distinct threshold target RSB36% [24]. The ecological constraint in Section 7.4 is instead the *probability of at least one discrete post-harvest crossing during a finite horizon.* Moreover, the denominator of an IPHC female *relative spawning biomass* measure is not the Gilpin–Ayala carrying capacity K, nor is X_t necessarily mature female spawning biomass. Without an explicit age/sex and reference-state mapping, identifying $X_t/K = 0.20$ with RSB20% is unjustified. All numerical critical stocks are therefore internal normalized model thresholds, not IPHC compliance criteria or fishery-management prescriptions.

Requirements for structural calibration. Identifying the mixed-noise model would require repeated stock or biomass-index observations at a frequency commensurate with the decision grid, harvest and effort records, observation-error models, and data on biological recruitment and environmental drivers. Estimation would need to distinguish persistent shocks (H), instantaneous Brownian variation (σ_W), fractional scale (σ_H), nonlinearity (r, K, θ), and regime changes, preferably with a state-space or likelihood/simulation-based method and an independent out-of-sample check. Conversion between total stock, female spawning biomass, landing mass, and the IPHC mortality-accounting categories would require an additional biological observation map. The present study neither identifies these components jointly nor equates normalized catch quotas with official fishery limits.

7 Numerical implementation and statistical protocol

7.1 Gaussian simulation and corrected time line

For every H, draw independent standard-normal arrays Z^W, Z^H. Let Γ_H be the $M \times M$ Toeplitz matrix $[\gamma_H(|j-k|)]_{jk}$ from (14), and set $L_H L_H^\top = \Gamma_H$. Generate

$$(\xi_0, \dots, \xi_{M-1})^\top = \Delta^H L_H Z^H, \quad \Delta W_j = \sqrt{\Delta} Z_j^W. \tag{46}$$

This is exact finite-dimensional Gaussian covariance simulation to floating-point accuracy. We construct the conditional next-increment predictor from the same triangular Gaussian decomposition. The transition uses (23), and *only* $t_1, \dots, t_{M-1}$ are harvest dates. For the constant trigger–effort rule (s, e),

$$\widehat{J}_N(s,e) = \frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{M-1} e^{-\rho t_j} \pi_e(X_j^{-,i}) \mathbf{1}_{\{X_j^{-,i} \geq s\}}. \tag{47}$$

This differs substantively from the earlier $j = 1, \ldots, M$ objective; old numerical returns are not presented as corrected estimates.

7.2 Two meaningful normalizations of fractional noise

Since $\mathrm{Var}(\sigma_H \xi_j) = \sigma_H^2 \Delta^{2H}$, fixed σ_H across H changes both dependence and stepwise marginal variance. We report: (A) fixed annual-intensity parameter $\sigma_H = \sigma_H^{(0)}$, and (B) a fixed one-step fractional variance,

$$\sigma_H^{\mathrm{adj}}(H) = \sigma_H^{(0)} \Delta^{H_0 - H}, \qquad \left(\sigma_H^{\mathrm{adj}}(H)\right)^2 \Delta^{2H} = (\sigma_H^{(0)})^2 \Delta^{2H_0}. \tag{48}$$

The second design holds marginal one-step variance fixed; it does *not* hold cumulative volatility or the entire multi-step conditional distribution fixed. In the constructive boundary illustration in Appendix C, $\Delta = 2 > 1$; in the main simulations $\Delta = 0.125 < 1$. The sign of $\partial_H \Delta^{2H} = 2\log(\Delta)\Delta^{2H}$ consequently differs between the two settings. Their numerical effects are not interchangeable.

7.3 Policy class, optimization and independent evaluation

The policy parameters are bounded and all transitions satisfy the no-depletion restriction. For the revised reference computation, $T = 12$, $M = 96$, $r = 0.30$, $K = 1$, $\theta = 1.5$, $\sigma_W = 0.20$, $a = 0.07$, $\beta = 0.20$, $c = 0.02$, $f = 0.04$, $q = 1$, $\rho = 0.05$, $X_0 = 0.5849$, and $\eta = 0.1$; constant candidate efforts are restricted to $[0.10, 0.40]$. This effort range is intentionally narrower than the theorem's constructive example ($[0.10, 0.85]$). The optimizer uses 17×17 coarse candidates over $s \in [0.55, 1.05]$, $e \in [0.10, 0.40]$, then a 21×21 local refinement around the best coarse candidate. Duplicates are removed; hence at most 730 unique candidates per run. The additional grid is chosen adaptively, so a fixed-730-family Bonferroni bound is not, by itself, valid for this search.

Selection uses $N_{\mathrm{train}} = 320$ paths, and economic holdout evaluation uses $N_{\mathrm{test}} = 1500$ independent paths. For ecological statements, an additional $N_{\mathrm{cert}} = 6000$ paths is generated independently and the selected policy remains fixed throughout certification. The seeds and algorithms are fully specified in the included `recompute_harvest.py`. All estimates and confidence intervals are conditional on this stylized simulator; they do not measure structural calibration uncertainty.

7.4 Ecological chance constraints: screening versus certification

For a normalized critical biomass $x_{\mathrm{crit}} > 0$ define

$$R_{x_{\mathrm{crit}}}(\varphi) = \mathbb{P}_\varphi\left(\min_{0 \le j \le M} X_j^+ < x_{\mathrm{crit}}\right), \quad X_M^+ := X_M^-. \tag{49}$$

The risk measure checks monitored post-harvest stock, not the continuous-time path infimum. Training *screens* candidate policies using the one-sided 95% Wilson upper limit $U_W(y, N_{\mathrm{train}}) \le \alpha$. Such screening is not a simultaneous guarantee after adaptive search. Conditional on a selected, frozen policy and a fresh independent certificate sample, an exact one-sided upper Clopper–Pearson bound is

$$U_{\mathrm{CP}}(y, n; 1-\delta) = \begin{cases} 1, & y = n, \\ \mathrm{Beta}^{-1}(1-\delta;\, y+1, n-y), & 0 \le y < n. \end{cases} \tag{50}$$

For a *single predeclared* policy and risk limit, certification at confidence $1 - \delta$ requires $U_{\mathrm{CP}} \le \alpha$. Multiple displayed certifications are pointwise unless a familywise adjustment is separately

implemented. If the certificate is used to select a new policy, an additional untouched certificate sample is required.

7.5 Nested rule families and interpretation

We distinguish constant and information-adaptive families

$$\mathcal{P}_0 : s_j = s_0, \quad E_j = e_0, \tag{51}$$
$$\mathcal{P}_1 : s_j = s_0 + s_t(t_j/T - 1/2), \quad E_j = e_0, \tag{52}$$
$$\mathcal{P}_2 : s_j = s_0 + s_t(t_j/T - 1/2) + s_\mu \mu_j(m_j), \quad E_j = e_0. \tag{53}$$

Decision rules harvest whenever pre-decision stock exceeds s_j. The coefficient choices $s_t = 0$ and $s_\mu = 0$ are always available; hence $\mathcal{P}_0 \subseteq \mathcal{P}_1 \subseteq \mathcal{P}_2$. At the population optimum and *under the same constraint set*, $V_0^* \leq V_1^* \leq V_2^*$. Selection using finite Monte Carlo may not attain those optima, and holdout returns need not satisfy the inequality. The increment $V_2^* - V_1^*$ pertains to the restricted *one-step predictor* and not to the full environmental history. In the supplemental pilot grid used here, effort remains constant across the three classes; this deliberate restriction makes the nested comparison transparent.

8 Verified calculations and corrected simulation results

The results in this section concern the original *quota-free* model; the separate, explicitly restricted numerical quota experiment is reported in Section 5.

8.1 Analytical and constructive checks

The fundamental constants for the constructive example in Appendix C are

$$x_{\text{low}} = 0.0716788173, \quad L_\pi = 0.1116089844, \quad a\beta(q\underline{e})^\beta B^{\beta-1} = 0.006386384003. \tag{54}$$

The last expression uses *minimum* effort $\underline{e} = 0.10$, whereas L_π uses $\overline{e} = 0.85$. The reproducibility script checks these constants and every row of Table 3 at tolerance 10^{-5} for exercise roots. The sufficient stock-space margins are positive in all three constructed histories.

Table 3: Constructed conditional threshold check ($H = 0.80,\ \Delta = 2$); not a fisheries calibration.

Previous raw increment ξ	$\mu = \kappa_H \xi$	Λ	δ	x^*
−0.50	−0.257858	0.003698523	0.002687861	0.374508
0.00	0	0.003894266	0.002492118	0.374783
0.50	0.257858	0.004100369	0.002286015	0.375072

Values reproduced by the supplied independent 64-node normal quadrature script. Rounding causes small discrepancies with the earlier working-paper values 0.374513, 0.374782, 0.375069.

8.2 Last-two-opportunity diagnostic at the economic monitoring frequency

For $r = 0.30$, $\rho = 0.05$, $\Delta = 0.125$, $H = 0.80$, $\sigma_H = 0.20$, and the other Appendix C primitives, $\Lambda_{M-2,1}(0) = 0.115212534$, while the global sufficient lower slope is 0.006386384, so

$$\delta_{M-2,1}(0) = -0.108826150 < 0. \tag{55}$$

Thus the conditions of Theorem 4.2 are *not* verified for this calibration at the decision covered by Corollary 4.5. The original constructive example uses $\rho = 2$ and $\Delta = 2$ and must not be used to certify the economic scenario. Nevertheless, applying (37) to the final-two-opportunity comparison gives the following *numerical* roots and actual derivatives, conditional on one preceding raw fractional increment. The three numerically located candidate roots are 0.786059, 0.875051 and 0.961138 for previous fractional increments -0.05, 0 and $+0.05$, respectively; the computed derivatives of D at these roots are positive. The full table, sampled derivative ranges, and figure are deferred to Appendix D. These checks use an explicitly reduced *two-opportunity* problem; they neither certify the derivative sign on an entire real interval nor characterize the original 96-interval decision problem with the full observed fractional history.

8.3 Recomputed economic and environmental sensitivity

Table 4 reports results from an independent reference implementation, with decisions only at $t_1, \ldots, t_{M-1}$. The output includes two normalization regimes and a density-shape sensitivity. These are *new* Monte Carlo estimates, not restored outputs from the original scripts. The corrected reference policy ($H = 0.75$, $\sigma_H = 0.10$, $\theta = 1.5$) selects $s = 0.746875$ and $e = 0.263125$, with holdout mean NPV 0.023688 (MC standard error 0.000296). At $\sigma_H = 0.20$, holding annual intensity fixed, the holdout NPV changes from 0.026610 at $H = 0.60$ to 0.027597 at $H = 0.90$. Under fixed marginal step variance relative to $H_0 = 0.75$, the corresponding values are 0.024628 and 0.032132; ecological crossing risk below 0.30 also changes markedly. These quantitative contrasts illustrate sensitivity to normalization, not a causal or universal sign theorem in H.

Table 4: Fresh corrected-grid economic policy experiments. Holdout NPV is reported with Monte Carlo standard error; ecological risk is the empirical probability of post-harvest hitting $x_{\text{crit}} = 0.30$.

Scenario	H	σ_H	θ	s^*	e^*	Holdout NPV (SE)
Baseline corrected	0.75	0.100	1.50	0.7469	0.2631	0.02369 (0.00030)
σ_H fixed	0.60	0.200	1.50	0.7594	0.2462	0.02661 (0.00036)
σ_H fixed	0.75	0.200	1.50	0.7844	0.2331	0.02684 (0.00046)
σ_H fixed	0.90	0.200	1.50	0.8188	0.2294	0.02760 (0.00060)
Step variance fixed	0.60	0.146	1.50	0.7844	0.2500	0.02463 (0.00032)
Step variance fixed	0.75	0.200	1.50	0.7844	0.2331	0.02684 (0.00046)
Step variance fixed	0.90	0.273	1.50	0.8625	0.2144	0.03213 (0.00082)
Shape sensitivity	0.75	0.100	0.70	0.7094	0.2631	0.01533 (0.00025)
Shape sensitivity	0.75	0.100	1.00	0.7375	0.2650	0.01869 (0.00027)
Shape sensitivity	0.75	0.100	2.00	0.7500	0.2613	0.02784 (0.00031)

Independent implementation with $M = 96$, $N_{\text{train}} = 320$, $N_{\text{test}} = 1500$, no harvest at t_0 or t_M. $H_0 = 0.75$ and $\sigma_H^{(0)} = 0.20$ for constant-step-variance experiments. Policy bounds and grids are specified in the text. The values are simulator-based and do not reproduce legacy tables.

8.4 Constrained policy search and independent risk check

The screened ecological-policy experiments in Table 5 separate search-based training admissibility from the independent Clopper–Pearson bound. A value "No" means that the particular selected policy is *not certified* at the stated confidence and tolerance; it does not

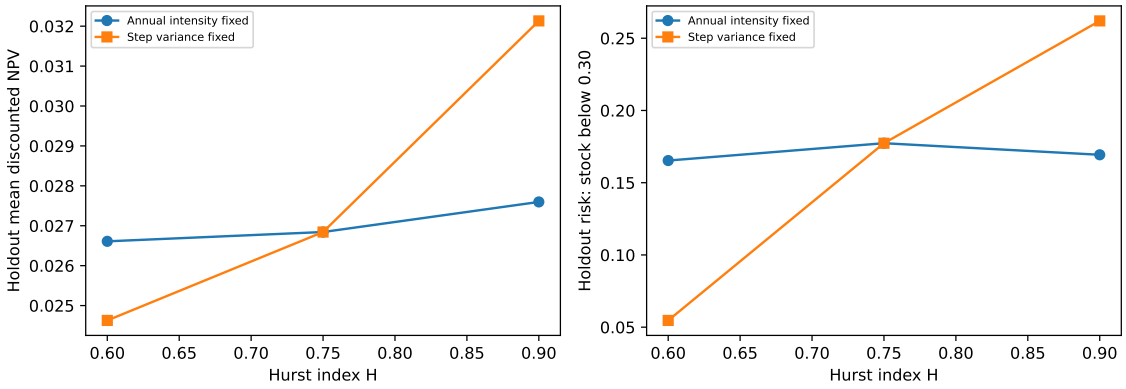

Figure 1: Corrected-grid model sensitivity under two normalizations. Each point is an independently selected constant rule evaluated on the same primitive validation shocks; the connecting segments are visual guides, not fitted response curves.

establish population-risk violation. For $x_{\text{crit}} = 0.30$ and $\alpha = 0.025$, training screening picks $(s, e) = (0.80, 0.25)$, but the independent certificate records 158 events in 6000 paths, with one-sided exact upper limit $0.029993 > 0.025$: certification fails. For the more permissive $\alpha = 0.05$, the selected $(0.746875, 0.263125)$ records 206 certificate events and upper limit $0.038455 < 0.05$, so that *single-policy pointwise* criterion is met. This contrast does not identify the true model-risk value without Monte Carlo uncertainty, and is not a statement of familywise certification across all rows.

Table 5: Ecological policy screening and independent pointwise certification; normalized biomass.

x_{crit}	Risk limit	s^*	e^*	Train events/320	Holdout risk	Cert events/6000	CP upper 95%	Certified?
0.25	2.50%	0.7469	0.2631	3	0.67%	54	1.13%	Yes
0.25	5.00%	0.7469	0.2631	3	0.67%	54	1.13%	Yes
0.25	10.00%	0.7469	0.2631	3	0.67%	54	1.13%	Yes
0.30	2.50%	0.8000	0.2500	3	2.87%	158	3.00%	No
0.30	5.00%	0.7469	0.2631	7	3.60%	206	3.85%	Yes
0.30	10.00%	0.7469	0.2631	7	3.60%	206	3.85%	Yes

Policy selection uses a one-sided Wilson screen on training paths; the independent 6000-path certificate uses the exact one-sided Clopper–Pearson upper bound. "Certified?" refers to one policy considered at a time; these six rows do not provide a simultaneous 95% familywise assertion. Non-certification does not establish a violation.

8.5 Nested information-use comparison

Table 6 compares nested pilot policy families on the same test paths. Since selection is finite and stochastic, an enlarged training-policy class need not yield a higher test-sample payoff. The table is evidence about this *implementable parametric comparison*, not an exact solution of the full history-dependent Bellman recursion. The holdout mean NPV rises from 0.027077 (constant rule) to 0.027447 (time-varying rule) and 0.028434 (time plus one-step fractional predictor). Paired differences relative to the immediately smaller class are 0.000370 with 95% MC interval $[0.000273, 0.000467]$ and 0.000987 with interval $[0.000865, 0.001109]$. However, the corresponding observed shortfall frequencies below 0.30 are 15.5%, 15.6%, and 19.0%. These are *unconstrained* information-flexibility comparisons and do not establish a positive value of predictor information when ecological risk is held fixed.

Table 6: Pilot nested policy-class comparison, with paired test-sample differences.

Rule class	Training NPV	Holdout NPV	s_t	s_μ	Paired holdout increment [95% MC CI]
P0-constant	0.02704	0.02708	0.00	0.00	–
P1-time	0.02736	0.02745	-0.12	0.00	0.000370 [0.000273, 0.000467]
P2-time-memory	0.02815	0.02843	-0.12	0.80	0.000987 [0.000865, 0.001109]

$H = 0.75$, $\sigma_H = 0.20$, $\theta = 1.5$. The grid contains 9, 45, and 225 nested candidates respectively; the displayed increment is P1−P0 or P2−P1 computed on the same 1500 shocks. These *unconstrained* pilot gains do not certify a common ecological risk limit: observed risks at $x_{\text{crit}} = 0.30$ are approximately 15.5%, 15.6%, and 19.0%, respectively. The last gain is attributable to inclusion of the one-step predictor in the policy class, not to full fractional-history optimality.

8.6 Interpretation and uncertainty limits

The sufficient global-slope condition is violated at the penultimate decision in the $r = 0.30$, $\rho = 0.05$ calibration: the measured negative margin is not evidence that no boundary exists, and the numerical two-opportunity roots of Table 7 do not certify full-horizon uniqueness. The local sufficient criterion in Proposition 4.6 would require verified derivative control on an interval, not just finite quadrature nodes.

First, holding one-step fractional variance constant changes the quantitative relationship between H and selected economic rules, so no universal comparative-static sign can be attributed to memory from fixed-σ_H experiments alone. Second, the numerical ecological limits concern hitting probabilities in an assumed biological simulator, not IPHC performance targets. Third, the pilot policy-family differences should be interpreted through paired test-sample uncertainty and possible optimization error rather than through in-sample returns alone. Finally, while covariance-factor Gaussian simulation is exact on its grid, the split-step biomass approximation has discretization error. A definitive empirical fisheries application would require convergence across monitoring grids, more independent trials, and identification of the biological and economic parameters. No management recommendation follows from the normalized example.

9 Conclusion

The combined Brownian–fractional Gilpin–Ayala model admits a positive closed-form representation and an exact finite-grid conditional control problem. Positive memory creates a genuinely history-dependent predictive state, yet a unique harvest threshold conditional on sampled history exists on a finite stock interval under checkable sufficient conditions. For the final two admissible decision dates, a favorable predictive fractional innovation raises the threshold when the stated sufficient single-crossing conditions hold. At the economic calibration these sufficient conditions fail, and the observed upward shift is a numerical illustration rather than an established uniqueness result.

The numerical analysis now uses the same admissible time line as the theory, evaluates both annual-intensity and matched one-step-variance normalizations, and reports independent ecological certification rather than treating screening as a regulatory guarantee. The controlled comparison of nested policies gives an operational, but necessarily partial, measure of the economic information value of one-step environmental prediction. The annual-quota constraint gives a parsimonious regulatory state and a monotone quota-value comparison; an independent, reoptimized pilot quantifies its restricted-policy costs, without identifying actual regulatory quota levels. Biological-model uncertainty, unobserved environmental histories, continuous-path ecological risk, and full structural fishery calibration remain outside the validated scope.

Declarations

Funding: This research received no external funding.
Competing interests: The author declares no competing interests.
Data and code availability: Public IPHC summaries are available on the IPHC website. The source archive provides Python scripts, numerical CSV outputs, tests and seeds. Original IPHC workbooks are not redistributed. No independently replicated structural estimation is claimed.
Author contributions: Sole-author conceptualization, analysis, interpretation, and manuscript preparation.

A Mixed Itô–Young derivation

For $Y = X^{-\theta}$, Itô–Young calculus yields $dY = (\alpha Y + \theta r/K^{\theta})dt - \theta\sigma_W Y\, dW - \theta\sigma_H Y\, dB^H$, with $\alpha = -\theta r + \frac{1}{2}\theta(\theta+1)\sigma_W^2$. The homogeneous factor F has logarithmic drift $\alpha - \frac{1}{2}\theta^2\sigma_W^2 = -\theta r + \frac{1}{2}\theta\sigma_W^2$. Multiplying by F^{-1} and integrating gives the representation in Proposition 3.1. For an interval starting at t_j, the same transformation gives (40). Neither derivation applies an Itô correction to the Young component.

B Gaussian conditioning and predictor implementation

Write $\widetilde{\Gamma}_{H,j}$ for the $j \times j$ leading block of the Toeplitz covariance, with L_H its full lower Cholesky factor. Given standard normals $Z_0^H, \dots, Z_j^H$ and raw fractional increments $\xi = \Delta^H L_H Z^H$, the observed increments $\xi_0, \dots, \xi_{j-1}$ uniquely determine $Z_0^H, \dots, Z_{j-1}^H$ because L_H is lower triangular with strictly positive diagonal. Hence

$$\mu_j(m_j) = \Delta^H \sum_{k<j} (L_H)_{jk} Z_k^H = \xi_j - \Delta^H (L_H)_{jj} Z_j^H. \tag{56}$$

The second equality is an algebraic device for simulator construction, not a look-ahead observation rule: when the policy acts at t_j, its predictor uses only terms $k < j$. In this way the predictor is exactly adapted to the sampled increment filtration. For a single preceding increment, $\mu = \kappa_H \xi$ and $v = \Delta^{2H}(1 - \kappa_H^2)$.

C Constructive non-vacuity and non-regression test

Take

$$(r, K, \theta) = (0.3, 1, 1.5), \quad (a, \beta, c, f) = (0.07, 0.2, 0.02, 0.04), \quad (\underline{e}, \overline{e}) = (0.10, 0.85), \tag{57}$$

$\sigma_W = \sigma_H = 0.2$, $H = 0.8$, $\Delta = 2$, $\rho = 2$, $\ell = 0.05$, $B = 1.5$, and $S \equiv 0$. Define $\kappa_H = 2^{2H-1} - 1$, $\mu = \kappa_H \xi$ and $v = \Delta^{2H}(1 - \kappa_H^2)$. Since the last admissible date has $V_{M-1,1}(x) = g(x)$, the preceding continuation slope bound becomes

$$\Lambda_{M-2,1}(\xi) = L_\pi \exp\big[(r - \rho)\Delta + \sigma_H \mu + \tfrac{1}{2}\sigma_H^2 v\big]. \tag{58}$$

The certified crossing margin is

$$\delta(\xi) = a\beta(q\underline{e})^\beta B^{\beta-1} - \Lambda_{M-2,1}(\xi). \tag{59}$$

Notice the *lower* effort $\underline{e}$ in the first term, in contrast to the *upper* effort $\overline{e}$ in L_π. The root is independently computed from

$$g(x) = e^{-\rho\Delta}\mathbb{E}\left[g\left(\phi_{\Delta/2}(e^Z\phi_{\Delta/2}(x))\right)\right], \quad Z \sim N\left(-\tfrac{1}{2}\sigma_W^2\Delta + \sigma_H\mu,\, \sigma_W^2\Delta + \sigma_H^2 v\right). \tag{60}$$

The companion `test_regression_annexe_C.py` recomputes the constants and roots using 64-point Gaussian quadrature, enforces the low-stock condition, and fails if reported roots differ by 10^{-5} or more. The deliberately high discount rate is a constructive mathematical example, not an empirical discount-rate estimate.

D Local-crossing numerical diagnostics and proof obligations

The following computations concern only the final two admissible intervention opportunities, a single remaining intervention, and conditioning on one preceding fractional innovation. In particular, the past-increment summary used here is not a sufficient predictive state for the full multi-period control problem.

Table 7: Two-opportunity diagnostic under $\rho = 0.05$, $\Delta = 0.125$; not a full-horizon uniqueness certificate.

Prior ξ	Root x^*	$g'(x^*)$	$C'(x^*)$	$D'(x^*)$
-0.05	0.786059	0.015340	0.014670	0.000670
0.00	0.875051	0.014154	0.013444	0.000710
0.05	0.961138	0.013192	0.012444	0.000748

The positive-profit onset is $b \simeq 0.371869$. A finite numerical scan over $(b, 1.5]$ gives positive sampled minimum D' of 0.000457, 0.000428, and 0.000398, respectively; the grid detects one economically relevant sign change in each displayed history. The global sufficient margin is deliberately not treated as an estimate of the actual derivative. Near the central numerical root, $g' = 0.014154$, $C' = 0.013444$, while $\Lambda = 0.115213$. Positive slopes *at* the three roots establish local transversality; the grid scan is supporting numerical evidence, not a proof that all economically relevant histories have unique boundaries.

Validated interval algorithm (not performed here). For a fixed history and $[b, B]$, construct a finite rational interval cover I_k and enclose $g'(I_k)$ using outward-rounded interval arithmetic, splitting wherever the optimal effort switches. For every interval I_k, enclose $\partial_x\Psi_\Delta(I_k, Z)$ and $g'(\Psi_\Delta(I_k, Z))$; integrate the Gaussian-weighted enclosure using a quadrature with a proven tail truncation and discretization remainder. Form the rigorous interval difference $\mathcal{D}'_k = g'(I_k) - \mathcal{C}'_k$. If $\inf \mathcal{D}'_k > 0$ for all k, and certified endpoint bounds have $D(b) < 0 < D(B)$, Proposition 4.6 applies to the entire continuum. Any interval with inconclusive sign must be bisected or reported as unresolved. A finite scan and ordinary floating-point quadrature, as used in this paper, do not execute this certificate.

What would constitute a validated computer-assisted proof. A rigorous verification of Proposition 4.6 for a specified history and bounded biomass domain $[b, B]$ would require: (i) an explicitly bounded error for each Gaussian expectation defining C and $\partial_x C$; (ii) interval enclosures for g' and for the split-step map and its derivative on a finite cover of (b, B); (iii) a strictly positive lower interval bound for $g' - \partial_x C$ on every subinterval, including neighborhoods of the effort-switching points and the positive-profit onset; and (iv) verified endpoint signs. The present Gaussian quadrature and finite stock scan do not provide items (i)–(iv). The

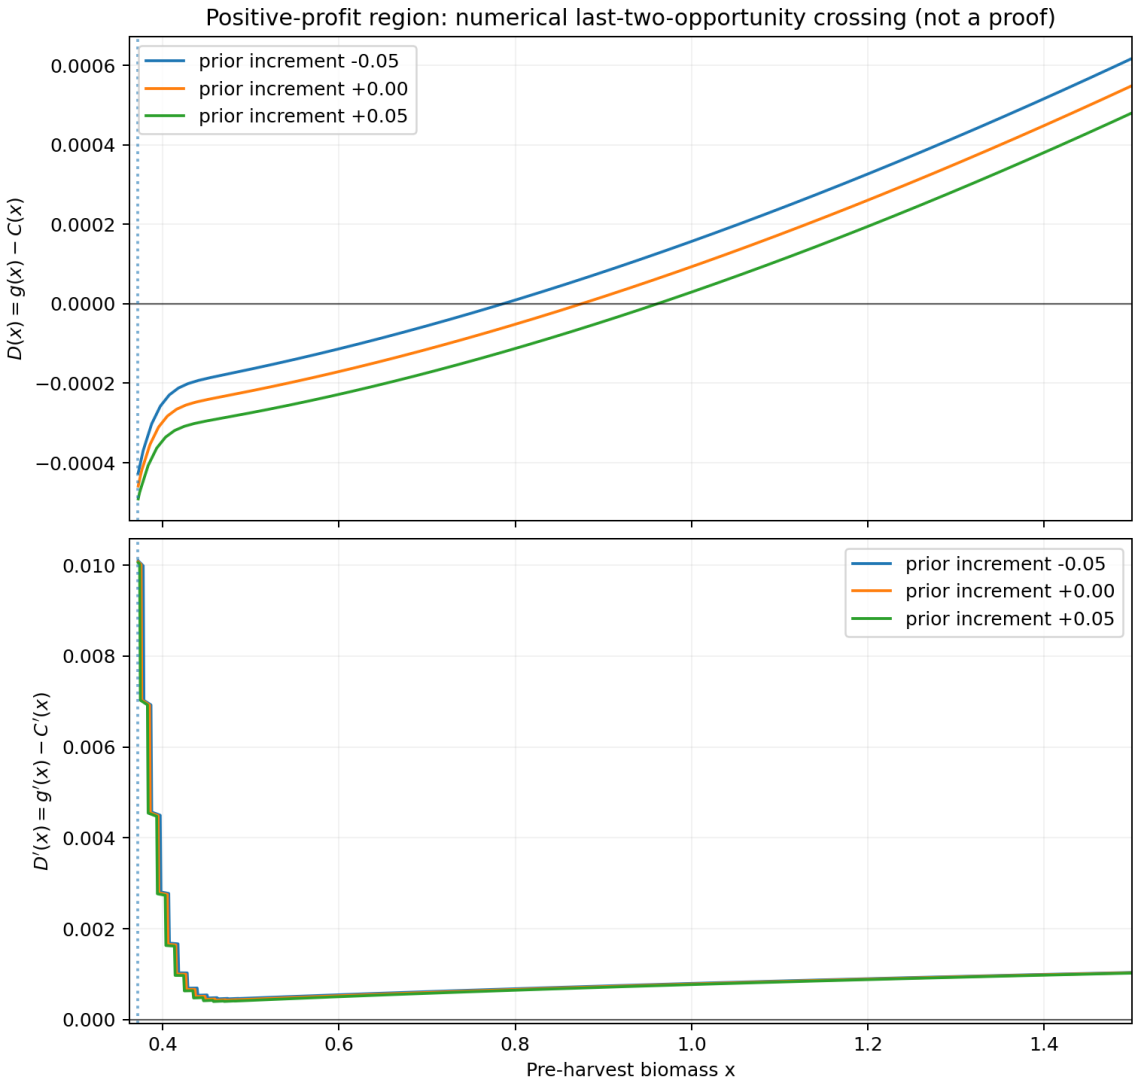

Figure 2: Numerical last-two-opportunity advantage and spatial slope on the positive-profit region. The positive sampled slopes do not certify (38) for every biomass point, nor do they characterize the full history-dependent 96-step control problem.

results in the table and figure must therefore be read as numerical diagnostics, not certified signs on the continuum.

Full-history limitation. At date t_j, a Gaussian fractional-increment predictor uses the entire observed innovation history through the relevant Toeplitz covariance blocks. The last-two-opportunity calculations do not solve the history-dependent Bellman recursion for all $j = 1, \ldots, 95$ and intervention balances n. There is no claim of a global unique boundary for the complete 96-interval problem, nor of a uniform monotonicity theorem in H.

E Simulation controls, certification and reproducibility

All reference calculations use NumPy's PCG64 generator through `default_rng`, SciPy Cholesky covariance factorization and exact beta quantiles for one-sided certification. Independent training, economic holdout and ecological certificate arrays use distinct seeds (20261002, 314159, 20261003). Gaussian primitives are shared across candidate policies and sensitivity regimes; different H values transform those primitives through their respective Cholesky factors. The non-regression test is independent of the long-horizon simulation code. All rates and losses are measured on normalized biomass, discounted NPV and model probabilities, not money-denominated fishery welfare.

Why a corrected model needs new tables. The previous objective included an intervention at terminal time. The present one excludes it and does not permit an intervention before the first propagation step. Consequently the old selected policies and their reported risks cannot be inserted as if they were recomputed under this chronology. The tables in Section 8 are produced by the companion corrected simulator. The supplied scripts remain the computational definition of the pilot experiments, and users should re-run them in the target computing environment before submission to verify platform-dependent rounding.

References

[1] Schaefer, M. B. (1954). Some aspects of the dynamics of populations important to management of commercial marine fisheries. *Bulletin of the Inter-American Tropical Tuna Commission*, 1, 27–56.

[2] Clark, C. W. (1973). Profit maximization and the extinction of animal species. *Journal of Political Economy*, 81, 950–961.

[3] Reed, W. J. (1979). Optimal escapement levels in stochastic and deterministic harvesting models. *Journal of Environmental Economics and Management*, 6, 350–363.

[4] Saphores, J.-D. (2003). Harvesting a renewable resource under uncertainty. *Journal of Economic Dynamics and Control*, 28, 509–529.

[5] Gilpin, M. E. and Ayala, F. J. (1973). Global models of growth and competition. *PNAS*, 70, 3590–3593.

[6] Sarkar, S. (2009). Optimal fishery harvesting rules under uncertainty. *Resource and Energy Economics*, 31, 272–286.

[7] Nyassoke Titi, G. C., Sadefo Kamdem, J., and Fono, L. A. (2020). Fishery management in a regime switching environment. *Results in Applied Mathematics*, 7, 100125.

[8] Nyassoke Titi, G. C., Sadefo Kamdem, J., and Fono, L. A. (2022). Optimal renewable resource harvesting model using price and biomass stochastic variations. *Mathematical Methods of Operations Research*, 95, 297–326.

[9] Nyassoke Titi, G. C., Sadefo Kamdem, J., and Fono, L. A. (2025). Optimal harvest under a Gilpin–Ayala model driven by the Hawkes process. *Results in Applied Mathematics*, 26, 100564.

[10] Sethi, G., Costello, C., Fisher, A., Hanemann, M., and Karp, L. (2005). Fishery management under multiple uncertainty. *JEEM*, 50, 300–318.

[11] Nøstbakken, L. (2006). Regime switching in a fishery with stochastic stock and price. *JEEM*, 51, 231–241.

[12] Singh, R., Weninger, Q., and Doyle, M. (2006). Fisheries management with stock growth uncertainty and costly capital adjustment. *JEEM*, 52, 582–599.

[13] MacCall, A. D. (2002). Fishery-management and stock-rebuilding prospects under low-frequency environmental variability. *Bulletin of Marine Science*, 70, 613–628.

[14] Vasseur, D. A. and Yodzis, P. (2004). The color of environmental noise. *Ecology*, 85, 1146–1152.

[15] Diebold, F. X. and Inoue, A. (2001). Long memory and regime switching. *Journal of Econometrics*, 105, 131–159.

[16] Mandelbrot, B. B. and van Ness, J. W. (1968). Fractional Brownian motions, fractional noises and applications. *SIAM Review*, 10, 422–437.

[17] Cheridito, P. (2001). Mixed fractional Brownian motion. *Bernoulli*, 7, 913–934.

[18] Djehiche, B., Hamadène, S., and Hdhiri, I. (2010). Stochastic impulse control of non-Markovian processes. *Applied Mathematics & Optimization*, 61, 1–26.

[19] Perninge, M. (2020). A finite horizon optimal switching problem with memory and application to controlled SDDEs. *MMOR*, 91, 465–500.

[20] Jönsson, J. and Perninge, M. (2023). Finite horizon impulse control of stochastic functional differential equations. *SIAM J. Control Optim.*, 61, 924–948.

[21] Perninge, M. (2023). Non-Markovian impulse control under nonlinear expectation. *Applied Mathematics & Optimization*, 88, 72.

[22] Liu, Z. and Zervos, M. (2025). The solution to an impulse control problem motivated by optimal harvesting. *Journal of Mathematical Analysis and Applications*, 542, 128809.

[23] van Dijk, D., Hendrix, E. M. T., Haijema, R., Groeneveld, R. A., and van Ierland, E. C. (2017). An adjustment restriction on fish quota: Resource rents, overcapacity and recovery of fish stock. *Environmental and Resource Economics*, 67, 203–230. `https://doi.org/10.1007/s10640-015-9983-0`.

[24] International Pacific Halibut Commission (2025a). *Harvest Strategy Policy*. IPHC-2025-HSP. `https://www.iphc.int/research-monitoring/harvest-strategy-policy/`.

[25] International Pacific Halibut Commission (2025b). *Assessment of the Pacific halibut stock at the end of 2025*. IPHC-2026-SA-01; `https://www.iphc.int/research/stock-assessment/`.

[26] International Pacific Halibut Commission (2026). *Time Series Datasets (TSD–007, TSD–009, TSD–010, TSD–029)*. `https://www.iphc.int/data/time-series-datasets/`.