

Dispensing with optimal control: a new approach for the pricing and management of share buyback contracts

Bastien BALDACCI*

Philippe BERGAULT†

Olivier GUÉANT‡

Abstract

This paper introduces a novel methodology for the pricing and management of share buyback contracts, overcoming the limitations of traditional optimal control methods, which frequently encounter difficulties with high-dimensional state spaces and the intricacies of selecting appropriate risk penalty or risk aversion parameter. Our methodology applies optimized heuristic strategies to maximize the contract's value. The computation of this value utilizes classical methods typically used for pricing path-dependent Bermudan options. Additionally, our approach naturally leads to the formulation of a hedging strategy.

Key words: share buyback contracts, optimal execution strategies, optimal stopping, stochastic optimal control.

Introduction

The analysis of payout policies, especially the distinction between dividends and share buybacks, is fundamental to corporate finance. Modigliani and Miller's theorem (see [9]), which asserted the irrelevance of payout policies in a perfect market, profoundly influenced this topic. Serving as a benchmark, this seminal result suggests that payout policy considerations only become relevant in the presence of market imperfections, including taxes, transaction costs, and informational asymmetry. These imperfections have a significant impact on a firm's payout strategy. Dividends are often viewed as indicators of consistent earnings and a commitment to future stability, attracting investors who seek regular and dependable income. Beyond their flexibility, share buybacks are strategically favored for several reasons, including signaling stock undervaluation and deterring potential takeovers. They are also appealing to companies aiming to recalibrate their capital structure. Each mechanism meets specific strategic requirements, reflecting a firm's operational situation and market perception (see for instance [1]).

Over the past decades, share buybacks have globally surged, even outpacing dividend increases, though they faced a downturn in 2023 in the US due to rising interest rates, as US companies frequently finance buybacks with borrowed funds. The shift from fixed-price tender offers and Dutch auctions to predominantly open market repurchases (OMRs) in the 90s, followed by the rise of Accelerated Share Repurchase (ASR) contracts and buyback mandates, historically marks a significant evolution in how buybacks are conducted. ASR setups are more common in the US. This approach involves a company rapidly repurchasing its own shares through a contract with an investment bank. Typically, the bank initially borrows the shares to deliver them to the company, then gradually closes its short position by buying shares in the open market. In contrast, the mandate format, which is more prevalent in Europe, allows the bank to slowly purchase shares on the market and progressively deliver them to the company.

*HSBC, 38 avenue Kléber, 75116 Paris, France.

†Université Paris Dauphine-PSL, Ceremade, Place du Maréchal de Lattre de Tassigny, 75116 Paris, France, bergault@ceremade.dauphine.fr.

‡Université Paris 1 Panthéon-Sorbonne, Centre d'Economie de la Sorbonne, 106 Boulevard de l'Hôpital, 75642 Paris Cedex 13, France, olivier.gueant@univ-paris1.fr.

In both ASRs and mandate share buyback contracts, the bank's compensation is often tied to a price benchmark, typically the Volume Weighted Average Price (VWAP) over a period decided dynamically by the bank which faces an optimal stopping time problem. Competitive pricing strategies in these contracts hinge on offering discounts relative to this benchmark. A company may request specific features in these contracts, and banks, particularly those advising the firm, can propose complex terms. Recent years have seen an increase in original features within contracts such as floating notional, lookback features, day exclusion upon insufficient daily volume or high price increase and subsequent maturity extensions or notional reduction. This trend emphasizes the importance of quantitative analysts in accurately pricing and managing these contracts.

Often seen as variants of optimal execution problems, the pricing and management of share buyback contracts have mainly been addressed with the tools of stochastic optimal control. Jaimungal and his coauthors in [7], followed by Guéant and his coauthors in [5], and then Guéant in [3] have indeed proposed several models relying on optimal control tools for both fixed quantity and fixed notional ASR contracts. These models ultimately involve solving Hamilton-Jacobi-Bellman equations, or simply Bellman equations, using grid or sometimes tree methods. While these methods are efficient in simple cases and have been adopted by several banks, they encounter three significant challenges. First, dealing with complex contracts or advanced price dynamics necessitates a high-dimensional state space, with mandates posing greater difficulties in this regard than ASR contracts. Although neural networks can potentially overcome the curse of dimensionality (see [4] and [6] for techniques inspired by reinforcement learning ideas), their opaque decision-making process and unpredictable behavior with extreme state values often lead to their rejection by practitioners. The second issue concerns the selection of an appropriate risk penalty and/or risk aversion parameter to mitigate contract execution risks. This choice significantly influences both optimal strategies and pricing. Moreover, the concept of pricing tied to optimal control methods relates to indifference pricing, seldom used by practitioners. Third, and related to the previous points, optimal control tools often conflict with the pricing and hedging frameworks that are commonly found in the libraries of most investment banks.

To address these limitations, Baldacci *et al.* introduced in [2] an alternative approach that simplifies the problem through a heuristic repurchase strategy. This strategy is designed to intuitively align with the expected directionality of parameter effects, reducing the complexity to the more manageable task of pricing exotic American or Bermudan options using conventional methods. Our research aims to expand upon this foundation, focusing on refining and optimizing heuristic strategies, thus offering a more detailed insight into the management of share buyback strategies and contracts.

Stochastic optimal control versus heuristic strategies

A simple framework

In practice, the problem faced by a trader entering a buyback contract is a discrete one: each day, the trader chooses the quantity to be bought. In this section, in order to simplify the mathematical presentation, we consider a continuous-time approximation of the problem, as often done in the literature.

Assume that the price process $(S_t)_t$ of the stock is given by

$$S_t = S_0 + \int_0^t \sigma S_u dW_u,$$

where $S_0, \sigma > 0$ are given and $(W_t)_t$ is a standard Brownian motion under the risk-neutral probability $\mathbb{Q}$. Then the process $(A_t)_t$ corresponding to the time-weighted average price process has the following dynamics:

$$dA_t = \frac{S_t - A_t}{t} dt.$$

In our first and simple framework, we consider the case of a theoretical buyback contract with notional $F > 0$, maturity $T > 0$, and no constraint whatsoever. We assume that there is no friction in the market,

i.e. the trader can buy or sell any number of shares at any time without transaction costs or market impact.¹

If the contract states that the bank gets paid the time-weighted average price upon delivery of the shares at time τ for an amount F spent on the market, then the payoff of the bank is $F \frac{A_\tau}{S_\tau} - F$ and the bank simply has to solve the following maximization problem:

$$\sup_{\tau \in \mathcal{T}_{0,T}} \mathbb{E}^\mathbb{Q} \left[\frac{A_\tau}{S_\tau} \right],$$

where for all $t \in [0, T]$, $\mathcal{T}_{t,T}$ is the set of stopping times taking values in $[t, T]$.

Defining the process $(Y_t)_t$ by $Y_t = \frac{A_t}{S_t}$, $\forall t \in (0, T]$ and $Y_0 = 1$, we clearly have that, for $t > 0$,

$$dY_t = \left(\frac{1 - Y_t}{t} + \sigma^2 Y_t \right) dt - \sigma Y_t dW_t. \quad (1)$$

In particular, $(Y_t)_t$ is Markovian, and our problem boils down to

$$\sup_{\tau \in \mathcal{T}_{0,T}} \mathbb{E}^\mathbb{Q} [Y_\tau]. \quad (2)$$

The PDE approach

Problem (2) is an optimal stopping problem that can be solved using standard tools from stochastic optimal control theory. Let us introduce the value function

$$u : (t, y) \in [0, T] \times \mathbb{R}_+^* \mapsto \sup_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}^\mathbb{Q} \left[Y_\tau \middle| Y_t = y \right]. \quad (3)$$

It is then well known that u satisfies the following quasi-variational inequality in the viscosity sense:

$$\begin{cases} \min \left\{ -\partial_t u(t, y) - \frac{1}{2} \sigma^2 y^2 \partial_{yy}^2 u(t, y) - \left(\sigma^2 y + \frac{1-y}{t} \right) \partial_y u(t, y), u(t, y) - y \right\} = 0 & \forall (t, y) \in [0, T) \times \mathbb{R}_+^*, \\ u(T, y) = y & \forall y \in \mathbb{R}_+^*. \end{cases} \quad (4)$$

Of course, this equation cannot be solved in closed-form. In order to obtain the pricing function and optimal strategy associated with Problem (2), a classical method consists in using an implicit Euler scheme to solve Equation (4) on a grid. Such schemes are widely used in practice, but they require to introduce some boundary conditions for the numerical computation, such as Neumann conditions, and there is of course a discretization error.

The Longstaff-Schwartz approach

For $N \in \mathbb{N}^*$, let us introduce a subdivision $\Delta^N = \{0 = t_0 < t_1 < \dots < t_N = T\}$ of the interval $[0, T]$.

The dynamic programming principle tells us that

$$u(t_n, y) = \max \left\{ \mathbb{E}^\mathbb{Q} \left[u(t_{n+1}, Y_{t_{n+1}}) \middle| Y_{t_n} = y \right], y \right\}, \quad \forall (n, y) \in \{0, \dots, N-1\} \times \mathbb{R}_+^*. \quad (5)$$

The Longstaff-Schwartz method is a Monte-Carlo method based on the above dynamic programming principle. One first needs to simulate $M \in \mathbb{N}^*$ sample paths of the process $(Y_{t_n}^1)_n, \dots, (Y_{t_n}^M)_n$. For each $m \in \{1, \dots, M\}$, we have the following terminal condition:

$$u(t_N, Y_{t_N}^m) = Y_{t_N}^m.$$

¹We diverge here from the traditional perspective commonly found in the academic literature, which treats buyback issues primarily as execution challenges. This departure is partly based on the idea that the market has already assimilated the informational content of trades. Moreover, we posit that the option component inherent in buyback contracts holds more significance than the aspects related to execution costs.

In order to proceed by backward induction, one needs to compute at each time $n \in \{0, \dots, N-1\}$ and for each sample path m the value of

$$\mathbb{E}^{\mathbb{Q}} \left[u(t_{n+1}, Y_{t_{n+1}}) \middle| Y_{t_n} = Y_{t_n}^m \right].$$

The idea of Longstaff and Schwartz [8] consists in approximating the above conditional expectation by regressing at each time step the values $\left(u(t_{n+1}, Y_{t_{n+1}}^m) \right)_m$ on simple functions of $(Y_{t_n}^m)_m$. This method can be very effective, but highly depends on the choice of the basis of functions used in the regression.

An optimized heuristic strategy

Instead of choosing boundary conditions for Equation (4), or a basis for the regression in order to approximate the conditional expectation in Equation (5), one can wonder if it is possible to choose directly a parametric form for the execution frontier and then optimize over the parameters.

More precisely, the idea consists first in choosing a family

$$\mathcal{F} := \{ f(\cdot, \cdot; \beta) : [0, T] \times \mathbb{R}_+^* \rightarrow \mathbb{R} \mid \beta \in B \},$$

where $B \subset \mathbb{R}^d$. We then define for each $\beta \in B$ the stopping time

$$\tau^\beta = \inf \{ t \in [0, T] \mid f(t, Y_t; \beta) > 0 \},$$

with the convention $\inf \emptyset = T$. We can approximate the value of $\mathbb{E}^{\mathbb{Q}}[Y_{\tau^\beta}]$ with a standard Monte-Carlo method, and then use a grid-search algorithm in order to optimize over $\beta \in B$. If the family $\mathcal{F}$ is rich enough, the approximate problem

$$\sup_{\beta \in B} \mathbb{E}^{\mathbb{Q}}[Y_{\tau^\beta}] \tag{6}$$

should be close to Problem (2).

Comparison of the different methods

In this part, we compare numerically the performance of the strategies obtained with the three methods described above, respectively. We consider the following parameters:

- Initial price $S_0 = 10$ €;
- Volatility $\sigma = 0.2$ year^{-1/2};
- Time horizon $T = \frac{1}{12}$ year.

For the first approach – hereafter denoted by PDE, we use an implicit Euler scheme for Equation (4) on $[0, T] \times [0.8, 1.2]$ with Neumann conditions at the boundaries.

For the second approach – hereafter denoted by LS, we approximate the conditional expectations with polynomials of order 2 in the y variable.

For the third approach – hereafter denoted by OHS, the family $\mathcal{F}$ is given by

$$\mathcal{F} := \{ f(\cdot, \cdot; \beta) : [0, T] \times \mathbb{R}_+^* \rightarrow \mathbb{R} \mid \beta \in \mathbb{R}^3 \},$$

where for all $(t, y) \in [0, T] \times \mathbb{R}_+^*$ and $\beta \in \mathbb{R}^3$,

$$f(t, y; \beta) = \beta_1 + \beta_2 \left(1 - \frac{t}{T} \right) + \beta_3 \left(1 - \frac{t}{T} \right)^2 - (y - 1).$$

We then play the strategies obtained with the different methods using a Monte-Carlo simulation with 10,000 trajectories. The average payoff and standard deviation for each method is reported in Table 1. The distribution of the values obtained for each method are plotted in Figures 1, 2, and 3, respectively. It is remarkable that the three strategies yield very close results both in terms of expected value and standard deviation. The distributions under $\mathbb{Q}$ of the resulting payoffs are spread below 1 and have a spike above 1. In all cases, the variability of the payoffs can be hedged away using standard mathematical finance techniques.

Method	PDE	LS	OHS
Expected value	1.0099	1.0112	1.0111
Standard deviation	0.0298	0.0241	0.0194

Table 1: Results for the different strategies.




Figure 1: Distribution of the values obtained with the PDE method.




Figure 2: Distribution of the values obtained with the LS method.




Figure 3: Distribution of the values obtained with the OHS method.

Back to buybacks

A high-dimensional problem

The three methods – PDE, Longstaff-Schwartz, and Optimized Heuristic Strategy – yield very similar results in terms of pricing for the simple contract discussed in the previous section. However, although the PDE and Longstaff-Schwartz methods can be easily implemented in the simple setting we considered earlier, they become impractical for more complex buyback contracts with numerous features. This impracticality arises because both methods suffer from the curse of dimensionality.

Many buyback contracts incorporate indeed constraints related to the daily volume that can be executed, which may vary based on factors such as the current asset price or the total market traded volume. Additionally, the computation of the average price may exclude certain days, for instance, if the trading volume is exceptionally low or the asset price is particularly high.

In what follows we consider a buyback contract with a floating maturity comprised between $N_{min} > 0$ and $N_{max} \geq N_{min}$, and a floating notional comprised between $F_{min} > 0$ and $F_{max} \geq F_{min}$. The trader can choose to exercise early at any time after a given date N_{ex} , with $0 \leq N_{ex} \leq N_{min}$.

Unlike in the previous section, we consider here a discrete-time model with daily decision making: on day n , the trader chooses the quantity $v_n \in [v_{min}, v_{max}]$ to be bought, with $0 \leq v_{min} < v_{max}$ specified in the contract.

There may be a cap S_{max} on the price such that, if $S_n > S_{max}$, the day is suspended, i.e. excluded from the computation of the average price.² For all n , let us denote by C_n the number of suspended days up to time n . Initially, the maturity of the contract is given by N_{min} . When a trading day is suspended, the maturity of the program is extended by one day up to the maximal maturity N_{max} . If the number of suspended days

²A barrier V_{min} might also be enforced such that, if the total market volume V_n on day n is such that $V_n < V_{min}$, the day n is suspended. However, this case is not treated in this paper, as the variability of V_n is not modeled and can hardly be hedged.

exceeds $N_{max} - N_{min}$, F_{min} is reduced pro-rata by the excess number of suspended days, i.e.

$$F_{min,n} := F_{min} \left(1 - \frac{\max(C_n - (N_{max} - N_{min}), 0)}{N_{max}} \right).$$

Buybacks as payoffs

Let us denote by q_n the number of shares bought by the trader and delivered to the company until day n , and X_n the associated cash spent by the trader up to day n . We have $q_0 = X_0 = 0$, and

$$\begin{cases} q_{n+1} = q_n + v_n, \\ X_{n+1} = X_n + v_n S_n, \end{cases} \quad (7)$$

for each $n \geq 0$.³ Moreover, S_0 is known, and we assume that

$$S_{n+1} = S_n e^{\bar{\sigma} \xi_{n+1}},$$

where $\bar{\sigma} > 0$ is the daily volatility of the price and $(\xi_n)_n$ are i.i.d. $\mathcal{N}(0, 1)$ random variables.

For each $n \geq 0$, let us denote by P_n the set of days that were not suspended up to day n , i.e.

$$P_n = \{k \in \{0, \dots, n\} \mid \text{day } k \text{ was not suspended}\}.$$

We also denote by $\#P_n$ the cardinal of P_n . We then introduce the process $(A_n)_n$ for the cumulative average of the price process $(S_n)_n$ excluding the suspended days, i.e.

$$A_n = \frac{1}{\#P_n} \sum_{k \in P_n} S_k.$$

In contrast to the simpler case discussed previously, it is not immediately evident that the buyback contract incorporating additional features can still be formulated as a payoff structure. Nevertheless, for a predetermined strategy $(v_n)_n$, it becomes apparent that if the trader decides to stop at time $\tilde{N} \geq N_{ex}$ (after N_{ex} , the trader has the right to stop whenever $F_{min, \tilde{N}} \leq X_{\tilde{N}} \leq F_{max}$), the resulting payoff can be precisely expressed as

$$q_{\tilde{N}} A_{\tilde{N}} - X_{\tilde{N}}.$$

To incentivize the trader to spend at least the minimum required notional,⁴ we slightly modify the payoff into

$$q_{\tau} A_{\tau} - \max(F_{min, \tau}, X_{\tau}) \quad (8)$$

where

$$\tau \in \mathcal{T}_{ex} = \{\tau \text{ stopping time} \mid N_{ex} \leq \tau \leq N_{min} + \min(C_{\tau}, N_{max} - N_{min}) \text{ a.s.}\}.$$

Therefore, the trader just solves the following problem

$$\sup_{(v_n)_n} \sup_{\tau \in \mathcal{T}_{ex}} \mathbb{E}^{\mathbb{Q}} [q_{\tau} A_{\tau} - \max(F_{min, \tau}, X_{\tau})]. \quad (9)$$

To solve this problem, we propose in the next section different heuristics for both the execution process $(v_n)_n$ and the stopping time τ .

³We impose that the cash spent remains below the maximum notional F_{max} by forbidding trades once it is reached.

⁴At the stopping time, we buy the required number of shares to reach the minimum notional if it has not been reached yet.

Numerical results and discussion

In this section, we present various numerical examples to illustrate the application of our method. We begin with a very simple buyback contract and then incrementally introduce additional features to demonstrate the method's adaptability and depth.

In all the examples, we choose $S_0 = 10 \text{ €}$ and $\bar{\sigma} = 0.2 \text{ year}^{-1/2}$.

We use three naive strategies as benchmarks:

- The linear policy: each day, choose $v_n = \frac{F_{max} - X_n}{S_n(T_n - n)}$ with $T_n = \min(N_{min} + C_n, N_{max})$, until the end (no early stopping).
- The minmaxtarget policy: each day, choose

$$v_n = \begin{cases} \frac{F_{max} - X_n}{S_n(T_n - n)} & \text{if } S_n \leq A_n, \\ \frac{F_{min} - X_n}{S_n(T_n - n)} & \text{otherwise,} \end{cases}$$

and stop when $S_n < A_n$ and $n \geq N_{ex}$.

- The no trade - no early stopping policy: wait until maturity and execute everything at the end.

A simple buyback

We first consider a buyback with the following parameters:

- Maturity: $N := N_{min} = N_{max} = 60$ days;
- Notional: $F := F_{min} = F_{max} = 200\text{M€}$;
- Early exercise: $N_{ex} = 40$ days;
- Bounds: $v_{min} = 0, v_{max} = +\infty$;
- Price cap: $S_{max} = +\infty$.

We propose a first heuristic strategy (hereafter denoted by Optuna- α, a), given as follows:

- Each day, choose

$$v_n = \frac{2}{1 + e^a} \times \frac{\bar{F} - X_n}{S_n(T_n - n)},$$

with $\bar{F} = \frac{F_{min} + F_{max}}{2}$.

- Stop when $A_n/S_n - 1 \geq \alpha$.

We can also try to improve the early stopping side of the strategy, and consider a second heuristic strategy (Optuna- α, β, γ, a) as follows:

- Each day, choose

$$v_n = \frac{2}{1 + e^a} \times \frac{\bar{F} - X_n}{S_n(T_n - n)},$$

with $\bar{F} = \frac{F_{min} + F_{max}}{2}$.

- Stop when $A_n/S_n - 1 \geq \alpha + \beta \frac{T_n - n}{T_n} + \gamma \frac{X_n}{F_{min}}$.

These two strategies are then optimized over α, β, γ , and a using the `optuna` library, and their respective performance are reported in the table below (corresponding to 2,000 simulations). We see in particular that the two strategies clearly outperform the benchmarks in terms of average PnL. We also observe that adding more parameters to the early execution strategy does not improve the performance.

	PnL in million	Mean (bp)	StdDev (bp)
Linear policy	0.33	16.34	15.07
Minmaxtarget policy	1.53	76.57	92.49
No trade - no early stopping	0.96	47.92	573.50
Optuna- α, a	2.01	100.23	544.67
Optuna- α, β, γ, a	1.98	99.19	548.31

Buyback with cap on the volume

We now add a cap on the volume that can be traded on each day, and consider a buyback with the following parameters:

- Maturity: $N := N_{min} = N_{max} = 60$ days;
- Notional: $F := F_{min} = F_{max} = 200\text{M€}$;
- Early exercise: $N_{ex} = 40$ days;
- Bounds: $v_{min} = 0, v_{max} = 0.2 \frac{\bar{V}}{S_0}$;
- Price cap: $S_{max} = +\infty$.

Here $\bar{V} = 400\text{M€}$ denotes the average daily volume.

In order to capture the increased complexity of the product, we introduce some new strategies.

The third heuristic strategy (Optuna- α, a, a_1) is given by:

- Each day, choose

$$v_n = \frac{2}{1 + \exp \left\{ a + a_1 \frac{T_n - n}{T_n} \right\}} \times \frac{\bar{F} - X_n}{S_n(T_n - n)},$$

$$\text{with } \bar{F} = \frac{F_{min} + F_{max}}{2}.$$

- Stop when $A_n/S_n - 1 \geq \alpha$.

The fourth heuristic strategy (Optuna- α, a, a_1, a_2) is given by:

- Each day, choose

$$v_n = \frac{2}{1 + \exp \left\{ a + a_1 \frac{T_n - n}{T_n} + a_2 \frac{\max(N_{ex} - n, 0)}{N_{ex}} \right\}} \times \frac{\bar{F} - X_n}{S_n(T_n - n)},$$

$$\text{with } \bar{F} = \frac{F_{min} + F_{max}}{2}.$$

- Stop when $A_n/S_n - 1 \geq \alpha$.

The fifth heuristic strategy (Optuna- α, a, a_1, b) is given by:

- Each day, choose

$$v_n = \frac{2}{1 + \exp \left\{ a + a_1 \frac{T_n - n}{T_n} + b \left(\frac{A}{S} - 1 \right) \right\}} \times \frac{\bar{F} - X_n}{S_n(T_n - n)},$$

$$\text{with } \bar{F} = \frac{F_{min} + F_{max}}{2}.$$

- Stop when $A_n/S_n - 1 \geq \alpha$.

And the sixth heuristic strategy (Optuna- $\alpha, a, a_1, b, b_1, c$) is given by:

- Each day, choose

$$v_n = \frac{2}{1 + \exp \left\{ a + a_1 \frac{T_n - n}{T_n} + \left(b + b_1 \frac{T_n - n}{T_n} \right) \left(\frac{A}{S} - 1 \right) + c \frac{F_{min} - X_n}{S_n v_{max}(T_n - n)} \right\}} \times \frac{\bar{F} - X_n}{S_n(T_n - n)},$$

$$\text{with } \bar{F} = \frac{F_{min} + F_{max}}{2}.$$

- Stop when $A_n/S_n - 1 \geq \alpha$.

The results are reported in the table below. Observe in particular that, as soon as we add a cap on the volume, the third benchmark strategy can never finish the program and therefore yields a negative payoff. Again, our optimized heuristic strategies outperform the benchmarks, and increasing the complexity may results in a slightly higher average PnL, but the difference is never significant: simple heuristic strategies seem to be already very close their more complex counterparts.

	PnL in million	Mean (bp)	StdDev (bp)
Linear policy	0.33	16.34	15.07
Minmaxtarget policy	1.52	76.22	91.33
No trade - no early stopping	-120.00	-5999.97	227.56
Optuna- α, a	1.62	80.87	60.34
Optuna- α, β, γ, a	1.63	81.41	63.21
Optuna- α, a, a_1	1.66	82.89	95.60
Optuna- α, a, a_1, a_2	1.66	83.23	106.10
Optuna- α, a, a_1, b	1.68	84.03	115.28
Optuna- $\alpha, a, a_1, b, b_1, c$	1.68	83.93	119.19

Buyback with cap on the volume and floating notional

We now consider a feature well-known to practitioners, sometimes called “flex size” or “Greenshoe”. More precisely, we consider a buyback program with the following parameters:

- Maturity: $N_{min} = N_{max} = 60$ days;
- Notional: $F_{min} = 200\text{M€}$, $F_{max} = 250\text{M€}$;
- Early exercise: $N_{ex} = 40$ days;
- Bounds: $v_{min} = 0$, $v_{max} = 0.2 \frac{\bar{V}}{S_0 \Delta t}$;
- Price cap: $S_{max} = +\infty$.

We compare the different strategies, and report the results in the table below. Notice first that allowing the trader to buy more shares always results in an increased PnL. Again, our six strategies clearly beat the benchmarks.

	PnL in million	Mean (bp)	StdDev (bp)
Linear policy	1.40	70.16	84.38
Minmaxtarget policy	2.21	110.66	97.32
No trade - no early stopping	-120.00	-5999.97	227.56
Optuna- α, a	2.84	142.19	94.25
Optuna- α, β, γ, a	2.88	143.75	98.39
Optuna- α, a, a_1	2.87	143.48	102.14
Optuna- α, a, a_1, a_2	2.87	143.62	99.09
Optuna- α, a, a_1, b	2.89	144.55	85.69
Optuna- $\alpha, a, a_1, b, b_1, c$	2.90	145.20	80.81

Buyback with cap on the volume, floating notional and price cap

Finally, we add a price cap to the program and consider the following parameters:

- Maturity: $N_{min} = 60$ days, $N_{max} = 65$ days;
- Notional: $F_{min} = 200\text{M€}$, $F_{max} = 250\text{M€}$;
- Early exercise: $N_{ex} = 40$ days;
- Bounds: $v_{min} = 0$, $v_{max} = 0.2 \frac{\bar{V}}{S_0 \Delta t}$;
- Price cap: $S_{max} = 1.2S_0 = 12\text{€}$.

We compare the different strategies and report the results in the table below.

	PnL in million	Mean (bp)	StdDev (bp)
Linear policy	1.36	68.21	83.49
Minmaxtarget policy	2.19	109.59	99.48
No trade - no early stopping	-118.77	-5938.64	537.30
Optuna- α, a	2.79	139.30	92.69
Optuna- α, β, γ, a	2.81	140.52	95.84
Optuna- α, a, a_1	2.80	140.11	98.13
Optuna- α, a, a_1, a_2	2.81	140.27	96.84
Optuna- α, a, a_1, b	2.83	141.61	88.10
Optuna- $\alpha, a, a_1, b, b_1, c$	2.86	143.21	77.23

Concluding remarks

In conclusion, our research marks a significant shift from traditional approaches that relied on stochastic optimal control tools for pricing and managing share buyback contracts, which were originally inspired by optimal execution methods. By refining and optimizing heuristic strategies, we successfully address the limitations posed by high-dimensional state spaces and the complex selection of risk parameters. Our method simplifies the issue to pricing and hedging a Bermudan payoff, making it tractable with conventional financial tools. This approach not only facilitates a practical bridge between theory and application in financial markets but also separates the strategies for delivering shares to the company from the hedging of the payoff, potentially elucidating execution patterns observed in the market.

Statement and acknowledgment

This research has been conducted with the support of the Research Initiative “Modélisation des marchés actions, obligations et dérivés” financed by HSBC France under the aegis of the Europlace Institute of Finance. The views expressed are those of the authors and do not necessarily reflect the views or the practices at HSBC.

References

- [1] Franklin Allen and Roni Michaely. Payout policy. *Handbook of the Economics of Finance*, 1:337–429, 2003.
- [2] Bastien Baldacci, Jérôme Lemue, and Lee Russell. A new method for the pricing of share buy-back programmes. *Risk Magazine*, 2024.
- [3] Olivier Guéant. Optimal execution of accelerated share repurchase contracts with fixed notional. *Journal of Risk*, 19(5), 2017.

- [4] Olivier Guéant, Iuliia Manziuk, and Jiang Pu. Accelerated share repurchase and other buyback programs: what neural networks can bring. *Quantitative Finance*, 20(8):1389–1404, 2020.
- [5] Olivier Guéant, Jiang Pu, and Guillaume Royer. Accelerated share repurchase: pricing and execution strategy. *International Journal of Theoretical and Applied Finance*, 18(03):1550019, 2015.
- [6] Mohamed Hamdouche, Pierre Henry-Labordere, and Huyên Pham. Policy gradient learning methods for stochastic control with exit time and applications to share repurchase pricing. *Applied Mathematical Finance*, 29(6):439–456, 2022.
- [7] Sebastian Jaimungal, Damir Kinzebulatov, and Dmitri Rubisov. Optimal accelerated share repurchase. *Available at SSRN 2360394*, 2013.
- [8] Francis A Longstaff and Eduardo S Schwartz. Valuing american options by simulation: a simple least-squares approach. *The review of financial studies*, 14(1):113–147, 2001.
- [9] Franco Modigliani and Merton H Miller. The cost of capital, corporation finance and the theory of investment. *The American economic review*, 48(3):261–297, 1958.