

Improving reinforcement learning algorithms: towards optimal learning rate policies

Othmane Mounjid*, Charles-Albert Lehalle †

Abstract

This paper investigates to what extent one can improve reinforcement learning algorithms. Our study is split in three parts. First, our analysis shows that the classical asymptotic convergence rate $O(1/\sqrt{N})$ is pessimistic and can be replaced by $O((\log(N)/N)^\beta)$ with $\frac{1}{2} \leq \beta \leq 1$ and N the number of iterations. Second, we propose a dynamic optimal policy for the choice of the learning rate $(\gamma_k)_{k \geq 0}$ used in stochastic approximation (SA). We decompose our policy into two interacting levels: the inner and the outer level. In the inner level, we present the **PASS** algorithm (for “PAst Sign Search”) which, based on a predefined sequence $(\gamma_k^o)_{k \geq 0}$, constructs a new sequence $(\gamma_k^i)_{k \geq 0}$ whose error decreases faster. In the outer level, we propose an optimal methodology for the selection of the predefined sequence $(\gamma_k^o)_{k \geq 0}$. Third, we show empirically that our selection methodology of the learning rate outperforms significantly standard algorithms used in reinforcement learning (RL) in the three following applications: the estimation of a drift, the optimal placement of limit orders and the optimal execution of large number of shares.

1 Introduction

We consider a discrete state space $\mathcal{Z} = \mathbb{N}$ or $\mathcal{Z} = \{1, \dots, d\}$ with $d \in \mathbb{N}^*$. We are interested in finding $q^* \in \mathcal{Q} \subset \mathbb{R}^{\mathcal{Z}}$ solution of

$$M(q, z) = \mathbb{E}[m(q, X(z), z)] = 0, \quad \forall z \in \mathcal{Z}, \quad (1)$$

with $X(z) \in \mathcal{X}$ a random variable with an unknown distribution and m a function from $\mathcal{Q} \times \mathcal{X} \times \mathcal{Z}$ to $\mathbb{R}$. Although the distribution of $X(z)$ is unspecified, we assume that we can observe some variables $(Z_n)_{n \geq 0}$ valued in $\mathcal{Z}$ and $(X_n(Z_{n-1}))_{n \geq 1}$ drawn from the same distribution of $X(Z_{n-1})$. Reinforcement learning (RL) addresses this problem through the following iterative procedure:

$$q_{n+1}(Z_n) = q_n(Z_n) - \gamma_n(Z_n)m(q_n, X_{n+1}(Z_n), Z_n), \quad \forall n \geq 0, \quad (2)$$

where q_0 is a given initial condition and each γ_n is a component-wise non-negative vector valued in $\mathbb{R}^{\mathcal{Z}}$. The connection between RL, problem (1) and Algorithm (2) is detailed in Section 2. It is possible to recover the classical SARSA, Q -learning and double Q -learning algorithms used in RL by taking a specific expression for m and X_{n+1} . Note that Algorithm (2) is different from the standard Robbins-Monro (RM) algorithm used in stochastic approximation (SA)

$$q_{n+1} = q_n - \gamma_n \bar{m}(q_n, X_{n+1}), \quad (3)$$

with $\bar{m} \in \mathbb{R}^{\mathcal{Z}}$ whose z -th coordinate is defined such that $\bar{m}(q, x)_z = m(q, x(z), z)$ for any $z \in \mathcal{Z}$ and $\gamma_n \geq 0$, mainly because, as it is frequent in RL, we do not observe the entire variable $(X_{n+1}(z))_{z \in \mathcal{Z}}$ but only its value according to the coordinate Z_n . Indeed, the way $(Z_n)_{n \geq 1}$ visits the set $\mathcal{Z}$ plays a key role in the convergence of Algorithm (2) which is not the case of Algorithm

*École Polytechnique, CMAP

†Capital Fund Management, Paris and Imperial College, London

(3). RM algorithm was first introduced by Robbins and Monro in [25]. After that, it was studied by many authors who prove the convergence of q_n towards q^* , see [3, 5, 6, 17]. The asymptotic convergence rate has also been investigated in many papers, see [3, 16, 26]. They show that this speed is in general proportional to $1/\sqrt{N}$ with N the number of iterations.

In this work, we give a special focus to RL problems. Nowadays RL cover a very wide collection of recipes to solve control problems in an exploration-exploitation context. This literature started in the seventies, see [31, 32], and became famous mainly with the seminal paper of Sutton, see [29]. It largely relied on the recent advances in the control theory developed in the late 1950s, see [2]. The key tool borrowed from this theory is the dynamic programming principle satisfied by the value function. This principle enables us to solve control problems numerically when the environment is known and the dimension is not too large. To tackle the curse of dimensionality, recent papers, see [28], use deep neural networks (DNN). For example, in [14], authors use DNN to derive optimal hedging strategies for finance derivatives and in [20] they use a similar method to solve a high dimensional optimal trading problem. To overcome the fact that environment is unspecified, it is common to use RM algorithm which estimates on-line quantities of interest. The combination of control theory and SA gave birth to numerous papers on RL.

Our contributions are as follows. We first conduct an error analysis to show that the classical asymptotic rate $O(1/\sqrt{N})$ is pessimistic and can be enhanced in many situations. It is indeed possible to get a $O((\log(N)/N)^\beta)$ asymptotic speed with $1/2 \leq \beta \leq 1$ and N the number of iterations. Then, we present our main result. It consists in proposing a dynamic policy for the choice of the step size $(\gamma_k)_{k \geq 0}$ used in (3). Our policy is decomposed into two interacting levels: the inner and the outer level. In the inner level, we propose the **PASS** algorithm, for “PAst Sign Search”. This algorithm builds a new sequence $(\gamma_k^i)_{k \geq 0}$, using a predefined sequence $(\gamma_k^o)_{k \geq 0}$ and the sign variations of $m(q_n, X_{n+1}(Z_n), Z_n)$. The error of $(\gamma_k^i)_{k \geq 0}$ decreases faster than the one of $(\gamma_k^o)_{k \geq 0}$. In the outer level, we propose an optimal dynamic policy for the choice of the predefined sequence $(\gamma_k^o)_{k \geq 0}$. These two levels are interacting in the sense that **PASS** influences the construction of $(\gamma_k^o)_{k \geq 0}$. Finally, we show that our selection methodology provides better convergence results than standard RL algorithms in three numerical examples: the drift estimation, the optimal placement of limit orders and the optimal execution of a large number of shares.

The structure of this paper is as follows: Section 2 describes the relation between RL and SA. Section 3 reformulates the problem (1) in terms of a minimisation issue and defines with accuracy the different sources of error. This enables us to exploit the most recent convergence results for each source of error to show that the slow convergence speed $O(1/\sqrt{N})$ can be replaced by $O((\log(N)/N)^\beta)$ with $1/2 \leq \beta \leq 1$ and N the number of iterations. Section 4 contains our main contribution. We start by defining the Algorithm **PASS**, comparing it with two other schemes and proving its convergence. Then, we describe our outer level policy and discuss its optimality. Finally, we explain our selection methodology of the learning rate $(\gamma_k^o)_{k \geq 0}$ which combines the **PASS** Algorithm with the outer level to enhance the convergence of (3). The last Section 5 provides numerical examples taken from the optimal trading literature: optimal placement of a limit order and the optimization of the trading speed of a liquidation algorithm. Proofs and additional results are relegated to an appendix.

2 Reinforcement learning

We detail in this section the relation between SA and RL since we are interested in solving RL problems. RL aims at estimating the Q -function which quantifies the value for the player to choose the action a when the system is at s . Let t be the current time, $U_t \in \mathcal{U}$ be a process defined on a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, \mathbb{P})$ which represents the current state of the system and $A_t \in \mathcal{A}$ the agent action at time t . We assume that the process (U_t, A_t) is Markov. The agent aims at maximizing

$$\mathbb{E}\left[\int_0^T \rho^s f(s, U_s, A_s) ds + \rho^T g(U_T)\right], \quad (4)$$

with g the terminal constraint, f the instantaneous reward, ρ a discount factor and T the final time. Let us fix a time step $\Delta > 0$ and allow the agent to take actions only at times³ $k\Delta$ with $k \in \mathbb{N}$. The Q -function satisfies

$$Q(t, u, a) = \sup_A \mathbb{E}_A \left[\int_t^T \rho^{(s-t)} f(s, U_s, A_s) ds + \rho^{(T-t)} g(U_T) | U_t = u, A_t = a \right], \quad \forall (t, u, a) \in \mathbb{R}_+ \times \mathcal{U} \times \mathcal{A},$$

with $A = \{A_t, t < T\}$ a possible control process for the agent. Note that the action of the agent depend on $s = (t, u)$ with t the current time and u the current state. We view the agent control A as a feedback process (i.e adapted to the filtration $\mathcal{F}_t$). The Q -function satisfies the classical dynamic programming principle (DPP)

$$Q(t, u, a) = \mathbb{E}[R_{t+\Delta} + \rho^\Delta \sup_{a' \in \mathcal{A}} Q(t + \Delta, U_{t+\Delta}, a') | U_t = u, A_t = a], \quad (5)$$

with $R_{t+\Delta} = \int_t^{t+\Delta} \rho^{(s-t)} f(s, U_s, A_s) ds$. Equation (5) reads that the optimal expected gain when the agent starts at s and chooses action a at time t is the sum of the next expected reward $R_{t+\Delta}$ plus the value of acting optimally starting from the new position $U_{t+\Delta}$ at time $t + \Delta$. By reformulating (5), we obtain that Q solves equation

$$\mathbb{E}[m(q, X_{t+\Delta}^z, z)] = 0, \quad \forall z = (t, u, a) \in \mathcal{Z} = [0, T] \times \mathcal{U} \times \mathcal{A}, \quad (6)$$

where $X_{t+\Delta}^z = (U_{t+\Delta}^z, R_{t+\Delta}^z) \in \mathcal{X} = \mathcal{U} \times \mathbb{R}$, U_s^z and R_s^z are respectively the conditional random variables U_s and R_s given the initial condition $(U_t^z, A_t^z) = (u, a)$ with $z = (t, u, a) \in \mathcal{Z}$ and m is defined as follows:

$$m(q, x, z^1) = H(q, x, z^1) - q(z^1), \quad H(q, x, z^1) = r + \rho^\Delta \sup_{a' \in \mathcal{A}} q(t^1 + \Delta, u, a'),$$

for any $x = (u, r) \in \mathcal{X}$ and $z = (t^1, u^1, a^1) \in \mathcal{Z}$. Thus, one can use stochastic approximation tools to solve (6).

Note that Equation (6) shows that one can study Q only on the time grid⁴ $D_T = \{n\Delta, n \leq T/\Delta\}$. Thus, we define A_k and U_k such that $A_k = A_{k\Delta}$ and $U_k = U_{k\Delta}$ for any $k \in \mathbb{N}$. The key variable to study is not the agent decision A_k but $Z_k = (k, U_k, A_k)$. Thus, the rest of the paper formulates the results in terms of Z_k only.

Actions of the agent. It is important in practice to visit the space $D_T \times \mathcal{U} \times \mathcal{A}$ sufficiently enough. Thus, to learn Q , it is common to not choose the maximising action⁵, but to encourage exploration by visiting the states where the error is large. We give in Appendix A examples of policies that promote exploration and others that maximize the Q -function.

3 Improvement of the asymptotic convergence rate

In [4, Part 2, Section 4], [16, Section 10] and [17, Section 7], the authors show a central limit theorem for the procedure (3) which ensures a convergence rate of $O(1/\sqrt{N})$ where N is the number of iterations. In this section, we extend such convergence rate to Algorithm (2) and aim at understanding how one can improve it. For this, we decompose our total error into two components: estimation error and optimization error.

3.1 Error decomposition

In this section, the space $\mathcal{Z} = \{1, \dots, d\}$ is finite with $d \in \mathbb{N}^*$. In such case, we view q and $M(q)$ as vectors of $\mathbb{R}^{\mathcal{Z}}$. Moreover, the process $(Z_n)_{n \geq 1}$ is Markov. We consider the following assumption.

Assumption 1 (Existence of a solution). *There exists a solution q^* of Equation (1).*

³We recall the following classical result: when Δ goes to zero, the value function and the optimal control of this problem converges towards the one where decisions are taken at any time.

⁴Here, we take $T = n^* \Delta$ with of $n^* \in \mathbb{N}^*$. Such approximation is not restrictive.

⁵The maximising action a^* for a state u is defined such that $a^* = \arg \max_{a \in \mathcal{A}} q(u, a)$.

Under Assumption 1, the function q^* is solution of the minimization problem

$$\min_{q \in \mathcal{Q}} g(q). \quad (7)$$

with $g(q) = \|M(q)\|^2$.

Remark 1. Note that, in the special case where M is the gradient of a given function f (i.e. $\nabla f = M$), the quantity q^* minimises a convex and differentiable cost

$$\tilde{g} = \sum_{z \in \mathcal{Z}} \mathbb{E}[L(q, X(z), z)],$$

and we can derive the following explicit expression for L

$$L(q, x, z) = \int_0^1 (q(z) - q^*(z)) \times m(q^* + r(q - q^*), x, z) dr, \quad \forall q \in \mathcal{Q}, \forall z \in \mathcal{Z}, \forall x \in \mathcal{X}.$$

with $\nabla \tilde{g} = M$. Thus, one can replace g by $\tilde{g}$ and all the results of this section hold. In the rest of this section, we use g instead of $\tilde{g}$.

In our context we do not have a direct access to the distribution of $X(z)$. Nevertheless, we assume that at time n we keep memory of a training sample of $n(z)$ independent variables $(X_i^z)_{i=1 \dots n(z)}$ drawn from the distribution $X(z)$ where $n(z)$ is the number of times the Markov chain Z_n visited z . We define q^n as a solution of

$$\min_{q \in \mathcal{Q}} g_n(q), \quad (8)$$

with $g_n(q) = \|M^n(q)\|^2$ and

$$M^n(q, z) = \mathbb{E}^n[m(q, X(z), z)] = \left(\sum_{j=1}^{n(z)} m(q, X_j(z), z) \right) / n(z),$$

expected value under the empirical measure $\mu = (\sum_{j=1}^{n(z)} \delta_{X_j(z)}) / n(z)$ (i.e. empirical risk). We finally define q_k^n as an approximate solution of the problem (8) returned by an optimization algorithm after k iterations. Thus, we can bound the error $g(q_k^n)$ by

$$0 \leq \mathbb{E} \left[\left(g(q_k^n) - g(q^*) \right) (z) \right] \leq \underbrace{\mathbb{E} \left[\left(g(q^n) - g(q^*) \right) (z) \right]}_{\text{estimation error}} + \underbrace{\mathbb{E} \left[\left| g(q^n) - g(q_k^n) \right| (z) \right]}_{\text{optimization error}},$$

since q^* minimizes g .

3.2 Convergence rate of the estimation error

3.2.1 Slow convergence rate

We have the following result.

Proposition 1. We assume that the Markov chain Z_n is irreducible. There exists $c_1 > 0$ such that

$$\mathbb{E}[\sup_{q \in \mathcal{Q}} |g(q) - g_n(q)| (z)] \leq c_1 \frac{1}{\sqrt{n}}, \quad \forall z \in \mathcal{Z}.$$

For sake of completeness, we give the proof of this result in Appendix D. Proposition 1 allows us to derive the following bound for the estimation error

$$\begin{aligned} \mathbb{E}[(g(q^n) - g(q^*)) (z)] &= \mathbb{E}[(g(q^n) - g_n(q^n)) (z)] + \underbrace{\mathbb{E}[(g_n(q^n) - g_n(q^*)) (z)]}_{\leq 0} + \mathbb{E}[(g_n(q^*) - g(q^*)) (z)] \\ &\leq 2\mathbb{E}[\sup_q |g(q) - g_n(q)| (z)] \leq 2c_1 \frac{1}{\sqrt{n}}. \end{aligned} \quad (9)$$

This bound is known to be pessimistic.

3.2.2 Fast convergence rate

We obtain the following fast statistical convergence rate.

Proposition 2. *Assume that the Markov chain Z_n is irreducible and*

$$\mathbb{E} \left[\sup_{q \in \mathcal{Q}} |g(q) - g_n(q)|(z) | n(z) \right] \leq c' \left(\frac{\log(\bar{n}(z))}{\bar{n}(z)} \right)^\beta, \quad (10)$$

with $\frac{1}{2} \leq \beta \leq 1$, $c' > 0$ and $\bar{n}(z) = n(z) \wedge 1$. Then, there exists $c_2 > 0$ such that

$$\mathbb{E}[|g(q^n) - g(q^*)|(z)] \leq c_2 \left(\frac{\log(n)}{n} \right)^\beta, \quad \forall z \in \mathcal{Z}.$$

The proof of this proposition is given in Appendix E. The condition (10) is established when

- The loss function g satisfies regularity conditions, of which the most important are : Lipschitz continuity and convexity, see [1, Section 4]. Moreover under the strong convexity assumption, the constant β is equal to 1.
- The data distribution satisfies some noise conditions, see for instance [1, Section 3] and [30].
- The function m has a bounded moment α (i.e. $\mathbb{E}[|m|^\alpha(q)] < \infty$) with $\alpha > 1$, see [11, Section 4].

It is also possible to get rid of the $\log(n)$ factor in (10), see the end of Section 5 in [8].

3.3 Convergence rate of the optimization error

We turn now to the optimization error. This means that expected value in (7) is replaced by the empirical risk which is known. In such case, one can use many algorithms to find q_n . We introduce the following assumption.

Assumption 2 (Pseudo strong convexity). *There exists $L > 0$ such that*

$$\langle M(q) - M(q'), q - q' \rangle \geq L \|q - q'\|^2, \quad \forall q \in \mathbb{R}^{\mathcal{Z}}, \forall q' \in \mathbb{R}^{\mathcal{Z}}.$$

The above assumption is natural since the gradient of any real valued strongly convex function f satisfies Assumption 2.

We present in the table below the most important properties of some gradient methods. Note that the results of Table 1 remain valid when strong convexity is replaced by Assumption 2.

Algorithm	Cost of one iteration	Iterations to achieve an ϵ precision		Time to reach an ϵ precision	
		Convex	Strongly convex	Convex	Strongly convex
GD	$O(d^2)$	$O\left(\frac{1}{\epsilon}\right)$	$O\left(\log\left(\frac{1}{\epsilon}\right)\right)$	$O\left(\frac{d^2}{\epsilon}\right)$	$O\left(d^2 \log\left(\frac{1}{\epsilon}\right)\right)$
SGD	$O(d)$	$O\left(\frac{1}{\epsilon^2}\right)$	$O\left(\frac{1}{\epsilon}\right)$	$O\left(\frac{d}{\epsilon^2}\right)$	$O\left(\frac{d}{\epsilon}\right)$
Proximal	$O(d)$	$O\left(\frac{1}{\epsilon}\right)$	$O\left(\log\left(\frac{1}{\epsilon}\right)\right)$	$O\left(\frac{d}{\epsilon}\right)$	$O\left(d \log\left(\frac{1}{\epsilon}\right)\right)$
Acc. prox.	$O(d)$	$O\left(\frac{1}{\sqrt{\epsilon}}\right)$	$O\left(\log\left(\frac{1}{\epsilon}\right)\right)$	$O\left(\frac{d}{\sqrt{\epsilon}}\right)$	$O\left(d \log\left(\frac{1}{\epsilon}\right)\right)$
SAGA	$O(d)$	$O\left(\frac{1}{\epsilon}\right)$	$O\left(\log\left(\frac{1}{\epsilon}\right)\right)$	$O\left(\frac{d}{\epsilon}\right)$	$O\left(d \log\left(\frac{1}{\epsilon}\right)\right)$
SVRG	$O(d)$		$O\left(\log\left(\frac{1}{\epsilon}\right)\right)$		$O\left(d \log\left(\frac{1}{\epsilon}\right)\right)$

Table 1: Asymptotic properties of some gradient methods. Note that d is the dimension of the state space $\mathcal{Z}$ and ϵ is a desired level of accuracy. Here ϵ corresponds to $1/n$. GD stands for Gradient Descent, SDG for Stochastic Gradient Descent, *Proximal* for Stochastic proximal gradient descent [10, 27], *Acc. prox.* for accelerated proximal stochastic gradient descent [23, 27], SAGA for Stochastic accelerated gradient approximation [12], and SVRG for stochastic variance reduced gradient [15].

3.4 Conclusion

Following the formalism of [7], we have decomposed our initial error into

- **Estimation error:** its convergence is $O(1/\sqrt{N})$ in pessimistic cases with N the number of iterations. In the other situations, the convergence is faster (i.e. $O((\log(N)/N)^\beta)$) with $1/2 \leq \beta \leq 1$.
- **Optimization error:** the convergence is exponential under suitable conditions. In unfavourable cases, the convergence rate is $O(1/N)$.

The comparison of these error sources shows that the estimation error is the dominant component. Thus, one can overcome the $O(1/\sqrt{N})$ asymptotic speed, in some situations, by improving the estimation error.

4 Optimal policy for the learning rate γ when is countable

In this section, we take $\mathcal{Z} = \mathbb{N}$ and consider the following type of algorithms:

$$q_{n+1}(Z_n) = q_n(Z_n) - \gamma_n(Z_n)m(q_n, X_{n+1}(Z_n), Z_n), \quad \forall n \in \mathbb{N}.$$

One can recover the classical SARSA, Q -learning and double Q -learning algorithms used in RL by considering a specific expression for m and X_{n+1} . In such algorithms the choice of γ_n is a crucial point. One can find in the literature general conditions that guarantee the convergence such that

$$\sum_{k \geq 0} \gamma_k(z) = \infty, \quad a.s., \quad \sum_{k \geq 0} \gamma_k^2(z) < \infty, \quad a.s., \quad \forall z \in \mathcal{Z}. \quad (11)$$

However, since the set of processes $(\gamma_n)_{n \geq 0}$ satisfying these conditions is large in general and may even be empty when $(Z_n)_{n \geq 0}$ is not recurrent, many authors suggest to take γ_n proportional to $1/n^\alpha$ for stochastic approximation (SA) algorithms to be more specific. The exponent α may vary from 0 to 1 depending on the algorithm used, see [13, 22]. Nonetheless, such a choice may be sub optimal. For example, Figure 1.a shows that the blue curve is a way higher than the orange one. Here, the blue curve represents the variation of the logarithm of the L^2 -error when $\gamma_n = \eta/n$, whereas the orange curve stands for a constant learning rate (i.e. $\gamma_n = \gamma$). We choose the constant η that ensures the fastest convergence for the blue curve.

In this paper, we propose to use a stochastic learning rate $(\gamma_k)_{k \geq 0}$; our learning policy is decomposed into two interacting levels: the inner and the outer level. In the inner level, we use the **PASS** algorithm, for “PAsT Sign Search”. This algorithm builds a new sequence $(\gamma_k^i)_{k \geq 0}$, based on a predefined sequence $(\gamma_k^o)_{k \geq 0}$ and the sign variations of $m(q_n, X_{n+1}(Z_n), Z_n)$, whose error decreases faster than the predefined one. In the outer level, we propose an optimal methodology for the selection of the predefined sequence $(\gamma_k^o)_{k \geq 0}$. These two levels are interacting in the sense that the **PASS** algorithm influences the construction of the sequence $(\gamma_k^o)_{k \geq 0}$.




Figure 1: L^2 -error for the estimation of the drift when γ_k is constant in orange and when $\gamma_k \propto \frac{1}{k}$ in blue.

4.1 The inner level

4.1.1 The algorithms

In this part, we introduce three algorithms. We start with our benchmark which is the standard algorithm used in RL. Then, we present a second algorithm inspired from SAGA [12], which is a method used to accelerate the convergence of the stochastic gradient descent. SAGA reduces the optimization error exponentially fast. Finally, we describe the **PASS** algorithm that modifies the learning rate $(\gamma_k)_{k \in \mathbb{N}}$ based on the sign variations of $m(q_n, X_{n+1}(Z_n), Z_n)$. The main idea is to increase γ_n as long as the sign of $m(q_n, X_{n+1}(Z_n), Z_n)$ remains unchanged. Then, we reinitialize or lower γ_n using a predefined sequence $(\gamma_k^0)_{k \in \mathbb{N}}$ when the sign of $m(q_n, X_{n+1}(Z_n), Z_n)$ switches. This algorithm can be seen as an adaptation of the line search strategy, which determines the maximum distance to move along a given search direction, to SA methods. Actually, the line search method requires a complete knowledge of the cost function because it demands to evaluate several times the difference $g(q_k + \gamma M(q^k)) - g(q_k)$ for different values of γ with g and M defined in Section 3.1. However, stochastic iterative models have neither access to g nor M . They can only compute $m(q_n, X_{n+1}(Z_n), Z_n)$ when the state $z = Z_k$ is visited. Moreover, to get a new observation they need to wait⁶ for the next visit of the state $z = Z_k$. Nevertheless, they have instantaneous access to the previously observed values. Thus, the main idea here is to use these past observations although it adds a small memory cost. Some theoretical properties of these algorithms are investigated in Section 4.3.

Algorithm 1 (RL). We start with an arbitrary $q_0 \in \mathcal{Q}$ and define by induction q_k

$$q_{k+1}(Z_k) = q_k(Z_k) - \gamma_k(Z_k)m(q_k, X_{k+1}(Z_k), Z_k).$$

Algorithm 2 (SAGA). We start with an arbitrary $q_0 \in \mathcal{Q}$, $M_0 = 0$ ⁷ and define by induction q_k and M_k

$$\begin{aligned} q_{k+1}(Z_k) &= q_k(Z_k) - \gamma_k(Z_k) \left[m(q_k, X_{k+1}(Z_k), Z_k) - M_k[Z_k, i] + \frac{(\sum_{j=1}^M M_k[Z_k, j])}{M} \right], \\ M_{k+1}[Z_k, i] &= m(q_k, X_{k+1}(Z_k), Z_k), \end{aligned}$$

with i picked from the distribution $p = (\sum_{i=1}^M \delta_i)/M$.

For the next algorithm, we give ourselves a predefined learning rate that we call $(\gamma_k)_{k \geq 0}$, a non-decreasing function $h : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ and a non-increasing function $l : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. The function h is used to increase the learning rate and hence to accelerate the descent, while the function l is used to go back to a slower pace.

Algorithm 3 (PASS). We start with an arbitrary q_0 and define by induction q_k and $\hat{\gamma}_k$

- If $m(q_n, X_{n+1}(Z_n), Z_n) \times m(q_{r_1^n}, X_{r_1^n+1}(Z_{r_1^n}), Z_{r_1^n}) \geq 0$, then do

$$\begin{aligned} q_{n+1}(Z_n) &= q_n(Z_n) - h(\hat{\gamma}_n(Z_n), \gamma_n(Z_n))m(q_n, X_{n+1}(Z_n), Z_n), \\ \hat{\gamma}_{n+1}(Z_n) &= h(\hat{\gamma}_n(Z_n), \gamma_n(Z_n)), \end{aligned}$$
 with r_1^n is the index of the last observation when the process X visits the state X_n .
- Else, do

$$\begin{aligned} q_{n+1}(X_n) &= q_n(X_n) - l(\hat{\gamma}_n(X_n), \gamma_n(X_n))m(q_n, X_n, X_{n+1}), \\ \hat{\gamma}_{n+1}(X_n) &= l(\hat{\gamma}_n(X_n), \gamma_n(X_n)). \end{aligned}$$

4.1.2 Assumptions

In this section, we present the assumptions needed to prove our main result about the convergence of Algorithms **RL**, **SAGA** and **PASS**. We assume that Assumption is in force. Hence, we denote by q^* a solution of (1) and write m^* for the vector $m^* = m(q^*) \in \mathcal{X} \times \mathcal{Z}$. Let us consider the following assumptions:

⁶This waiting time may be very long depending on the dimension of the state space $\mathcal{Z}$ and the properties of the process $(Z_k)_{k \geq 0}$.

⁷Here M_0 is the zero function in the sense that $M_0[z, i] = 0$ for any $z \in \mathcal{Z}$ and $i \in \{1, \dots, M\}$.

Assumption 3 (Pseudo strong convexity 2). *There exists a constant $L > 0$ such that*

$$(\mathbb{E}_k[m(q_k, X_{k+1}(Z_k), Z_k) - m^*(X_{k+1}(Z_k), Z_k)])(q_k(Z_k) - q^*(Z_k)) \geq L(q_k(Z_k) - q^*(Z_k))^2,$$

with q^* of (1) and $\mathbb{E}_k[X] = \mathbb{E}[X|\mathcal{F}_k]$ for any random variable X .

Recall that Assumption 3 is natural in the deterministic framework. For instance, if we take a strongly convex function f and call m its gradient (i.e $m = \nabla f$). Then, m satisfies Assumption 3. Additionally, the pseudo-gradient property (PG) considered in [5, Section 4.2] is close to Assumption 3. However, Assumption 3 is slightly less restrictive than PG since it involves only the norm of the component $(q_k - q^*)(Z_k)$ instead of the norm of the vector $(q_k - q^*)$. To get tighter approximations, we will also need the quantity L_k defined as follows:

$$L_k = \begin{cases} \frac{\mathbb{E}_k[m(q_k, X_{k+1}(Z_k), Z_k) - m^*(X_{k+1}(Z_k), Z_k)]}{q^k(Z_k) - q^*(Z_k)}, & \text{If } q^k(Z_k) - q^*(Z_k) \neq 0, \\ 0, & \text{otherwise.} \end{cases}$$

Note that $L_k \geq 0$ under Assumption 3. It is also the biggest constant that satisfies Assumption 3 for a fixed k . In particular, this means that $L_k \geq L$.

Assumption 4 (Lipschitz continuity of m). *There exists a positive constant $B > 0$ such that for any random variables X and X' valued in $\mathcal{X}$ we have*

$$\mathbb{E}_k[(m(q_k, X, Z_k) - m^*(X', Z_k))^2] \leq B\{1 + (q_k(Z_k) - q^*(Z_k))^2 + \mathbb{E}_k[(X - X')^2]\},$$

with $\mathbb{E}_k[X] = \mathbb{E}[X|\mathcal{F}_k]$ for any random variable X .

Assumption 4 guarantees that m is Lipschitz. Authors in [5, Section 4.2] use a similar condition. To get better bounds, we introduce B_k

$$B_k = \frac{\mathbb{E}_k[(m(q_k, X, Z_k) - m^*(X', Z_k))^2]}{1 + (q_k(Z_k) - q^*(Z_k))^2 + \mathbb{E}_k[(X - X')^2]}.$$

We have $B_k \leq B$ since B_k is the smallest constant satisfying Assumption 4 for a fixed k . We finally need an assumption on the learning $(\gamma_k)_{k \geq 0}$ that describes indirectly how the process Z communicates with its different states.

Assumption 5 (Learning rate explosion). *For any $z \in \mathcal{Z}$, we have*

$$\sum_{k \geq 1} \gamma_k(z) = \infty, \quad a.s.$$

When the process Z is Markov and $\gamma_k(z)$ is bounded, Assumption 5 ensures that Z is recurrent. To see this, we first remark that the boundedness assumption of $\gamma_k(z)$ is not restrictive and can always be fulfilled. After that, we use the boundedness of $\gamma_k(z)$ to exhibit a positive constant A such that $\gamma_k(z) \leq A$, a.s for all $k \geq 1$. Thus, we get $\sum_{k \geq 1} \mathbb{E}[\gamma_k(z) \mathbf{1}_{Z_k=z}] \leq A \sum_{k \geq 1} \mathbb{P}[Z_k = z]$. Since the left hand side of the previous inequality diverges under Assumption 5, we have

$$\sum_{k \geq 1} \mathbb{P}[Z_k = z] = \infty,$$

which proves that Z is recurrent. Moreover, in the Markov setting one can replace Assumption 5 by the more general condition (11).

4.1.3 Main results

In this section, we compare the algorithms **RL**, **SAGA** and **PASS** and prove the convergence of **PASS**. Let c be a positive constant and $k \in \mathbb{N}$. We define the error function for the different algorithms as follows:

$$e^k(z) = \begin{cases} (q_k(z) - q^*(z))^2, & \text{for algorithms **RL** and **PASS**,} \\ \frac{\sum_{j=1}^M (M^k[z, j] - m^*(z))^2}{M} + c(q_k(z) - q^*(z))^2, & \text{for algorithm **SAGA**,} \end{cases}$$

for all $z \in \mathcal{Z}$ and $j \in \{1, \dots, M\}$. We write E^k for the total error $E^k = \|e^k\|$. We also use the following notations:

$$p(x) = 2Lx - Bx^2, \quad p_k(x) = 2L_kx - B_kx^2, \quad \bar{\gamma}_k = \arg \sup_{l \in \mathbb{R}} p_k(l) = \frac{L_k}{B_k}, \quad \forall x \in \mathbb{R}.$$

Proposition 3. Let $z \in \mathcal{Z}$. Under Assumptions 1, 3 and 4 and when there exists $r_1 \geq 1$ such that

$$\gamma_2 \leq l(\gamma_1, \gamma_2), \quad \text{and} \quad h(\gamma_1, \gamma_2) \leq r_1 \gamma_2, \quad \forall (\gamma_1, \gamma_2) \in \mathbb{R}^2,$$

we have

$$\mathbf{1}_A \mathbb{E}_k[e^{k+1}(z)] \leq \mathbf{1}_A[\alpha_k e^k(z) + M_k], \quad (12)$$

with $A = \{Z_k = z\}$. The values of the constant α_k and M_k vary from an algorithm to another as follows:

$$\alpha_k(z_1) = \begin{cases} [1 - p(\gamma_k(z))], & \text{for algorithm RL,} \\ \max\left(1 - p(\gamma_k(z)) + \frac{B}{Mc}, 1 - \left(\frac{1}{M} - 6\gamma_k^2(z)c\right)\right), & \text{for algorithm SAGA,} \\ \{1 - p_k(\underline{\gamma}_k^*) + d^1 \gamma_k^2(z) \mathbf{1}_{c_k \geq 1}\}, & \text{for algorithm PASS,} \end{cases} \quad (13)$$

and

$$M_k = \begin{cases} B\gamma_k^2(z)(2 + v_k), & \text{for algorithm RL,} \\ 3B\gamma_k^2(z)(2 + v_k), & \text{for algorithm SAGA,} \\ B(c_k \bar{\gamma}_k)^2(2 + v_k), & \text{for algorithm PASS,} \end{cases} \quad (14)$$

with $c_k = \frac{\mathbb{E}_k[\hat{\gamma}_k(z)m(q^k, X_{k+1}(z), z)]}{\bar{\gamma}_k(z)\mathbb{E}_k[m(q^k, X_{k+1}(z), z)]}$, $\underline{\gamma}_k^* = c_k(z)\bar{\gamma}_k \vee \bar{\gamma}_k$, $d^1 = (r_1 - 1)^2 B_k$ and $v_k = \text{Var}(Z_k)$.

Equation (12) reveals that the performances of algorithms RL, SAGA and PASS depend on the interaction between two competing terms:

- On the one hand the slope α_k controls the decrease of the error from one step to the other,
- On the other hand the quantity M_k gathers two sources of imprecision: the estimation error and the optimization error. Both sources of imprecision have a term in the variance v_n (because the distribution of Z is unknown) and a positive constant (coming from the noise generated by the noisy nature of observations).

There is a competition between these two terms: to decrease M_k we need to send γ_k towards zero while the reduction of α_k requires a relatively small but still non-zero value of γ_k . Thus, γ_k should satisfy a trade-off in order to ensure the convergence of the algorithms. The RM conditions (11) are a way to address this trade-off. Now, in order to analyse the properties of each algorithm, we compare for a given γ_k their respective values for α_k and M_k in Table 2. For sake of clarity, we choose to present the variable $(1 - \alpha_k)$ instead of α_k in this table; note that a large value of $1 - \alpha_k$ means that α_k is small and thus induces a fast convergence.

Algorithms	$(1 - \alpha_k)$		M_k	
	Value	Comparison with Algo 1	Value	Comparison with Algo 1
RL (Algo 1)	$2\gamma_k(z_1)L - B\gamma_k^2(z_1)$	—	$B\gamma_k^2(2 + v_k)$	—
SAGA (Algo 2)	$\left(2\gamma_k(z_1)L - B\gamma_k^2(z_1) - \frac{B}{Mc}\right) \vee \left(\frac{1}{M} - 6\gamma_k^2(z_1)c\right)$	smaller	$3B\gamma_k^2(2 + v_k)$	larger
PASS (Algo 3)	$2\gamma_k^* L_k - B_k(\gamma_k^*)^2 + d^1 \gamma_k^2 \mathbf{1}_{c_k \geq 1}$	larger	$B(c_k \bar{\gamma}_k)^2(2 + v_k)$	larger

Table 2: Comparison of the algorithms RL, SAGA and PASS.

Let $n \in \mathbb{N}^*$, $z_1 \in \mathcal{Z}$, and $\tau_{\mathcal{Z}_1}^n = \inf\{l > 0, Z_{l+n} = z_1\}$. We need to introduce the following notations $\tilde{a}_j^n = \mathbb{P}[\tau_{\mathcal{Z}_1}^{n-j} \geq j | Z_{n-j} = z_1]$, $\tilde{b}_j = \mathbb{P}[Z_j = z_1]$, and $\mu_j = \frac{\mathbb{E}_0[\mathbf{1}_{Z_j=z_1} \alpha_j e^j(z_1)]}{\tilde{b}_j \mathbb{E}_0[e^j(z_1)]}$ for all $j < n$. We write $\bar{\mu}_j^n = \prod_{k=n-j+1}^n \mu_k$ and $\tilde{a}_n^{*\infty} = \lim_{m \rightarrow \infty} \tilde{a}_n^{*m}$. The sequence $(\tilde{a}_k^{*m})_{k \geq 1}$ is defined recursively such that $\tilde{a}_k^{*1} = a_k$ and $\tilde{a}_k^{*(m+1)} = \sum_{l=1}^k \tilde{a}_{k+1-l} \tilde{b}_l \tilde{a}_l^{*m}$ for all $k \geq 1$ and $m \geq 1$. The result below holds only for Algorithms 1 and 3.

Theorem 1. *Let the Assumptions 1, 3, 4 and 5 be in force. Algorithms 1 and 3 verify*

- When $\sum_{k \geq 0} \mathbb{E}[\gamma_k^2](z) < \infty$ for all $z \in \mathbb{Z}$, we have

$$\mathbb{E}_0[E^n] \xrightarrow{n \rightarrow \infty} 0. \quad (15)$$

- There exists a constant $B' \geq 0$ such that

$$\mathbb{E}_0[e^n(z_1)] \leq B' \sum_{(l,j,i) \in (\mathbb{N}^*)^3, l+j+i=n} \tilde{\epsilon}_j^1 \bar{\mu}_l^j a_i^{*\infty}, \quad (16)$$

$$\text{with } \tilde{\epsilon}_n = e^1(z_1) \mathbb{P}[\tau_{Z_1}^0 \geq n] + \sum_{j=1}^{n-1} \tilde{a}_{n-j}^n \mathbb{E}[M_j].$$

The first part of the above result ensures the convergence of the error E^n towards zero and the second part describes the behaviour of the error by giving an upper bound. This upper bound shows how the different terms: $\tilde{\epsilon}_n$, $\bar{\mu}_k^n$, and $\tilde{a}_n^{*\infty}$ interact together in order to guarantee the decrease of the error. We recall that $\tilde{\epsilon}_n$ is a noise term that gather the model uncertainty, $\bar{\mu}$ represents the influence of the slope factor α , and $\tilde{a}_n^{*\infty}$ is related to the probability distribution of the process $(Z_k)_{k \geq 0}$. Note that the right-hand side of (16) is not trivial and it converges towards 0 when there exists constant $L' \geq 0$ such that $\mathbb{E}_0[\mathbf{1}_{Z_j=z_1} \alpha_j e^j(z_1)] \leq L' \tilde{b}_j \mathbb{E}_0[e^j(z_1)] \mathbb{E}[\alpha_j]$ for all $j \geq 2$, see Appendix G.

4.2 The upper level

In practice, to apply **PASS** we need an appropriate predefined sequence $(\gamma_k)_{k \in \mathbb{N}}$. It is possible to take γ_k proportional to $1/k^\alpha$ with $\alpha \in (0, 1]$ as proposed in [22, 25]. However, in this section, we present an optimal dynamic policy for the choice of the learning rate $(\gamma_k)_{k \in \mathbb{N}}$. To do so, we assume that

$$e^{n+1}(z) = \mathbf{1}_A(\alpha_n e^n(z) + M_n + S_n) + \mathbf{1}_{A^c} e^n(z), \quad (17)$$

with $A = \{Z_n = z\}$, S_n does not depend on the learning rate and verifies $\mathbb{E}[\mathbf{1}_A S_n] \leq 0$. Equation (17) is consistent with Proposition 3. Moreover, we force e^n to stay below the upper bound $x_2 = \arg \sup_{x \in \mathbb{R}_+} g(x)$ with

$$g(x) = x - \frac{L^2 x^2}{2B(x + (2 + v))}, \forall x \in \mathbb{R}.$$

Such constraint is not that restrictive since we know that the error e^n converges towards 0. Since α_n and M_n are both functions of the learning rate, the idea is to choose the learning rate γ_n such that

$$\gamma_n = \arg \min_{\gamma} (\alpha_n(\gamma) e^n(z) + M_n(\gamma)), \quad (18)$$

which gives

$$\gamma_n = \begin{cases} \frac{L}{B} \frac{e^n(z)}{e^n(z) + (2 + v_n)}, & \text{for algorithm RL,} \\ \frac{L_n}{B_n} \frac{e^n(z)}{e^n(z) + d_1/B_n \mathbf{1}_{c_n \geq 1} + c_n^2(2 + v_n)}, & \text{for algorithm PASS,} \end{cases}$$

The constants L , L_n , B , B_n , d_1 and c_n are defined in Proposition 3. We write Γ for the set of processes $\tilde{\gamma} = (\tilde{\gamma}_n)_{n \geq 0}$ adapted to the filtration generated by the observed errors $(e^n)_{n \geq 0}$. We have the following result.

Proposition 4. *The sequence $\gamma = (\gamma_n)_{n \geq 0}$ defined in (18) satisfies*

$$e_\gamma^n \leq e_{\tilde{\gamma}}^n, \quad a.s., \quad \forall n \in \mathbb{N}, \forall \tilde{\gamma} \in \Gamma.$$

We write e_γ^n to point out the dependence of the error e^n on the chosen control γ .

The proof of the above result is given in Appendix H. Proposition 4 shows that γ ensures that fastest convergence speed of the average error and in particular

$$\mathbb{E}[e_\gamma^n] = \inf_{\tilde{\gamma} \in \Gamma} \mathbb{E}[e_{\tilde{\gamma}}^n], \quad \forall n \in \mathbb{N}.$$

This guarantees its optimality. Note that e^n is not known in practice because we do not have access to q^* . However, one can take the average value of $m(q_n, X_{n+1}(Z_n), Z_n)$ over the last $p \in \mathbb{N}^*$ visit times as a proxy of e^n since we aim at solving (1) and Assumption 3 guarantees that $e^n(Z_n) = 0$ as soon as $\mathbb{E}_n[m(q_n, X_{n+1}(Z_n), Z_n)] = 0$. Moreover, the constants L and B are also unknown in practice. To tackle this, a first solution consists in starting with arbitrary values for B and L and generating a sequence of learning rates. If the error $m(q_n, X_{n+1}(Z_n), Z_n)$ increases, we take a larger value for B and a smaller one for L otherwise he keeps B and L unchanged. A second solution consists in considering a piece-wise constant (PC) policy for the learning $(\gamma_k)_{k \geq 0}$. To do so, one can track the average error of $m(q_n, X_{n+1}(Z_n), Z_n)$ over the lasts p visit times with $p \in \mathbb{N}^*$ fixed by the controller. If this average error does not decrease, we reduce the step size by a factor α .

4.3 Extension

The results of this section still hold when the descent sequence (2) is replaced by

$$q_{n+1}(z) = q_n(z) - \gamma_n(z)m(q_n, X_{n+1}(z), z), \quad \forall z \in \mathcal{Z}, \forall n \in \mathbb{N}. \quad (19)$$

When $\gamma_n(z) = 0$ if $z \neq Z_n$, we recover (2). Thus, Equation (19) is slightly more general. Moreover, Algorithm (19) appears in many contexts: stochastic iterative algorithms, gradient methods, fixed point iterative techniques, etc. The results of this section can be directly transposed to (19). We present in Appendix B an adaptation of the Algorithm PASS for (19).

5 Some examples

5.1 Methodology

The code and numerical results presented in this section can be found in https://github.com/othmaneM/RL_adap_stepsize. Here, we compare four algorithms. The two first ones are two different versions of RL. In the first version, the learning rate $\gamma_k(z)$ is taken such that $\gamma_k = \eta/n_k(z)$ with $\eta > 0$ selected to provide the best convergence results and $n_k(z)$ the number of visits to the state z . In the second version, the step size follows the piece-wise constant policy (PC) described at the end of Section 4.2. The third algorithm is SAGA where the step size is derived from PC policy. Finally, we use the PASS algorithm presented in the previous sections with a predefined learning rate following the PC policy. We consider three numerical examples to compare the convergence speed of these algorithms: drift estimation, optimal placement of limit orders and the optimal liquidation of shares.

5.2 Drift estimation

Formulation of the problem. We observe a process $(S_n)_{n \geq 0}$ which satisfies

$$S_{n+1} = S_n + f_{n+1} + W_n, \quad (20)$$

with W_n a centred noise with finite variance. We want to estimate the quantities f_i with $i \in \{1, \dots, n_{max}\}$. Using (20) and $\mathbb{E}[W_t] = 0$, we get

$$\mathbb{E}[S_{i+1} - S_i - f_{i+1}] = 0, \quad \forall i \in \{0, \dots, n_{max} - 1\}.$$

Thus, we can estimate f_i using stochastic iterative algorithms. The pseudo-code of our implementation of PASS for this problem can be found in the Appendix C under the name Implementation 1.

Numerical results. Figure 2 shows the variation of the L^2 -error when the number of iterations increases. We can see that the algorithm PASS outperforms standard stochastic approximation algorithms. Moreover, other algorithms behave as expected: the standard RL decreases very slowly (but we know it will drive the asymptotic error to zero), the constant learning rate and SAGA provides

better results than RL, while **PASS** seems to have captured the best of the two worlds for this application: very fast acceleration at the beginning and the asymptotic error goes to zero.




Figure 2: The L^2 -error between f^k and f for different numerical methods averaged over 1000 simulated paths.

5.3 Optimal placement of a limit order

Formalisation of the problem. We consider an agent who aims at buying a unit quantity using limit orders and market orders (see [19] for detailed explanations). In such case, the agents wonder how to find the right balance between fast execution and avoiding trading costs associated to the bid-ask spread. The agent state at time t is modelled by $X_t = (Q^{Before}, Q^{After}, P)$ with Q^{Before} the number of shares placed before the agent’s order, Q^{After} the queue size after the agent’s order and P_t the mid price, see Figure 3. The agents wants to minimise the quantity

$$\mathbb{E}[F(X_{\tau \vee T^{\text{exec}}}) + \int_0^{\tau \vee T^{\text{exec}}} c \, ds]$$

where

- $T^{\text{exec}} = \inf\{t \geq 0, P_t = 0\}$ the first time when the limit order gets a transaction.
- τ the first time when a market order is sent.
- $X = (Q^{Before}, Q^{After}, P)$ the state of the order book.
- $F(u)$ is the price of the transaction (i.e. $F(u) = p + \psi$ when the agents crosses the spread and $F(u) = p$ otherwise).

We show in Section 2 that the Q -function is solution of (6), see details in https://github.com/othmaneM/RL_adap_stepsize. Thus, we can use Algorithms **RL**, **SAGA** and **PASS** to estimate it. The pseudo-code of our implementation of **PASS** is available as Implementation 2 in Appendix C.

Numerical results. Figure 4 shows three control maps: the x-axis reads the quantity on “same side” (i.e. $Q^{\text{same}} = Q^{Before} + Q^{After}$) and the y-axis reads the position of the limit order in the queue, i.e. Q^{Before} . The color and numbers gives the control associated to a pair $(Q^{\text{same}}, Q^{Before})$: 1 (blue) means “stay in the book”, while 0 (red) means “cross the spread” to obtain a transaction. The panel (at the left) gives the reference optimal controls obtained with a finite difference scheme, the middle panel the optimal controls obtained for a **RL** algorithm where the step-size $(\gamma_k)_{k \geq 0}$ is derived from the upper level policy, and the right panel the optimal control obtained with our optimal policy (i.e. upper level and inner level combined). It shows that after few iterations our optimal policy already found the optimal controls. Figure 5 compares the log of the L^2 error, averaged over 100 trajectories, between the different algorithms. We see clearly that our methodology improves basic stochastic approximation algorithms. Again, the other algorithms behave as expected: SAGA




Figure 3: The state space of our limit order control problem.

is better than a constant learning rate that is better than the standard RL (at the beginning, since we know that asymptotically RL will drive the error to zeros whereas a constant learning rate does not).




Figure 4: Comparison optimal control after 300 iteration for different methods: left is the optimal control, middle is **RL** with a step size derived from the upper level and right is our optimal policy for the step size (i.e. upper level and inner level combined).




Figure 5: The log L^2 -error against the number of iterations averaged over 1000 simulated paths.

5.4 Optimal execution

Formalisation of the problem. An investor wants to buy a given quantity q_0 of a tradable instrument (see [9] and [18] for details about this application). The price S_t of this instrument satisfies the

following dynamic:

$$dS_t = \alpha dt + \sigma dB_t,$$

where $\alpha \in \mathbb{R}$ is the drift and σ is the price volatility. The state of the investor is described by two variables its inventory Q_t and its wealth X_t at time t . The evolution of these two variables reads

$$\begin{cases} dQ_t = \nu_t dt, \\ dW_t = -\nu_t(S_t + \kappa \nu_t) dt, \end{cases} \quad \begin{matrix} Q_0 = q_0, \\ W_0 = 0, \end{matrix} \quad (21)$$

with ν_t the trading speed of the agent and $\kappa > 0$. The term $\kappa \nu_t$ corresponds to the temporary price impact. The investor wants to maximize the following quantity

$$W_T + Q_T(S_T - AQ_T) - \phi \int_t^T Q_s^2 ds,$$

it represents its final wealth X_T at time T , plus the value of liquidating its inventory minus a running quadratic cost. The value function V is defined such that

$$V(t, w, q, s) = \sup_{\nu} \mathbb{E}[W_T + Q_T(S_T - AQ_T) - \phi \int_t^T Q_s^2 ds | W_t = w, Q_t = q, S_t = s].$$

We remark that $v(t, w, q, s) = V(t, w, q, s) - w - qs$ verifies

$$v(t, w, q, s) = \sup_{\nu} \mathbb{E}[\underbrace{(W_T - W_t) + (Q_T S_T - Q_t S_t)}_{=M_T^t} - AQ_T^2 - \phi \int_t^T Q_s^2 ds | W_t = w, Q_t = q, S_t = s].$$

Using (21), we can see that the variable M_T^t is independent of the initial values W_t , S_t and Q_t . This means that v is a function of only two variables the time t and the inventory q . The dynamic programming principle ensures that v satisfies

$$v(t, q) = \sup_{\nu} \mathbb{E}[M_{t+\Delta}^t - \phi \int_t^{t+\Delta} Q_s^2 ds + v(t + \Delta, Q_{t+\Delta}) | Q_t = q]. \quad (22)$$

We fix a maximum inventory $\bar{q}$. Let $k = (k_T, k_q) \in (\mathbb{N}^*)^2$, $\Delta = T/k_T$, $D_T = \{t_i^{k_T}; i \leq k_T\}$ and $D_q = \{q_i^{k_q}; i \leq k_q\}$ with $t_i^{k_T} = i\Delta$ and $q_i^{k_q} = -\bar{q} + 2i\bar{q}/k_q$. To estimate v we use the numerical scheme $(v_n^k)_{n \geq 1, k \in (\mathbb{N}^*)^2}$ defined such that

$$v_{n+1}^k(Z_n) = v_n^k(Z_n) + \gamma_n(Z_n) \left[\sup_{\nu \in A(Z_n)} \{M_{n+1}^\nu - \phi \Delta Q_n^2 + v_n^k(Z_{n+1}^\nu) - v_n^k(Z_n)\} \right],$$

with $Z_n = (n\Delta, Q_{n\Delta})$ and $A(Z_n) \in D_q$ is the set of admissible actions⁸. When the final time T is reached (i.e. $n = k_T$), we pick a new initial inventory from the set D_q and start again its liquidation. At a first sight, it is not clear that v_n^k approximates v . However, we show in Appendix I that v_n^k converges point-wise to v on $D_T \times D_q$ when $n \rightarrow \infty$ and $k \rightarrow \infty$. see Appendix C for a detailed implementation of the algorithm with the corresponding pseudo-code (as Implementation 3).

Numerical results. Figure 6 shows the value function v for different values of the elapsed time t and the remaining inventory Q_t . The panel (at the left) gives the reference value function. It is computed by following the same approach of [24]. The middle panel the value function obtained after 120 000 iterations for RL algorithm where the step-size $(\gamma_k)_{k \geq 0}$ is derived from the upper level of our optimal policy, and the right panel the value function obtained with our optimal policy (i.e upper level and inner level combined). It shows that our optimal strategy leads to better performance results. We also plot, in Figure 7, a simulated path for the variations of the log L^2 error for different algorithms. Here again, we notice that our methodology improves the basic RL algorithm and that the ordering of other approaches is similar to the one of the “drift estimation” approximation (i.e. SAGA and the constant learning rate are very similar).

⁸We do not allow controls that lead to states where the inventory exceeds $\bar{q}$.

a) Theoretical value function b) step_ cste value function c) PASS value function




Figure 6: Comparison value function between methods.




Figure 7: The log L^2 -error against the number of iterations averaged over 1000 simulated paths.

A Actions of the agent

We present here policies that encourage exploration and others that maximize the agent's decisions.

Exploration policies: Since it is important to visit the state space sufficiently enough, we propose to set the conditional distribution of the random variable A_k such that

$$\mathbb{P}[A_k = a | \mathcal{F}_k] = \frac{e^{\bar{\epsilon}_k(Z_k^a)}}{\sum_{a'} e^{\bar{\epsilon}_k(Z_k^{a'})}}, \quad \forall a \in \mathcal{A}, \quad (23)$$

with $Z_k^a = (k, U_k, a)$, $\bar{\epsilon}_k(Z_k^a) = \beta(Z_k^a) \sum_{a'} \epsilon_k(U_{k+1}^{Z_k^a}, a')$ and

$$\epsilon_k(z) = \begin{cases} |m(q_{r_1(z)}, X_{r_1(z)+1}(z), z)| & \text{when the state } z \text{ already visited,} \\ b & \text{otherwise,} \end{cases}$$

where $b > 0$ encourages the exploration, q_k satisfies (2) and $r_1(z)$ is the last observation time of the state z .

Optimal policies. To give more importance to the maximizing action, one may consider the following policy:

$$\mathbb{P}[A_k = a | \mathcal{F}_k] = \frac{e^{\beta_k(U_k)q_k(Z_k^a)}}{\sum_{a'} e^{\beta_k(U_k)q_k(Z_k^{a'})}}, \quad \forall a \in \mathcal{A}. \quad (24)$$

Any mixture of these two procedures can also be considered.

B Extension of the PASS Algorithm

First, we can adapt Algorithm PASS to (19) by considering a component-wise version which consists in applying Algorithm 3 introduced in Section 4.1.1 to each visited coordinate of q_n . Next, we also propose a more direct approach with the algorithm below.

Algorithm 4 (PASS for (19)). *We start with an arbitrary q_0 and define by induction q_k and $\hat{\gamma}_k$*

- If $\langle m(q_n, X_{n+1}), m(q_{n-1}, X_n) \rangle \geq 0$, then do

$$\begin{aligned} q_{n+1} &= q_n - h(\hat{\gamma}_n, \gamma_n) m(q_n, X_{n+1}), \\ \hat{\gamma}_{n+1} &= h(\hat{\gamma}_n, \gamma_n). \end{aligned}$$

- Else, do

$$\begin{aligned} q_{n+1} &= q_n - l(\hat{\gamma}_n, \gamma_n) m(q_n, X_{n+1}), \\ \hat{\gamma}_{n+1} &= l(\hat{\gamma}_n, \gamma_n). \end{aligned}$$

C Implementations

We give here the pseudo code used for each one of the three numerical examples considered in Section 5.

Drift estimation. We consider the following expression for the functions h and l :

$$h(\gamma, \gamma_{base}) = \min(\gamma + \gamma_{base}, 3\gamma_{base}), \quad l(\gamma, \gamma_{base}) = \max(\gamma - \gamma_{base}, \gamma_{base}), \quad \forall (\gamma, \gamma_{base}) \in \mathbb{R}_+.$$

We use Implementation 1 for the numerical experiments.

Optimal placement of limit orders. We consider the following expression for the functions h and l :

$$h(\gamma, \gamma_{base}) = \max(\min(\gamma + 2/3\gamma_{base}, 3\gamma_{base}), \gamma_{base}), \quad l(\gamma, \gamma_{base}) = \max(\gamma - 2/3\gamma_{base}, \gamma_{base}), \quad \forall (\gamma, \gamma_{base}) \in \mathbb{R}_+. \quad (25)$$

We use Algorithm 2 for the numerical experiments. Note that we do not need to send a market order to know our expected future gain.

Implementation 1 PAST Sign Search (PASS) for (RL) drift estimation problem

```
1: Algorithm parameters: step size  $(\gamma^o)_{n \geq 0} \in (0, 1]$ , number of episodes  $n$   
   initial guess  $q_0$ , past error value  $E_{past}$   
   Initialise  $\hat{\gamma}_0 = \gamma_0^o$   
2: for episode in  $1 : n$  do  
3:   for  $t \in \{0, \dots, n_{\max} - 1\}$  do  
4:     Observe  $\Delta X_{next} = S_{t+1} - S_t$   
5:     if the first visit time to  $t$  then  
6:        $q_0(t) \leftarrow q_0(t) - \hat{\gamma}_0(t)m(q_0, \Delta X_{next}, t)$   
7:     else if  $m(q_0, t, \Delta X_{next}) \times E_{past}(t) \geq 0$  then  
8:        $\hat{\gamma}_0(t) \leftarrow h(\hat{\gamma}_0(t), \gamma^o(t))$   
9:        $q_0(t) \leftarrow q_0(t) - \hat{\gamma}_0(t)m(q_0, \Delta X_{next}, t)$   
10:    else if  $m(q_0, t, \Delta X_{next}) \times E_{past}(t) < 0$  then  
11:       $\hat{\gamma}_0(t) \leftarrow l(\hat{\gamma}_0(t), \gamma^o(t))$   
12:       $q_0(t) \leftarrow q_0(t) - \hat{\gamma}_0(t)m(q_0, \Delta X_{next}, t)$   
13:    end if  
14:     $E_{past}(t) \leftarrow m(q_0, t, \Delta X_{next})$   
15:  end for  
16:  Save the norm  $\|E\|$  of the vector  $E_{past}(t)$ .  
17:  if the average value of  $\|E\|$  over the last  $w = 5$  episodes is not reduced by  $p = 1\%$  then  
18:     $\gamma^o(t) \leftarrow \max(\gamma^o(t)/2, 0.01)$  (this is done each  $w$  episodes)  
19:  end if  
20: end for
```

Implementation 2 PAST Sign Search (PASS) for (RL) optimal placement problem

```
1: Algorithm parameters: step size  $(\gamma^o)_{n \geq 0} \in (0, 1]$ , number of episodes  $n$   
   initial guess  $q_0$ , past error value  $E_{past}$   
   Initialise  $\hat{\gamma}_0 = \gamma_0^o$   
2: for episode in  $n$  do  
3:   Select initial state  $X_0$   
4:   for each step within episode do  
5:     Take the action stay in the order book  
6:     Observe the new order book state  $X_{next}$   
7:     for  $a \in \{0, 1\}$  do  
8:       if the first visit time to  $X_{next}$  then  
9:          $q_0(X_0, a) \leftarrow q_0(X_0, a) - \hat{\gamma}_0(X_0, a)m^a(q_0, X_{next}, X_0)$   
10:      else if  $m^a(q_0, X_0, X_{next}) \times E_{past}(X_0, a) \geq 0$  then  
11:         $\hat{\gamma}_0(X_0, a) \leftarrow h(\hat{\gamma}_0(X_0, a), \gamma^o(X_0, a))$   
12:         $q_0(X_0, a) \leftarrow q_0(X_0, a) - \hat{\gamma}_0(X_0, a)m^a(q_0, X_{next}, X_0)$   
13:      else if  $m^a(q_0, X_0, X_{next}) \times E_{past}(X_0, a) < 0$  then  
14:         $\hat{\gamma}_0(X_0, a) \leftarrow l(\hat{\gamma}_0(X_0, a), \gamma^o(X_0, a))$   
15:         $q_0(X_0, a) \leftarrow q_0(X_0, a) - \hat{\gamma}_0(X_0, a)m^a(q_0, X_{next}, X_0)$   
16:      end if  
17:       $E_{past}(X_0, a) \leftarrow m^a(q_0, X_{next}, X_0)$   
18:    end for  
19:     $X_0 \leftarrow X_{next}$   
20:  end for  
21:  Save the norm  $\|E\|$  of the vector  $E_{past}(t)$ .  
22:  if the average value of  $\|E\|$  over the last  $w = 40$  episodes is not reduced by  $p = 5\%$  then  
23:     $\gamma^o(t) \leftarrow \max(\gamma^o(t)/2, 0.01)$  (this is done each  $w$  episodes)  
24:  end if  
25: end for
```

Optimal execution of a large number of shares. To solve this problem we use the same functions h and l considered in the previous problem, see (25). Then, we apply Algorithm 3. In this problem, it is crucial to select actions according to the policy (23) in order to encourage exploration. The coefficient $\bar{\beta}$ used by the agent to select its actions is taken constant equal to $\bar{\beta} = 5$. We consider the same policy for all the tested algorithms.

Implementation 3 PAST Sign Search (PASS) for (RL) optimal execution problem

```

1: Algorithm parameters: step size  $(\gamma^o)_{n \geq 0} \in (0, 1]$ , number of episodes  $n$ 
   initial guess  $q_0$ , past error value  $E_{past}$ 
   Initialise  $\hat{\gamma}_0 = \gamma_0^o$ 
2: for episode in  $n$  do
3:   Select the initial inventory  $Q_0$ 
4:   for  $t \in \{0, \dots, n_T - 1\}$  do
5:     Observe the new price state  $S_{next}$  and set  $X_0 = (t, Q_0)$ 
6:     Observe the new price state  $S_{next}$ 
7:     if the first visit time to  $X_0$  then
8:        $q_0(X_0) \leftarrow q_0(X_0) - \hat{\gamma}_0(X_0)m(q_0, S_{next}, X_0)$ 
9:     else if  $m(q_0, S_{next}, X_0) \times E_{past}(X_0) \geq 0$  then
10:       $\hat{\gamma}_0(X_0) \leftarrow h(\hat{\gamma}_0(X_0), \gamma^o(X_0))$ 
11:       $q_0(X_0) \leftarrow q_0(X_0) - \hat{\gamma}_0(X_0)m(q_0, S_{next}, X_0)$ 
12:    else if  $m(q_0, S_{next}, X_0) \times E_{past}(X_0) < 0$  then
13:       $\hat{\gamma}_0(X_0) \leftarrow l(\hat{\gamma}_0(X_0), \gamma^o(X_0))$ 
14:       $q_0(X_0) \leftarrow q_0(X_0) - \hat{\gamma}_0(X_0)m(q_0, S_{next}, X_0)$ 
15:    end if
16:     $E_{past}(X_0) \leftarrow m(q_0, S_{next}, X_0)$ 
17:    Select an action  $A$  and observe  $Q_{next}$ 
18:     $Q_0 \leftarrow Q_{next}$ 
19:  end for
20:  Save the norm  $\|E\|$  of the vector  $E_{past}(t)$ .
21:  if the average value of  $\|E\|$  over the last  $w = 300$  episodes is not reduced by  $p = 1\%$  then
22:     $\gamma^o(t) \leftarrow \max(\gamma^o(t) - 0.01, 0.01)$  (this is done each  $w$  episodes)
23:  end if
24: end for

```

D Proof of Proposition 1

Proof of Proposition 1. Let $z \in \mathcal{Z}$. Standard uniform convergence results ensure that

$$\mathbb{E}[\sup_q |g(q) - g_n(q)|(z)|n(z)] \leq c \frac{1}{\sqrt{n(z) \wedge 1}}, \quad a.s.$$

with $c > 0$ a positive constant. Since the Markov chain $(Z_n)_{n \geq 1}$ is irreducible and the set $\mathcal{Z}$ is finite, the sequence $(Z_n)_{n \geq 1}$ is positive recurrent and we have

$$\frac{n(z)}{n} = \frac{\sum_{k=1}^n \mathbf{1}_{Z_k=z}}{n} \xrightarrow{n \rightarrow \infty} \mathbb{P}_\mu[Z_n = z] > 0 = p(z), \quad a.s.,$$

with μ the unique invariant distribution of $(Z_n)_{n \geq 1}$. Thus, we have

$$u_n(z) = \mathbb{E} \left[\sqrt{\frac{n}{n(z) \wedge 1}} \right] \xrightarrow{n \rightarrow \infty} \frac{1}{\sqrt{p(z)}} > 0.$$

This shows that $u_n(z)$ is bounded by a constant called $u_\infty(z)$ and ensures that

$$\mathbb{E}[\sup_q |g(q) - g_n(q)|(z)] \leq c_1(z) \frac{1}{\sqrt{n}},$$

with $c_1(z) = cu_\infty(z)$. The constant c_1 can be taken independent of z since $\mathcal{Z}$ is finite. $\square$

E Proof of Proposition 2

Proof of Proposition 2. Let $z \in \mathcal{Z}$. We follow the same approach used in the proof of Proposition 1 to get

$$v_n(z) = \mathbb{E} \left[\left(\frac{\log(\bar{n}(z))/\log(n)}{\bar{n}(z)/n} \right)^\beta \right] \leq \mathbb{E} \left[\left(\frac{1}{\bar{n}(z)/n} \right)^\beta \right] \xrightarrow{n \rightarrow \infty} \frac{1}{p(z)^\beta} > 0.$$

This shows that $v_n(z)$ is bounded by a constant $v_\infty(z)$ and ensures that

$$\mathbb{E}[\sup_q |g(q) - g_n(q)|(z)] \leq c_2(z) \frac{1}{\sqrt{n}},$$

with $c_2(z) = c' v_\infty(z)$. The constant c_2 can be taken independent of z since $\mathcal{Z}$ is finite. Using the same manipulations of (9), we complete the proof. $\square$

F Proof of Proposition 3

Proof of Proposition 3. Let $k \geq 0$, A the set $A = \{Z_k = z\}$ and $m(q_k) \in \mathbb{R}^{\mathcal{Z}}$ such that $m(q_k)(z') = m(q_k, X_{k+1}(z'), z')$ for any $z' \in \mathcal{Z}$. We split the proof in three steps. In each one of these steps, we prove (12) for a given algorithm.

Step (i): In this step, we prove (12) for Algorithm RL. We have

$$\begin{aligned} \mathbf{1}_A \mathbb{E}_k[(q_{k+1} - q^*)^2(z)] &= \mathbf{1}_A \mathbb{E}_k[(q_k - \gamma_k m(q_k) - q^*)^2(z)] \\ &= \mathbf{1}_A \left\{ (q_k - q^*)^2(z) - \underbrace{2\gamma_k \mathbb{E}_k[m(q_k)](q_k - q^*)(z)}_{=(i)} + \gamma_k^2 \underbrace{\mathbb{E}_k[m(q_k)^2(z)]}_{=(ii)} \right\}. \end{aligned}$$

Using Assumption 3 and $\mathbb{E}_k[m^*] = 0$, we get (i) $\leq -2L\gamma_k|q_k - q^*|^2(z)$. Since $\mathbb{E}_k[m^*] = 0$, Assumption 4 gives

$$\begin{aligned} (ii) &= \mathbb{E}_k[(m(q_k) - \mathbb{E}_k[m(q_k)])^2(z)] + (\mathbb{E}_k[m(q_k) - m^*])^2(z) \\ &\leq B(1 + v_k) + B(1 + (q_k - q^*)^2(z)). \end{aligned}$$

Thus, we deduce that

$$\mathbf{1}_A \mathbb{E}_k[(q_{k+1} - q^*)^2(z)] \leq \mathbf{1}_A \left\{ \underbrace{(1 - 2\gamma_k L + B\gamma_k^2)}_{=-p(\gamma_k)} (q_k - q^*)^2(z) + \underbrace{\gamma_k^2(z)B(2 + v_k)}_{=M_k} \right\},$$

which shows (12) for Algorithm RL.

Step (ii): Here we show (12) for Algorithm SAGA. Let $\bar{M}^k(z) = (\sum_{j=1}^M M^k[z, j])/M$. Using $\mathbb{E}_k[m^*] = 0$ and $\mathbb{E}_k[M^k[z, i]] = \bar{M}^k(z)$, we have

$$\begin{aligned} \mathbf{1}_A \mathbb{E}_k[(q_{k+1} - q^*)^2(z)] &= \mathbf{1}_A \left\{ (q_k - q^*)^2(z) + 2(\mathbb{E}_k[q_{k+1}] - q_k)(q_k - q^*)(z) + \mathbb{E}_k[(q_{k+1} - q_k)^2(z)] \right\} \\ &= \mathbf{1}_A \left\{ (q_k - q^*)^2(z) - 2(\gamma_k \mathbb{E}_k[m(q_k) - m^*])(q_k - q^*)(z) \right. \\ &\quad \left. + \gamma_k^2 \mathbb{E}_k[(m(q_k)(z) - M^k[z, i] + \bar{M}^k(z))^2] \right\} \\ &\stackrel{\text{Assumption 3}}{\leq} \mathbf{1}_A \left\{ (1 - 2L\gamma_k)(q_k - q^*)^2(z) + \gamma_k^2 \underbrace{\mathbb{E}_k[(m(q_k)(z) - M^k[z, i] + \bar{M}^k(z))^2]}_{=(1)} \right\}. \end{aligned} \tag{26}$$

We first dominate the term (1). Since $\mathbb{E}_k[m^*](z) = 0$ and

$$\mathbb{E}_k[(M^k[z, i] - m^*)^2(z)] = 1/M \sum_j \mathbb{E}_k[(M^k[z, j] - m^*(z))^2],$$

we have

$$\begin{aligned}
(1) &= \mathbb{E}_k \left| (m(q_k) - \mathbb{E}_k[m^*])(z) - (M^k[z, i] - m^*(z)) + (1/M \sum_j (M^k[z, j] - m^*(z))) \right|^2 \\
&\leq 3 \left[\mathbb{E}_k [m(q_k) - m^*]^2(z) + \mathbb{E}_k [(M^k[z, i] - m^*(z))^2] + \mathbb{E}_k \left[(1/M \sum_j (M^k[z, j] - m^*(z)))^2 \right] \right] \\
&\stackrel{\text{Jensen's inequality}}{\leq} 3 \left[\mathbb{E}_k [(m(q_k) - m^*)^2(z)] + \mathbb{E}_k [(M^k[z, i] - m^*)^2(z)] + 1/M \sum_j \mathbb{E}_k [(M^k[z, j] - m^*(z))^2] \right] \\
&= 3 \left[\mathbb{E}_k [(m(q_k) - \mathbb{E}_k[m(q_k)])^2(z)] + (\mathbb{E}_k[m(q_k) - m^*](z))^2 + 2/M \sum_j \mathbb{E}_k [(M^k[z, j] - m^*(z))^2] \right] \\
&\stackrel{\text{Assumption 4}}{\leq} 3 \left[B(2 + v_k + (q_k - q^*)^2(z)) + 2/M \sum_j \mathbb{E}_k [(M^k[z, j] - m^*(z))^2] \right] \\
&= 3B(2 + v_k) + 3B(q_k - q^*)^2(z) + 6/M \sum_j \mathbb{E}_k [(M^k[z, j] - m^*(z))^2]. \tag{27}
\end{aligned}$$

By combining (26) and (27), we get

$$\begin{aligned}
\mathbf{1}_A \mathbb{E}_k [|q_{k+1} - q^*|^2(z)] &\leq \mathbf{1}_A \left\{ (1 - 2\gamma_k(z)L + 3B\gamma_k^2(z))(q_k - q^*)^2(z) \right. \\
&\quad \left. + 6/M \sum_j \mathbb{E}_k [(M^k[z, j] - m^*(z))^2] + 3B\gamma_k^2(v_k + 2)(z) \right\}. \tag{28}
\end{aligned}$$

Moreover, we have

$$\begin{aligned}
\mathbf{1}_A 1/M \mathbb{E}_k \left[\sum_{j=1}^M (M^{k+1}[z, j] - m^*(z))^2 \right] &= \mathbf{1}_A \left\{ \frac{1}{M} \mathbb{E}_k [(m(q_k) - m^*)^2(z)] \right. \\
&\quad \left. + (1 - \frac{1}{M}) \frac{1}{M} \sum_{j=1}^M \mathbb{E}_k [(M^k[z, j] - m^*(z))^2] \right\} \\
&\leq \mathbf{1}_A \left\{ \frac{B}{M} (1 + (q_k - q^*)^2(z)) \right. \\
&\quad \left. + (1 - \frac{1}{M}) \frac{1}{M} \sum_{j=1}^M \mathbb{E}_k [(M^k[z, j] - m^*(z))^2] \right\}. \tag{29}
\end{aligned}$$

Thus using (28) and (29), we conclude

$$\begin{aligned}
\mathbf{1}_A \mathbb{E}_k [e^{k+1}(z)] &\leq \mathbf{1}_A \left\{ \underbrace{(1 - 2\gamma_k L + 3B\gamma_k^2 + \frac{B}{Mc})}_{=\alpha_1} c(q_k - q^*)^2(z) \right. \\
&\quad \left. + \underbrace{(1 - \frac{1}{M} + 6\gamma_k^2(z)c)}_{\alpha_2} \frac{1}{M} \sum_j \mathbb{E}_k [(M^k[z, j] - m^*(z))^2] \right\} + 3B\gamma_k^2(v_k + 2)(z) \\
&\leq \mathbf{1}_A \alpha e^k(z) + 3B\gamma_k^2(v_k + 2)(z),
\end{aligned}$$

with $\alpha = \max(\alpha_1, \alpha_2) \in [0, 1)$.

Step (iii): In this final step, we show (12) for Algorithm PASS. We have

$$\begin{aligned} \mathbf{1}_A \mathbb{E}_k[(q_{k+1} - q^*)^2(z)] &= \mathbf{1}_A \mathbb{E}_k[(q_k - \hat{\gamma}_k m(q_k) - q^*)^2(z)] \\ &= \mathbf{1}_A \left\{ (q_k - q^*)^2(z) - \underbrace{2\mathbb{E}_k[\hat{\gamma}_k m(q_k)](q_k - q^*)}_{=(i)}(z) + \underbrace{\mathbb{E}_k[\hat{\gamma}_k^2 (m(q_k))^2(z)]}_{=(ii)} \right\}. \end{aligned}$$

For the term (i), using Assumption 3 and $\mathbb{E}_k[m^*] = 0$, we have $(i) \leq -2c_k \bar{\gamma}_k (q_k - q^*)^2(z)$ with $c_k = \frac{\mathbb{E}_k[\hat{\gamma}_k m(q_k)(z)]}{\bar{\gamma}_k \mathbb{E}_k[m(q_k)(z)]}$. Using Assumption 4 and $\mathbb{E}_k[m^*] = 0$, we get

$$\begin{aligned} (ii) &= c_k^2 \bar{\gamma}_k^2 \mathbb{E}_k[m(q_k)^2(z)] = c_k^2 \bar{\gamma}_k^2 \left(\mathbb{E}_k[(m(q_k) - \mathbb{E}_k[m(q_k)])^2(z)] + (\mathbb{E}_k[m(q_k) - m^*])^2(z) \right) \\ &\leq c_k^2 \bar{\gamma}_k^2 \left(B_k(1 + v_k) + B_k(1 + (q_k - q^*)^2(z)) \right). \end{aligned}$$

Thus, we deduce that

$$\mathbf{1}_A \mathbb{E}_k[(q_{k+1} - q^*)^2(z)] \leq \mathbf{1}_A \left\{ \underbrace{(1 - 2c_k \bar{\gamma}_k L_k + B_k(c_k \bar{\gamma}_k)^2)}_{=-p_k(c_k \bar{\gamma}_k)} (q_k - q^*)^2(z) + \underbrace{c_k^2 \bar{\gamma}_k^2(z) B_k(2 + v_k)}_{=M_k} \right\}. \quad (30)$$

We write $\underline{\gamma}_k$ for the quantity $\underline{\gamma}_k = c_k \bar{\gamma}_k \wedge \bar{\gamma}_k$. Since $\hat{\gamma}_k \in [\gamma_k, r_1 \gamma_k]$, we have $c_k \bar{\gamma}_k \in [\gamma_k, r_1 \gamma_k]$. When $c_k \bar{\gamma}_k \in [\gamma_k, \bar{\gamma}_k]$, we have $p_k(\underline{\gamma}_k) = p_k(c_k \bar{\gamma}_k) > p_k(\gamma_k) \geq 0$. When $c_k \bar{\gamma}_k \in [\bar{\gamma}_k, r_1 \gamma_k]$ (i.e $c_k \geq 1$), we use the following canonical decomposition of the function p_k :

$$p_k(x) = B_k(x - \bar{\gamma}_k)^2 - \left(\frac{L_k}{B_k} - 1 \right),$$

and $\underline{\gamma}_k = \bar{\gamma}_k$ to get

$$|p_k(c_k \bar{\gamma}_k) - p_k(\underline{\gamma}_k)| = |p_k(c_k \bar{\gamma}_k) - p_k(\bar{\gamma}_k)| \leq B_k(c_k \bar{\gamma}_k - \bar{\gamma}_k)^2 \leq B_k(r_1 \gamma_k - \gamma_k)^2 = \gamma_k^2 d^1,$$

with $d^1 = (r_1 - 1)^2 B_k$. Thus, using (30), we conclude

$$\mathbf{1}_A \mathbb{E}_k[(q^{k+1} - q^*)^2(z)] \leq \mathbf{1}_A \left\{ (1 - p_k(\underline{\gamma}_k) + \gamma_k^2 d^1 \mathbf{1}_{c_k \geq 1}) (q_k - q^*)^2(z) + \underbrace{c_k^2 \bar{\gamma}_k^2(z) B_k(2 + v_k)}_{=M_k} \right\}.$$

This completes the proof. $\square$

G Proof of Theorem 1

G.1 Preparation for the proof of Theorem 1

We introduce the following notations. Let $j \in \mathbb{N}^*$ and $(\mu_n, a_n, b_n)_{n \geq 1}$ be a sequence valued in $\mathbb{R}_+^3$. We write $(\mu, b)^j = (\mu_n^j, b_n^j)_{n \geq 1}$ for the delayed sequence $\mu_n^j = \mu_{j+n}$ and $b_n^j = b_{n+(j-1)}$ with $n \geq 1$. Additionally, we define recursively the sequence $(a_n^{\mu, b})_{n \geq 1}$ as follows:

$$a_1^{\mu, b} = 1, \quad \text{and} \quad a_{n+1}^{\mu, b} = \mu_n \sum_{l=1}^n a_{n+1-l} b_l a_l^{\mu, b}, \quad \forall n \geq 1. \quad (31)$$

Lemma 1. *By convention, an empty sum is equal to zero. Let $(v_n)_{n \geq 1}$ be the sequence defined as follows:*

$$v_n = \epsilon_n + \mu_n \left(\sum_{j=1}^{n-1} a_{n-j} b_j v_j \right), \quad \forall n \geq 1,$$

where $(\epsilon_n)_{n \geq 1}$ is a sequence. Then, we have

$$v_n = \sum_{j=1}^n a_{n+1-j}^{\mu, b} \epsilon_j, \quad \forall n \geq 1. \quad (32)$$

Proof of Lemma 1. Let us prove the result by induction on $n \geq 1$. By definition, Equation (32) is satisfied for $n = 1$. By applying the induction hypothesis (32) to all $j \leq n$, we get

$$\begin{aligned}
v_{n+1} &= 1 \times \epsilon_{n+1} + \mu_{n+1} \left(\sum_{j=1}^n a_{n+1-j} b_j v_j \right) = a_1^{(\mu,b)^{n+1}} \epsilon_{n+1} + \mu_{n+1} \left(\sum_{j=1}^n a_{n+1-j} b_j \sum_{l=1}^j a_{j+1-l}^{(\mu,b)^l} \epsilon_l \right) \\
&= a_1^{(\mu,b)^{n+1}} \epsilon_{n+1} + \left[\sum_{l=1}^n \mu_{n+1} \left(\sum_{j=l}^n a_{n+1-j} b_j a_{j+1-l}^{(\mu,b)^l} \right) \epsilon_l \right] \\
&= a_1^{(\mu,b)^{n+1}} \epsilon_{n+1} + \left[\sum_{l=1}^n \mu_{n+1} \left(\sum_{j=1}^{n-l+1} a_{n-l+2-j} b_{j+l-1} a_j^{(\mu,b)^l} \right) \epsilon_l \right] \\
&= a_1^{(\mu,b)^{n+1}} \epsilon_{n+1} + \sum_{l=1}^n a_{n+2-l}^{(\mu,b)^l} \epsilon_l = \sum_{j=1}^{n+1} a_{n+2-j}^{(\mu,b)^l} \epsilon_j.
\end{aligned}$$

□

Lemma 2. Let $n \in \mathbb{N}$, $(a_n^{\mu,b})_{n \geq 0}$ be the sequence defined in (31), $(\mu_n)_{n \geq 0}$ be a positive non-decreasing sequence, and $r_n = 1 - \mu_n$.

- When $\sum_{n \geq 0} r_n = +\infty$, we have

$$a_n^{\mu,b} \xrightarrow{n \rightarrow \infty} 0. \quad (33)$$

- There exists a non-negative constant B such that

$$a_n^{\mu,b} \leq B \sum_{k=1}^{n-1} \bar{\mu}_k^n a_{n-k}^{*\infty}, \quad \forall n \geq 2, \quad (34)$$

with $\bar{\mu}_k^n = e^{-\sum_{j=n-k+1}^n r_j}$ and $a_k^{*\infty} = \lim_{n \rightarrow \infty} a_k^{*n}$. The sequence $(a_k^{*n})_{k \geq 1}$ is defined recursively such that $a_k^{*1} = a_k$ and $a_k^{*(n+1)} = \sum_{l=1}^k a_{k+1-l} b_l a_l^{*n}$ for all $k \geq 1$.

Proof of Lemma 2. • *Step 1.* We first prove by induction the following subsidiary inequality:

$$a_n^{\mu,b} \leq \sum_{k=1}^n \bar{\mu}_k^n \tilde{a}_{n-k}^n, \quad \forall n \geq 2, \quad (35)$$

where $(\tilde{a}_k^n)_{k \geq 0}$ is a non-negative sequence that does not depend on $(\mu_n)_{n \geq 0}$. We also prove that Inequality (35) is optimal since it becomes an equality when the sequence $(\mu_n)_{n \geq 0}$ is constant. Inequality (35) holds clearly for $n = 2$ with $\tilde{a}^2 = a$. Using (31) and the induction assumption, we get

$$\begin{aligned}
a_{n+1}^{\mu,b} &= \mu_n \sum_{l=1}^n a_{n+1-l} b_l a_l^{\mu,b} \leq \mu_n b_1 a_n + \mu_n \sum_{l=2}^n a_{n+1-l} b_l \sum_{k=1}^{l-1} \bar{\mu}_k^l \tilde{a}_{l-k}^l \\
&\leq \mu_n b_1 a_n + \mu_n \sum_{l=2}^n a_{n+1-l} b_l \sum_{k=1}^{l-1} \bar{\mu}_k^n \tilde{a}_{l-k}^l = (i) + (ii), \quad (36)
\end{aligned}$$

where the inequality $\bar{\mu}_k^l \leq \bar{\mu}_k^n$ for any $l \leq n$ comes from the monotonicity assumption on $(\mu_n)_{n \geq 0}$. Note that $\bar{\mu}_k^l = \bar{\mu}_k^n$ when the sequence $(\mu_n)_{n \geq 0}$ is constant. In such case, we can replace all the previous inequalities by equalities. The inequality (36) gives

$$\begin{aligned}
(ii) &= \sum_{k=1}^{n-1} \mu^{n+1} \bar{\mu}_k^n \sum_{l=k+1}^n a_{n+1-l} b_l \tilde{a}_{l-k}^l = \sum_{k=1}^{n-1} \bar{\mu}_{k+1}^{n+1} \sum_{l=1}^{n-k} a_{n-k+1-l} b_{l+k} \tilde{a}_l^{k+l} = \sum_{k=1}^{n-1} \bar{\mu}_{k+1}^{n+1} \tilde{a}_{n-k}^{n+1} \\
&= \sum_{k=2}^n \bar{\mu}_k^{n+1} \tilde{a}_{n+1-k}^{n+1},
\end{aligned}$$

with

$$\tilde{a}_k^{n+1} = \sum_{l=1}^k a_{k+1-l} b_{n-k+l} \tilde{a}_l^{n-k+l}, \quad \text{when } k < n, \quad \tilde{a}_n^{n+1} = a_n, \quad (37)$$

and $\tilde{a}_k^{n+1} = 0$ otherwise. This proves (35).

• *Step 2.* Here, we show that $\sum_{m \geq 1} \tilde{a}_m^n \leq 1$ for all $n \geq 2$. Let us take $\mu_n^* = 1$ for all $n \geq 0$. In such case, Inequality (35) is optimal which gives $a_n^{\mu^*} = \sum_{k=1}^n \tilde{a}_{n-k}^n$. One can also show using a direct induction and (31) that $a_n^{\mu^*} \leq 1$. Thus, we have $\sum_{m \geq 1} \tilde{a}_m^n \leq 1$.

• *Step 3.* Here, we prove (33). Note that the condition $\sum_{n \geq 0} r_n = +\infty$ ensures that $\bar{\mu}_k^n \xrightarrow{n \rightarrow \infty} 0$. Since $\sum_{m \geq 1} \tilde{a}_m^n \leq 1$, we can then apply the monotone convergence theorem to (35) and get (33).

• *Step 4.* Here, we prove by induction that

$$\tilde{a}_k^\infty = a_k^{*\infty}, \quad \forall k \geq 1, \quad (38)$$

with $\tilde{a}_k^\infty = \lim_{n \rightarrow \infty} \tilde{a}_k^n$ and $a_k^{*\infty}$ defined in (34). Using (37) and the definition of a_1^{*n} , we check directly that $\tilde{a}_1^n = \tilde{a}_1^{*n}$ for any $n \geq 2$. This proves $\tilde{a}_1^\infty = a_1^{*\infty}$ and shows (38) for $k = 1$. Since $\sum_{l \leq n} a_{n+1-l} b_l \leq 1$ and $(\tilde{a}_l^n)_{n \geq 2, l \geq 1}$ is bounded by 1, see Step 2, we can apply the dominated convergence theorem to (37) and get

$$\tilde{a}_k^\infty = \sum_{l=1}^k a_{k+1-l} b_l \tilde{a}_l^\infty + a_1 b_k \tilde{a}_k^\infty. \quad (39)$$

When $a_1 b_k = 1$, we have $a_j b_{k'} = 0$ for any $k' \neq k$ or $j \neq 1$. Thus, a direct induction applied to (37) shows that $\tilde{a}_k^n = 0$ for any $k \geq 2$. By sending n to $+\infty$, we find $\tilde{a}_k^\infty = 0$ for any $k \geq 2$ which guarantees (38). We can then assume that $a_1 b_k < 1$ for all $k \geq 2$. Under this assumption, Equation (39) reads

$$\tilde{a}_k^\infty = \sum_{l=1}^{k-1} a_{k+1-l}^* b_l \tilde{a}_l^\infty, \quad (40)$$

with $a_l^* = a_l / (1 - a_1 b_k)$. The same lines of arguments show that $(a_k^{*\infty})_{k \geq 1}$ also satisfies (40). We can then apply the induction assumption to complete the proof (38).

• *Step 5.* Finally, we show (34). Using (38), we have $\lim_{n \rightarrow \infty} \bar{\mu}_{n-k}^n \tilde{a}_k^n / \bar{\mu}_{n-k}^n a_k^{*\infty} = 1$ for all $k \geq 1$. This means that $(1 - \epsilon) \bar{\mu}_{n-k}^n a_k^{*\infty} \leq \bar{\mu}_{n-k}^n \tilde{a}_k^n \leq (1 + \epsilon) \bar{\mu}_{n-k}^n a_k^{*\infty}$ for any $\epsilon > 0$ when n becomes large enough. By summing the previous inequality over all the possible values of k we get

$$\lim_{n \rightarrow \infty} \frac{\sum_{k=1}^n \bar{\mu}_{n-k}^n \tilde{a}_k^n}{\sum_{k=1}^n \bar{\mu}_{n-k}^n a_k^{*\infty}} = 1. \quad (41)$$

We combine (35) and (41), to show $\lim_{n \rightarrow \infty} \frac{a_n^\mu}{\sum_{k=1}^n \bar{\mu}_{n-k}^n \tilde{a}_k^{*\infty}} \leq \lim_{n \rightarrow \infty} \frac{\sum_{k=1}^n \bar{\mu}_{n-k}^n \tilde{a}_k^n}{\sum_{k=1}^n \bar{\mu}_{n-k}^n \tilde{a}_k^{*\infty}} = 1$.

Thus, there exists $B \geq 0$ such that $a_n^\mu \leq B \sum_{k=1}^n \bar{\mu}_{n-k}^n \tilde{a}_k^{*\infty}$ for all $n \geq 2$ which proves (34). $\square$

G.2 Propagation of the error

We introduce the following notations. Let $n \in \mathbb{N}^*$ and $(\tilde{z}^1, \dots, \tilde{z}^n) \in \mathcal{Z}^n$. For any $i \in \{1, \dots, n\}$, we denote by A^i the set $A^i = \{Z_1 = \tilde{z}^1, \dots, Z_i = \tilde{z}^i\}$. Moreover, we write $\tau_{\tilde{z}^1}^n = \inf\{l > 0, Z_{l+n} = z_1\}$ and $a_k^n = \mathbb{P}[\tau_{\tilde{z}^1}^{n-k} \geq k | \mathcal{F}_n]$. We have the following result.

Proposition 5. *Let $z_1 \in \mathcal{Z}$ and $n \in \mathbb{N}^*$. Under Assumptions 1, 3, 4 and 5, we have*

$$\mathbb{E}_0[e^n(z_1) \mathbf{1}_{A^n}] \leq \epsilon_n + \sum_{j=1}^{n-1} a_{n-j}^n b_j \alpha_j \mathbb{E}_0[e^j(z_1) \mathbf{1}_{A^j}],$$

with $\epsilon_n = e^1(z_1) a_n^0(z_1) \mathbb{E}[\mathbf{1}_{A^n}] + \sum_{j=1}^{n-1} a_{n-j}^n \mathbb{E}[M_j \mathbf{1}_{A^j}]$ and $b_j = \mathbf{1}_{\tilde{z}^j = z_1}$. The variables α_j and M_j are given by (13) and (14).

Proof of Proposition 5. Using the last-exit decomposition, see Section 8.2.1 in [21], we have

$$\begin{aligned}
\mathbb{E}_z[e^n(z_1)\mathbf{1}_{A^n}] &= \mathbb{E}_z[e^n(z_1)\mathbf{1}_{\tau_{Z_1}^0 \geq n, A^n}] + \sum_{j=1}^{n-1} \mathbb{E}_z[e^n(z_1)\mathbf{1}_{\{\tau_{Z_1}^0 \geq n-j, Z_j=z_1, A^n\}}] \\
&= \mathbb{E}_z[e^1(z_1)\mathbf{1}_{\tau_{Z_1}^0 \geq n, A^n}] + \sum_{j=1}^{n-1} \mathbb{E}_z[e^{j+1}(z_1)\mathbf{1}_{\{\tau_{Z_1}^0 \geq n-j, Z_j=z_1, A^n\}}] \\
&\leq \mathbb{E}_z[e^1(z_1)\mathbf{1}_{\tau_{Z_1}^0 \geq n, A^n}] + \sum_{j=1}^{n-1} \mathbb{P}[\tau_{Z_1}^{n-k} \geq k|\mathcal{F}_n]\mathbf{1}_{Z_j=z_1} \mathbb{E}_z[e^{j+1}(z_1)\mathbf{1}_{A^j}] \\
&\stackrel{\text{Proposition 3}}{\leq} \mathbb{E}_z[e^1(z_1)\mathbf{1}_{\tau_{Z_1}^0 \geq n, A^n}] + \sum_{j=1}^{n-1} \mathbb{P}[\tau_{Z_1}^{n-k} \geq k|\mathcal{F}_n]\mathbf{1}_{Z_j=z_1} \mathbb{E}_z[(\alpha_j e^j(z_1) + M_j)\mathbf{1}_{A^j}] \\
&= \epsilon_n + \sum_{j=1}^{n-1} a_{n-j}^n b_j \alpha_j \mathbb{E}_z[e^j(z_1)\mathbf{1}_{A^j}],
\end{aligned}$$

with ϵ_n , a_j^n , and b_j defined in Proposition 5. The variables α_n and M_n are defined in Proposition 3. In the second equality, we use that $e^n(z_1)$ does not change as long as the state z_1 is not reached. This completes the proof. $\square$

G.3 Proof of Theorem 1

Proof of Theorem 1. We split the proof in two steps. In Step (i), we show (15) and then in Step (ii) we show (16).

Step (i): In this part, we first prove (15) when the space $\mathcal{Z}$ is finite. Then, we show how to extend (15) to the general case.

Sub-step (i-1): Here we demonstrate (15) when the space $\mathcal{Z}$ is finite. Let $z_1 \in \mathcal{Z}$ and $n \in \mathbb{N}^*$. For any sequence $\mathbf{z} = (z_k)_{k \in \mathbb{N}^*} \in \mathcal{Z}$ and set $A_{\mathbf{z}}^n = \{Z_1 = \mathbf{z}_1, \dots, Z_n = \mathbf{z}_n\}$, we define v_n by $v_n = \alpha_n \mathbb{E}[e^n(z_1)\mathbf{1}_{A_{\mathbf{z}}^n}]$. Using Proposition 5, the sequence v_n verifies

$$v_n \leq \bar{\epsilon}_n + \alpha_n \sum_{j=1}^{n-1} a_{n-j}^n b_j v_j(z_1),$$

with $\bar{\epsilon}_n(z_1) = \alpha_n \epsilon_n$. Thus, using Lemma 1, we get

$$v_n = \sum_{j=0}^n a_{n-j}^{(\mu, b)^j} \bar{\epsilon}_j,$$

with $(a_n^{(\mu, b)})_{n \geq 1}$ defined in (31). We can assimilate the sequence $\bar{\epsilon}$ to the measure $\mu = \sum_{k \geq 1} \bar{\epsilon}_k \delta_k$ with δ_k the Dirac measure at k . Since $\epsilon_n = e^1 a_n^0 + \sum_{k=1}^{n-1} a_{n-k}^n b_k \mathbb{E}[M_k]$, $\sum_k a_k < \infty$, $\sum_k \mathbb{E}[\gamma_k^2] < \infty$, and $\mathbb{E}[M_k] = O(\mathbb{E}[\gamma_k^2])$ for all the algorithms, we get $\sum_k \epsilon_k < \infty$ using the properties of the Cauchy product between two sequences. This property ensures that the measure μ has a finite mass. Additionally, using Lemma 2, we have $a_n^{(\mu, b)^j} \xrightarrow{n \rightarrow \infty} 0$ for any $j \geq 1$. Thus, the dominated convergence theorem ensures that $v_n \xrightarrow{n \rightarrow \infty} 0$. Since α_n converges to 1, it also means that $\mathbb{E}[e^n(z_1)\mathbf{1}_{A_{\mathbf{z}}^n}] \xrightarrow{n \rightarrow \infty} 0$. Finally, we use the dominated convergence theorem to get

$$\lim_{n \rightarrow \infty} \mathbb{E}[e^n] = \lim_{n \rightarrow \infty} \mathbb{E}\left[\sum_{\mathbf{z} \in S} e^n \mathbf{1}_{A_{\mathbf{z}}^n}\right] \stackrel{\text{Fubini}}{=} \lim_{n \rightarrow \infty} \sum_{\mathbf{z} \in S} \mathbb{E}[e^n \mathbf{1}_{A_{\mathbf{z}}^n}] \stackrel{\text{dominate convergence}}{=} \sum_{\mathbf{z} \in S} \lim_{n \rightarrow \infty} \mathbb{E}[e^n \mathbf{1}_{A_{\mathbf{z}}^n}] = 0,$$

where S is the set of sequences valued in $\mathcal{Z}$ and $A_{\mathbf{z}}^n = \{Z_1 = \mathbf{z}_1, \dots, Z_n = \mathbf{z}_n\}$ for any $\mathbf{z} \in S$. Since the space $\mathcal{Z}$ is finite, we deduce that $\mathbb{E}[E^n] \xrightarrow{n \rightarrow \infty} 0$.

Sub-step(i-2): In this second step, we show (15) when the space $\mathcal{Z}$ is countable. Let $\epsilon > 0$. Since $E^1 = \sum_{k \geq 1} \mathbb{E}[e^1(z_k)] < \infty$, there exists $k_0 \in \mathbb{N}$ such that

$$\sum_{k \geq k_0} \mathbb{E}[e^1(z_k)] < \frac{\epsilon}{2}.$$

We write A_{k_0} for the set $A_{k_0} = \{z_k, k \leq k_0\}$. Since A_{k_0} is finite, we use **Sub-step(i-1)** to show the existence of $k_1 \in \mathbb{N}$ such that

$$\sum_{k \leq k_0} \mathbb{E}[e^{k_1}(z_k)] < \frac{\epsilon}{2}, \quad \forall k \geq k_1, \forall z_k \in A_{k_0}.$$

We take now $k \geq k_1$. Using $(\mathbb{E}[e^l(z)])_{l \geq 1}$ is non-increasing for any $z \in \mathcal{Z}$, we get $E^k = \sum_{k' \geq k_0} \mathbb{E}[e^k(z_{k'})] + \sum_{k' < k_0} \mathbb{E}[e^k(z_{k'})] \leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$.

Step (ii): In this step, we show (16). Using the same approach as in Proposition 5, we have

$$\mathbb{E}[e^n(z_1)] = \tilde{\epsilon}_n + \sum_{j=1}^{n-1} \tilde{a}_{n-j}^n \mathbb{E}_0[\mathbf{1}_{Z_j=z_1} \alpha_j e^j(z_1)] = \tilde{\epsilon}_n + \sum_{j=1}^{n-1} \tilde{a}_{n-j}^n \tilde{b}_j \bar{\alpha}_j \mathbb{E}_0[e^j(z_1)],$$

with $\tilde{\epsilon}_n = e^1(z_1) \mathbb{P}[\tau_{\mathcal{Z}_1}^0 \geq n] + \sum_{j=1}^{n-1} \tilde{a}_{n-j}^n \mathbb{E}[M_j]$, $\tilde{a}_j^n = \mathbb{P}[\tau_{\mathcal{Z}_1}^{n-j} \geq j | Z_{n-j} = z_1]$, $\tilde{b}_j = \mathbb{P}[Z_j = z_1]$, and $\bar{\alpha}_j = \frac{\mathbb{E}_0[\mathbf{1}_{Z_j=z_1} \alpha_j e^j(z_1)]}{\tilde{b}_j \mathbb{E}_0[e^j(z_1)]}$. We define $(\tilde{v}_n)_{n \geq 1}$, $(\tilde{\epsilon}_n^1)_{n \geq 1}$, and $(\mu_n)_{n \geq 1}$ respectively as follows $\tilde{v}_n = \bar{\alpha}_n \mathbb{E}[e^n(z_1)]$, $\tilde{\epsilon}_n^1 = \bar{\alpha}_n \tilde{\epsilon}_n$, and $\mu_n = \bar{\alpha}_n$ for all $n \geq 1$. In such case, we have $\tilde{v}_n \leq \tilde{\epsilon}_n^1 + \mu_n \sum_{j=1}^{n-1} \tilde{a}_{n-j}^n \tilde{b}_j \tilde{v}_j$. Thus, we can use Lemma 1 and Lemma 2 to get

$$\tilde{v}_n \leq B \sum_{j=1}^n \sum_{l=1}^j \bar{\mu}_l^j \tilde{a}_{l-j}^{*\infty} \tilde{\epsilon}_j^1,$$

with $\bar{\mu}_j^l$, $\tilde{a}_{l-j}^{*\infty}$, and B are introduced in (34). Since α_n converges towards 1 almost surely, the variable $\bar{\alpha}_j$ is bounded from below and thus there exists a constant $B' \geq 0$ such that

$$\mathbb{E}[e^n(z_1)] \leq B' \sum_{l+j+i=n} \tilde{\epsilon}_j^1 \bar{\mu}_l^j a_i^{*\infty}. \quad (42)$$

Note that the right-hand side of (42) is not trivial. To see this, we assume that there exists a constant $L' \geq 0$, such that $\mathbb{E}[\alpha_j \mathbf{1}_{Z_j=z_1} e^j(z_1)] \leq L' \mathbb{E}[\alpha_j] \mathbb{E}[\mathbf{1}_{Z_j=z_1}] \mathbb{E}[e^j(z_1)]$. Thus, we have $\mu_j \leq L \mathbb{E}[\alpha_j]$. Given that $\sum_k \mathbb{E}[\gamma_k^2] \leq \infty$, we also have $\mathbb{E}[B \gamma_k^2] \leq \mathbb{E}[L \gamma_k]$ and $\mathbb{E}[B_k \gamma_k^2] \leq \mathbb{E}[L_k \gamma_k]$ for k large enough where the constant B , B_k , L , and L_k are introduced in Assumptions 3 and 4. We can then find L'' such that $r_j = (1 - \mu_j) \geq L'' \mathbb{E}[\gamma_j]$. Since $\sum_{j \geq 1} \gamma_j = +\infty$ almost surely, we deduce that $\sum_{j \geq 1} r_j = +\infty$. This ensures that the right-hand side of (42) converges towards 0. $\square$

H Proof of Proposition 4

Proof of Proposition 4. For simplicity, we prove the result for Algorithm 1 however the same arguments hold for Algorithm 3. Let us prove by induction on $n \in \mathbb{N}$ that

$$e_\gamma^n \leq e_\gamma^n, \quad a.s., \quad \forall \gamma \in \Gamma. \quad (43)$$

For $n = 0$, Inequality (43) is directly satisfied since the initial error e^0 does not depend on the choice of the learning rate. Let $\tilde{\gamma} \in \Gamma$. We assume now that $e_\gamma^n \leq e_\gamma^n, a.s.$ Using Equation (17) and the definition of $(\gamma_n)_{n \geq 0}$, we have

$$e_\gamma^{n+1} = \mathbf{1}_A(g(e_\gamma^n) + S_n) + \mathbf{1}_{A^c} e_\gamma^n, \quad (44)$$

with $g(x) = x - \frac{L^2 x^2}{2B(x + (2 + v_n))}$. The previous expression of g is obtained by minimising the function $y \mapsto (1 - 2Ly + By^2)e_\gamma^n + B(2 + v_n) \times y^2$. This means that

$$g(x) \leq (1 - 2Ly + By^2)x + M_n y^2, \quad \forall y \in \mathbb{R}, \forall x \in \mathbb{R}_+. \quad (45)$$

A study of the function g shows that it is non-decreasing on the interval $[0, x_2[$ with $x_2 = \arg \sup_{x \in \mathbb{R}_+} g(x)$ and $g(x) \leq x$ for any $x \in \mathbb{R}$. Thus, we have necessarily $\mathbf{1}_A e_\gamma^n \leq g(\mathbf{1}_A e_\gamma^{n-1}) \leq g(x_2) \leq x_2$ for any $n \geq 1$ and $\tilde{\gamma} \in \Gamma$. Note that when $x_2 = +\infty$, we do not need to assume that $\mathbf{1}_A S_n \leq 0$. Using that $e_\gamma^n \leq e_\gamma^n$ a.s., the monotonicity of g on $[0, x_2[$ and Equation (45), we get

$$\begin{aligned} \mathbf{1}_A e_\gamma^{n+1} &= \mathbf{1}_A g(\mathbf{1}_A e_\gamma^n) + \mathbf{1}_A S_n \leq \mathbf{1}_A g(\mathbf{1}_A e_\gamma^n) + \mathbf{1}_A S_n \\ &\leq (1 - 2L\tilde{\gamma}_n + B\tilde{\gamma}_n^2) \mathbf{1}_A e_\gamma^n + \mathbf{1}_A M_n \tilde{\gamma}_n^2 + \mathbf{1}_A S_n \\ &= \mathbf{1}_A e_\gamma^{n+1}, \quad \text{a.s.} \end{aligned} \quad (46)$$

We complete the proof by combining (44), (46) and the induction assumption. $\square$

I Proof of the convergence of $(v_n^k)_{n \geq 1, k \in (\mathbb{N}^*)^2}$

In this section we prove the following result.

Proposition 6. *The sequence $(v_n^k)_{n \geq 1, k \geq 1}$ converges point-wise towards v on $D_T \times D_q$ when $n \rightarrow \infty$ and $k \rightarrow \infty$.*

Proof of Proposition 6. We prove this result in three steps. First, we show that v can be approximated by a numerical scheme $\bar{v}^k$. Then, we replace $\bar{v}^k$ by another scheme v^k that also converges towards v . Finally, we show that v_n^k tends to v^k when $n \rightarrow \infty$.

Step (i): We start with our initial control problem where the agents may choose its trading speed at any time. It was studied by many authors, see for example [24], who show that the optimal trading speed verifies

$$\nu(t, q) = \frac{h_1(t) - qh_2(t)}{2\kappa}, \quad (47)$$

with $h_1 : [0, T] \rightarrow \mathbb{R}$ and $h_2 : [0, T] \rightarrow \mathbb{R}_+$ a positive function. This means that the optimal inventory follows:

$$Q'(t) = \frac{h_1(t) - Qh_2(t)}{2\kappa}, \quad \forall t \in [0, T],$$

with $Q(0) = q_0$ given. Thus Q verifies

$$Q(t) = q_0 e^{-\int_0^t \frac{h_2(s)}{2\kappa} ds} + e^{-\int_0^t \frac{h_2(s)}{2\kappa} ds} \int_0^t \frac{h_1(s)}{2\kappa} e^{-\int_0^s \frac{h_2(u)}{2\kappa} du} ds. \quad (48)$$

Using Equation (48) and basic inequalities, one can show that if we take a large enough $\bar{q} \in \mathbb{R}_+$ and place ourselves in $S_{\bar{q}} = [-\bar{q}, \bar{q}]$ then

$$Q(t) \in S_{\bar{q}},$$

when $Q(0) \in S_{\bar{q}}$. Hence, we rewrite the dynamic programming principle as follows:

$$v(t, q) = \sup_{\nu \in \bar{A}(t, q)} \mathbb{E} \left[M_{t+\Delta}^t - \phi \int_t^{t+\Delta} Q_s^2 ds + v(t + \Delta, Q_{t+\Delta}) \mid Q_t = q \right]. \quad (49)$$

where $\bar{A}(t, q) \subset S_{\bar{q}}$ is the set of admissible actions⁹. We can focus on the set $\bar{A}(t, q)$ instead of $\mathbb{R}$ since other controls are not optimal. Now, we approximate this problem in a classical way using the numerical scheme $\bar{v}^k$ defined such that

$$\bar{v}^k(n_t, n_q) = \sup_{\nu \in D_a} \mathbb{E} \left[M_{n_t+1}^{n_t} - \phi Q_{n_t \Delta_t}^2 \Delta + \bar{v}^k(n_t + 1, n_{q+1}^\nu) \mid Q_{n_t \Delta_t} = n_q \Delta_q \right], \quad \forall (n_t, n_q) \in D_T \times D_q,$$

with $M_{n_t+1}^{n_t} = M_{\Delta_t(n_t+1)}^{n_t \Delta_t}$, n_{q+1}^ν the index such that $Q_{(n_t+1)\Delta_t}^\nu = n_{q+1}^\nu \Delta_q$ and $\bar{A} = A \cap \{i \Delta_q, i \in \mathbb{Z}\}$. The convergence of $(\bar{v}^k)_{k \geq 1}$ towards v on the set $D_T \times D_q$ when $k \rightarrow \infty$ is standard.

⁹We only allow controls that lead to states where the inventory stays in $S_{\bar{q}}$.

Step (ii): We denote by v^k the numerical scheme

$$v^k(n_t, n_q) = \mathbb{E} \left[\sup_{\nu \in D_a} \{M_{n_t+1}^{n_t} - \phi Q_{n_t \Delta_t}^2 \Delta + v^k(n_t + 1, n_{q+1}^\nu)\} | Q_{n_t \Delta_t} = n_q \Delta_q \right], \quad \forall (n_t, n_q) \in D_T \times D_q.$$

Let us show that $\bar{v}^k$ and v^k have the same limit. For this, we use a backward recurrence. For the moment, we assume that $|\sup_\nu \mathbb{E}[M_{n_t+1}^{n_t}] - \mathbb{E}[\sup_\nu M_{n_t+1}^{n_t}]| \leq K\Delta_t^2$ and we will prove it at the end of Step (ii). We want to show that $|\bar{v}^k(n_t, n_q) - v^k(n_t, n_q)| \leq K(T-t)\Delta_t$ for all $(n_t, n_q) \in D_T \times D_q$. At the terminal time $\bar{v}^k$ and v^k coincide. We move now to the induction part. We have

$$\begin{aligned} |\bar{v}^k(n_t, n_q) - v^k(n_t, n_q)| &= \left| \sup_\nu \mathbb{E}[M_{n_t+1}^{n_t} - \phi q^2 \Delta_t + \bar{v}^k(n_t + 1, n_{q+1}^\nu)] \right. \\ &\quad \left. - \mathbb{E}[\sup_\nu M_{n_t+1}^{n_t} - \phi q^2 \Delta_t + v^k(n_t + 1, n_{q+1}^\nu)] \right| \\ &\stackrel{\substack{\leq \\ \bar{v} \text{ and } v \text{ are not random}}}{\leq} \left| \sup_\nu \mathbb{E}[M_{n_t+1}^{n_t}] - \mathbb{E}[\sup_\nu M_{n_t+1}^{n_t}] \right| + \left| \sup_\nu \{ \bar{v}^k(n_t + 1, n_{q+1}^\nu) - v^k(n_t + 1, n_{q+1}^\nu) \} \right| \\ &\leq \mathbb{E} \left[\sup_\nu \{ K\Delta_t^2 + K(T-t-\Delta_t)\Delta_t \} \right] \\ &= K(T-t)\Delta_t. \end{aligned}$$

In the third inequality, we use the induction assumption to complete the proof. Now, let us show that $|\sup_\nu \mathbb{E}[M_{n_t+1}^{n_t}] - \mathbb{E}[\sup_\nu M_{n_t+1}^{n_t}]| \leq K\Delta_t^2$. We write $t = n_t \Delta_t$, $\Delta Q = Q_{t+\Delta_t} - Q_t = \nu \Delta_t$, $\Delta S = S_{t+\Delta_t} - S_t$ and $\Delta \bar{S} = \int_t^{t+\Delta_t} (S_s - S_t) ds$. Thus, we have

$$\begin{aligned} M_{n_t+1}^{n_t} &= (W_{t+\Delta} - W_t) + (Q_{t+\Delta} S_{t+\Delta} - Q_t S_t) \\ &= -\nu \Delta \bar{S} - \Delta Q S_t - \kappa \nu^2 \Delta_t + Q_t \Delta S + \Delta Q S_t + \Delta Q \Delta S \\ &= -\nu \Delta \bar{S} - \kappa \nu^2 \Delta_t + Q_t \Delta S + \nu \Delta_t \Delta S. \end{aligned} \tag{50}$$

The above equation shows $\sup_\nu M_{n_t+1}^{n_t} = \frac{(\Delta_t \Delta S - \Delta \bar{S})^2}{4\kappa \Delta_t} + Q_t \Delta S$. Using $\mathbb{E}[\Delta S] = \alpha \Delta_t$ and

$$\mathbb{E}[\Delta \bar{S}] = \int_t^{t+\Delta_t} \alpha(s-t) ds = \alpha \Delta_t^2 / 2,$$

we get

$$\sup_\nu \mathbb{E}[M_{n_t+1}^{n_t}] = \frac{\alpha^2 \Delta_t^3}{16\kappa} + \alpha Q_t \Delta_t, \quad \text{and} \quad \mathbb{E}[\sup_\nu M_{n_t+1}^{n_t}] = \frac{\alpha^2 \Delta_t^3}{16\kappa} + \frac{\sigma^2 \Delta_t^2}{12\kappa} + \alpha Q_t \Delta_t.$$

Thus we deduce that $|\sup_\nu \mathbb{E}[M_{n_t+1}^{n_t}] - \mathbb{E}[\sup_\nu M_{n_t+1}^{n_t}]| \leq K\Delta_t^2$ with $K = \frac{\sigma^2}{12\kappa}$.

Step (iii): Theorem 1, proves that v_n^k converges towards v^k . Thus by composition we have v_n^k converges point-wise towards v when $n \rightarrow \infty$ and $k \rightarrow \infty$ which completes the proof. $\square$

References

- [1] Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [2] Richard Bellman and Stuart Dreyfus. Functional approximations and dynamic programming. *Mathematical Tables and Other Aids to Computation*, pages 247–251, 1959.
- [3] Albert Benveniste, Michel Métivier, and Pierre Priouret. Algorithmes adaptatifs et approximations stochastiques: théorie et applications à l'identification, au traitement du signal et à la reconnaissance des formes. 1987.
- [4] Albert Benveniste, Michel Métivier, and Pierre Priouret. *Adaptive algorithms and stochastic approximations*, volume 22. Springer Science & Business Media, 2012.
- [5] Dimitri P Bertsekas and John N Tsitsiklis. *Neuro-dynamic programming*, volume 5. Athena Scientific Belmont, MA, 1996.

- [6] Julius R Blum et al. Approximation methods which converge with probability one. *The Annals of Mathematical Statistics*, 25(2):382–386, 1954.
- [7] Léon Bottou and Olivier Bousquet. The tradeoffs of large scale learning. In *Advances in neural information processing systems*, pages 161–168, 2008.
- [8] Olivier Bousquet, Stéphane Boucheron, and Gábor Lugosi. Introduction to statistical learning theory. In *Summer School on Machine Learning*, pages 169–207. Springer, 2003.
- [9] Álvaro Cartea, Sebastian Jaimungal, and José Penalva. *Algorithmic and High-Frequency Trading (Mathematics, Finance and Risk)*. Cambridge University Press, October 2015.
- [10] Patrick L Combettes and Jean-Christophe Pesquet. Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, pages 185–212. Springer, 2011.
- [11] Corinna Cortes, Spencer Greenberg, and Mehryar Mohri. Relative deviation learning bounds and generalization with unbounded loss functions. *Annals of Mathematics and Artificial Intelligence*, 85(1):45–70, 2019.
- [12] Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in neural information processing systems*, pages 1646–1654, 2014.
- [13] Sébastien Gadat and Fabien Panloup. Optimal non-asymptotic bound of the ruppert-polyak averaging without strong convexity. *arXiv preprint arXiv:1709.03342*, 2017.
- [14] James M Hutchinson, Andrew W Lo, and Tomaso Poggio. A nonparametric approach to pricing and hedging derivative securities via learning networks. *The Journal of Finance*, 49(3):851–889, 1994. bibtex*[publisher=Wiley Online Library].
- [15] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in neural information processing systems*, pages 315–323, 2013.
- [16] Harold Kushner and G George Yin. *Stochastic approximation and recursive algorithms and applications*, volume 35. Springer Science & Business Media, 2003.
- [17] Harold Joseph Kushner and Dean S Clark. *Stochastic approximation methods for constrained and unconstrained systems*. Springer, 1978.
- [18] Charles-Albert Lehalle, Sophie Laruelle, Romain Burgot, Stéphanie Pelin, and Matthieu Lasnier. *Market Microstructure in Practice*. World Scientific publishing, 2018.
- [19] Charles-Albert Lehalle and Othmane Mounjid. Limit Order Strategic Placement with Adverse Selection Risk and the Role of Latency. *Market Microstructure and Liquidity*, 3, 2017.
- [20] Iuliia Manziuk and Olivier Guéant. Deep reinforcement learning for market making in corporate bonds: beating the curse of dimensionality, June 2019.
- [21] Sean P Meyn and Richard L Tweedie. *Markov chains and stochastic stability*. Springer Science & Business Media, 2012.
- [22] Eric Moulines and Francis R Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems*, pages 451–459, 2011.
- [23] Atsushi Nitanda. Stochastic proximal gradient descent with acceleration techniques. In *Advances in Neural Information Processing Systems*, pages 1574–1582, 2014.
- [24] Cardialaguet Pierre and Lehalle Charles-Albert. Mean Field Game of Controls and An Application To Trade Crowding. November 2016.
- [25] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- [26] Jerome Sacks. Asymptotic distribution of stochastic approximation procedures. *The Annals of Mathematical Statistics*, 29(2):373–405, 1958.
- [27] Mark Schmidt, Nicolas L Roux, and Francis R Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. In *Advances in neural information processing systems*, pages 1458–1466, 2011.

- [28] Justin Sirignano and Konstantinos Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375:1339–1364, 2018. bibtex*[publisher=Elsevier].
- [29] Richard S Sutton, Andrew G Barto, et al. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.
- [30] Alexander B Tsybakov et al. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.
- [31] Christopher John Cornish Hellaby Watkins. Learning from delayed rewards. 1989.
- [32] Paul J Werbos. Building and understanding adaptive systems: A statistical/numerical approach to factory automation and brain research. *IEEE Transactions on Systems, Man, and Cybernetics*, 17(1):7–20, 1987.