



Accelerated Share Repurchase and other buyback programs: what neural networks can bring

Olivier Guéant, Iuliia Manziuk, Jiang Pu

► To cite this version:

Olivier Guéant, Iuliia Manziuk, Jiang Pu. Accelerated Share Repurchase and other buyback programs: what neural networks can bring. 2020. hal-02987889

HAL Id: hal-02987889

<https://hal.science/hal-02987889>

Preprint submitted on 4 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Accelerated Share Repurchase and other buyback programs: what neural networks can bring*

Olivier Guéant[†], Iuliia Manziuk[‡], Jiang Pu[§]

Abstract

When firms want to buy back their own shares, they have a choice between several alternatives. If they often carry out open market repurchase, they also increasingly rely on banks through complex buyback contracts involving option components, e.g. accelerated share repurchase contracts, VWAP-minus profit-sharing contracts, etc. The entanglement between the execution problem and the option hedging problem makes the management of these contracts a difficult task that should not boil down to simple Greek-based risk hedging, contrary to what happens with classical books of options. In this paper, we propose a machine learning method to optimally manage several types of buyback contract. In particular, we recover strategies similar to those obtained in the literature with partial differential equation and recombining tree methods and show that our new method, which does not suffer from the curse of dimensionality, enables to address types of contract that could not be addressed with grid or tree methods.

Key words: ASR contracts, Optimal stopping, Stochastic optimal control, Deep learning, Recurrent neural networks, Reinforcement learning.

*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 authors would like to thank Philippe Bergault (Université Paris 1 Panthéon-Sorbonne), Marc Chataigner (Université Evry Val-d’Essonne), Dan Ederly (HSBC), Nicolas Grandchamp des Raux (HSBC), Greg Molin (HSBC), and Kamal Omari (HSBC) for the discussions they had on the topic. The authors also would like to thank two anonymous referees for their relevant remarks and questions that allowed to improve the paper. The readers should nevertheless be aware that the views, thoughts, and opinions expressed in the text belong solely to the authors.

[†]Université Paris 1 Panthéon-Sorbonne. Centre d’Economie de la Sorbonne. 106, boulevard de l’Hôpital, 75013 Paris, France. Corresponding author. email: olivier.gueant@univ-paris1.fr

[‡]Université Paris 1 Panthéon-Sorbonne. Centre d’Economie de la Sorbonne. 106, boulevard de l’Hôpital, 75013 Paris, France.

[§]Institut Europlace de Finance. 28, place de la Bourse, 75002 Paris, France.

1 Introduction

Payout policy has been a major research topic in corporate finance since the payout irrelevance proposition of Modigliani and Miller [32] stating the equivalence of dividend payment and share buyback in an idealised market without taxes, frictions, and information asymmetries. When taxes, frictions, and information asymmetries enter the scene, there could be reasons to prefer share buybacks over dividend payments, or *vice versa*. In practice, in addition to fiscal motives in some regions, share buybacks are often favoured for signalling stock price undervaluation, for deterring takeover, or for offsetting the dilution effect associated with stock options (see [2, 15] for a review on payout policy).

Share buybacks can be carried out using several methods. Until the end of the 80s, share repurchases were predominantly made via fixed-price tender offers and Dutch auctions.¹ Then, in the 90s, open market repurchases (OMRs) took over and represented the vast majority of share buyback programs (see [35]). However, as reported for instance in [8], after a share repurchase announcement, a substantial number of companies usually do not commit to it. In order to make a credible commitment, an increasing number of firms started, from the early 2000s, to sign contracts with investment banks to delegate buyback programs in the form of VWAP-minus programs. The main examples of such contracts are Accelerated Share Repurchase (ASR) contracts.

In a nutshell, ASR contracts work as follows. Upon signature of an ASR contract between a firm and an investment bank, the latter delivers shares to the former by borrowing them from shareholders (typically institutional investors). Subsequently the bank has a short position and needs to buy shares in the open market to return them back to the lenders. The contract typically involves an option component to determine either the price per share paid by the firm, the number of shares it receives, or both. This option component is usually of the Asian type with Bermudan exercise features, or even more complex in the case of profit-sharing programs (see Section 2 for more details).

In addition to higher credibility (see [8]), the motives of firms carrying out buyback through accelerated programs are numerous. An important segment of the academic literature deals with the financial reporting advantages and the immediate boost of earnings per share (EPS) provided by ASR contracts. For instance, [29] and [30] find evidence of EPS enhancement as a motive of ASR adoption,² but this finding has to be put in perspective because of other studies such as [1, 8, 12, 25] finding little evidence. The literature also discusses the signalling content of ASR over OMR programs, as the commitment associated with ASRs reinforces the classical undervaluation signal of share buyback programs (see [12, 25, 29]).

The economic literature on ASRs also deals with the short- and long-term effects of ASR announcement on the firm stock price. Many papers suggest indeed an immediate increase in the

¹Privately negotiated repurchases also existed and continue to exist.

²The literature discusses for instance the incentive of management to sign ASR contracts to boost EPS for increasing performance-based compensation, see [30].

stock price, although the amplitude of this effect is debated (see for instance [1, 6, 25, 31, 36]). Some also discuss price manipulations of firms willing to reduce the price of stocks before the announcement of ASR programs (see [13, 14]). Market microstructure changes around ASR announcements are also discussed in [28].

In spite of an extensive economic literature on ASR contracts, the pricing and management of complex buyback contracts has rarely been tackled. Pioneer works on the subject include that of Jaimungal *et al.* [24] and papers by Guéant *et al.* [19, 20]. They all show that ASR contracts should not be managed like traditional equity derivatives, i.e. managed with Greeks, because the execution problem at the heart of these contracts cannot be disentangled from the option component. The payoff of the option constitutes indeed, in most cases, a partial hedge for the execution process. Moreover, the volumes to be executed are often very large and execution costs must be taken into account. Furthermore, there are often participation constraints in buyback programs preventing to buy more than a given proportion of the daily volume, or even forbidding the use of stock selling.

In [24], the authors focus on ASR contracts with fixed number of shares and American exercise. They propose a continuous-time model where the stock price is modeled as a geometric Brownian motion with a drift reflecting permanent market impact, and add quadratic execution costs as in Almgren-Chriss models (see [3, 4]). The strategy they propose is optimal for a bank maximising its expected profit and loss (PnL) and penalising inventory (a penalty that can also be regarded as a form of ambiguity aversion as far as the stock price process is concerned). Jaimungal *et al.* derive the dynamic programming equation associated with the problem: a degenerate quasi-variational inequality in dimension 4 (time + 3 state variables). Interestingly, in an attempt to tame the degeneracy of the equation, they introduce the ratio between the stock price and its average value since inception and subsequently reduce the dimensionality of the problem. In addition to obtaining a new quasi-variational inequality – this time in dimension 3, including time –, they show that the exercise boundary only depends on the time to maturity and the above ratio.

The case of ASR contracts with fixed number of shares is also dealt with in the paper [20] by Guéant, Pu, and Royer who proposed a discrete-time model with a general execution cost function, and an expected utility objective function. As in [24], they show that the problem boils down to a set of equations with 3 variables; here time to maturity, the number of shares to be bought, and the difference between the current stock price and the average price since inception (and not the ratio because of the different assumptions regarding price dynamics). The case of ASR contracts with fixed notional is dealt with in [19] and it must be noted that there is no similar dimensionality reduction in that case. It is noteworthy that more complex VWAP-minus programs, such as profit-sharing programs, are not dealt with in the literature.

Because of its high-dimensional nature, it is natural to try solving the problem of pricing and managing ASRs and other (more complex) VWAP-minus programs with the help of neural networks instead of grids or trees as in the above literature. This paper proposes a machine learning approach involving recurrent neural networks to find the optimal execution strategy associated with different types of VWAP-minus programs: ASRs with fixed number of shares,

ASRs with fixed notional, and profit-sharing contracts.

In recent years, following the craze regarding neural networks, several research papers have encouraged the idea that neural network techniques could be a way to tackle financial issues suffering from the curse of dimensionality. In particular, several papers written by Jentzen and collaborators – see for instance [9, 21, 34] – proposed new methods, based on neural networks, to approximate the solutions of linear and nonlinear parabolic partial differential equations (PDE). In particular, [34] solves linear PDEs including that of Black and Scholes with correlated noises and that of Heston, and [21] solves the nonlinear equation associated with the Black-Scholes model when different interest rates are considered for borrowing and lending.³ A group of researchers around Pham (see [7, 23]) recently proposed other methods based on neural networks to solve optimal control problems with applications to energy issues, and proved results of convergence. In finance, papers on the hedging of options with (deep) neural network techniques include the famous “Deep Hedging” paper (see [11]) written by Buehler *et al.* that uses a semi-recurrent neural network. The case of American and Bermudan payoffs is also addressed in [10] with an interesting idea that we also use, though in a slightly different manner: the relaxation of the optimal stopping decision.

Our approach is innovative in that, in addition to looking for the best execution strategy using a recurrent neural network, we do not look directly for the optimal stopping time, but rather for the optimal probability to stop at each step, given the current state. This relaxation allows to go from a discrete decision problem to a continuous one, and therefore enables the use of gradient descent tools. In practice, we use a second neural network for modelling the probability to stop. Our approach recovers results similar to those of [19, 20] in the case of ASR contracts. Compared to the approaches based on the dynamic programming principle, our approach has a number of advantages: it does not require one to solve non-linear PDEs in a high-dimensional space, and thus allows to handle more sophisticated contracts – see our treatment of VWAP-minus profit-sharing contracts – and allows essentially any price dynamics unlike what happens with the grid or tree approaches developed in the literature.

In Section 2, we describe the three different types of buyback contracts addressed in the paper: two types of ASR contracts and one VWAP-minus profit-sharing contract. In Section 3, we propose a discrete-time model similar to that of [18, 19, 20] and define the objective functions. In Section 3, we also describe the architecture of our deep recurrent neural network to approximate the optimal strategy for managing the different contracts. In Section 4, we provide numerical results and discuss our findings. An appendix is dedicated to neural networks in order to provide the readers with ideas that are still seldom used in mathematical finance.

³Interestingly, these papers do not approximate directly the solution of the PDEs, but their (space-)gradient (related to the actions in the vocabulary of reinforcement learning). In other words, prices must be deduced from Greeks and not the other way round as with the classical tools of mathematical finance.

2 Buyback contracts

In this paper we consider three different types of buyback contract: the two types of ASR contract tackled in [18, 19, 20, 24], and one VWAP-minus profit-sharing contract never addressed in the academic literature. The termsheets of these contracts can be summarised as follows:

I. ASR contract with fixed number of shares:

- 1) At time $t = 0$, the bank borrows Q shares from the firm's shareholders (usually institutional investors) and delivers these shares to the firm in exchange for the current Mark-to-Market (MtM) value of these assets (QS_0).⁴
- 2) The bank has to progressively buy Q shares in the open market to give them back to the initial shareholders and return to a flat position on the stock.
- 3) The final settlement of the contract is associated either with the early exercise of an option or with the expiry of the contract (at time T). If the bank decides to early exercise the option at time $\tau \in \mathcal{T}$, where $\mathcal{T} \subset (0, T)$ is the set of possible early exercise dates specified in the contract, then the firm pays to the bank the difference between the average market price between 0 and τ (in this section, we denote by A_t the average price between 0 and t) and the price at inception S_0 . This can be regarded as the bank being long a Bermudan option with Asian payoff $Q(A_\tau - S_0)$. If the contract goes to expiry the final payoff is instead $Q(A_T - S_0)$.

II. ASR contract with fixed notional:

- 1) At time $t = 0$, the firm pays to the bank a fixed amount of cash F . In return, the bank delivers to the firm Q shares borrowed from the firm's shareholders, where $Q = \zeta \frac{F}{S_0}$ (ζ is usually around 80%).
- 2) The bank has to progressively buy back Q shares in the open market to give them back to the initial shareholders.
- 3) The final settlement of the contract is associated either with the early exercise of an option or with the expiry of the contract (at time T). If the bank decides to early exercise the option at time $\tau \in \mathcal{T}$, where $\mathcal{T} \subset (0, T)$ is the set of possible early exercise dates specified in the contract, then there is a transfer of $\frac{F}{A_\tau} - Q$ shares from the bank to the firm, so that the actual number of shares acquired by the firm is $\frac{F}{A_\tau}$. If the contract goes to expiry, then there is a transfer of $\frac{F}{A_T} - Q$ shares from the bank to the firm.

Remark 1. *In practice, for both types of ASR, there is often a discount proposed to the firm: the bank gives back part of the option value in the form of a discount on the average price – hence the expression VWAP-minus used for most of these programs. Considering this discount does not raise any difficulty when using our approach, unlike what would happen with classical methods.*⁵

⁴Here we consider the case of a pre-paid ASR contract. The case of a post-paid ASR contract is the same, if funding and interest rate are ignored.

⁵For instance, the dimensionality reduction obtained through a change of variables in [20] does not work anymore in presence of a multiplicative discount.

III. VWAP-minus profit-sharing contract:

- 1) At time $t = 0$, there is no initial transaction.
- 2) The bank has to buy shares in the open market on behalf of the client either until an amount of cash equal to F has been spent or until the expiry of the contract (at time T). For this type of contract, selling is prohibited.
- 3) If the contract expires before the required amount of cash is spent, the contract is settled by the payment of a penalty by the bank to the firm.⁶ Otherwise, once an amount of cash equal to F has been spent (we denote by time τ_0 the occurrence of that event), the bank becomes long a Bermudan option with expiry date T and payoff $\alpha(q(A_\tau - \kappa S_0) - F)_+$, where:

- q is the number of shares bought by the bank on behalf of the firm against the amount F ;
- $\tau \in \mathcal{T} \cap [\tau_0, T]$ designates a stopping time (as in all Bermudan/American options), where $\mathcal{T} \subset (0, T]$ is the set of possible exercise dates specified in the contract;
- α is the proportion of profit sharing (typically 25%);
- κ is a hurdle rate required by the firm (typically below 1%).

In other words, the bank is incentivised to carry out the execution at a better price than the average price minus a discount.

Remark 2. *These contracts are common in the brokerage and corporate derivatives industry, but it is not clear that they really give the bank an incentive to carry out a good execution in all situations. If, indeed, the beginning of the execution process is poor, and if the bank subsequently realises that the option will be worth almost nothing, then it has no reason to provide the best possible execution to the client. For this reason banks, in order to give the best service to the client, should manage the option as if the payoff was $\alpha(q(A_\tau - \kappa S_0) - F)$ or $\alpha(q(A_\tau - \kappa S_0) - F)_+ - \beta(q(A_\tau - \kappa S_0) - F)_-$ (where $\beta \in [0, \alpha)$) instead of $\alpha(q(A_\tau - \kappa S_0) - F)_+$.*

3 The model

3.1 Mathematical setting

3.1.1 Dynamics of the state variables

We consider a discrete-time model where each period of time of length δt corresponds to one day. In other words, given a contract with maturity date T corresponding to N days ($T = N\delta t$), we consider the subdivision $(t_n = n\delta t)_{0 \leq n \leq N}$ of the interval $[0, T]$. We denote

⁶This should never happen as T is chosen to ensure the possibility of the delivery.

by $\mathcal{N} = \{n \in \{0, \dots, N\} | t_n \in \mathcal{T}\}$ the set of indices corresponding to the possible (early) exercise dates.

We consider a probability space (Ω, \mathbb{P}) and a listed firm whose stock price is modelled by a stochastic process $(S_n)_n$. We denote by $(\mathcal{F}_n)_n$ the completed filtration generated by $(S_n)_n$ (i.e. we assume that \mathcal{F}_0 contains all the \mathbb{P} -null sets).

Remark 3. *It is noteworthy that we do not set a particular model for the price dynamics.*

For $n \in \{1, \dots, N\}$, the running average price of the stock over $\{t_1, \dots, t_n\}$ is denoted by

$$A_n = \frac{1}{n} \sum_{k=1}^n S_k.$$

Let us consider a bank in charge of buying shares of that firm. We assume that the bank executes an order each day, and we denote by $(v_n \delta t)_n$ the daily volumes of transactions: $v_0 \delta t$ for the first day, $v_1 \delta t$ for the second day, etc. Subsequently, the number of shares $(q_n)_n$ bought by the bank in the market is given by

$$\begin{cases} q_0 = 0 \\ q_{n+1} = q_n + v_n \delta t. \end{cases}$$

For each share bought over the n -th day the bank pays $S_n + g\left(\frac{v_n}{V_{n+1}}\right)$, where g is a nonnegative function modelling execution costs per share and $(V_n)_n$ is the market volume process, assumed to be deterministic. In other words, the trader pays the reference price for the n -th day plus execution costs depending on her participation rate to the market over the n -th day.

Following [18], we consider the function $L : \rho \in \mathbb{R} \mapsto \rho g(\rho)$ and assume that g is such that L verifies the following assumptions:

- L is strictly convex on \mathbb{R} , increasing on \mathbb{R}_+ , and decreasing on \mathbb{R}_- ;
- L is asymptotically superlinear, i.e.:

$$\lim_{\rho \rightarrow +\infty} \frac{L(\rho)}{\rho} = +\infty.$$

The resulting cumulative cash spent by the bank modelled by $(X_n)_n$ has the following dynamics:

$$\begin{cases} X_0 = 0 \\ X_{n+1} = X_n + v_n S_{n+1} \delta t + g\left(\frac{v_n}{V_{n+1}}\right) v_n \delta t = X_n + v_n S_{n+1} \delta t + L\left(\frac{v_n}{V_{n+1}}\right) V_{n+1} \delta t, \end{cases}$$

In the following, we first compute the profit and loss associated with each type of contract. Then, we introduce the set of admissible controls and propose an objective function that could be used by the bank to carry out optimisation.

3.1.2 Profit and Loss

I. ASR contract with fixed number of shares:

No matter if the bank chooses to early exercise on day $n \in \mathcal{N}$ or if the contract expires on day $n = N$, the bank has to acquire $Q - q_n$ shares. We assume that these remaining shares could be purchased at price S_n plus execution costs. The resulting amount of cash spent by the bank at time n is $(Q - q_n)S_n + \ell(Q - q_n)$, where $\ell : \mathbb{R} \mapsto \mathbb{R}_+$ satisfies the same properties as the execution cost function L .

At exercise date or at expiry (day n) the bank receives from the firm an amount of cash equal to QA_n . The resulting profit and loss of the bank is

$$\text{PnL}_n^Q = QA_n - X_n - (Q - q_n)S_n - \ell(Q - q_n).$$

II. ASR contract with fixed notional:

No matter if the bank chooses to early exercise on day $n \in \mathcal{N}$ or if the contract expires on day $n = N$, the bank has to acquire $\frac{F}{A_n} - q_n$ shares. We assume that these remaining shares could be purchased at price S_n plus execution costs. The resulting amount of cash spent by the bank at time n is $\left(\frac{F}{A_n} - q_n\right)S_n + \ell\left(\frac{F}{A_n} - q_n\right)$, where $\ell : \mathbb{R} \mapsto \mathbb{R}_+$ is as above.

At exercise date or at expiry (day n) the bank receives from the firm an amount of cash equal to F . The resulting profit and loss of the bank is

$$\text{PnL}_n^F = F - X_n - \left(\frac{F}{A_n} - q_n\right)S_n - \ell\left(\frac{F}{A_n} - q_n\right).$$

III. VWAP-minus profit-sharing contract:

If the bank manages to spend the amount F before expiry, then its profit and loss is

$$\text{PnL}_n^S = F - X_n + \alpha(q_n(A_n - \kappa S_0) - F)_+,$$

where n corresponds to the date of exercise of the option.⁷ Otherwise, we assume that the profit and loss at expiry date is just a penalty.

In our approach, we consider (i) that the option can be exercised even if the amount F has not been spent and (ii) that once an amount of cash F has been spent the bank stops trading. Moreover, we consider the modification of the profit and loss discussed in Remark 2. This results in the following modified profit and loss formula:

$$\text{PnL}_n^S = -\ell(F - X_n) + \alpha(q_n(A_n - \kappa S_0) - F)_+ - \beta(q_n(A_n - \kappa S_0) - F)_-,$$

where $\ell : \mathbb{R} \mapsto \mathbb{R}_+$ is as above.

If the bank exercises the option before the amount F has been spent, then the penalty associated with ℓ is paid and we assume that it is large enough to compensate the profit sharing term (should it be positive) if X_n is far below F . Otherwise, the payoff is just the same as above, except when it comes to the additional β term.

⁷In practice, X_n should be equal to F .

3.1.3 Objective function

Before introducing the objective function let us first define the set of admissible controls. We consider minimal and maximal market participation rates. In other words, we impose the market participation constraints $\underline{\rho}V_{n+1} \leq v_n \leq \bar{\rho}V_{n+1}$, where $\bar{\rho}$ is positive and $\underline{\rho}$ can be of either sign.⁸

Therefore the set of admissible strategies of the bank can be represented as follows:

$$\mathcal{A} = \left\{ (v, n^*) \mid v = (v_n)_{0 \leq n \leq n^*-1} \text{ is } \mathcal{F}\text{-adapted, } \underline{\rho}V_{n+1} \leq v_n \leq \bar{\rho}V_{n+1}, 0 \leq n \leq n^* - 1, \right. \\ \left. \text{and } n^* \text{ is a } \mathcal{F}\text{-stopping time taking values in } \mathcal{N} \cup \{N\} \right\}.$$

To be consistent with [19, 20], we consider that the bank is willing to maximise the expected CARA utility of its PnL. Therefore, the optimisation problem faced by the bank is the following:

$$\sup_{(v, n^*) \in \mathcal{A}} \mathbb{E}[-\exp(-\gamma \text{PnL}_{n^*})]$$

where γ is the risk aversion parameter of the bank and PnL is either PnL^Q , PnL^F or PnL^S .

Remark 4. We assume that the dynamics of the stock is chosen so that the above problem has a solution, i.e.

$$\sup_{(v, n^*) \in \mathcal{A}} \mathbb{E}[-\exp(-\gamma \text{PnL}_{n^*})] \neq -\infty$$

3.2 Relaxation and mean-variance approximation: towards a machine learning approach

3.2.1 Relaxation of the optimal stopping problem

Given the structure of the problem, the optimal number of shares to be bought on day $n + 1$ can be written as a closed-loop control $v(n, S_n, A_n, X_n, q_n)$. Similarly, the optimal decision to exercise the option can be written as: $\mathbf{1}_{\{n^*=n\}} = p(n, S_n, A_n, X_n, q_n)$.

Since the function p takes values in $\{0, 1\}$, this problem is not suitable for the optimisation methods commonly associated with neural networks, *e.g.* stochastic gradient descent. In this regard, we extend the set of admissible controls to allow stochastic stopping decisions.

More precisely, an admissible strategy is determined by:

- the number of shares to be bought on each day, modelled (up to the δt multiplicative term) by a \mathcal{F} -adapted process $(v_n)_n$;
- the stochastic stopping policy $(p_n)_n$, which is a \mathcal{F} -adapted process that takes values in the interval $[0, 1]$ with $p_n = \mathbf{1}_{n=N}$ if $n \notin \mathcal{N}$.

⁸Constraints of this type are sometimes specified explicitly in the contract.

In order to sample effective stopping decisions based on the stochastic stopping policy $(p_n)_n$, we introduce an extended σ -algebra $\mathcal{G} \supset \mathcal{F}_N$ and i.i.d random variables $(\tilde{\epsilon}_n)_n$ defined on $(\Omega, \mathbb{P}, \mathcal{G})$, uniform on $[0, 1]$, and assumed to be independent of \mathcal{F}_N .

The effective stopping time n^* is then defined as $\min\{n \in \mathcal{N} \cup \{N\} | \tilde{\epsilon}_n \leq p_n\}$, so that the stopping decision \hat{p}_n defined by $\hat{p}_n = \mathbf{1}_{\tilde{\epsilon}_n < p_n}$ is conditionally distributed as a Bernoulli with parameter p_n given \mathcal{F}_n .

Therefore the PnL of the strategy is given by:

$$\text{PnL} = \sum_{n=1}^N \prod_{k=1}^{n-1} (1 - \hat{p}_k) \hat{p}_n \text{PnL}_n.$$

We search for the optimal strategy v in the form of $v_\theta(n, S_n, A_n, X_n, q_n)$ for $n \in \{0, \dots, N-1\}$, and p in the form of $p_n = p_\phi(n, S_n, A_n, X_n, q_n)$ for $n \in \mathcal{N}$, both of them lying in a finite-dimensional set of functions parameterised by θ and ϕ respectively.

Remark 5. *For the fixed number of shares and fixed notional ASR contracts, the optimal strategy does not depend on the cash variable when using a CARA utility framework (see [19, 20]). Therefore, the cash variable is absent of v_θ and p_ϕ in these cases.*

In our relaxed setting, the objective function has then a differentiable dependency on the parameters of the neural networks.⁹

3.2.2 Neural networks

For the neural networks to be robust with respect to scaling effects, we ensure that the variables that are the inputs of the neural networks are dimensionless and centered. Using the finding of [19, 20], we give as an input $\frac{A-S}{S_0}$ instead of A , as the strategy has a strong dependency on the spread between the spot price and the running average. Likewise, the outputs of the networks are designed as perturbations of naive strategies (see below for details).

I. ASR contract with fixed number of shares:

We parameterise the rate of share repurchase v_θ by:

$$v_\theta(n, S, A, X, q) = Q \cdot \min \left(\left(1 + \tilde{v}_\theta \left(\frac{n}{N} - \frac{1}{2}, \frac{S}{S_0} - 1, \frac{A-S}{S_0}, \frac{q}{Q} - \frac{1}{2} \right) \right) \cdot \frac{n+1}{N}, 1 \right) - q,$$

where \tilde{v}_θ is a neural network consisting of 4 inputs, a hidden layer of 50 neurons with ReLU activation function and 1 output.

It is noteworthy that if \tilde{v}_θ is equal to 0, then the portfolio to reach at step $n+1$ is $\frac{n+1}{N}Q$, which corresponds to the trading schedule of a trader buying the same amount of shares each day until maturity.

The stochastic stopping policy p_ϕ is represented by:

$$p_\phi(n, S, A, X, q) = \mathbf{1}_{n \in \mathcal{N}} \cdot \mathcal{S} \left(\nu_\phi \cdot \left(\frac{q}{Q} - \tilde{p}_\phi \left(\frac{n}{N} - \frac{1}{2}, \frac{S}{S_0} - 1, \frac{A-S}{S_0} \right) \right) \right) + \mathbf{1}_{n=N},$$

⁹A similar idea is used in [10] to handle American options with neural networks.

where \tilde{p}_ϕ is a neural network consisting of 3 inputs, a hidden layer of 50 neurons with ReLU activation function and 1 output, ν_ϕ is a scaling parameter, and \mathcal{S} is the activation function defined by:

$$\mathcal{S} : x \mapsto \min \left(\max \left(\frac{2}{1 + e^{-x}} - \frac{1}{2}, 0 \right), 1 \right).$$

We use the activation function \mathcal{S} that is a modified version of the logistic function (rescaled and bounded to $[0, 1]$) to allow the values 0 and 1 to be reached.

The output of the network p_ϕ can be interpreted as the frontier in terms of the ratio $\frac{q}{Q}$, depending on n , S and A , above which we exercise the option.

II. ASR contract with fixed notional:

We parameterise the rate of share repurchase v_θ by:

$$v_\theta(n, S, A, X, q) = \frac{F}{A} \cdot \frac{n+1}{N} \left(1 + \tilde{v}_\theta \left(\frac{n}{N} - \frac{1}{2}, \frac{S}{S_0} - 1, \frac{A-S}{S_0}, \frac{qA}{F} - \frac{1}{2} \right) \right) - q,$$

where \tilde{v}_θ is a neural network consisting of 4 inputs, a hidden layer of 50 neurons with ReLU activation function and 1 output.

It is noteworthy that if \tilde{v}_θ is equal to 0, then the portfolio to reach at step $n+1$ is $\frac{n+1}{N} \frac{F}{A}$, which corresponds to a natural naive trading schedule.

The stochastic stopping policy p_ϕ is represented by:

$$p_\phi(n, S, A, X, q) = \mathbf{1}_{n \in \mathcal{N}} \cdot \mathcal{S} \left(\nu_\phi \cdot \left(\frac{qA}{F} - \tilde{p}_\phi \left(\frac{n}{N} - \frac{1}{2}, \frac{S}{S_0} - 1, \frac{A-S}{S_0} \right) \right) \right) + \mathbf{1}_{n=N},$$

where \tilde{p}_ϕ is a neural network consisting of 3 inputs, a hidden layer of 50 neurons with ReLU activation function and 1 output, and ν_ϕ is a scaling parameter.

The output of the network p_ϕ here can be interpreted as the frontier in terms of the ratio $\frac{qA}{F}$, depending on n , S and A , above which we exercise the option.

III. VWAP-minus profit-sharing contract:

We parameterise the rate of share repurchase v_θ by:

$$v_\theta(n, S, A, X, q) = \mathbf{1}_{X < F} \cdot \frac{F - X}{S} \max \left(\min \left(\frac{1}{N - n} (1 + \tilde{v}_\theta(\dots)), 1 \right), 0 \right),$$

where $\tilde{v}_\theta(\dots)$ stands for $\tilde{v}_\theta \left(\frac{n}{N} - \frac{1}{2}, \frac{S}{S_0} - 1, \frac{A-S}{S_0}, \frac{X}{F} - \frac{1}{2}, \frac{qS_0}{F} - \frac{1}{2} \right)$, and \tilde{v}_θ is a neural network consisting of 5 inputs, a hidden layer of 50 neurons with ReLU activation function and 1 output.

It is noteworthy that if \tilde{v}_θ is equal to 0, then the cash spent is equal (up to the execution costs) to the ratio of the remaining cash to spend to the remaining number of days until maturity, which is a natural naive strategy. The max and min functions prevent us from selling and over-buying.

The stochastic stopping policy p_ϕ is represented by:

$$p_\phi(n, S, A, X, q) = \mathbf{1}_{n \in \mathcal{N}} \cdot \mathcal{S} \left(\nu_\phi \cdot \left(\frac{X}{F} - \tilde{p}_\phi \left(\frac{n}{N} - \frac{1}{2}, \frac{S}{S_0} - 1, \frac{A-S}{S_0}, \frac{qS_0}{F} - \frac{1}{2} \right) \right) \right) + \mathbf{1}_{n=N},$$

where \tilde{p}_ϕ is a neural network consisting of 4 inputs, a hidden layer of 50 neurons with ReLU activation function and 1 output, and ν_ϕ is a scaling parameter.

The output of the network p_ϕ here can be interpreted as the frontier in terms of the ratio $\frac{X}{F}$, depending on n , S and A , above which we exercise the option.

3.2.3 Objective function approximation

We could use a stochastic gradient descent or a mini-batched gradient descent on the expected CARA utility objective function to approximate an optimal trading and an optimal stopping strategy. However the very fact that the utility is exponential typically causes numerical issues. For that reason, we consider the classical Arrow-Pratt (see for instance [33]) approximation of the expected CARA utility objective function by a mean-variance objective function:¹⁰

$$-\frac{1}{\gamma} \log \mathbb{E} [\exp(-\gamma \text{PnL})] \approx \mathbb{E} [\text{PnL}] - \frac{\gamma}{2} \mathbb{V} [\text{PnL}].$$

In our relaxed setting, we have

$$\begin{aligned} \mathbb{E} [\text{PnL}] &= \mathbb{E} [\mathbb{E} [\text{PnL} | \mathcal{F}_N]] \\ &= \mathbb{E} \left[\mathbb{E} \left[\sum_{n=1}^N \prod_{k=1}^{n-1} (1 - \hat{p}_k) \hat{p}_n \text{PnL}_n \middle| \mathcal{F}_N \right] \right] \\ &= \mathbb{E} \left[\sum_{n=1}^N \mathbb{E} \left[\prod_{k=1}^{n-1} (1 - \hat{p}_k) \hat{p}_n \middle| \mathcal{F}_N \right] \text{PnL}_n \right] \\ &= \mathbb{E} \left[\sum_{n=1}^N \prod_{k=1}^{n-1} (1 - p_k) p_n \text{PnL}_n \right], \end{aligned}$$

since

$$\mathbb{E} \left[\prod_{k=1}^{n-1} (1 - \hat{p}_k) \hat{p}_n \middle| \mathcal{F}_N \right] = \mathbb{E} \left[\prod_{k=1}^{n-1} (1 - \mathbf{1}_{\tilde{\epsilon}_k < p_k}) \mathbf{1}_{\tilde{\epsilon}_n < p_n} \middle| \mathcal{F}_N \right] = \prod_{k=1}^{n-1} (1 - p_k) p_n,$$

where we used the fact that (p_1, \dots, p_n) is \mathcal{F}_N -measurable, and that $\tilde{\epsilon}_1, \dots, \tilde{\epsilon}_n$ are i.i.d. and independent of \mathcal{F}_N .

Similarly,

$$\mathbb{E} [\text{PnL}^2] = \mathbb{E} \left[\sum_{n=1}^N \prod_{k=1}^{n-1} (1 - p_k) p_n \text{PnL}_n^2 \right].$$

¹⁰It is important to note that we could have chosen a mean-variance objective function from the very beginning. The reason why we started with an exponential utility is to relate our paper to [19, 20].

It is also important to recall that the mean-variance approximation of the certainty equivalent associated with a CARA utility function (i) turns out to be exact in the case of Gaussian risks and (ii) corresponds to a first order Taylor expansion in the risk aversion parameter (around 0).

Subsequently,

$$\begin{aligned}
-\frac{1}{\gamma} \log \mathbb{E} [\exp(-\gamma \text{PnL})] &\approx \mathbb{E} \left[\sum_{n=1}^N \prod_{k=1}^{n-1} (1 - p_k) p_n \text{PnL}_n \right] \\
&\quad - \frac{\gamma}{2} \left[\mathbb{E} \left[\sum_{n=1}^N \prod_{k=1}^{n-1} (1 - p_k) p_n (\text{PnL}_n)^2 \right] \right. \\
&\quad \left. - \left(\mathbb{E} \left[\sum_{n=1}^N \prod_{k=1}^{n-1} (1 - p_k) p_n \text{PnL}_n \right] \right)^2 \right].
\end{aligned}$$

Therefore, using a Monte-Carlo approximation with I trajectories of prices $(S_n^i)_{0 \leq n \leq N, 1 \leq i \leq I}$, and the resulting stopping policy $(p_n^i)_{1 \leq n \leq N, 1 \leq i \leq I}$ and profit and losses $(\text{PnL}_n^i)_{1 \leq n \leq N, 1 \leq i \leq I}$, we can consider the following approximation

$$\begin{aligned}
-\frac{1}{\gamma} \log \mathbb{E} [\exp(-\gamma \text{PnL})] &\approx \frac{1}{I} \sum_{i=1}^I \sum_{n=1}^N \prod_{k=1}^{n-1} (1 - p_k^i) p_n^i \text{PnL}_n^i \\
&\quad - \frac{\gamma}{2} \left[\frac{1}{I} \sum_{i=1}^I \sum_{n=1}^N \prod_{k=1}^{n-1} (1 - p_k^i) p_n^i (\text{PnL}_n^i)^2 \right. \\
&\quad \left. - \left(\frac{1}{I} \sum_{i=1}^I \sum_{n=1}^N \prod_{k=1}^{n-1} (1 - p_k^i) p_n^i \text{PnL}_n^i \right)^2 \right].
\end{aligned}$$

Given the sampled trajectories $(S_n^i)_{0 \leq n \leq N, 1 \leq i \leq I}$, the right-hand side of the above equation depends only on θ and ϕ . Therefore using automatic differentiation tools we can perform gradient descent on this proxy of the objective function.

4 Numerical results

In this section we illustrate the practical use of our method. We consider the reference case described below which corresponds to rounded values for the stock Total SA, deliberately chosen to be the same as in [20] in order to show that the strategies obtained with our method are similar to those derived in [20] for an ASR contract with fixed number of shares.

For the same comparison purpose, we train the neural networks with arithmetic Brownian motion price trajectories $S_{n+1} = S_n + \sigma \sqrt{\delta t} \epsilon_{n+1}$, where $(\epsilon_n)_n$ are i.i.d. $\mathcal{N}(0, 1)$ random variables. Contrary to what happens with the method presented in [20], our method can be used with almost any price dynamics or even historical data.¹¹

More precisely, we consider the following market model:

- $S_0 = 45$ €;
- $\sigma = 0.6$ €·day^{-1/2}, corresponding to an annual volatility approximately equal to 21%;

¹¹Unfortunately, historical time series are often not long enough. A new practice consists in using data simulated with generative models calibrated to historical data (e.g. generative adversarial networks – see [17]).

- $T = 63$ trading days. The set of possible early exercise dates is $\mathcal{N} = [22, 62] \cap \mathbb{N}$;
- $\forall n \in \{1, \dots, N\}, V_n = V = 4\,000\,000 \text{ shares} \cdot \text{day}^{-1}$;
- $L(\rho) = \eta |\rho|^{1+\phi}$ with $\eta = 0.1 \text{ €} \cdot \text{share}^{-1} \cdot \text{day}^{-1}$ and $\phi = 0.75$.

4.1 ASR contract with fixed number of shares

For this contract, we consider the following characteristics:

- $Q = 20\,000\,000$ shares;
- $\ell : q \mapsto Cq^2$ for the terminal penalty, where $C = 2 \cdot 10^{-7} \text{ €} \cdot \text{share}^{-2}$;
- $\underline{\rho} = -\infty, \bar{\rho} = +\infty$, meaning that there is no participation constraints.

Our choice for the risk aversion parameter is $\gamma = 2.5 \cdot 10^{-7} \text{ €}^{-1}$.

Let us consider three different trajectories for the price in order to exhibit several features of the optimal strategy of the bank.

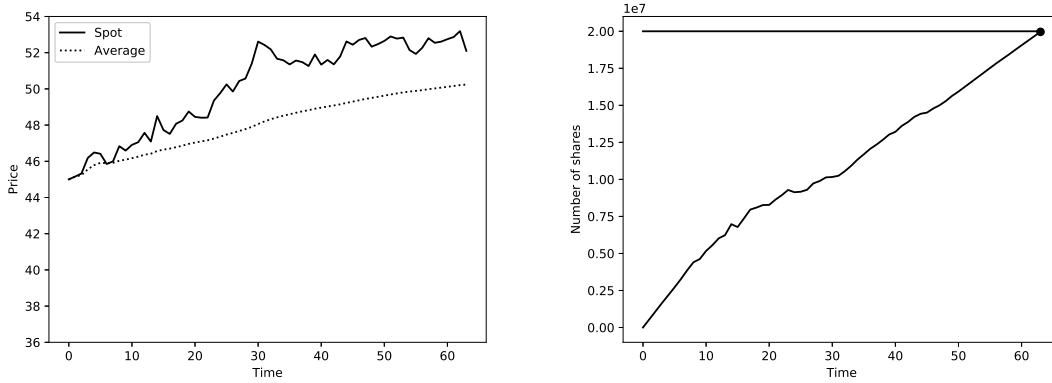


Figure 1: Price trajectory 1 and corresponding strategy for the ASR with fixed number of shares

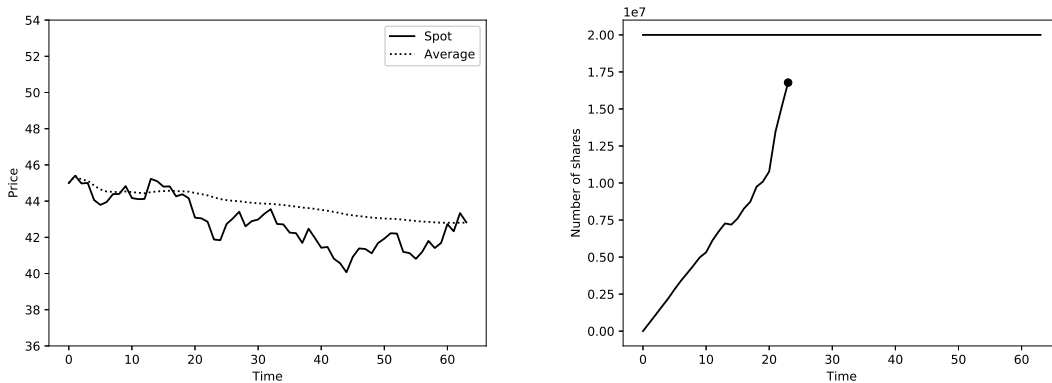


Figure 2: Price trajectory 2 and corresponding strategy for the ASR with fixed number of shares

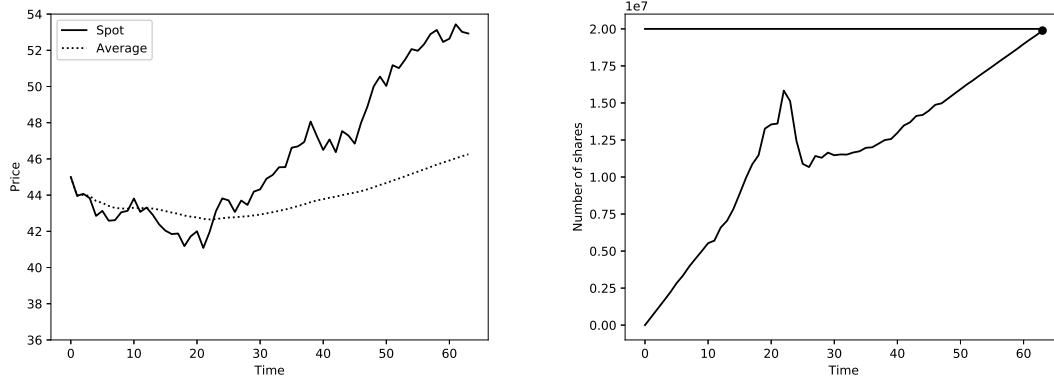


Figure 3: Price trajectory 3 and corresponding strategy for the ASR with fixed number of shares

The first price trajectory exhibits an upward trend. In this case, the optimal strategy of the bank consists in buying the shares slowly to minimise execution costs, as illustrated in Figure 1.

The second price trajectory exhibits a downward trend. In that case, the bank has an incentive to exercise rapidly, even without all the required shares being bought (see Figure 2). Indeed, as the price stays below its average, the latter is pulled down over time, making it less profitable to postpone the exercise of the option.

The third price trajectory we consider corresponds to the price decreasing and then increasing. As in both preceding examples, we see in Figure 3 that the behaviour of the bank is strongly linked to the relative position of the price to its running average. At the beginning of the contract, when the price is below its running average, the bank is acquiring shares at a high pace. Afterwards, when the price goes above its running average, it is not profitable anymore for the bank to accelerate execution. Instead, the bank is incentivised to delay the exercise of the option and it is even selling shares in order to stay close to the strategy $q_n = \frac{n}{N}Q$ because the risk associated with that strategy is hedged by the payoff of the ASR contract.

Now, as in [20], we study the effects of the parameters on the optimal strategy. More precisely, we focus on the execution cost parameter η and the risk aversion parameter γ .

Let us focus first on execution costs and more precisely on the liquidity parameter η . We consider our reference case with 4 values for the parameter η : 0.01, 0.1, 0.2, and 0.5. As we can see in Figure 4, corresponding to the third price trajectory, the less liquid the stock, the smoother the optimal strategy to avoid abrupt trading and round trips.

The optimal values of the mean-variance criterion for the different values of the parameter η are presented in the table below:

η	0.01	0.1	0.2	0.5
$\frac{\text{MeanVar}}{QS_0}$	1.13%	1.05%	0.99%	0.81%

As expected, the more liquid the stock, the more profitable the contract for the bank.

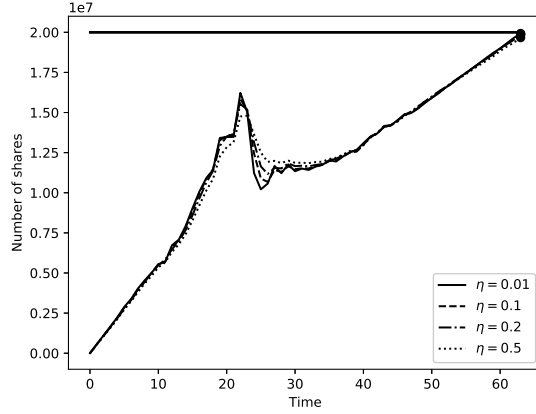


Figure 4: Effect of execution costs

Let us come now to risk aversion. We consider our reference case with 4 values for the parameter γ : 0, $2.5 \cdot 10^{-9}$, $2.5 \cdot 10^{-7}$, and $5.0 \cdot 10^{-7}$. Figure 5 shows the influence of γ on the optimal strategy. We see that the more risk averse the bank, the closer to the naive strategy (*i.e.* $q_n = \frac{n}{N}Q$) its strategy. This is intuitive as the risk associated with this strategy is perfectly hedged by the payoff of the ASR contract. At the other end of the spectrum, when $\gamma = 0$, the corresponding strategy is much more aggressive at the beginning of the contract in order to be able to benefit from the optionality as soon as possible (because the function ℓ we have chosen provides a very strong incentive to have only a few shares to buy at the time of early exercise). What prevents the bank from buying instantaneously is just execution costs. An interesting point is also that, when $\gamma = 0$, the optimal strategy does not involve any stock selling.¹²

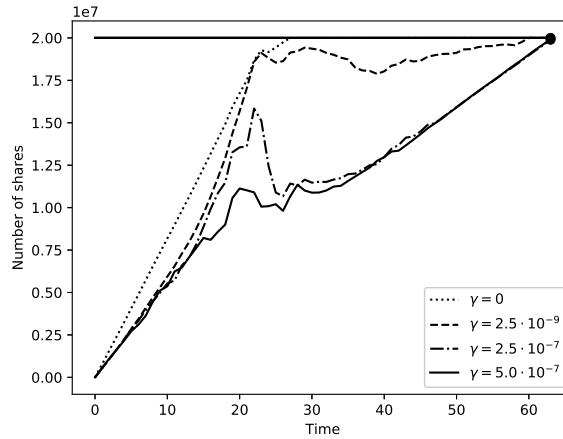


Figure 5: Effect of risk aversion

¹²Since neural networks are just approximations, sometimes we can see some small deviations from the optimal buy-only strategy in the no risk aversion case.

The mean-variance values for the different values of the risk aversion parameter γ are presented in the table below:

γ	0	$2.5 \cdot 10^{-9}$	$2.5 \cdot 10^{-7}$	$5 \cdot 10^{-7}$
$\frac{\text{MeanVar}}{QS_0}$	1.35%	1.32%	1.05%	0.86%

Unsurprisingly, the more risk averse the bank, the lower the optimal value of the mean-variance criterion.

As our optimisation problem is not convex, the optimisation procedure might lead to a local optimum. Because of the random initialisation of the neural networks weights and because of Monte Carlo sampling, the learning process is not always the same. Figure 6 illustrates two very different learning curves associated with two different instances of the learning procedure with $\gamma = 5 \cdot 10^{-7} \text{ €}^{-1}$. We see that the optimisation process for the second instance stalls in a suboptimal state (with a mean-variance score slightly below 0), whereas the first instance manages to reach a state with a significantly higher score.

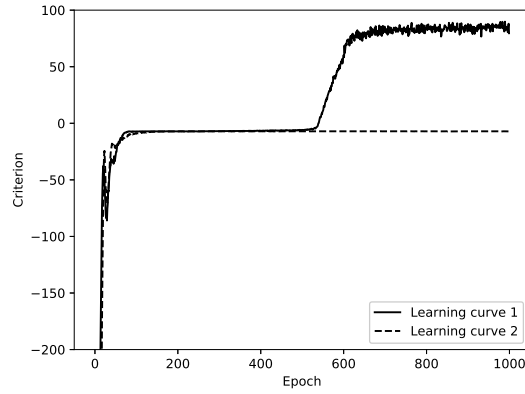


Figure 6: Training curve (the unit of the y-axis is $\frac{\text{MeanVar}}{QS_0}$ expressed in basis points)

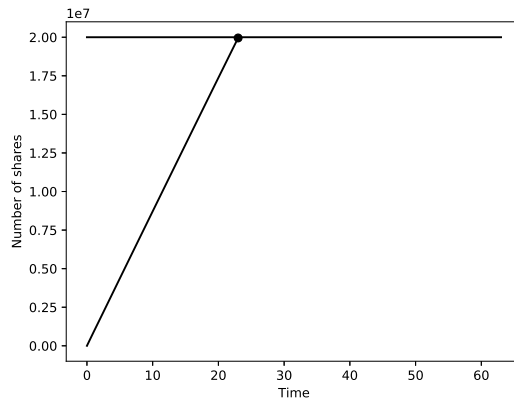


Figure 7: Locally optimal strategy

Interestingly, the suboptimal strategy associated with the second learning instance consists in buying shares at a constant pace until having the required quantity Q of shares and exercising immediately the option, regardless of the price trajectory (see Figure 7). It is not surprising that this strategy could be a local optimum as the option payoff provides a perfect hedge for the execution process.

In order to deter the learner from being caught in the domain of attraction of the type of local optimum described above, we can modify the objective function by setting γ to 0 over the first training epochs in order to remove the incentive to hedge. We refer to this procedure as pretraining.

We illustrate in Figure 8 the learning curve associated with the learning procedure where we performed pretraining over the first 100 epochs, and we compare it to the two above examples without pretraining. From this graph, we see that pretraining the network helps to avoid this type of local optimum. Moreover, when pretraining is used, we see in Figure 8 that the learning curve does not exhibit an intermediate plateau.

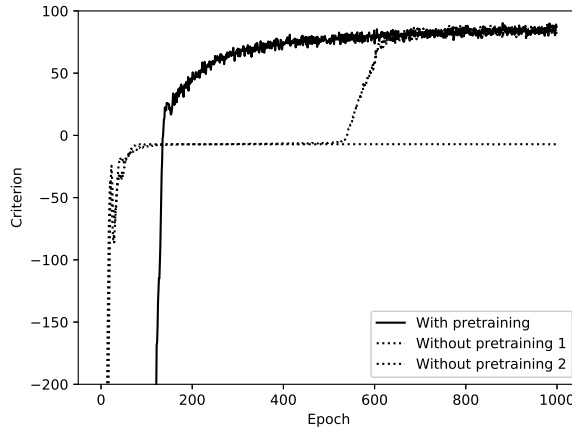


Figure 8: Comparison between learning curves with and without pretraining.

Remark 6. *It should be mentioned that our pretraining procedure enables to avoid a trivial local optimum, but does not theoretically ensure the convergence towards a global optimum (this is a common issue in machine learning, especially with neural networks). However, the strategies we obtain are in line with the results of our previous works [19, 20] based on solving Bellman equations. We believe therefore that our method succeeds in reaching a global optimum.*

4.2 ASR contract with fixed notional

For this contract, we consider the following characteristics:

- $F = 900\,000\,000$ €;
- $\ell : q \mapsto Cq^2$ as terminal penalty, where $C = 2 \cdot 10^{-7}$ € · share⁻²;
- $\underline{\rho} = -\infty, \bar{\rho} = +\infty$, meaning that there is no participation constraints.

We choose the risk aversion parameter $\gamma = 2.5 \cdot 10^{-7} \text{ €}^{-1}$.

In Figures 9, 10 and 11, we plot the strategies obtained for the fixed notional ASR contract (for the same three price trajectories as above). The targeted number of shares is represented by a solid line (it is not constant due to the stock price change).

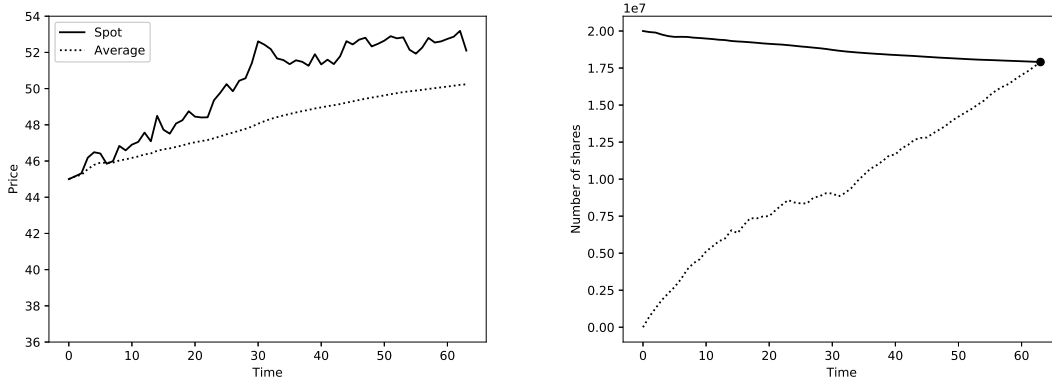


Figure 9: Price trajectory 1 and corresponding strategy for the ASR with fixed notional

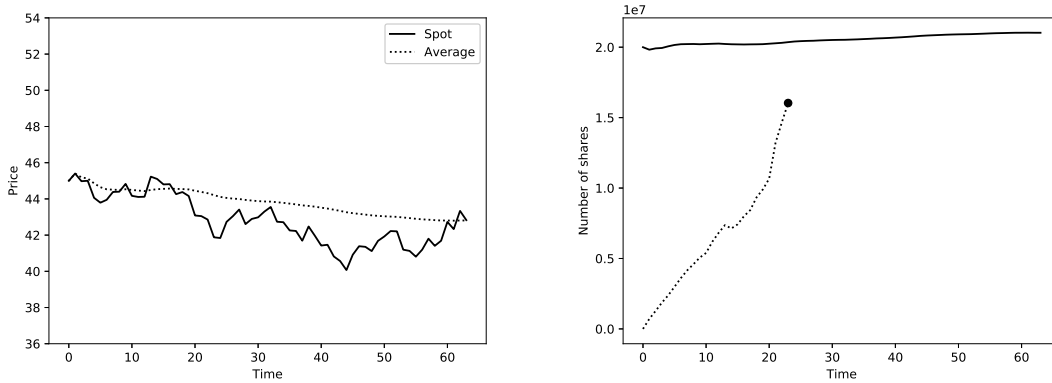


Figure 10: Price trajectory 2 and corresponding strategy for the ASR with fixed notional

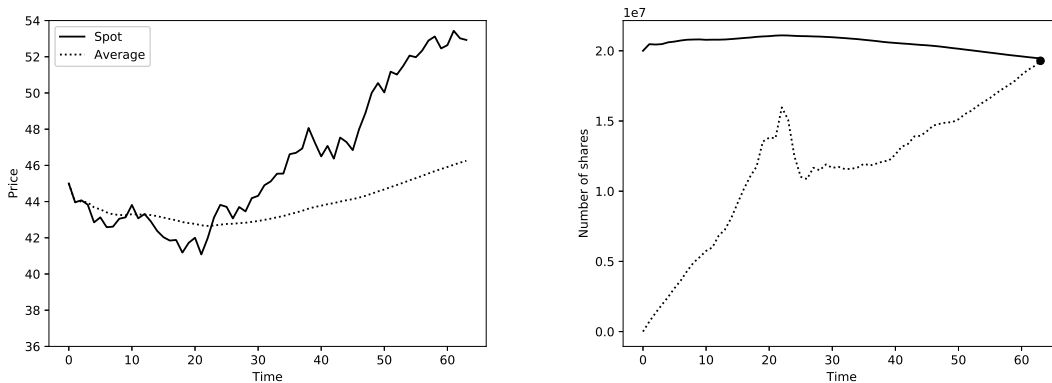


Figure 11: Price trajectory 3 and corresponding strategy for the ASR with fixed notional

Now, let us compare the strategies associated with the two types of ASR contract. In Figures 12, 13 and 14, we see that we obtain similar strategies for both types of ASR: accelerating purchase when the difference between the average price and the price is positive and decelerating purchase or even selling when that difference is negative.

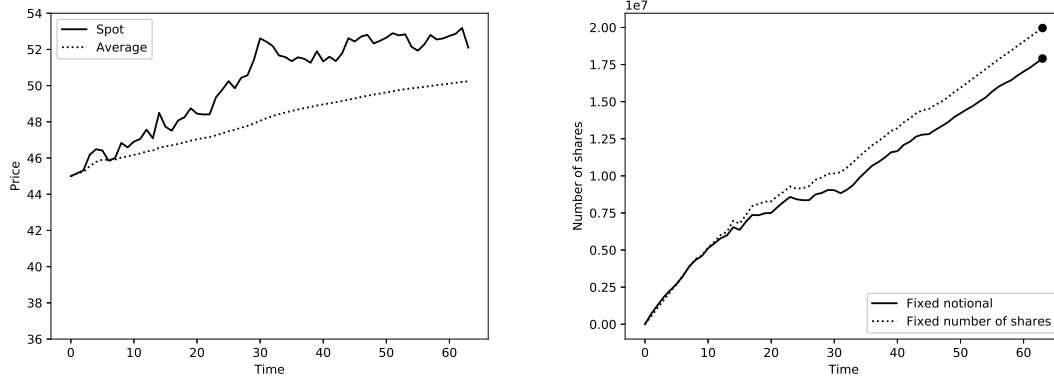


Figure 12: Price trajectory 1 and comparison of the strategies for the two types of ASR

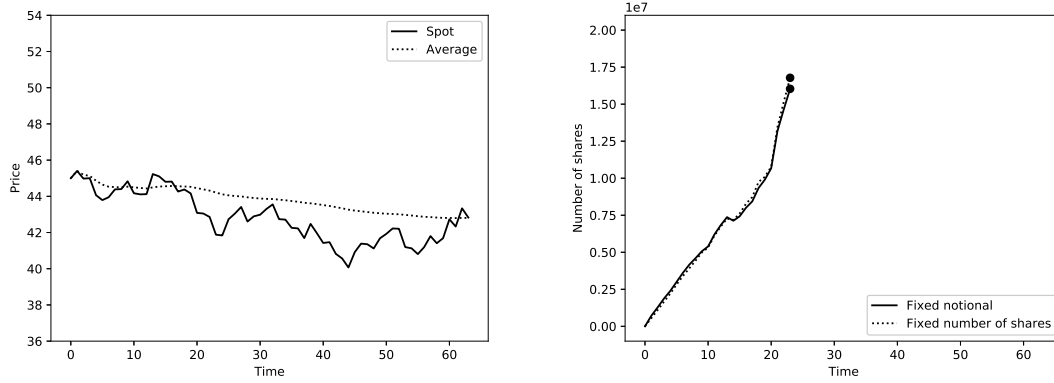


Figure 13: Price trajectory 2 and comparison of the strategies for the two types of ASR

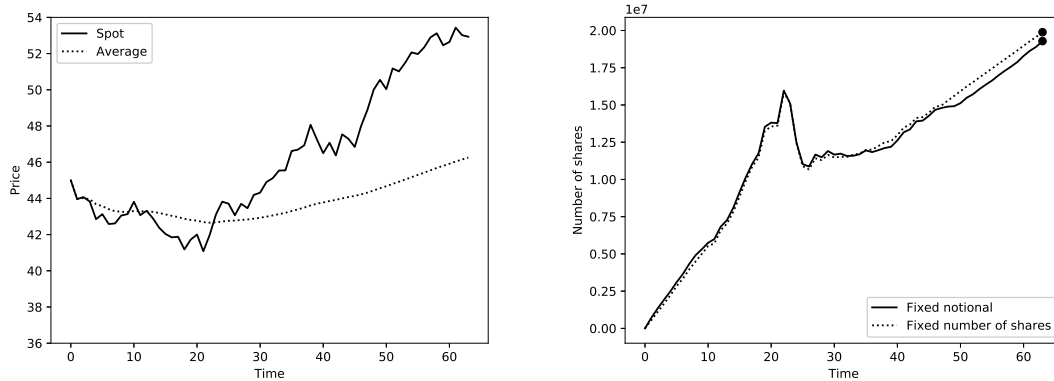


Figure 14: Price trajectory 3 and comparison of the strategies for the two types of ASR

4.3 VWAP-minus profit-sharing contract

For this contract, we consider the following characteristics:

- $F = 900\,000\,000\text{ €}$;
- $\alpha = 25\%$;
- $\kappa = 0.5\%$;
- $\ell : x \mapsto Cx^2$ as terminal penalty, where $C = 2 \cdot 10^{-9}\text{ €}^{-1}$;
- $\underline{\rho} = 0, \bar{\rho} = +\infty$, reflecting the prohibition to sell.

To manage the contract, we consider the modified payoff described in Remark 2 and in Section 3.1.2 with $\beta = 5\%$: the bank does not only get part of the profit but also part of the loss.

We choose $\gamma = 10^{-6}\text{ €}^{-1}$ so that $\alpha\gamma = 2.5 \cdot 10^{-7}$.

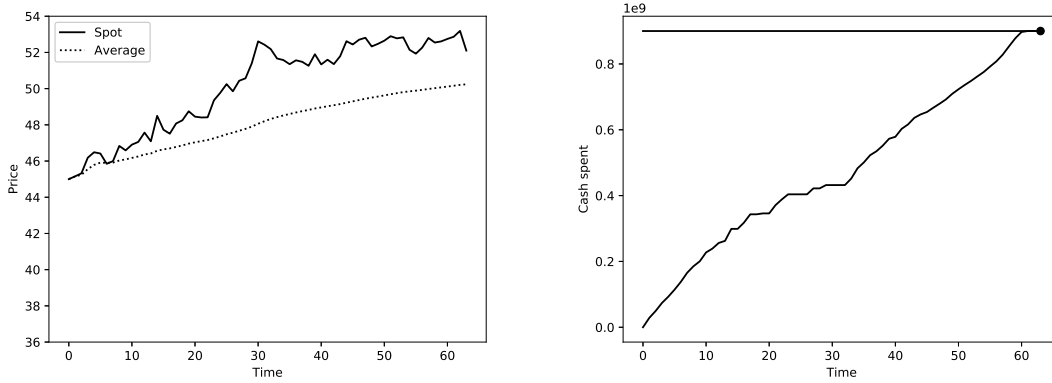


Figure 15: Price trajectory 1 and corresponding strategy for the VWAP-minus profit-sharing contract

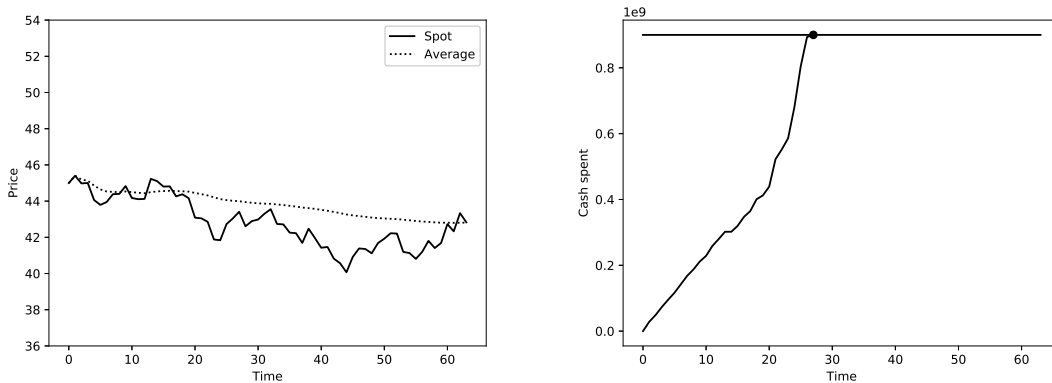


Figure 16: Price trajectory 2 and corresponding strategy for the VWAP-minus profit-sharing contract

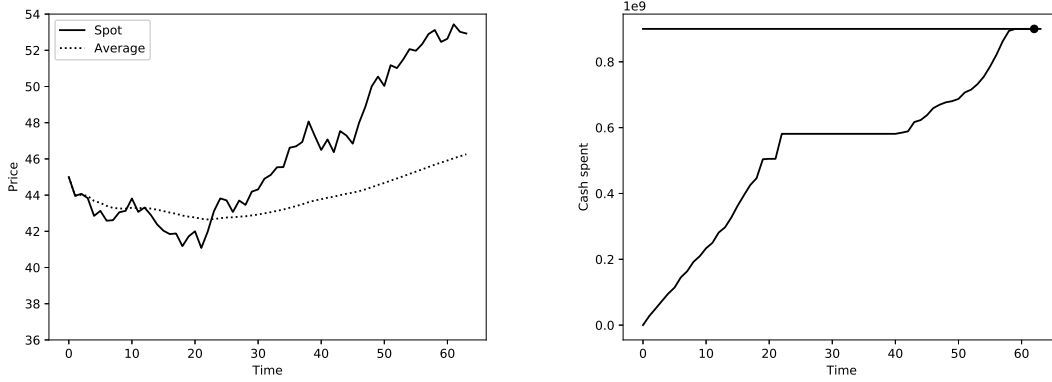


Figure 17: Price trajectory 3 and corresponding strategy for the VWAP-minus profit-sharing contract

The strategies obtained with our neural network algorithm for this type of contract are plotted in Figures 15, 16 and 17.¹³ Here we represent our strategy in terms of the cash spent in repurchasing because cash is the crucial variable for this contract.

We see that the strategy consists in accelerating the purchase process when the price goes below its average and decelerating it when the price increases above its average. In the case of this contract, there is no round trip as selling is prohibited. This explains in particular the shape of the execution strategy in the case of the third price trajectory.

It is interesting to notice (see Figures 18, 19 and 20) that this strategy is similar to that of an ASR contract with fixed number of shares (with the same trading constraints) when one compares the proportion of the cash spent in the case of the former contract with the proportion of shares bought in the case of latter contract.

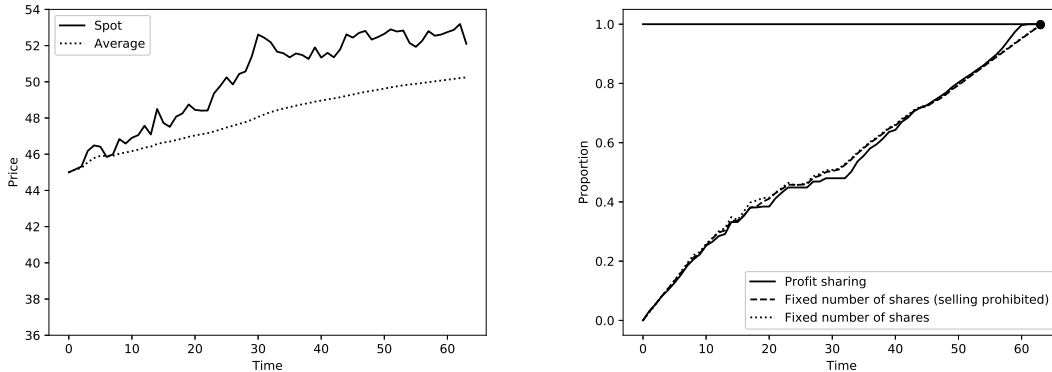


Figure 18: Price trajectory 1 and comparison of the strategies

¹³It must be mentioned that, in the case of this type of contract, we used the value $C = \frac{0.01}{F}$ for the final penalty function during the pretraining phase.

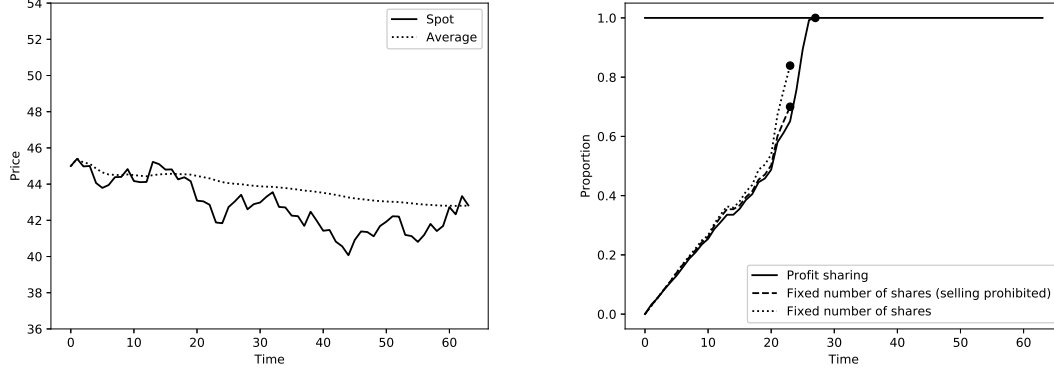


Figure 19: Price trajectory 2 and comparison of the strategies

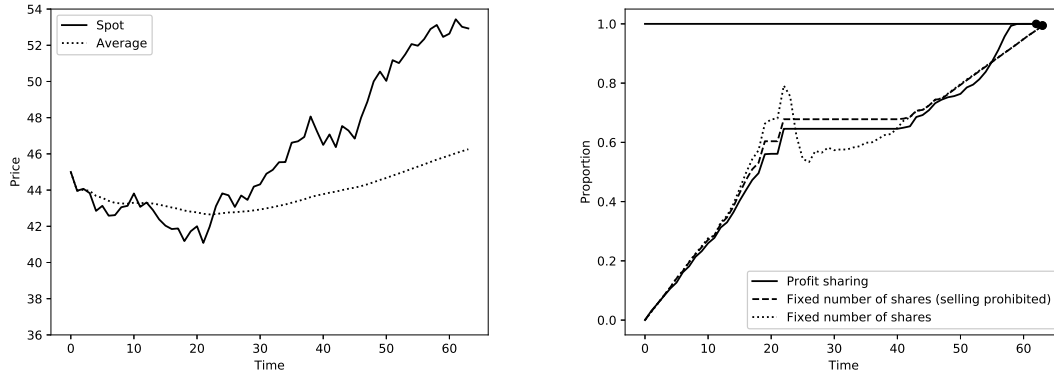


Figure 20: Price trajectory 3 and comparison of the strategies

Conclusion

In this paper, we propose a machine learning approach involving recurrent neural networks to find the optimal strategy associated with different types of VWAP-minus program: ASRs with fixed number of shares, ASRs with fixed notional, and profit-sharing contracts. The results we obtain are in line with both intuition and previous studies. The interest of our method lies in the fact that almost any price dynamics can be considered and that new types of contract can be handled. In particular, we manage to handle contracts for which classical methods usually fail because of (i) high dimensionality and (ii) the very complexity of some contracts that cannot be written as payoffs.

Appendix

In this appendix, we propose a brief introduction to neural networks and their use for approximating functions. We also briefly expose their interest for solving optimisation problems. Finally, we describe the networks used in this paper with a focus on their recurrent structure.

A bit of history

In the 1950s, decisive results were obtained regarding representations of real continuous functions of several variables through addition and composition of continuous functions depending on a smaller number of variables. First, Kolmogorov obtained in [26] a representation with functions of three variables. Then, Arnold, his 19-year-old student at that time, obtained a representation with functions of two variables (see [5]), thus providing an answer to the continuous (as opposed to algebraic) version of Hilbert’s thirteenth problem. Finally, Kolmogorov derived in [27] his celebrated superposition theorem stating the existence of a representation by superpositions of continuous functions of one variable:

Theorem 1 (Kolmogorov’s superposition theorem). *Let $n \geq 2$ be an integer. There exist real continuous functions $(\phi_{p,q})_{1 \leq p \leq n, 1 \leq q \leq 2n+1}$ defined on $[0, 1]$ such that for any real continuous function f defined on $[0, 1]^n$, there exist real continuous functions $(\chi_q)_{1 \leq q \leq 2n+1}$ defined on $[0, 1]$ such that*

$$\forall (x_1, \dots, x_n) \in [0, 1]^n, f(x_1, \dots, x_n) = \sum_{q=1}^{2n+1} \chi_q \left(\sum_{p=1}^n \phi_{p,q}(x_p) \right).$$

This result has then been improved in many ways, from reducing the number of necessary functions, to imposing monotonicity, Lipschitz, or Hölder conditions on the functions. However, despite constructive proofs of superposition theorems, real computation of representations remains almost always impossible because constructions always involve limits, hence infinite loops. Furthermore, the functions involved are often unreasonably complicated because the goal is to obtain a representation rather than an approximation.

Function approximation with neural networks

In fact, an important strand of research has been dedicated to obtain approximate representations with simple functions of a single variable. In particular, feedforward neural networks are often regarded as good candidates for approximating nonlinear functions.

A feedforward neural network is a function of the form Ψ^L (for $L \geq 1$ an integer) where $(\Psi^l)_{0 \leq l \leq L}$ are defined recursively¹⁴ by $\Psi^0 : x \in \mathbb{R}^{d_0} \mapsto x$ and

$$\Psi^l : x \in \mathbb{R}^{d_{l-1}} \mapsto g^l(A^l \Psi^{l-1}(x) + b^l), \quad \forall l \in \{1, \dots, L\},$$

¹⁴Layer 0 is called the input layer while layer L is called the output layer.

where for each layer $l \in \{1, \dots, L\}$, $g^l : (x_1, \dots, x_{d_l}) \in \mathbb{R}^{d_l} \mapsto (h^l(x_1), \dots, h^l(x_{d_l})) \in \mathbb{R}^{d_l}$ applies elementwise either the identity function or a nonlinear function called activation function (in the latter case h^l is typically a sigmoid function, the hyperbolic tangent function, a softmax function, or the rectified linear unit (ReLU) function – *i.e.* $x \in \mathbb{R} \mapsto x \mathbf{1}_{x \geq 0}$), $A^l \in \mathbb{R}^{d_l \times d_{l-1}}$ is a matrix of weights, and $b^l \in \mathbb{R}^{d_l}$ a vector of weights also called bias. A feedforward neural network is often denoted by Ψ_θ where $\theta = (A_l, b_l)_{1 \leq l \leq L}$ stacks all the parameters.

In addition to the link with neurons and the functioning of the animal/human brain (on which we shall remain silent throughout this short appendix), the main interest of feedforward neural networks lies in universal approximation theorems such as the one proved by Hornik in [22]. In a nutshell, these theorems state, under various – and usually mild – hypotheses, that any real function of several variables (here d_0 variables) can be approximated by a feedforward neural network of the above form, provided that there are sufficiently many neurons (*i.e.* d_j for $j \in \{1, \dots, J\}$ large enough). More precisely, Hornik [22] proved the following universal approximation theorem for continuous functions:¹⁵

Theorem 2 (Hornik’s universal approximation theorem). *Let K be a compact set of \mathbb{R}^{d_0} . Let h be a real-valued continuous, bounded, and nonconstant function defined on \mathbb{R} . Then*

$$\left\{ x \in K \mapsto \sum_{j=1}^{d_1} A_{1,j}^2 h \left(\sum_{k=1}^{d_0} A_{j,k}^1 x_k + b_k^1 \right) \mid d_1 \in \mathbb{N}^*, A^1 \in \mathbb{R}^{d_1 \times d_0}, b^1 \in \mathbb{R}^{d_1}, A^2 \in \mathbb{R}^{1 \times d_1} \right\}$$

is dense in the set $C(K)$ of real continuous functions defined on K .

Most of early universal approximation theorems state that one hidden layer with an activation function and a second output layer is sufficient to obtain a good approximation. However, and this is why deep learning is so important, it is often more efficient to approximate a function with several layers consisting of a few neurons than with one layer with a lot of neurons.

Remark 7. *It is noteworthy that, in this paper, neural networks with one hidden layer and a few dozens of neurons were sufficient to obtain satisfactory results.*

Optimisation with neural networks

Feedforward neural networks allow to approximate a large class of functions but so do other families of functions. The reasons why they are often favored have to do with (i) the definition (or characterisation) of the functions to be approximated in machine learning, (ii) the possibility to compute $\Psi_\theta(x)$ and its gradient with respect to θ in an efficient way using forward and backward propagation respectively (see for instance [16] for an introduction), and (iii) the fact that approximations with feedforward neural networks of functions defined on a high-dimensional space are often parsimonious, hence the advantage of neural networks over other families of functions in front of the curse of dimensionality.

¹⁵Hornik also proved a L^p version of the universal approximation theorem. Many other versions of this theorem exist, for instance to handle ReLU activation functions which are unbounded.

In machine learning indeed, be it for supervised, unsupervised, or reinforcement learning, the problem often boils down to approximating a function defined as the solution to an optimisation problem. Therefore, the ability to efficiently differentiate the function with respect to its parameters is essential to use all the classical techniques of gradient descent.¹⁶

Recurrent network structure

In this paper we use neural networks to approximate two functions: the optimal trading strategy of the trader and its optimal stopping decision in the form of a probability to exercise the option. These functions depend, in general, on the time to maturity, the cash spent since inception, the current price of the stock, its running average since the beginning of the contract, and the number of shares already acquired. The problem is discretised in time and decisions made at any time step influence the state at all future time steps including the final one. In particular, the learning procedure to optimise the mean-variance objective function requires to take into account all subsequent effects when computing the derivatives with respect to the weights of the neural networks involved in the decisions taken at a time period $n < N - 1$.

Formally speaking, the trading decision made at a period $n < N - 1$, *i.e.* $v_\theta(n, S_n, A_n, X_n, q_n)$ where θ stands as always for the parameters of the neural network, allows to compute the next value of the inventory process, *i.e.* $q_{n+1} = q_n + v_\theta(n, S_n, A_n, X_n, q_n)\delta t$, which enters as an input in the network to compute the next trading decision $v_\theta(n + 1, S_{n+1}, A_{n+1}, X_{n+1}, q_{n+1})$, and so on (in the case of the profit-sharing contract the same problem occurs with the cash variable). As a consequence, the networks involved in our approach are recurrent. This does not however raise any technical difficulty for computing gradients because we use automatic differentiation.¹⁷

References

- [1] Ali Akyol, Jin S. Kim, and Chander Shekhar. The causes and consequences of accelerated stock repurchases. *International Review of Finance*, 14(3):319–343, 2014.
- [2] Franklin Allen and Roni Michaely. Payout policy. In *Handbook of the Economics of Finance*, volume 1, pages 337–429. Elsevier, 2003.
- [3] Robert Almgren and Neil Chriss. Value under liquidation. *Risk*, 12(12):61–63, 1999.
- [4] Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–40, 2001.
- [5] Vladimir I Arnold. On functions of three variables. collected works: Representations of functions. *Celestial Mechanics and KAM Theory*, 1965:5–8, 1957.

¹⁶The examples of this paper have been computed using the open-source package TensorFlow, distributed by Google.

¹⁷Recurrent structures are sometimes prone to vanishing gradient problems. However, we never noticed any such problem while using our methods.

- [6] Ladshiya Atisoothanan, Balasingham Balachandran, Huu Nhan Duong, and Michael Theobald. Informed trading in option markets around accelerated share repurchase announcements. In *27th Australasian Finance and Banking Conference*, 2014.
- [7] Achref Bachouch, Côme Huré, Nicolas Langrené, and Huyen Pham. Deep neural networks algorithms for stochastic control problems on finite horizon, part 2: numerical applications. 2018.
- [8] Leonce Barger, Manoj Kulchania, and Shawn Thomas. Accelerated share repurchases. *Journal of Financial Economics*, 101(1):69–89, 2011.
- [9] Christian Beck, Sebastian Becker, Philipp Grohs, Nor Jaafari, and Arnulf Jentzen. Solving stochastic differential equations and kolmogorov equations by means of deep learning. 2018.
- [10] Sebastian Becker, Patrick Cheridito, and Arnulf Jentzen. Deep optimal stopping. *Journal of Machine Learning Research*, 20(74):1–25, 2019.
- [11] Hans Buehler, Lukas Gonon, Josef Teichmann, and Ben Wood. Deep hedging. *Quantitative Finance*, pages 1–21, 2019.
- [12] Thomas J. Chemmanur, Yingmei Cheng, and Tianming Zhang. Why do firms undertake accelerated share repurchase programs? 2010.
- [13] Kai Chen. News management and earnings management around accelerated share repurchases. 2017.
- [14] Yung-Chin Chiu and Woan-Lih Liang. Do firms manipulate earnings before accelerated share repurchases? *International Review of Economics & Finance*, 37:86–95, 2015.
- [15] Joan Farre-Mensa, Roni Michaely, and Martin Schmalz. Payout policy. *Annual Review of Financial Economics*, 6(1):75–134, 2014.
- [16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [18] Olivier Guéant. *The Financial Mathematics of Market Liquidity: From optimal execution to market making*, volume 33. CRC Press, 2016.
- [19] Olivier Guéant. Optimal execution of ASR contracts with fixed notional. *Journal of Risk*, 19(5):77–99, 2017.
- [20] 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.
- [21] Jiequn Han, Arnulf Jentzen, and E Weinan. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- [22] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.

- [23] Côme Huré, Huyên Pham, Achref Bachouch, and Nicolas Langrené. Deep neural networks algorithms for stochastic control problems on finite horizon, part i: convergence analysis. 2018.
- [24] Sebastian Jaimungal, Damir Kinzebulatov, and Dmitri Rubisov. Optimal accelerated share repurchases. *Applied Mathematical Finance*, 24(3):216–245, 2017.
- [25] Tao-Hsien Dolly King and Charles E. Teague. Accelerated share repurchases: Value creation or extraction. 2017.
- [26] Andreï Kolmogorov. The representation of continuous functions of several variables by superpositions of continuous functions of a smaller number of variables.
- [27] Andreï Kolmogorov. The representation of continuous functions of several variables by superpositions of continuous functions of one variable and addition.
- [28] Manoj Kulchania. Market microstructure changes around accelerated share repurchase announcements. *Journal of Financial Research*, 36(1):91–114, 2013.
- [29] Ahmet C. Kurt. Managing eps and signaling undervaluation as a motivation for repurchases: The case of accelerated share repurchases. *Review of Accounting and Finance*, 17(4):453–481, 2018.
- [30] Carol A Marquardt, Christine Tan, and Susan M Young. Accelerated share repurchases, bonus compensation, and CEO horizons. In *2012 Financial Markets & Corporate Governance Conference*, 2011.
- [31] Allen Michel, Jacob Oded, and Israel Shaked. Not all buybacks are created equal: The case of accelerated stock repurchases. *Financial Analysts Journal*, 66(6):55–72, 2010.
- [32] Merton Miller and Franco Modigliani. Dividend policy, growth, and the valuation of shares. 1961.
- [33] John W Pratt. Risk aversion in the small and in the large. In *Uncertainty in Economics*, pages 59–79. Elsevier, 1978.
- [34] E Weinan, Jiequn Han, and Arnulf Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in Mathematics and Statistics*, 5(4):349–380, 2017.
- [35] J Fred Weston and Juan A Siu. Changing motives for share repurchases. 2003.
- [36] Ken C. Yook and Partha Gangopadhyay. The wealth effects of accelerated stock repurchases. *Managerial Finance*, 40(5):434–453, 2014.