

On the Relevance of Stochastic Controls

Orlando Gomes*

Lisbon School of Accounting and Administration (ISCAL-IPL) and Business Research Unit of the Lisbon University Institute (BRU/LUI), Portugal

Overview

Many scientifically relevant problems share two salient features. First, they are dynamic; second, they involve uncertainty. Hence, it is natural that researchers would be concerned with the study of events taking place in stochastic dynamic settings. In these settings, the decision-maker will typically want to select optimal paths for an array of control variables in order to maximize or minimize the current value of a sequence of future expected outcomes. In this article, we defend the argument that exploring techniques and applications in the field of stochastic optimal control theory is vital for the advancement of applied science. Although solid steps have been taken, in the last few years, to consolidate the theory of stochastic controls and to make this an adequate tool to address many important problems in multiple fields of knowledge, further work is still necessary to accomplish new insights and to unveil new results in an area of extreme complexity, where the search for efficient paths is many times hampered by the high degree of underlying uncertainty.

The benchmark optimization problem

Stochastic optimal control problems are concerned with the inter-temporal optimization (maximization or minimization) of an objective function subject to one or more constraints that, in continuous time, acquire the form of stochastic differential equations. The objective function typically respects to the expected value of a sequence of utility levels that range from the initial date $t=0$ to some future horizon. In the case of economic problems, the horizon is commonly assumed as infinite and the future is discounted at a constant rate $\rho > 0$. Taking the autonomous case, for which time is not an explicit argument of the problem's functional, the objective function acquires the following shape,

$$U(0) = E_0 \left[\int_0^{+\infty} f(x(t), u(t)) \exp(-\rho t) dt \right] \quad (1)$$

Function $f: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ is assumed to be a real-valued continuous and differentiable function. Two categories of variables constitute the arguments of f , namely the state variables, $x(t) \in \mathbb{R}^n$ and the control variables, $u(t) \in \mathbb{R}^m$. State variables are those that have their laws of motion determined by the differential equations respecting to the problem's constraints; control variables are the ones that the decision-maker is able to control in order to pursue the specified dynamic goal.

The constraints underlying the optimization problem are, as mentioned above, stochastic differential equations. A generic specification is the following,

$$dx(t) = g(x(t), u(t))dt + \sigma(x(t), u(t))dB(t), x(0) = x_0 \text{ given } (2)$$

In equation (2), $g: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ and $\sigma: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^{n \times m}$. The

In control problems involving a constant discount of future rewards and an infinite horizon are designated as reinforcement learning problems.

term $B(t)$ is an m -dimensional stochastic process defined in a filtered probability space $(\dot{U}, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, P)$. Frequently, the stochastic process takes the form of a Wiener process or Brownian motion. The formal definition of Brownian motion is as follows:

If, for all $0 \leq s < t$, $B(t) - B(s)$ is independent of the stopping time σ -algebra \mathcal{F}_s and is normally distributed with mean 0 and variance-covariance matrix $(t-s)I$, with I an $m \times m$ identity matrix, then $B(t)$ is a Brownian motion. The two mentioned properties are presentable in the form:

$$E(B(t) - B(s) | \mathcal{F}_s) = 0, P - \text{a.s.} \quad (3)$$

$$E((B(t) - B(s))(B(t) - B(s))^T | \mathcal{F}_s) = (t-s)I, P - \text{a.s.} \quad (4)$$

Roughly speaking, one might say that a Brownian motion is a normally distributed stochastic process with stationary independent increments. By adding the Brownian motion to an otherwise multi-dimensional deterministic differential equation, the time trajectories of the problem's endogenous variables will no longer correspond to deterministic paths; instead, the trajectories will exhibit persistent fluctuations around the non-stochastic trend. Instead of a Brownian motion, one may consider other types of stochastic processes in order to build a stochastic differential equation. A popular alternative is a Poisson process. A Poisson process, also known as counting process or pure jump process, is a stochastic process $q(t)$ such that

$$dq(t) = \begin{cases} 0, & \text{with probability } 1 - \lambda dt \\ 1, & \text{with probability } \lambda dt \end{cases} \quad (5)$$

Parameter $\lambda > 0$ is designated arrival rate. The Poisson process implies that the state variables in the stochastic differential equations are subject to jumps at random dates. This kind of process is useful to model phenomena as technological progress, where uncertainty concerning the arrival date of new innovations and about the extent of such innovations exists. Other types of, more sophisticated, stochastic processes might also be included in the constraints of the optimal control problem, namely Lévy processes, which combine features of the Brownian motion and of the Poisson process².

For a systematic and thorough analysis of Lévy processes. Lévy processes are also designated by the term jump diffusions.

*Corresponding author: Orlando Gomes, Lisbon School of Accounting and Administration (ISCAL-IPL) and Business Research Unit of the Lisbon University Institute (BRU/LUI), Portugal, Tel: 351-933420915; E-mail: omgomes@iscal.ipl.pt

Received August 18, 2014; Accepted August 20, 2014; Published August 26, 2014

Citation: Gomes O (2014) On the Relevance of Stochastic Controls. J Appl Computat Math 3: 181. doi:10.4172/2168-9679.1000181

Copyright: © 2014 Gomes O. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Independently of the stochastic process that best describes the type of randomness associated to the evolution of the assumed state variables, the optimal control stochastic problem will typically correspond to the maximization of $U(0)$ as given by expression (1) subject to a series of n stochastic constraints, as those in equation (2). In its generic form, this is not a simple problem to solve. Nevertheless, under some simplifications and constraints, this framework is able to deliver meaningful results that answer challenging questions that are posed by events taking place in nature, in the society and in the economy.

A Brief Note on General Solution Techniques

The most common technique to approach and solve stochastic optimal control problems consists in the construction of the respective Hamilton-Jacobi-Bellman (HJB) equation, which applies to the optimization problem in its generic form.

Let

$$V(x) = \max_{u(t)} E_0 \left[\int_0^{+\infty} f(x(t), u(t)) \exp(-\rho t) dt \right] \quad (6)$$

Given $V(x)$, the HJB equation takes the form of a non-linear ordinary differential equation,

$$\rho V(x) = \max_{u(t)} E_0 \left[f(x(t), u(t)) + g(x(t), u(t)) V'(x) + \frac{1}{2} \sigma(x(t), u(t))^2 V''(x) \right] \quad (7)$$

HJB equations are not always tractable from an analytical point of view, namely when the dimensionality of the underlying system is high or when the problem involves nonlinear constraints. The general stochastic control problem, as presented above, is, in fact, intractable to solve, requiring an unreasonable quantity of computational resources. Common algorithms demand millions of iterations for a task to be learned in the context of such generic problem.

Recent relevant advancements in the treatment of stochastic controls have allowed, on one hand, for a deeper understanding of the implications of the general problem and, on the other hand, for a rigorous and detailed analysis of some meaningful particular cases. For instance, Horowitz [1] takes a more insightful look at the HJB equation and develops algorithms that are capable of dealing directly with the optimal solutions of high dimensional nonlinear systems. Previously, Kappen et al., Todorov et al., and Theodorou et al., [2-5] have discovered that particular assumptions on the structure of the dynamic system make it possible to transform the HJB equation into a linear equation, allowing for analytical tractability. Specifically, the work by Kappen [2,3] indicates that an efficient solution is attainable when the control problem is defined in a finite horizon, the control term is linear and additive and the cost of the control is quadratic. Another form of approaching stochastic optimal control problems in order to deliver important insights is through state space discretization, what transforms the problem into a Markov decision process. In discrete time, the problem becomes easier to deal with, since it avoids the consideration of partial differential equations, which are typically difficult to analyze and often imply the impossibility of determining explicit analytical results [6].

Some Specific Applications and Contributions

Stochastic control techniques may be applied to any inter-temporal optimization problem for which the decision-maker has the faculty of controlling partially the environment and chooses to attain a specific goal. In Kappen [7], control problems are associated with animal behavior. Living organisms, including human beings, employ cognitive resources to take decisions. Frequently, these are not static decisions,

what implies the need to revisit the problem at successive time periods. Therefore, there is a recurrent process of adaptation and learning. Control theory studies how living organisms optimize a sequence of actions to attain inter-temporal goals. Because future events are not known with certainty, the optimization process implies the assumption of a probabilistic model of the expected outcomes. At each time step, given the uncertain outcome, the agent must re-estimate the trajectories of the control variables, and in this way one might state that there is a close link between adaptation and learning, on one hand, and stochastic optimization, on the other. Besides individual decision-making that takes place at the level of the brain, stochastic control might also be applied to scenarios where multiple agents have to solve a task. Besides optimization in time, this setting requires also the optimal coordination of the agents' actions. An example of this type of problem and the exploration of the respective outcomes is offered by Wiegierinck et al. [8]. The problem is one in which there is a common goal but where agents act in a decentralized way, choosing the paths that desirably lead to an optimal distribution of the agents across a given number of targets.

The prototypical example concerning multi-agent optimal control presented in the mentioned paper respects to a problem of firemen allocation across a number of active fires. Stochastic control techniques allow encountering a solution that is optimal from the social point of view, i.e., a solution such that firemen do not tend to concentrate in the same fires; on the contrary, they are able to distribute themselves in the direction of different final locations just by observing the trajectories followed by others. This is an interesting problem because it can be straightforwardly adapted to many situations that arise in society and in the economy: although agents behave on their own interest, their ultimate goal can only be accomplished by maximizing the performance of the group or the team. For instance, one could extend this framework in order to explain how market relations are organized or how different individuals choose different transportation to commute in an urban area. In fact, optimal control allows dealing with every stochastic environment in which agents have to distribute themselves efficiently over a number of targets. Again, the complexity of the problem emerges from the underlying uncertainty: in a stochastic environment, a configuration that is apparently optimal at $t=0$ may no longer be optimal at a future date.

In the field of economics, one of the most popular applications of stochastic optimal control is the one proposed by Merton [9]. In Merton's model, a representative consumer selects efficient consumption and portfolio investment strategies over a long time horizon. Uncertainty is, in this case, associated to a risky asset for which the expected return and the volatility are estimated. This model basically consists in the adaptation of the Ramsey model, i.e., of the utility maximization problem of the representative agent facing a resource constraint, to a stochastic environment.

A related area where stochastic controls have been applied with success is neoclassical growth theory. In a stochastic growth model, the standard benchmark framework is modified in order to account for uncertainty in technological progress. As in the deterministic case, growth can be explained through the analysis of a capital accumulation differential equation, but now this becomes a stochastic differential equation Brock et al., and Merton et al., for the analytical treatment of this model [10,11].

Why Stochastic Controls?

As highlighted by Horowitz [1], in control theory the ultimate

objective is to direct the system under appreciation towards a specified goal. The system involves laws of motion for the state variables, which constrain the choices of the agent solving the problem. The agent has, however, at her disposal, a set of control variables that constitute signals that she may manipulate in order to achieve the intended goals. Control theory is, then, concerned with designing efficient and robust solutions for the problem at hand, i.e., it is concerned with defining optimal time trajectories for the available control variables. At a first glance, optimal control problems do not seem too difficult to approach. In theory, one knows the initial state of the system, the inter-temporal goal to fulfill and the constraints that are faced by the problem solver. The fundamental point, though, is that one is dealing with the future; the plan is set now, at $t=0$, for an horizon that starts at $t=0$ and extends to a pre-defined future date. Therefore, as emphasized by Kappen [7], the control problem is stochastic in nature. There is uncertainty associated with future outcomes and the best the agent can do is to compute the optimal trajectory of some control variable(s) contingent on how the system is supposed to evolve. It is the stochastic element that sophisticates the problem and makes it hard to approach.

A deterministic plan can be solved at the initial period and never revised, because nothing in the environment will presumably change. Stochastic controls require a permanent reevaluation of the dynamic conditions as these potentially depart from expected values. When evolving from determinism to stochasticity one loses in tractability but one surely gains in realism. The relevance of approaching plans through the lenses of stochastic optimal control is not exhausted in the determination of optimal paths. It is also important to guarantee their stability. Stabilization methods in the presence of uncertainty must also be an issue in mind when addressing this type of problems. This is also a concern that optimal control theorists have, as highlighted by Horowitz [1].

In synthesis, stochastic dynamic optimization problems are a significant part of the decision problems one finds in nature, in the society and in the economy. As a consequence, it is undoubtedly of extreme relevance to continue to build the tools and to explore the applications that somehow fall in the extremely rich scientific domain that is stochastic controls.

Literature

The modern theory of optimal control starts with Bellman et al. [12]; it is also Bellman [13] who presents the first results on stochastic controls. Since then, many important results have come to light. Today, there is a vast literature on stochastic optimal control, which is synthesized in some volumes that constitute the main references in this area. These include, just to cite a few, Kamien et al. [14], Yong and Zhou [15], Kendrick [16], Bertsekas [17,18] and Oksendal and Sulem [19]. Walde [20-22] offers a comprehensive presentation of dynamic optimization models in economics. In four parts, the book addresses deterministic optimization in discrete and continuous time and stochastic optimization in discrete and continuous time. Concerning stochastic optimal control problems in continuous time, it is explained, in this manuscript, in a simple mode, what a stochastic process is, how a stochastic process might integrate a stochastic differential equation and how the shape of the stochastic differential equation depends on the type of stochastic process (Brownian motion, Poisson process or Lévy process). A similar discussion may be found in Brito [23].

References

- Horowitz MB (2014) Efficient Methods for Stochastic Optimal Control. PhD thesis, Pasadena: California Institute of Technology.

- Kappen HJ (2005a) Linear Theory for Control of Nonlinear Stochastic Systems. *Physical Review Letters* 95: 200201.
- Kappen HJ (2005b) Path Integrals and Symmetry Breaking for Optimal Control Theory. *Journal of Statistical Mechanics: Theory and Experiment* 2005: P11011.
- Todorov E (2009) Efficient Computation of Optimal Actions. *Proceedings of the National Academy of Sciences* 28: 11478-11483.
- Theodorou EA, Todorov E (2012) Relative Entropy and Free Energy Dualities: Connections to Path Integral and KL Control. *Proceedings of the 51st IEEE Conference on Decision and Control (CDC)* 1466-1473.
- Todorov E (2006) Linearly-solvable Markov Decision Problems. *Advances in Neural Information Processing* 19: 1369-1376.
- Kappen HJ (2007) An Introduction to Stochastic Control Theory, Path Integrals and Reinforcement Learning. *Cooperative Behavior in Neural Systems* 887: 149-181.
- Wiegerinck W, Broek BVD, Kappen HJ (2012) Stochastic Optimal Control in Continuous Space-time Multi-agent Systems.
- Merton R (1971) Optimum Consumption and Portfolio Rules in a Continuous Time Model. *Journal of Economic Theory* 3: 373-413.
- Brock WA, Mirman LJ (1972) Optimal Economic Growth and Uncertainty: the Discounted Case. *Journal of Economic Theory* 4: 479-513.
- Merton R (1975) An Asymptotic Theory of Growth under Uncertainty. *Review of Economic Studies* 42: 375- 393.
- Bellman R (1952) On the Theory of Dynamic Programming. *Proceedings of the National Academy of Sciences* 38: 716-719.
- Bellman R (1958) Dynamic Programming and Stochastic Control Process. *Information & Control* 1: 228-239.
- Kamien MI, Schwartz NL (1991) Dynamic Optimization. 2nd edition, Amsterdam: North-Holland.
- Yong J, Zhou XY (1999) Stochastic Controls: Hamiltonian Systems and HJB Equations. *Applications of Mathematics -Stochastic Modelling and Applied Probability*, New York: Springer-Verlag.
- Kendrick DA (2002) Stochastic Control for Economic Models. 2nd edition, Austin, Texas: University of Texas.
- Bertsekas DP (2005) Dynamic Programming and Optimal Control. 3rd edition, Nashua, NH: Athena Scientific.
- Bertsekas DP (2012) Dynamic Programming and Optimal Control: Approximate Dynamic Programming 4th edition, Nashua, NH: Athena Scientific.
- Oksendal B, Sulem A (2009) Applied Stochastic Control of Jump Diffusions. 3rd edition, Berlin: Springer.
- Walde K (2011) Applied Intertemporal Optimization. Mainz University Gutenberg Press.
- Walde K (1999) Optimal Saving under Poisson Uncertainty. *Journal of Economic Theory* 87: 194-217.
- Walde K (2005) Endogenous Growth Cycles. *International Economic Review* 46: 867-894.
- Brito P (2008) Introduction to Dynamic Programming Applied to Economics. Lisbon: Lisbon Technical University.