Global and Stochastic Analysis
Vol. 13 No. 4 (2026)

DOI: 10.64837/GSA.13.4.2

OPTIMAL SWITCHING UNDER COMPETITION: A DISCRETE-TIME OPTIMAL STOPPING APPROACH

MANUEL L. ESQUÍVEL AND NADEZHDA P. KRASII

ABSTRACT. Motivated by markets in which new entrants challenge established monopolies, we study the problem of determining an optimal switching time for consumers facing competing stochastic price dynamics. In the particular case of the sequences of prices—or premiums—being sub-martingales, an optimal stopping time is described by a simple relation between the increasing processes of the Doob decompositions of these sub-martingales. We present a statistical methodology allowing the application of the results including a Bayesian approach for the case of scarce data. We present a simulation study in the case of average linear growth of the prices increments and we develop two efficient estimation procedures of the parameters of these models allowing accurate estimation of the optimal stopping time. In the particular cases presented, the simulated smallest optimal stopping time coincides with the optimal stopping time studied.

Introduction

While optimal stopping theory has been extensively applied to financial markets—e.g. American options—and corporate investment—Real Options—its application to consumer behaviour in competitive insurance markets is less explored. Standard economic models of switching focus on equilibrium costs; in contrast, we propose a stochastic approach that treats the timing of the switch as an optimal stopping problem driven by the predictable trends—in Doob decomposition—of the competing price processes.

Optimal stopping in a finite time horizon builds on the Snell envelope. We follow [LL08, p. 38–43] on the presentation of the main results we will use next. Let $(\Omega, \mathscr{F}, \mathbb{P})$ be a complete probability space and $\mathbb{F} = (\mathscr{F}_n)_{n\geq 0}$ a filtration on this space; we will consider that $\mathscr{F}_0 = \{\emptyset, \Omega\}$. We consider a finite time horizon N and a $\mathbb{F}$ adapted sequence $(Z_n)_{0\leq n\leq N}$. The **Snell envelope** of this sequence is the smallest super-martingale that dominates the sequence $(Z_n)_{0\leq n\leq N}$, that is, the sequence $(U_n)_{0\leq n\leq N}$ defined by backward induction by:

$$\begin{cases} U_N := Z_N \\ U_n := \max\left(Z_n, \mathbb{E}\left[U_{n+1} \mid \mathscr{F}_n\right]\right) , \quad n = 0, \ldots, N-1 \ . \end{cases} \tag{0.1}$$

Date: Date of Submission May 08, 2026 ; Date of Acceptance June 30, 2026 ; Communicated by Yuri E. Gliklikh.

2020 *Mathematics Subject Classification.* 60G40, 60G42, 62F15, 91B42,91B24, 62M10.

Key words and phrases. optimal stopping, snell envelope, doob decomposition, submartingales, consumer choice, monopoly breaking, 1-SLA rule, bayesian estimation, real options.

This work is funded by national funds through the FCT – Fundação para a Ciência e a Tecnologia, I.P., under the scope of the projects UID/00297/2025 (`https://doi.org/10.54499/UID/00297/2025`) and UID/PRR/00297/2025 (`https://doi.org/10.54499/UID/PRR/00297/2025`) (Center for Mathematics and Applications – NOVA Math).

We denote by τ or ν a generic stopping time taking values in $\{0,1,\ldots,N\}$. Given the sequence $(Z_n)_{0\leq n\leq N}$ and a stopping time τ, we define Z_τ by:

$$Z_\tau := \sum_{n=0}^{N} Z_n \mathbb{I}_{\{\tau=n\}} .$$

A stopping time ν is an **optimal stopping time** for the sequence $(Z_n)_{0\leq n\leq N}$ if and only if,

$$\mathbb{E}\left[Z_\nu\right] = \sup_\tau \mathbb{E}\left[Z_\tau\right] .$$

The fundamental characterization of optimal stopping times of an adapted sequence $(Z_n)_{0\leq n\leq N}$ in a finite time horizon N uses the Snell envelope $(U_n)_{0\leq n\leq N}$.

Theorem 0.1. *The stopping time ν is optimal for the sequence $(Z_n)_{0\leq n\leq N}$ iff the two following conditions on the Snell envelope $(U_n)_{0\leq n\leq N}$ are satisfied:*

$$\begin{cases} Z_\nu = U_\nu \\ (U_{\nu\wedge n})_{0\leq n\leq N} \quad \textit{is a martingale.} \end{cases} \tag{0.2}$$

Finally, we have the following characterization of the smallest optimal stopping time.

Theorem 0.2. *The stopping time τ° defined by:*

$$\tau^\circ = \inf\{k\geq 0 : U_k = Z_k\} = \inf\{k\geq 0 : \max\left(Z_k, \mathbb{E}\left[U_{k+1} \mid \mathscr{F}_k\right]\right) = Z_k\} \tag{0.3}$$

is the smallest optimal stopping time.

According to [CRS91, p. 51], conditions (0.3) for the determination of the smallest optimal stopping time *can rarely be computed analytically by the backward induction which defines them*. So, instead of τ°, let us consider the following stopping time:

$$\tau_{1-SLA} := \tau^\star = \inf\{0\leq k< N : Z_k \geq \mathbb{E}[Z_{k+1}|\mathscr{F}_k]\} \tag{0.4}$$

As condition $U_k = Z_k$ is sufficient to ensure that $(U_{\tau^\star\wedge n})_{0\leq n\leq N}$ is a martingale (see Proposition 2.2.1 in [LL08, p. 39]) and obviously, $Z_{\tau^\star} = U_{\tau^\star}$ (as well as $Z_{\tau^\star+1} = U_{\tau^\star+1}$ if $\tau^\star < N$), we have that conditions (0.2) are verified and so $\tau^\star$ is an optimal stopping time for $(Z_n)_{0\leq n\leq N}$. Clearly, as

$$\{0\leq k<N : U_k = Z_k \wedge U_{k+1} = Z_{k+1}\} \subseteq \{k\geq 0 : U_k = Z_k\} ,$$

we have that $\tau^\circ \leq \tau^\star$ in accordance with theorem 0.2.

Other major references for optimal stopping—in particular, in discrete time—are the following [Shi08], [Nev75] and [PS06]. Optimal stopping in discrete time received some attention recently, the main motivation being problems in Mathematical Finance; see [DLU09] and references therein; in the context of the so called disorder problems an important reference is [Shi19]. A special algorithm for solving optimal stopping problems of Markov chains is presented in [Pre11]. An interesting application of optimal stopping in continuous time to risk theory is given in [Sch98]. The study of changes in the Snell envelope adequate for computability or other purposes is carried in [DHOR11].

We must refer that there is a connection of this work with articles dealing with consumer switching, competition, real options, dynamic pricing, and optimal stopping, for instance, [Kle87, FK07].

Let us also report that in corporate finance, firms use optimal stopping to decide when to enter a market in the context of real options theory (see [Sch05, LM22, ET22]). In this work we have, in a way, addressed a real option optimal stopping for the consumer model.

The main contribution of the present work is not the development of new optimal stopping theory per se, but rather the derivation of an operational stopping framework adapted to competitive pricing dynamics. In particular, we show that under a monotonicity condition on the predictable components of the competing price processes, the one-step look-ahead rule coincides with the smallest optimal stopping time. This characterisation allows a direct integration between stochastic optimal stopping and statistical estimation procedures suitable for practical implementation.

1. The optimal stopping model when monopoly breaks

Let us suppose that company A detains a monopoly in some market and that at time $t=0$ a company B enters the market practicing a price p_0^B such that $p_0^B < p_0^A$, p_0^A representing the price practiced by the previously monopolistic company A. Let us consider the price sequences—reported daily, weekly, monthly, etc—of the two companies $(p_n^B)_{n\geq 0}$ and $(p_n^A)_{n\geq 0}$. It is natural to assume at first that, in some sense to be clarified in the following,

- p_n^A decreases towards p_n^B on account of the competition between A and B;
- p_n^B increases towards p_n^A as B increases its market share.

For the *new* consumer, as long as $p_n^A - p_n^B > 0$ there is a favorable choice. We will consider the optimal stopping problem - in a finite time horizon N - that consists in determining an optimal stopping time $\tau^\star$ such that:

$$\mathbb{E}\left[p_{\tau^\star}^A - p_{\tau^\star}^B\right] = \sup_{0\leq\tau\leq N} \mathbb{E}\left[p_\tau^A - p_\tau^B\right] , \tag{1.1}$$

where the supremum is taken among all non negative stopping times bounded by N.

1.1. Optimal Stopping Criterion. Let $(p_n^A)_{n\geq 0}$ and $(p_n^B)_{n\geq 0}$ be two sub-martingales adapted to the filtration $\mathbb{F} = (\mathscr{F}_n)_{n\geq 0}$, representing the price sequences of Company A and Company B, respectively. By the Doob decomposition, we write:

$$p_n^A = M_n^A + I_n^A \quad \text{and} \quad p_n^B = M_n^B + I_n^B, \tag{1.2}$$

where M denotes a martingale and I denotes a predictable, non-decreasing process with $I_0 = 0$.

Remark 1.1. Although both price processes are modeled as submartingales, the competitive effect is expressed relatively rather than absolutely. In particular, the incumbent's prices may continue increasing on average, but at a slower predictable rate than those of the entrant. Thus, the gap process

$$Y_n = p_n^A - p_n^B ,$$

may decrease even when both individual price processes are submartingales.

We seek the optimal stopping time τ^* to minimize the expected future cost, or equivalently, to maximize the relative benefit of switching strategies. We define the local comparison set S as the set of times where the expected incremental growth of Company B exceeds that of Company A:

$$S = \{k \in \{0,\dots,N-1\} : \Delta I^B_{k+1} \geq \Delta I^A_{k+1}\}, \tag{1.3}$$

where $\Delta I_{k+1} = I_{k+1} - I_k$.

Lemma 1.2. *Let $(\Omega, \mathscr{F}, \mathbb{P})$ be a probability space and $\mathbb{F} = (\mathscr{F}_n)_{n=0}^N$ a filtration. Let $(p^A_n)_{n=0}^N$ and $(p^B_n)_{n=0}^N$ be integrable, $\mathbb{F}$-adapted submartingales with Doob decompositions*

$$p^A_n = M^A_n + I^A_n, \qquad p^B_n = M^B_n + I^B_n,$$

where M^A, M^B are $\mathbb{F}$-martingales and I^A, I^B are predictable, nondecreasing processes (with $I^A_0 = I^B_0 = 0$). Define the gap process $Y_n := p^A_n - p^B_n$ and the one-step look-ahead conditional expectation $\mathbb{E}[Y_{k+1} \mid \mathscr{F}_k]$ for $0 \leq k \leq N-1$. Then for every $k \in \{0,\dots,N-1\}$ we have the almost sure identity

$$\mathbb{E}[Y_{k+1} \mid \mathscr{F}_k] = Y_k + \left(\Delta I^A_{k+1} - \Delta I^B_{k+1}\right), \tag{1.4}$$

where $\Delta I^A_{k+1} := I^A_{k+1} - I^A_k$ and $\Delta I^B_{k+1} := I^B_{k+1} - I^B_k$. In particular,

$$Y_k \geq \mathbb{E}[Y_{k+1} \mid \mathscr{F}_k] \iff \Delta I^B_{k+1} \geq \Delta I^A_{k+1}, \tag{1.5}$$

a.s.

Proof. For fixed $k \in \{0,\dots,N-1\}$ write the Doob decompositions at times k and $k+1$:

$$Y_{k+1} = (M^A_{k+1} - M^B_{k+1}) + (I^A_{k+1} - I^B_{k+1}), \qquad Y_k = (M^A_k - M^B_k) + (I^A_k - I^B_k).$$

Take conditional expectation with respect to $\mathscr{F}_k$. The martingale increments satisfy

$$\mathbb{E}\left[M^A_{k+1} - M^A_k \mid \mathscr{F}_k\right] = 0, \qquad \mathbb{E}\left[M^B_{k+1} - M^B_k \mid \mathscr{F}_k\right] = 0,$$

so $\mathbb{E}[M^A_{k+1} - M^B_{k+1} \mid \mathscr{F}_k] = M^A_k - M^B_k$. Because I^A and I^B are predictable (hence I^A_{k+1}, I^B_{k+1} are $\mathscr{F}_k$-measurable), we have

$$\mathbb{E}[I^A_{k+1} - I^B_{k+1} \mid \mathscr{F}_k] = I^A_{k+1} - I^B_{k+1}.$$

Combining these equalities yields

$$\mathbb{E}[Y_{k+1} \mid \mathscr{F}_k] = (M^A_k - M^B_k) + (I^A_{k+1} - I^B_{k+1}) = Y_k + \left(\Delta I^A_{k+1} - \Delta I^B_{k+1}\right),$$

which proves (1.4). The equivalence (1.5) follows immediately by rearranging the terms. □

Theorem 1.3 (Monotone case; first-entry 1-SLA optimal)**.** *Under the same hypotheses as Lemma 1.2, define the gap process $Y_n := p^A_n - p^B_n$ and the one-step look-ahead set*

$$S := \{\, k \in \{0,\dots,N-1\} : \; Y_k \geq \mathbb{E}[Y_{k+1} \mid \mathscr{F}_k] \,\}.$$

Assume the following monotonicity *property:*

$$k \in S \implies j \in S \text{ for all } j = k, k+1, \dots, N-1.$$

Let $U = (U_n)_{n=0}^N$ be the Snell envelope of $(Y_n)_{n=0}^N$, defined by backward induction

$$U_N := Y_N, \qquad U_n := \max\left\{Y_n, \, \mathbb{E}[U_{n+1} \mid \mathscr{F}_n]\right\}, \quad n = N-1, \dots, 0.$$

Define stopping times

$$\tau^\circ := \inf\{\, n \geq 0 : \; U_n = Y_n \,\} \quad \text{and} \quad \tau^{1\mathrm{SLA}} := \inf\{\, 0 \leq k < N : \; k \in S \,\},$$

with the convention that the infimum of the empty set equals N. Then $\tau^\circ = \tau^{1\mathrm{SLA}}$ a.s.; in particular $\tau^{1\mathrm{SLA}}$ is optimal for maximizing $\mathbb{E}[Y_\tau]$ over all stopping times $\tau \in \{0,\dots,N\}$.

Proof. We prove by backward induction on $k = N, N-1, \dots, 0$ that for every k the following two statements hold almost surely:

(i) If $k \in S$ then $U_k = Y_k$.
(ii) If $k \notin S$ then $U_k = \mathbb{E}[U_{k+1} \mid \mathscr{F}_k]$.

The result $\tau^\circ = \tau^{1\mathrm{SLA}}$ will follow immediately, since $U_n = \max\{Y_n, \mathbb{E}[U_{n+1} \mid \mathscr{F}_n]\}$ and the first n with $U_n = Y_n$ is therefore the first n for which the stopping inequality holds; by the monotonicity assumption this coincides with first entry into S.

Base case. At $k = N$ we have $U_N = Y_N$ by definition; and $N \notin S$ by construction of $S \subset \{0, \dots, N-1\}$. Thus the claims hold at $k = N$.

Induction step. Fix $k \in \{0, \dots, N-1\}$ and assume (i)–(ii) hold for all times j with $k < j \le N$. We show they hold at time k.

(A) Case $k \in S$. By the monotonicity assumption, $j \in S$ for every $j \ge k$ with $j \le N-1$. By the induction hypothesis (i) we therefore have $U_j = Y_j$ for all $j \ge k$. In particular $U_{k+1} = Y_{k+1}$, and hence

$$\mathbb{E}[U_{k+1} \mid \mathscr{F}_k] = \mathbb{E}[Y_{k+1} \mid \mathscr{F}_k].$$

Since $k \in S$ we have $Y_k \ge \mathbb{E}[Y_{k+1} \mid \mathscr{F}_k]$. Thus

$$U_k = \max\{Y_k,\ \mathbb{E}[U_{k+1} \mid \mathscr{F}_k]\} = \max\{Y_k,\ \mathbb{E}[Y_{k+1} \mid \mathscr{F}_k]\} = Y_k,$$

which proves (i) at time k.

(B) Case $k \notin S$. By definition $Y_k < \mathbb{E}[Y_{k+1} \mid \mathscr{F}_k]$. Observe that for any $m \ge k+1$ we have $U_m \ge Y_m$, hence by monotonicity and the tower property

$$\mathbb{E}[U_{k+1} \mid \mathscr{F}_k] \ge \mathbb{E}[Y_{k+1} \mid \mathscr{F}_k] > Y_k.$$

Therefore the Snell recursion at time k picks the continuation value and

$$U_k = \max\{Y_k,\ \mathbb{E}[U_{k+1} \mid \mathscr{F}_k]\} = \mathbb{E}[U_{k+1} \mid \mathscr{F}_k],$$

which is (ii) at time k.

By backward induction both (i) and (ii) hold for all k. Consequently the stopping region $\{n : U_n = Y_n\}$—and, in particular, its first entrance time τ°—coincides with the first entry into S. The Snell envelope characterization of optimal stopping times yields that τ° is optimal and minimal, hence $\tau^{1\mathrm{SLA}}$ is the smallest optimal stopping time. This completes the proof. □

Remark 1.4. Lemma 1.2 makes explicit why predictability of I^A, I^B is essential: the one-step look-ahead inequality can be checked by comparing the predictable increments ΔI^A_{k+1} and ΔI^B_{k+1}, which are $\mathscr{F}_k$-measurable. The monotonicity assumption is strong: economically it says that once the predictable drift advantage of B over A appears it persists for all later times (so the 1-stage rule remains valid thereafter). If monotonicity fails, the 1-SLA rule need not be optimal and the full Snell-envelope computation is required.

Remark 1.5 (Economic interpretation of the stopping rule). The stopping problem studied in this work should be interpreted as follows. Initially, company B enters the market with prices lower than those of the previously monopolistic company A, making company B the economically preferable choice for the consumer. As competition evolves, the relative price dynamics may progressively reduce this advantage. The stopping time τ^* therefore represents the optimal moment at which the consumer should stop benefiting from the entrant's temporary price advantage and reconsider switching back to the incumbent company.

1.2. Optimal Stopping with Discounting. In financial applications, future gains must be discounted to reflect the time value of money. We introduce a discount factor $\beta \in (0,1)$, related to an interest rate r by $\beta = \frac{1}{1+r}$.

We define the discounted payoff process as $\tilde{Y}_n = \beta^n(p_n^A - p_n^B)$. The problem is to find τ^* that maximizes $\mathbb{E}[\tilde{Y}_\tau]$.

Lemma 1.6 (Discounted one-step decomposition)**.** *Let $(\Omega, \mathscr{F}, \mathbb{P})$ be a probability space equipped with a filtration $\mathbb{F} = (\mathscr{F}_n)_{n=0}^N$ (finite horizon $N \in \mathbb{N}$). Let $(p_n^A)_{n=0}^N$ and $(p_n^B)_{n=0}^N$ be integrable, $\mathbb{F}$-adapted submartingales with Doob decompositions*

$$p_n^F = M_n^F + I_n^F, \qquad F \in \{A, B\}, \tag{1.6}$$

where M^F are $\mathbb{F}$-martingales and I^F are predictable, nondecreasing processes with $I_0^F = 0$. Define the gap $Y_n := p_n^A - p_n^B$ and, for a discount factor $\beta \in (0,1]$, the discounted process

$$Z_n := \beta^n Y_n, \qquad 0 \le n \le N.$$

Then for every $k \in \{0, \dots, N-1\}$ we have the identity (a.s.)

$$\mathbb{E}[Z_{k+1} \mid \mathscr{F}_k] = \beta^{k+1}\big(Y_k + \Delta I_{k+1}^A - \Delta I_{k+1}^B\big), \tag{1.7}$$

where $\Delta I_{k+1}^F := I_{k+1}^F - I_k^F$. Consequently, the one-step look-ahead inequality $Z_k \ge \mathbb{E}[Z_{k+1} \mid \mathscr{F}_k]$ is equivalent to

$$(1-\beta)Y_k \ \ge\ \beta\big(\Delta I_{k+1}^A - \Delta I_{k+1}^B\big), \tag{1.8}$$

and equivalently

$$\Delta I_{k+1}^B \ \ge\ \Delta I_{k+1}^A - \frac{1-\beta}{\beta} Y_k. \tag{1.9}$$

Proof. By (1.6) we may write

$$Y_{k+1} = (M_{k+1}^A - M_{k+1}^B) + (I_{k+1}^A - I_{k+1}^B).$$

Taking conditional expectation with respect to $\mathscr{F}_k$, the martingale increments have zero conditional mean and the predictable increments I_{k+1}^F are $\mathscr{F}_k$-measurable. Hence

$$\mathbb{E}[Y_{k+1} \mid \mathscr{F}_k] = Y_k + \Delta I_{k+1}^A - \Delta I_{k+1}^B.$$

Multiplying both sides by β^{k+1} yields (1.7). The equivalence $Z_k \ge \mathbb{E}[Z_{k+1} \mid \mathscr{F}_k]$ becomes

$$\beta^k Y_k \ge \beta^{k+1}\big(Y_k + \Delta I_{k+1}^A - \Delta I_{k+1}^B\big),$$

which, after dividing by $\beta^k > 0$, is exactly (1.8); rearranging gives (1.9). □

Theorem 1.7 (Discounted monotone case; 1-SLA optimal)**.** *Under the hypotheses of Lemma 1.6 fix $\beta \in (0,1]$ and define*

$$Z_n := \beta^n Y_n, \qquad n = 0, \dots, N.$$

Let the one-step look-ahead set for the discounted process be

$$S_\beta := \{\, k \in \{0, \dots, N-1\} : \ Z_k \ge \mathbb{E}[Z_{k+1} \mid \mathscr{F}_k] \,\}.$$

Assume the following right-monotonicity property for S_β:

$$k \in S_\beta \quad \Longrightarrow \quad j \in S_\beta \ \text{ for all } j = k, k+1, \dots, N-1.$$

Let $U^Z = (U_n^Z)_{n=0}^N$ be the Snell envelope of Z,

$$U_N^Z := Z_N, \qquad U_n^Z := \max\{Z_n,\, \mathbb{E}[U_{n+1}^Z \mid \mathscr{F}_n]\}, \quad n = N-1, \dots, 0,$$

and define the minimal optimal stopping time

$$\tau_\beta^\circ := \inf\{n \geq 0 : \; U_n^Z = Z_n\},$$

(with infimum of empty set taken as N). Then

$$\tau_\beta^\circ \; = \; \inf\{0 \leq k < N : \; k \in S_\beta\},$$

i.e. the first entry time into S_β equals the smallest optimal stopping time for maximizing $\mathbb{E}[Z_\tau] = \mathbb{E}[\beta^\tau Y_\tau]$. In particular the first-entry 1-step look-ahead rule for the discounted process is optimal.

Proof. We prove by backward induction on $k = N, N-1, \dots, 0$ that for each k the two assertions hold a.s.:

(i) If $k \in S_\beta$ then $U_k^Z = Z_k$.
(ii) If $k \notin S_\beta$ then $U_k^Z = \mathbb{E}[U_{k+1}^Z \mid \mathscr{F}_k]$.

The theorem follows because the first time n with $U_n^Z = Z_n$ is therefore the first time the one-step stopping inequality holds, and by right-monotonicity this is the first entry into S_β.

Base case. By definition $U_N^Z = Z_N$, and $N \notin S_\beta$, so (i)–(ii) hold at $k = N$.

Induction step. Fix $k \in \{0, \dots, N-1\}$ and assume (i)–(ii) hold for all j with $k < j \leq N$. We show they hold at k.

Case (A): $k \in S_\beta$. By right-monotonicity, every $j \geq k$ with $j \leq N-1$ belongs to S_β. By the induction hypothesis (i), $U_j^Z = Z_j$ for all such j; in particular $U_{k+1}^Z = Z_{k+1}$ and $\mathbb{E}[U_{k+1}^Z \mid \mathscr{F}_k] = \mathbb{E}[Z_{k+1} \mid \mathscr{F}_k]$. Since $k \in S_\beta$ we have $Z_k \geq \mathbb{E}[Z_{k+1} \mid \mathscr{F}_k]$. Hence

$$U_k^Z = \max\{Z_k,\, \mathbb{E}[U_{k+1}^Z \mid \mathscr{F}_k]\} = Z_k,$$

so (i) holds at k.

Case (B): $k \notin S_\beta$. Then $Z_k < \mathbb{E}[Z_{k+1} \mid \mathscr{F}_k]$. As $U_{k+1}^Z \geq Z_{k+1}$ we obtain

$$\mathbb{E}[U_{k+1}^Z \mid \mathscr{F}_k] \geq \mathbb{E}[Z_{k+1} \mid \mathscr{F}_k] > Z_k,$$

and therefore the Snell recursion yields

$$U_k^Z = \max\{Z_k,\, \mathbb{E}[U_{k+1}^Z \mid \mathscr{F}_k]\} = \mathbb{E}[U_{k+1}^Z \mid \mathscr{F}_k],$$

establishing (ii) at k.

Since both cases are covered, the induction closes and (i)–(ii) hold for every k. Thus the stopping region $\{n : U_n^Z = Z_n\}$ coincides with S_β (modulo the terminal time), and the minimal optimal stopping time τ_β° equals the first entry into S_β. By the Snell envelope optimality theorem this stopping time is optimal and minimal, completing the proof. □

Remark 1.8. Lemma 1.6 shows that discounting introduces a dependence on the current gap Y_k in the one-step inequality: the condition is no longer a pure comparison of predictable increments but contains the term $\frac{1-\beta}{\beta} Y_k$. Consequently the monotonicity assumption for S_β is typically stronger than in the undiscounted case: once the inequality holds at time k it must remain true at all later times, which requires both favorable behavior of the predictable increment differences and limited growth of Y_j thereafter. As $\beta \uparrow 1$

the extra term vanishes and the result reduces to the undiscounted theorem. High interest rates—ow β—raise the bar for continuing: company B's growth must not just beat A's, but beat it by enough to cover the interest income lost by not stopping today.

Remark 1.9 (On discrete-time modelling). Discrete-time formulations are particularly natural in insurance and consumer pricing applications, where premiums and prices are typically revised at monthly, quarterly, or yearly frequencies and decisions are implemented at discrete observation times. For this reason, the discrete-time Snell-envelope framework developed in the present work provides an operationally convenient formulation while retaining a rigorous optimal stopping interpretation.

2. Estimation, prediction, and stopping under uncertainty

This section explains how the increment processes of firms A and B can be estimated from data, how predictive distributions are obtained, and how these estimates feed into the—discounted—one-step look-ahead stopping rule. Our goal is to connect statistical inference with optimal stopping in a clear and operational manner.

We observe per-period price increments

$$X_t^F, \qquad F \in \{A,B\}, \quad t = 1,\dots,n,$$

and form cumulative prices and their gap:

$$p_k^F := \sum_{t=1}^{k} X_t^F, \qquad Y_k := p_k^A - p_k^B.$$

A stopping decision is made at time k by comparing the current gap with its expected one-step-ahead value:

$$Y_k \;\geq\; \mathbb{E}[Y_{k+1} \mid \mathscr{F}_k], \qquad \text{equivalently } \beta^k Y_k \;\geq\; \mathbb{E}[\beta^{k+1} Y_{k+1} \mid \mathscr{F}_k], \tag{2.1}$$

where $\beta \in (0,1]$ is the discount factor and $\mathscr{F}_k$ is the information available at time k. Thus, statistical inference on increments directly enters the stopping decision.

2.1. Increment model and predictive quantities. To keep the analysis tractable, we adopt a simple exponential-family model.

Definition 2.1 (Increment model)**.** For each firm $F \in \{A,B\}$,

$$X_t^F \mid \lambda_F \overset{\text{iid}}{\sim} \text{Exponential}(\lambda_F), \qquad \lambda_F > 0.$$

The corresponding mean increment is

$$\mu_F = \frac{1}{\lambda_F}.$$

This specification yields closed-form frequentist and Bayesian estimators.

2.2. Frequentist estimation and 1-SLA decisions. A natural estimator for the exponential rate is the inverse sample mean.

$$\widehat{\lambda}_F = \frac{1}{\overline{X}^F}, \qquad \overline{X}^F = \frac{1}{n}\sum_{t=1}^{n} X_t^F, \qquad \widehat{m}^F = \frac{1}{\widehat{\lambda}_F}.$$

The plug-in predictive mean for the next increment is therefore $\widehat{m}^F$.
Using the approximation $\mathbb{E}[X_{k+1}^F \mid \mathscr{F}_k] \approx \widehat{m}^F$, the one-step condition (2.1) becomes

$$\widehat{m}^B \;\geq\; \widehat{m}^A + \frac{1-\beta}{\beta} Y_k.$$

In the undiscounted case ($\beta = 1$),

$$\widehat{m}^B \;\geq\; \widehat{m}^A.$$

Thus, the frequentist 1-SLA rule amounts to comparing the estimated mean increments of the two firms.

2.3. Bayesian Estimation and Prediction. A Bayesian approach incorporates parameter uncertainty by averaging over the posterior distribution of each firm's rate.

Definition 2.2 (Prior)**.** Each rate parameter has an independent Gamma prior:

$$\lambda_F \sim \mathrm{Gamma}(a_F, b_F), \qquad a_F, b_F > 0.$$

Proposition 2.3 (Posterior)**.** *Given observed increments $x_{1:n}^F$,*

$$\lambda_F \mid x_{1:n}^F \sim \mathrm{Gamma}(a_F', b_F'), \qquad a_F' = a_F + n, \qquad b_F' = b_F + \sum_{t=1}^{n} x_t^F.$$

Proof. Recall the prior density and the Exponential likelihood (rate parametrisation). For firm F the prior is

$$\pi(\lambda_F) = \frac{b_F^{a_F}}{\Gamma(a_F)} \lambda_F^{a_F - 1} e^{-b_F \lambda_F},$$

and the likelihood for the observed increments $x_{1:n}^F$ is

$$L(\lambda_F; x_{1:n}^F) = \prod_{t=1}^{n} \lambda_F e^{-\lambda_F x_t^F} = \lambda_F^n e^{-\lambda_F \sum_{t=1}^{n} x_t^F}.$$

The posterior density is proportional to prior $\times$ likelihood:

$$\pi(\lambda_F \mid x_{1:n}^F) \propto \lambda_F^{a_F - 1} e^{-b_F \lambda_F}\, \lambda_F^n e^{-\lambda_F \sum_{t=1}^{n} x_t^F} = \lambda_F^{a_F + n - 1} e^{-(b_F + \sum_{t=1}^{n} x_t^F)\lambda_F}.$$

Recognising this as the kernel of a Gamma density with shape $a_F + n$ and rate $b_F + \sum_{t=1}^{n} x_t^F$ finishes the proof. $\square$

Proposition 2.4 (Posterior predictive distribution)**.** *Integrating out λ_F yields the Lomax (Pareto II) predictive density:*

$$p(x \mid x_{1:n}^F) = a_F' \, b_F'^{\,a_F'} (b_F' + x)^{-a_F' - 1}, \qquad x \geq 0.$$

If $a_F' > 1$, the predictive mean exists and satisfies

$$m^F = \frac{b_F'}{a_F' - 1}.$$

Using predictive means (or a Monte Carlo approximation),

$$Y_k \;\geq\; \beta(Y_k + m^A - m^B) \quad\iff\quad m^B \;\geq\; m^A + \frac{1-\beta}{\beta} Y_k.$$

Proof. The posterior predictive density for a new increment X_{new} is

$$p(x \mid x^F_{1:n}) = \int_0^\infty f(x \mid \lambda_F)\, \pi(\lambda_F \mid x^F_{1:n})\, d\lambda_F,$$

where $f(x \mid \lambda_F) = \lambda_F e^{-\lambda_F x}$ is the Exponential likelihood for the new observation. Substituting the Gamma posterior $\text{Gamma}(a'_F, b'_F)$ with $a'_F = a_F + n$ and $b'_F = b_F + \sum_{t=1}^n x^F_t$ gives

$$\begin{aligned} p(x \mid x^F_{1:n}) &= \int_0^\infty \lambda_F e^{-\lambda_F x} \frac{b_F'^{a'_F}}{\Gamma(a'_F)} \lambda_F^{a'_F - 1} e^{-b'_F \lambda_F}\, d\lambda_F \\ &= \frac{b_F'^{a'_F}}{\Gamma(a'_F)} \int_0^\infty \lambda_F^{a'_F} e^{-(b'_F + x)\lambda_F}\, d\lambda_F. \end{aligned}$$

The integral is a standard Gamma integral:

$$\int_0^\infty \lambda^{a'_F} e^{-(b'_F + x)\lambda}\, d\lambda = \frac{\Gamma(a'_F + 1)}{(b'_F + x)^{a'_F + 1}}.$$

Using $\Gamma(a'_F + 1) = a'_F \Gamma(a'_F)$ yields

$$p(x \mid x^F_{1:n}) = \frac{b_F'^{a'_F}}{\Gamma(a'_F)} \cdot \frac{\Gamma(a'_F + 1)}{(b'_F + x)^{a'_F + 1}} = a'_F\, b_F'^{a'_F} (b'_F + x)^{-a'_F - 1},$$

which is the stated Lomax (Pareto type II) density. Finally, when $a'_F > 1$ the mean of this Lomax distribution exists and equals $\dfrac{b'_F}{a'_F - 1}$, as claimed. □

2.4. Posterior Sampling of the Stopping Time. When uncertainty in λ_A and λ_B matters—especially in the early life of a new entrant—it is natural to sample from the posterior distribution of the stopping time itself. The following algorithm provides a complete Bayesian procedure based on posterior draws of rate parameters.

Proposition 2.5 (Posterior inference for stopping decisions)**.** *The sample $\{\tau_\beta^{(i)}\}_{i=1}^M$ yields:*

- *posterior mean and median of τ_β,*
- *credible intervals,*
- *the posterior probability of stopping immediately,*
- *posterior continuation probabilities to any horizon.*

Proof. Let $\{\tau_\beta^{(i)}\}_{i=1}^M$ be the sample of stopping times produced by Algorithm 1 by drawing i.i.d. posterior samples of the parameters and, for each draw, computing the random 1-SLA stopping time using either analytic or Monte Carlo conditional expectations.

By construction each $\tau_\beta^{(i)}$ is a draw from the posterior predictive distribution of the stopping time (i.e. the distribution of the stopping time under the posterior on the model parameters and the predictive law for future increments). Therefore standard Monte Carlo law-of-large-numbers arguments apply:

Algorithm 1 Posterior stopping-time sampler (posterior 1-SLA)

Require: Observed increments $X^F_{1:n}$; posterior draws $\{(\lambda_A^{(i)}, \lambda_B^{(i)})\}_{i=1}^M$; discount $\beta \in (0,1]$; optional horizon $H \geq 1$; method for computing the conditional expectation (analytic or Monte Carlo).

Ensure: Posterior sample $\{\tau_\beta^{(i)}\}_{i=1}^M$

Compute observed cumulative prices p_k^F for $k \leq n$.
for $i = 1$ to M **do**
    Draw predictive future increments $X^{F,(i)}_{n+1:n+H} \sim p(\cdot \mid \lambda_F^{(i)})$.
    Construct extended cumulative prices $p_k^{F,(i)}$ and gaps $Y_k^{(i)}$.
    Set $\tau_\beta^{(i)} \leftarrow \infty$.
    for $k = n$ to $n+H-1$ **do**
        if analytic **then**
            Compute the closed-form predictive moment of $\beta^{k+1} Y_{k+1}$.
        else
            Approximate the expectation via L Monte Carlo one-step draws.
        end if
        if $\beta^k Y_k^{(i)} \geq \widehat{\mathbb{E}}[\beta^{k+1} Y_{k+1}^{(i)}]$ **then**
            $\tau_\beta^{(i)} \leftarrow k$; **break**.
        end if
    end for
    if $\tau_\beta^{(i)} = \infty$ **then**
        Set $\tau_\beta^{(i)} \leftarrow n+H$ (or leave ∞).
    end if
end for
return Posterior sample and summaries (mean, median, quantiles, etc.)

- The sample mean $\frac{1}{M}\sum_{i=1}^M \tau_\beta^{(i)}$ converges almost surely to the posterior expectation $\mathbb{E}[\tau_\beta \mid \text{data}]$ as $M \to \infty$.
- Empirical quantiles of the sample converge to the corresponding posterior quantiles, and hence the Monte Carlo credible intervals converge to posterior credible intervals.
- The empirical frequency $\frac{1}{M}\sum_{i=1}^M \mathbf{1}\{\tau_\beta^{(i)} \leq t\}$ converges to the posterior probability $\Pr(\tau_\beta \leq t \mid \text{data})$ for each fixed t.

When the algorithm uses a Monte Carlo inner loop to approximate the conditional expectation $\mathbb{E}[\beta^{k+1} Y_{k+1} \mid \mathscr{F}_k]$, the outer Monte Carlo sample still targets the correct posterior predictive law provided that the inner MC estimator is unbiased (or that its bias is controlled and vanishes as the inner sample size grows). In practice one chooses the inner sample size L large enough that the numerical error is negligible compared with posterior sampling variability.

Thus the collection $\{\tau_\beta^{(i)}\}_{i=1}^M$ yields consistent Monte Carlo estimators of posterior summaries: mean, median, credible intervals, and event probabilities such as the posterior probability of immediate stopping. This proves the proposition. □

Remark 2.6 (Summary). The framework presented above provides:

(1) estimation of increment distributions either frequentist or Bayesian;
(2) predictive means or full predictive distributions;
(3) discounted one-step stopping decisions;
(4) posterior uncertainty quantification for the stopping time.

3. Numerical illustration: a linear-drift synthetic market

We now illustrate the estimation—prediction—stopping framework developed in Section 2 on a stylised synthetic market. The purpose of this section is twofold: (i) to validate the procedure in a controlled setting where the growth pattern of increments is known, and (ii) to show how frequentist estimators behave under both long and short observation windows.

Throughout this section the "true" increments are generated from a parametric model with deterministic linear drift, but estimation and stopping decisions are computed using only data-driven procedures, without access to the true parameters.

3.1. Model specification. We fix a horizon of $N = 120$ months (ten years). For constants $\beta, \alpha^A, \alpha^B > 0$ we impose the following conditional expectations:

$$\mathbb{E}[X_k^A] = \beta + \alpha^A \frac{k}{N}, \qquad \mathbb{E}[X_k^B] = \alpha^B \frac{k}{N}, \qquad k = 1, \ldots, N. \tag{3.1}$$

This specification induces an initial gap of magnitude β and average incremental growth rates α^A/N and α^B/N for firms A and B. The condition $\alpha^B > \alpha^A$ ensures that the increments of firm B eventually dominate those of firm A in expectation.

A simple deterministic bound for the expected-time crossing produced by Formula (3.1) is

$$k_0 > \frac{\beta N}{\alpha^B - \alpha^A}, \tag{3.2}$$

which provides a benchmark against which to compare the data-driven stopping rules derived below.
Noise specification. To avoid deterministic increments, we draw X_k^A and X_k^B from Gamma distributions whose means equal the right-hand sides of Formula (3.1) and whose shapes are fixed (shape $= 2$), resulting in moderate dispersion and increasing variance over time. This choice is purely illustrative and does not affect the qualitative behaviour of the estimation procedures.

Figures 1 and 2 below illustrate typical samples.

3.2. Parameter choices and qualitative interpretation. We take:

$$\beta = 25, \qquad \alpha^A = 9, \qquad \alpha^B = 59,$$

motivated by the real case of an insurance monopoly break-up in the Angolan market. Here β represents the initial premium difference, α^A the average incremental growth of the incumbent's prices, and α^B the larger growth rate of the entrant.

The deterministic bound (3.2) yields

$$k_0 > 60,$$

suggesting that (in expectation) the entrant undercuts the incumbent only after approximately five years. The simulation results below confirm this prediction when full data are available and show how uncertainty increases when only early data are used.

3.3. Increment simulation. Figure 1 displays the simulated densities of increments for firms A and B at selected times. As k increases, both means and variances increase linearly, with the increments of firm B eventually dominating those of firm A.

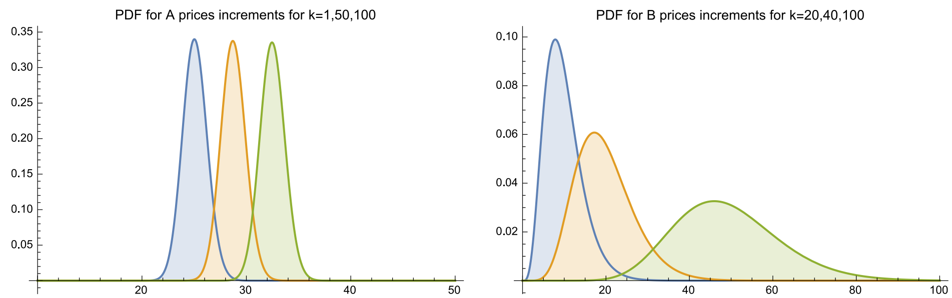

FIGURE 1. Simulated densities of increments X_k^A (left) and X_k^B (right) for selected months.

3.4. Frequentist estimation under a long observation window. When a long time series is available (here the full $N = 120$ months), the parameters $\beta, \alpha^A, \alpha^B$ in Formula (3.1) can be estimated by ordinary least squares regression:

$$X_k^A = \beta + \alpha^A \frac{k}{N} + \varepsilon_k^A, \qquad X_k^B = \alpha^B \frac{k}{N} + \varepsilon_k^B.$$

For the simulated sample, the regression fits, shown in Figure 2, produce accurate parameter estimates. Plugging the estimates into (3.2) yields

$$\widehat{\tau}^\star \approx 60.047,$$

in close agreement with the true value $k_0 > 60$.

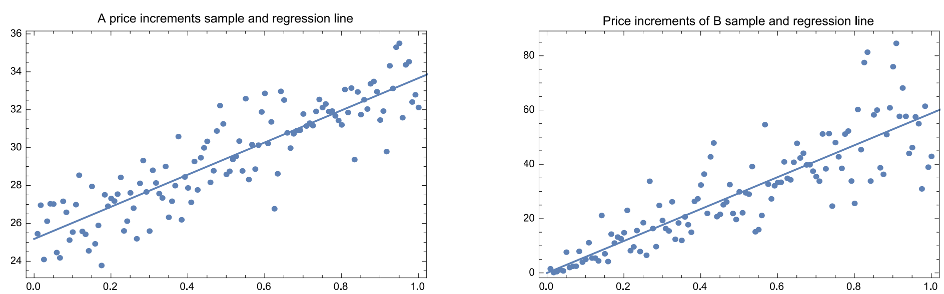

	Estimate	Standard Error	t-Statistic	P-Value
1	25.1751	0.240985	104.467	3.9194×10^{-118}
x	8.48876	0.414807	20.4644	1.26055×10^{-40}

	Estimate	Standard Error	t-Statistic	P-Value
x	58.7995	1.55555	37.7998	3.86251×10^{-68}

FIGURE 2. Simulated increments (dots) and linear regressions for firms A and B. The inferred growth rates closely match the true values.

3.5. Early-phase estimation: pooled small samples. When the new entrant has been active for only a short period, frequentist estimators may become unstable. To reproduce this setting, we simulate $M = 500$ independent one-year samples:

$$X_k^{F,j}, \qquad k = 1, \ldots, 12, \quad j = 1, \ldots, M.$$

From the moment identities

$$\mathbb{E}[X^A_{k+1} - X^A_k] = \frac{\alpha^A}{N}, \qquad \mathbb{E}[X^B_{k+1} - X^B_k] = \frac{\alpha^B}{N}, \qquad \mathbb{E}\left[\frac{X^A_k}{\alpha^A} - \frac{X^B_k}{\alpha^B}\right] = \frac{\beta}{\alpha^A},$$

moment estimators for $\beta, \alpha^A, \alpha^B$ are obtained by averaging across the M samples, giving

$$\widehat{\alpha}^A = 9.19476, \qquad \widehat{\alpha}^B = 59.4274, \qquad \widehat{\beta} = 24.7811.$$

These estimates again imply a stopping time $\widehat{\tau}^\star \approx 59.20$, consistent with the theoretical value.

Remark 3.1 (Economical interpretation). Economically, larger values of $\alpha_B - \alpha_A$ correspond to more aggressive entrant pricing dynamics, thereby accelerating the optimal switching time. Conversely, larger initial gaps β delay the stopping decision because the incumbent's price disadvantage persists longer.

3.6. Smallest optimal stopping time: Monte Carlo evidence. We compute the empirical analogue of:

$$V(\tau) := \mathbb{E}[p^A_\tau - p^B_\tau]$$

by simulating 500 replications of length 120. For each $V(\tau)$ we select the top three maximisers $\widehat{\tau}^\circ$. These serve as Monte Carlo estimates of the smallest optimal stopping time. The results in Table 1 show excellent agreement with the theoretical benchmark in Fornula (3.2) across all parameter choices.

Parameters	Theoretical $\tau^\star$	1st: $\widehat{\tau}^\circ; V(\widehat{\tau}^\circ)$	2nd: $\widehat{\tau}^\circ; V(\widehat{\tau}^\circ)$	3rd: $\widehat{\tau}^\circ; V(\widehat{\tau}^\circ)$
$\beta = 25$, $\alpha^A = 9$, $\alpha^B = 59$	60	60; 733.135	59; 733.044	58; 732.591
$\beta = 15$, $\alpha^A = 9$, $\alpha^B = 49$	45	44; 328.713	43, 328.539	42; 328.516
$\beta = 36$, $\alpha^A = 12$, $\alpha^B = 68$	77.14	77; 1365.59	76; 1365.42	75; 1364.95

TABLE 1. Monte Carlo estimates of the smallest optimal stopping time τ° and corresponding values of $V(\tau)$.

3.7. Bayesian estimation for scarce data. In realistic scenarios only a small number of increments from the entrant may be available. Using the results of Section 2, in order to stabilise inference we can adopt a Bayesian update for the rate parameter of a Gamma increment model with known shape parameter. With prior $\lambda \sim \text{Gamma}(a_0, b_0)$, observing n increments yields the posterior,

$$\lambda \mid X_{1:n} \sim \text{Gamma}(a_0 + n\alpha,\ b_0 + \sum_{i=1}^{n} X_i).$$

The posterior mean of the predictive increment is:

$$\widehat{\mu}_B^{(n)} = \frac{\alpha(b_0 + S^B_n)}{a_0 + n\alpha - 1},$$

provided $a_0 + n\alpha > 1$. This estimator smoothly interpolates between prior beliefs and data and produces stable stopping decisions when n is small.

The main results of the implementation of Algorithm 1 are presented in Figure 3.

It is to be noticed that the smoothed PDF of the stopping time τ has a mean $\overline{\tau} = 177.561$ and a median $\tau_{\text{med}} = 180$, thus corresponding after subtraction of the initial segment of

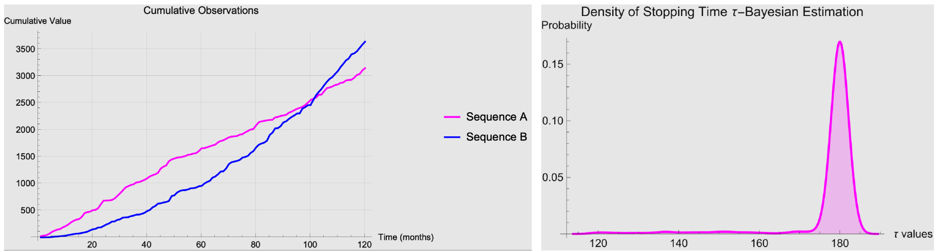

FIGURE 3. Simulated accumulated increments (left) and smoothed distribution of the stopping time τ obtained from the simulation (right).

simulated data of length 120 to the a value $\widehat{\tau}^{\star}_{\text{Mean-Bayes}} \approx 57.56$ and $\widehat{\tau}^{\star}_{\text{Median-Bayes}} \approx 60.$, also consistent with the theoretical value in Formula 3.2.

3.8. Discussion. Across all experiments the estimated stopping times are in close agreement with the theoretical benchmark implied by Formula (3.2). Frequentist regression performs well with long samples; moment estimators remain accurate across many short independent replications; and Bayesian estimation provides stabilised drift estimates in true small-sample situations. This example demonstrates the internal coherence of the estimation–prediction–stopping pipeline and confirms that the one-step rule provides accurate practical stopping decisions in settings with structured but noisy increment dynamics.

3.9. Continuous-Time Analogue and Relation with Perpetual Stopping Problems. The discrete-time optimal stopping framework developed above admits a natural continuous-time analogue related to perpetual stopping and real-options models. Such formulations are classical in Mathematical Finance and allow one to highlight the role of volatility and correlation effects in the switching decision.

Suppose that the price processes of the incumbent company A and the entrant company B evolve according to correlated geometric Brownian motions:

$$dP_t^A = \mu_A P_t^A\,dt + \sigma_A P_t^A\,dW_t^A,$$

$$dP_t^B = \mu_B P_t^B\,dt + \sigma_B P_t^B\,dW_t^B,$$

where

$$d\langle W^A, W^B\rangle_t = \rho\,dt,$$

$\mu_A, \mu_B \in \mathbb{R}$ denote the drift parameters, $\sigma_A, \sigma_B > 0$ the volatilities, and $\rho \in [-1,1]$ the instantaneous correlation coefficient.

The consumer seeks an optimal stopping time τ maximizing the expected discounted price advantage:

$$V(P^A, P^B) = \sup_{\tau} \mathbb{E}\left[e^{-r\tau}(P_\tau^A - P_\tau^B)\right],$$

where $r > 0$ is the discount rate.

Since the payoff function is positively homogeneous of degree one, the dimensionality of the problem can be reduced by introducing the ratio variable

$$x = \frac{P^A}{P^B},$$

and writing

$$V(P^A, P^B) = P^B v(x).$$

Standard stochastic calculus arguments then show that, in the continuation region, the reduced value function v satisfies the second-order ordinary differential equation

$$\frac{1}{2}\sigma^2 x^2 v''(x) + (\mu_A - \mu_B) x v'(x) - (r - \mu_B) v(x) = 0,$$

where

$$\sigma^2 = \sigma_A^2 + \sigma_B^2 - 2\rho\sigma_A\sigma_B.$$

The corresponding characteristic equation is

$$\frac{1}{2}\sigma^2\beta^2 + \left(\mu_A - \mu_B - \frac{1}{2}\sigma^2\right)\beta - (r - \mu_B) = 0.$$

Depending on the relative positions of the drift and discount parameters, different geometric structures of the stopping region may arise. In the classical regime

$$r > \max(\mu_A, \mu_B),$$

one obtains a single-threshold stopping rule analogous to standard perpetual American-option problems. However, for certain parameter regimes satisfying

$$r < \mu_A < \mu_B,$$

the continuation region may become disconnected, producing two continuation regions separated by an intermediate stopping region, a phenomenon related to double continuation regions studied in perpetual stopping problems and real-options theory.

The continuous-time analogue highlights an important structural distinction relative to the present discrete-time framework. In perpetual diffusion models, optimal stopping boundaries are jointly determined by drift, volatility, discounting, and correlation effects through a free-boundary problem. By contrast, the discrete-time one-step look-ahead rule developed in the present work is primarily driven by comparisons between predictable increments in the Doob decompositions of the competing price processes. This distinction makes the discrete-time framework particularly suitable for operational implementations based on sequentially observed market data, such as monthly or yearly insurance premium updates.

Finally, the continuous-time formulation also clarifies the role of correlation. Unlike the discrete-time one-step look-ahead inequalities, which depend mainly on predictable trend comparisons, the correlation coefficient ρ affects the optimal switching policy through the effective volatility σ^2 of the price ratio process, thereby modifying the geometry of the stopping regions.

Remark 3.2. The continuous-time formulation discussed above connects the present work with the literature on perpetual optimal stopping, real options, and free-boundary problems. In particular, perpetual switching and entry problems frequently arise in corporate finance and investment timing under uncertainty; see, for example, [DP94, MS86, PS06, Sch05]. The possibility of disconnected continuation regions is also related to certain perpetual American-option problems studied in the optimal stopping literature.

4. Conclusions

We introduced a model for determining an optimal stopping time for buying to competition breaking monopoly. We present a simulation study in the case of independent price increments with linear growth and develop two estimation procedures for the parameters of this model. With a parameter choice—consistent with the information we have on the real life case—we could estimate with remarkable accuracy the optimal stopping time studied. We computed by simulation the smallest optimal stopping time for three different choices of the parameters and observed that the computed values are remarkably close to the values of the particular optimal stopping time studied. The remarkable agreement between the simulated and theoretical values is fully explained by Theorem 1.3, as the linear growth model used in our simulations satisfies the monotonicity condition required for the global optimality of the one-step look-ahead rule.

The framework developed in this work treats the price dynamics of the two companies as exogenous stochastic processes. This modelling choice is intentional and allows the derivation of tractable optimal stopping rules based on the predictable components of the associated Doob decompositions.

In realistic markets, however, competitive interactions may generate feedback effects between consumer behaviour and future price dynamics. For example, a large migration of consumers toward the entrant company may induce the incumbent to accelerate price reductions, while capacity constraints or market-share objectives may lead the entrant to modify its pricing policy dynamically. Such mechanisms may produce endogenous drift adjustments, regime changes, or strategic interactions between firms.

Incorporating these effects would naturally lead toward stochastic game formulations, mean-field interaction models, or adaptive regime-switching frameworks in which the price dynamics themselves depend on the aggregate switching behaviour of consumers. Although such extensions lie beyond the scope of the present work, they constitute a promising direction for future research linking optimal stopping, market competition, and dynamic pricing theory.

Aknowledgements

Research Financial Support: This work was partially funded—for M.L.E. and N.P.K.—through FCT (the Fundação para a Ciência e a Tecnologia, I.P., Portugal) under the scope of projects UID/297/2025 and UID/PRR/297/2025 (Center for Mathematics and Applications - NOVA Math, Universidade Nova de Lisboa).

Acknowledgement Computational Work: The computations and graphics in this work were done with Wolfram's Mathematica ™ 14.3.0.0 in a Mac OS X ARM (64-bit).

Acknowledgement AI-I: This work made use of OpenAI's ChatGPT to assist with (i) rephrasing and clarifying mathematical arguments, (ii) generating draft sections of explanatory text, and (iii) producing preliminary versions of simulation and computational code. All mathematical proofs, final algorithms, theoretical results, simulations, and interpretations were independently verified and validated by the authors. No confidential or proprietary data were supplied to the AI system.

Acknowledgement AI-II: During the later stages of the preparation of the latest version of this work, the authors used the large language model Gemini to verify the logical flow of the mathematical derivations and for assistance with LaTeX formatting. The authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

[CRS91] Yuan Shih Chow, Herbert Robbins, and David Siegmund. *The theory of optimal stopping. Corrected reprint of the 1971 original.* New York: Dover Publications, Inc., corrected reprint of the 1971 original edition, 1991.

[DHOR11] Pierre Del Moral, Peng Hu, Nadia Oudjane, and Bruno Rémillard. On the robustness of the Snell envelope. *SIAM J. Financ. Math.*, 2:587–626, 2011.

[DLU09] G. Deligiannidis, H. Le, and S. Utev. Optimal stopping for processes with independent increments, and applications. *J. Appl. Probab.*, 46(4):1130–1145, 2009.

[DP94] Avinash K. Dixit and Robert S. Pindyck. *Investment Under Uncertainty*. Princeton University Press, Princeton, New Jersey, 1994.

[ET22] Christian Oliver Ewald and Bart Taub. Real options, risk aversion and markets: A corporate finance perspective. *Journal of Corporate Finance*, 72:102164, 2022.

[FK07] Joseph Farrell and Paul Klemperer. Chapter 31 coordination and lock-in: Competition with switching costs and network effects. In M. Armstrong and R. Porter, editors, *Handbook of Industrial Organization*, volume 3, pages 1967–2072. Elsevier, 2007.

[Kle87] Paul Klemperer. Markets with consumer switching costs. *The Quarterly Journal of Economics*, 102(2):375–394, 1987.

[LL08] Damien Lamberton and Bernhard Lapeyre. *Introduction to stochastic calculus applied to finance. 2nd ed.* Boca Raton, FL: Chapman & Hall/CRC, 2nd ed. edition, 2008.

[LM22] Zhenya Liu and Yuhao Mu. Optimal stopping methods for investment decisions: A literature review. *International Journal of Financial Studies*, 10(4), 2022.

[MS86] Robert McDonald and Daniel Siegel. The value of waiting to invest. *The Quarterly Journal of Economics*, 101(4):707–727, 1986.

[Nev75] J. Neveu. *Discrete-parameter martingales. Translated by T. P. Speed.* North-Holland Mathematical Library. Vol. 10. Amsterdam - Oxford: North- Holland Publishing Company; New York: American Elsevier Publishing Company, Inc. VIII., 1975.

[Pre11] Ernst Presman. A new approach to the solution of optimal stopping problem in a discrete time. *Stochastics*, 83(4-6):467–475, 2011.

[PS06] Goran Peskir and Albert Shiryaev. *Optimal stopping and free-boundary problems.* Basel: Birkhäuser, 2006.

[Sch98] A. Schöttl. Optimal stopping of a risk reserve process with interest and cost rates. *Journal of Applied Probability*, 35(1):pp. 115–123, 1998.

[Sch05] Marcus Schulmerich. *Real options valuation*, volume 559 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, Berlin, 2005. The importance of interest rate modelling in theory and practice, With a foreword by Stewart C. Myers and Ulrich Hommel.

[Shi08] Albert N. Shiryaev. *Optimal stopping rules. Reprint of the 1978 3rd ed.* Berlin: Springer, reprint of the 1978 3rd ed. edition, 2008.

[Shi19] Albert N. Shiryaev. *Stochastic disorder problems*, volume 93 of *Probability Theory and Stochastic Modelling*. Springer, Cham, 2019. With a foreword by H. Vincent Poor.

MANUEL L. ESQUÍVEL AND NADEZHDA P. KRASII: DEPARTAMENTO DE MATEMÁTICA, FACULDADE DE CIÊNCIAS E TECNOLOGIA DA UNIVERSIDADE NOVA DE LISBOA, QUINTA DA TORRE, 2829-516, CAPARICA, PORTUGAL & NOVA MATH, UNL.

Email address: `mle@fct.unl.pt`