

Numerical method for optimal stopping of Piecewise deterministic Markov processes under partial observations

ANR-09-SEGI-004

Extra report of team CQFD on task 3.1



The aim of this task is to propose a numerical method for the optimal stopping of piecewise deterministic Markov processes when the only observations available are partial and noisy. The partners involved are INRIA CQFD. People involved are A. Brandejsky (CQFD), B. de Saporta (CQFD) and F. Dufour (CQFD).

1 Context

The aim of this paper is to investigate an optimal stopping problem under partial observation for piecewise-deterministic Markov processes (PDMP) both from the theoretical and numerical points of view. PDMP's have been introduced by Davis [10] as a general class of stochastic models. They form a family of Markov processes involving deterministic motion punctuated by random jumps. The motion depends on three local characteristics, the flow Φ , the jump rate λ and the transition measure Q , which selects the post-jump location. Starting from the point x , the motion of the process $(X_t)_{t \geq 0}$ follows the flow $\Phi(x, t)$ until the first jump time T_1 , which occurs either spontaneously in a Poisson-like fashion with rate $\lambda(\Phi(x, t))$ or when the flow hits the boundary of the state space. In either case, the location of the process at T_1 is selected by the transition measure $Q(\Phi(x, T_1), \cdot)$ and the motion restarts from X_{T_1} . We define similarly the time until the next jump and the next post-jump location and so on. One important property of a PDMP, relevant for the approach developed in this paper, is that its distribution is completely characterized by the discrete time Markov chain $(Z_n, S_n)_{n \in \mathbb{N}}$ where Z_n is the n -th post-jump location and S_n is the n -th inter-jump time. A suitable choice of the state space and local characteristics provides stochastic models covering a large number of applications such as operations research [10, section 33], reliability [12], neurosciences [19], internet traffic [9], finance [4]. This list of examples and references is of course not exhaustive.

In this paper, we consider an optimal stopping problem for a partially observed PDMP $(X_t)_{t \geq 0}$. Roughly speaking, the observation process $(Y_t)_{t \geq 0}$ is a point process defined through the embedded discrete time Markov chain $(Z_n, S_n)_{n \in \mathbb{N}}$. The inter-arrival times are given by $(S_n)_{n \in \mathbb{N}}$ and the marks by a noisy function of $(Z_n)_{n \in \mathbb{N}}$. For a given reward function g and a computation horizon $N \in \mathbb{N}$, we study the following optimal stopping problem

$$\sup_{\sigma \leq T_N} \mathbf{E} [g(X_\sigma)],$$

where T_N is the N -th jump time of the PDMP $(X_t)_{t \geq 0}$, σ is a stopping time with respect to the natural filtration $\mathfrak{F}^Y = (\mathfrak{F}_t^Y)_{t \geq 0}$ generated by the observations $(Y_t)_{t \geq 0}$. In some applications, it may be more appropriate to consider a fixed optimization horizon t_f rather than the random horizon T_N . This is a difficult problem with few references in the literature, see for instance [13] where the underlying process is not piecewise deterministic. Regarding PDMP's, this problem could be addressed using the same ideas as in [7]. It involves the time-augmented process (X_t, t) . Although this process is still a PDMP, its local characteristics may not have the same good properties as those of the original process leading to several new technical difficulties.

2 Approach

A general methodology to solve such a problem is to split it into two sub-problems. The first one consists in deriving the filter process given by the conditional expectation of X_t with respect to the observed information \mathfrak{F}_t^Y . Its main objective is to transform the initial problem into a completely observed optimal stopping problem where the new state variable is the filter process. The second

step consists in solving this reformulated problem, the new difficulty being its infinite dimension. Indeed, the filter process takes values in a set of probability measures.

Our work is inspired by [20] which deals with an optimal stopping problem under partial observation for a Markov chain with finite state space. The authors study the optimal filtering and convert their original problem into a standard optimal stopping problem for a continuous state space Markov chain. Then they propose a discretization method based on a quantization technique to approximate the value function. However, their method cannot be directly applied to our problem for the following main reasons related to the specificities of PDMPs.

Firstly, PDMPs are continuous time processes. Although the dynamics can be described by the discrete-time Markov chain $(Z_n, S_n)_{n \in \mathbb{N}}$, this optimization problem remains intrinsically a *continuous-time* optimization problem. Indeed, the performance criterion is maximized over the set of stopping times defined with respect to the *continuous-time* filtration $(\mathfrak{F}_t^Y)_{t \geq 0}$. Consequently, our problem cannot be converted into a fully discrete time problem.

Secondly, the distribution of a PDMP combines both absolutely continuous and singular components. This is due to the existence of forced jumps when the process hits the boundary of the state space. As a consequence the derivation of the filter process is not straightforward. In particular, the absolute continuity hypothesis **(H)** of [20] does not hold.

Thirdly, in our context the reformulated optimization problem is not standard, unlike in [20]. As already explained, this reformulated optimization problem combines *continuous-time* and *discrete-time* features. Consequently, this problem does not correspond to the classical optimal stopping problem of a discrete-time Markov chain. Moreover, it is different from the optimal stopping problem of a PDMP under complete observation mainly because the new state variables given by the Markov chain $(\Pi_n, S_n)_{n \geq 0}$ are not the underlying Markov chain of some PDMP. Therefore the results of the literature [11, 15, 20] cannot be used.

Finally, a natural way to proceed with the numerical approximation is then to follow the ideas developed in [11, 20] namely to replace the filter Π_n and the inter-jump time S_n by some finite state space approximations in the dynamic programming equation. However, a noticeable difference from [11] lies in the fact that the dynamic programming operators therein were Lipschitz continuous whereas our new operators are only Lipschitz continuous between some points of discontinuity. We overcome this drawback by splitting the operators into their restrictions onto their continuity sets. This way, we obtain not only an approximation of the value function of the optimal stopping problem but also an ϵ -optimal stopping time with respect to the filtration $(\mathfrak{F}_t^Y)_{t \geq 0}$ that can be computed in practice.

Our approximation procedure for random variables is based on quantization. There exists an extensive literature on this method. The interested reader may for instance consult [14, 18] and the references within. The quantization of a random variable X consists in finding a finite grid such that the projection \hat{X} of X on this grid minimizes some L^p norm of the difference $X - \hat{X}$. Roughly speaking, such a grid will have more points in the areas of high density of X . As explained for instance in [18, section 3], under some Lipschitz-continuity conditions, bounds for the rate of convergence of functionals of the quantized process towards the original process are available, which makes this technique especially appealing. Quantization methods have been developed recently in numerical probability or optimal stochastic control with applications in finance, see e.g. [18, 2, 3].

3 Results

We have given a recursive form for the filter of the state space of a PDMP. We have derived the dynamic programming equations for the partially observed optimal stopping problem for PDMPs

and we have proposed a numerical approximation for the value function and close to optimal stopping time. Our results were tested on an academic example.

4 Dissemination of results

This work was presented in an invited conference [6] and published in an international peer-reviewed journal [5].

References

- [1] E. Arjas, P. Haara, and I. Norros. Filtering the histories of a partially observed marked point process. *Stochastic Process. Appl.*, 40(2):225–250, 1992.
- [2] V. Bally and G. Pagès. A quantization algorithm for solving multi-dimensional discrete-time optimal stopping problems. *Bernoulli*, 9(6):1003–1049, 2003.
- [3] V. Bally, G. Pagès, and J. Printems. A quantization tree method for pricing and hedging multidimensional American options. *Math. Finance*, 15(1):119–168, 2005.
- [4] N. Bäuerle and U. Rieder. *Markov decision processes with applications to finance*. Universitext. Springer, Heidelberg, 2011.
- [5] Adrien Brandejsky, Benoîte de Saporta, and François Dufour. Optimal stopping for partially observed piecewise-deterministic markov processes. *Stochastic Processes and their Applications*, 123(8):3201–3238, 2013.
- [6] Adrien Brandejsky, Benoîte de Saporta, and François Dufour. Optimal stopping for partially observed piecewise deterministic markov processes. In *XIème Colloque Franco-Roumain de Mathématiques Appliquées*, Bucarest, Romania, 2012.
- [7] A. Brandejsky, B. de Saporta, and F. Dufour. Numerical method for expectations of piecewise-deterministic markov processes. *CAMCoS*, 7(1):63–104, 2012.
- [8] P. Brémaud. *Point processes and queues*. Springer Series in Statistics. Springer-Verlag, New York, 1981.
- [9] D. Chafaï, F. Malrieu, and K. Paroux. On the long time behavior of the TCP window size process. *Stochastic Process. Appl.*, 120(8):1518–1534, 2010.
- [10] M. H. A. Davis. *Markov models and optimization*, volume 49 of *Monographs on Statistics and Applied Probability*. Chapman & Hall, London, 1993.
- [11] B. de Saporta, F. Dufour, and K. Gonzalez. Numerical method for optimal stopping of piecewise deterministic Markov processes. *Ann. Appl. Probab.*, 20(5):1607–1637, 2010.
- [12] B. de Saporta, F. Dufour, H. Zhang, and C. Elegbede. Optimal stopping for the predictive maintenance of a structure subject to corrosion. *Journal of Risk and Reliability*, 226(2):169–181, 2012.
- [13] E. Edoli and W. J. Runggaldier. On optimal investment in a reinsurance context with a point process market model. *Insurance Math. Econom.*, 47(3):315–326, 2010.

- [14] R. M. Gray and D. L. Neuhoff. Quantization. *IEEE Trans. Inform. Theory*, 44(6):2325–2383, 1998. Information theory: 1948–1998.
- [15] U. S. Gugerli. Optimal stopping of a piecewise-deterministic Markov process. *Stochastics*, 19(4):221–236, 1986.
- [16] M. Ludkovski and S. O. Sezer. Finite horizon decision timing with partially observable Poisson processes. *Stoch. Models*, 28(2):207–247, 2012.
- [17] J.-Y. Ouvrard. *Probabilités 2*. Enseignement des mathématiques. Cassini, 2004.
- [18] G. Pagès, H. Pham, and J. Printems. Optimal quantization methods and applications to numerical problems in finance. In *Handbook of computational and numerical methods in finance*, pages 253–297. Birkhäuser Boston, Boston, MA, 2004.
- [19] K. Pakdaman, M. Thieullen, and G. Wainrib. Fluid limit theorems for stochastic hybrid systems with application to neuron models. *Adv. in Appl. Probab.*, 42(3):761–794, 2010.
- [20] H. Pham, W. Runggaldier, and A. Sellami. Approximation by quantization of the filter process and applications to optimal stopping problems under partial observation. *Monte Carlo Methods Appl.*, 11(1):57–81, 2005.