



KTH Engineering Sciences

Contributions to the Stochastic Maximum Principle

DANIEL ANDERSSON

Doctoral Thesis
Stockholm, Sweden 2009

TRITA-MAT-09-MS-12
ISSN 1401-2286
ISRN KTH/MAT/DA 09/08-SE
ISBN 978-91-7415-436-8

KTH Matematik
SE-100 44 Stockholm
SWEDEN

Akademisk avhandling som med tillstånd av Kungl Tekniska högskolan framlägges till offentlig granskning för avläggande av teknologie doktorsexamen i matematisk statistik fredagen den 30 oktober 2009 klockan 13.00 i sal F3, Kungl Tekniska högskolan, Valhallavägen 79, Stockholm.

© Daniel Andersson, september 2009

Tryck: Universitetsservice US AB

Abstract

This thesis consists of four papers treating the maximum principle for stochastic control problems.

In the first paper we study the optimal control of a class of stochastic differential equations (SDEs) of mean-field type, where the coefficients are allowed to depend on the law of the process. Moreover, the cost functional of the control problem may also depend on the law of the process. Necessary and sufficient conditions for optimality are derived in the form of a maximum principle, which is also applied to solve the mean-variance portfolio problem.

In the second paper, we study the problem of controlling a linear SDE where the coefficients are random and not necessarily bounded. We consider relaxed control processes, i.e. the control is defined as a process taking values in the space of probability measures on the control set. The main motivation is a bond portfolio optimization problem. The relaxed control processes are then interpreted as the portfolio weights corresponding to different maturity times of the bonds. We establish existence of an optimal control and necessary conditions for optimality in the form of a maximum principle, extended to include the family of relaxed controls.

The third paper generalizes the second one by adding a singular control process to the SDE. That is, the control is singular with respect to the Lebesgue measure and its influence on the state is thus not continuous in time. In terms of the portfolio problem, this allows us to consider two investment possibilities – bonds (with a continuum of maturities) and stocks – and incur transaction costs between the two accounts.

In the fourth paper we consider a general singular control problem. The absolutely continuous part of the control is relaxed in the classical way, i.e. the generator of the corresponding martingale problem is integrated with respect to a probability measure, guaranteeing the existence of an optimal control. This is shown to correspond to an SDE driven by a continuous orthogonal martingale measure. A maximum principle which describes necessary conditions for optimal relaxed singular control is derived.

Acknowledgements

First of all I wish to thank my supervisor Professor Boualem Djehiche for his support and for constantly suggesting new problems for me to work on. I also would like to thank my assistant supervisor Professor Lars Holst for support and for introducing me to the theory of weak convergence. I am also grateful to Professor Tomas Björk for providing me with very useful feedback in connection with the presentation of my licentiate thesis. Finally, many thanks to Professor Brahim Mezerdi for insightful comments on the first and second paper of the thesis.

Contents

Contents	v
Introduction	1
Optimal control theory	1
The maximum principle	2
Stochastic control	4
The stochastic maximum principle	5
Relaxed control	6
Summary of the papers	9
Paper I	9
Paper II	10
Paper III	11
Paper IV	12
Bibliography	13

List of papers

I: A maximum principle for SDEs of mean-field type (2009). *Joint work with Boualem Djehiche.*

II: A maximum principle for relaxed stochastic control of linear SDEs with application to bond portfolio optimization (2007). *Joint work with Boualem Djehiche.*

III: A mixed relaxed singular maximum principle for linear SDEs with random coefficients (2008).

IV: The relaxed general maximum principle for singular optimal control of diffusions (2009). *Systems & Control Letters, Volume 58, Issue 1, January 2009, 76–82.*

Introduction

In this section we give some background on optimal control theory, stochastic control and in particular the stochastic maximum principle. We also introduce and discuss the notion of relaxed control. We conclude by providing a summary of the four papers.

Optimal control theory

Optimal control theory can be described as the study of strategies to optimally influence a system x with dynamics evolving over time according to a differential equation. The influence on the system is modeled as a vector of parameters, u , called the control. It is allowed to take values in some set U , which is known as the action space. For a control to be optimal, it should minimize a cost functional (or maximize a reward functional), which depends on the whole trajectory of the system x and the control u over some time interval $[0, T]$. The infimum of the cost functional is known as the value function (as a function of the initial time and state). This minimization problem is infinite dimensional, since we are minimizing a functional over the space of functions $u(t), t \in [0, T]$. Optimal control theory essentially consists of different methods of reducing the problem to a less transparent, but more manageable problem. The two main methods are dynamic programming and the maximum principle.

As for dynamic programming, it is essentially a mathematical technique for making a sequence of interrelated decisions. By considering a family of control problems with different initial times and states and establishing relationships between them one obtains a nonlinear first-order partial differential equation known as the Hamilton-Jacobi-Bellman (HJB) equation. Solving this equation gives the value function, after which a finite dimensional maximization problem can be solved.

On the other hand, the maximum principle gives necessary conditions for optimality by perturbing an optimal control on a small time interval of length θ . Performing a Taylor expansion with respect to θ and then sending θ to zero one obtains a variational inequality. By duality the maximum principle is obtained. It states that any optimal control must solve the Hamiltonian system associated with the control problem. The Hamiltonian system is a linear differential equation

with terminal conditions, called the adjoint equation, and a (finite dimensional) maximization problem.

Optimal control theory and the maximum principle originate from the works of Pontryagin and co-authors (see e.g. Pontryagin et al. (1962)). The method of dynamic programming was developed by Bellman (see e.g. Bellman (1957)). For an introduction to optimal control theory, see e.g. Sontag (1990).

The maximum principle

As an illustration, we provide a sketch of how the maximum principle for a deterministic control problem is derived. In this setting, the state of the system is given by the differential equation

$$\begin{cases} dx(t) &= b(x(t), u(t)) dt, \\ x(0) &= x_0, \end{cases} \quad (1)$$

where $u(t) \in U$ for all $t \in [0, T]$, and the action space U is some subset of \mathbb{R} . The objective is to minimize some cost function

$$J(u) = \int_0^T h(x(t), u(t)) dt + g(x(T)). \quad (2)$$

That is, the function h inflicts a running cost and the function g inflicts a terminal cost. We now assume that there exists a control $\hat{u}(t)$ which is optimal, i.e.

$$J(\hat{u}) = \inf_u J(u).$$

We denote by $\hat{x}(t)$ the solution to (1) with the optimal control $\hat{u}(t)$. We are going to derive necessary conditions for optimality by analyzing what happens when we make a small perturbation of the optimal control. Therefore we introduce a so-called spike variation, i.e. a control which is equal to \hat{u} except on some small time interval:

$$u^\theta(t) = \begin{cases} v & \text{for } \tau - \theta \leq t \leq \tau, \\ \hat{u}(t) & \text{otherwise.} \end{cases} \quad (3)$$

We denote by $x^\theta(t)$ the solution to (1) with the control $u^\theta(t)$. We see that $\hat{x}(t)$ and $x^\theta(t)$ are equal up to $t = \tau - \theta$ and that

$$\begin{aligned} x^\theta(\tau) - \hat{x}(\tau) &= (b(x^\theta(\tau), v) - b(\hat{x}(\tau), \hat{u}(\tau)))\theta + o(\theta) \\ &= (b(\hat{x}(\tau), v) - b(\hat{x}(\tau), \hat{u}(\tau)))\theta + o(\theta), \end{aligned} \quad (4)$$

where the second equality holds since $x^\theta(\tau) - \hat{x}(\tau)$ is of order θ . Next, we look at the Taylor expansion of the state with respect to θ . Let

$$z(t) = \left. \frac{\partial}{\partial \theta} x^\theta(t) \right|_{\theta=0},$$

i.e. the Taylor expansion of $x^\theta(t)$ is

$$x^\theta(t) = \widehat{x}(t) + z(t)\theta + o(\theta). \quad (5)$$

Then, by (4),

$$z(\tau) = b(\widehat{x}(\tau), v) - b(\widehat{x}(\tau), \widehat{u}(\tau)). \quad (6)$$

Moreover, we can derive the following differential equation for $z(t)$.

$$\begin{aligned} dz(t) &= \left. \frac{\partial}{\partial \theta} dx^\theta(t) \right|_{\theta=0} \\ &= \left. \frac{\partial}{\partial \theta} b(x^\theta(t), u^\theta(t)) dt \right|_{\theta=0} \\ &= b_x(x^\theta(t), u^\theta(t)) \left. \frac{\partial}{\partial \theta} x^\theta(t) dt \right|_{\theta=0} \\ &= b_x(x^\theta(t), u^\theta(t)) z(t) dt, \end{aligned}$$

where b_x denotes the derivative of b with respect to x . If we for the moment assume that $h = 0$, the optimality of $\widehat{u}(t)$ leads to the inequality

$$\begin{aligned} 0 &\leq \left. \frac{\partial}{\partial \theta} J(u^\theta) \right|_{\theta=0} = \left. \frac{\partial}{\partial \theta} g(x^\theta(T)) \right|_{\theta=0} = g_x(x^\theta(T)) \left. \frac{\partial}{\partial \theta} x^\theta(T) \right|_{\theta=0} \\ &= g_x(\widehat{x}(T)) z(T). \end{aligned}$$

We shall use duality to obtain a more explicit necessary condition from this. To this end we introduce the adjoint equation:

$$\begin{cases} dp(t) &= -b_x(\widehat{x}(t), \widehat{u}(t)) p(t) dt, \\ p(T) &= g_x(\widehat{x}(T)). \end{cases}$$

Then it follows that

$$d(p(t)z(t)) = 0,$$

i.e. $p(t)z(t) = \text{constant}$. By the terminal condition for the adjoint equation we have

$$p(t)z(t) = g_x(\widehat{x}(T)) z(T) \geq 0, \quad \text{for all } 0 \leq t \leq T.$$

In particular, by (6)

$$p(\tau) (b(\widehat{x}(\tau), v) - b(\widehat{x}(\tau), \widehat{u}(\tau))) \geq 0.$$

Since τ was chosen arbitrarily, this is equivalent to

$$p(t)b(\widehat{x}(t), \widehat{u}(t)) = \inf_v p(t)b(\widehat{x}(t), v) \quad \text{for all } 0 \leq t \leq T.$$

This specifies a necessary condition for $\hat{u}(t)$ to be optimal when $h = 0$. To account for the running cost h one can construct an extra state $dx^0(t) = h(x(t), u(t)) dt$, which allows us to write the cost function in terms of two terminal costs:

$$J(u) = x^0(T) + g(x(T)).$$

By repeating the calculations above for this two-dimensional system, one can derive the necessary condition

$$H(\hat{x}(t), \hat{u}(t), p(t)) = \inf_v H(\hat{x}(t), v, p(t)), \text{ for all } 0 \leq t \leq T, \quad (7)$$

where H is the so-called Hamiltonian (sometimes defined with a minus sign which turns the minimum condition above into a maximum condition):

$$H(x, u, p) = h(x, u) + pb(x, u),$$

and the adjoint equation is given by

$$\begin{cases} dp(t) &= - (h_x(\hat{x}(t), \hat{u}(t)) + b_x(\hat{x}(t), \hat{u}(t)) p(t)) dt, \\ p(T) &= g_x(\hat{x}(T)). \end{cases} \quad (8)$$

The minimum condition (7) together with the adjoint equation (8) specifies the Hamiltonian system for our control problem.

Stochastic control

Stochastic control is the extension of optimal control to problems where it is of importance to take into account some uncertainty in the system. One possibility is then to replace the differential equation by an SDE:

$$dx(t) = b(t, x(t), u(t)) dt + \sigma(t, x(t)) dB(t), \quad (9)$$

where b and σ are deterministic functions and the last term is an Itô integral with respect to a Brownian motion B defined on a probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$. More generally, the diffusion coefficient σ may have an explicit dependence on the control:

$$dx(t) = b(x(t), u(t)) dt + \sigma(x(t), u(t)) dB(t). \quad (10)$$

The cost function for the stochastic case is the expected value of the cost function (2), i.e. we want to minimize

$$\mathbb{E} \left(\int_0^T h(x(t), u(t)) dt + g(x(T)) \right).$$

When solving these types of stochastic control problems with dynamic programming, the obtained HJB equation is a nonlinear second-order partial differential equation, see e.g. Fleming and Soner (2006). For an extensive survey of the developments in stochastic control theory see e.g. Yong and Zhou (1999) and Fleming and Soner (2006).

The stochastic maximum principle

The earliest paper on the extension of the maximum principle to stochastic control problems is Kushner and Schweppe (1964). One major difficulty that arises in such an extension is that the adjoint equation (c.f. (8)) becomes an SDE with terminal conditions. In contrast to a deterministic differential equation, one cannot simply reverse the time since the control process, and consequently the solution to the SDE, is required to be adapted to the filtration. Bismut solved this problem by introducing conditional expectations and obtained the solution to the adjoint equation from the martingale representation theorem, see e.g. Bismut (1978) and also Haussmann (1986). An extensive study of these so-called backward SDEs can be found in e.g. Ma and Yong (1999).

For the case (9) the adjoint equation is given by

$$\begin{cases} dp(t) = -(h_x(\hat{x}(t), \hat{u}(t)) + b_x(\hat{x}(t), \hat{u}(t))p(t) + \sigma_x(\hat{x}(t))q(t))dt \\ \quad - q(t)dB(t), \\ p(T) = g_x(\hat{x}(T)). \end{cases} \quad (11)$$

A solution to this kind of backward SDE is a pair $(p(t), q(t))$ which fulfills (11). The Hamiltonian is

$$H(x, u, p, q) = h(x, u) + pb(x, u) + q\sigma(x),$$

and the maximum principle reads

$$H(\hat{x}(t), \hat{u}(t), p(t), q(t)) = \inf_v H(\hat{x}(t), v, p(t), q(t)), \text{ for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s.} \quad (12)$$

For the stochastic maximum principle, there is a major difference between the cases (9) and (10). As for (10), when performing the expansion with respect to the perturbation θ (c.f. (5)), the fact that the perturbed Itô integral turns out to be of order $\sqrt{\theta}$ (rather than θ as with the ordinary Lebesgue integral) poses a problem. In fact, one needs to take into account both the first-order and second-order terms in the Taylor expansion (5). This ultimately leads to a maximum principle containing two adjoint equations, both in the form of linear backward SDEs. The Hamiltonian is replaced by an extended Hamiltonian:

$$\begin{aligned} \mathcal{H}^{\hat{x}(t), \hat{u}(t)}(t, x, v) = \\ H\left(t, x, v, p(t), q(t) - P(t)\sigma(t, \hat{x}(t), \hat{u}(t))\right) - \frac{1}{2}\sigma^2(t, \hat{x}(t), v)P(t), \end{aligned}$$

where $(p(t), q(t))$ is the solution to the first order adjoint equation (11) and $(P(t), Q(t))$ is the solution to the second order adjoint equation – see Peng (1990) where the first proof of this general stochastic maximum principle is given. The optimal control is in this case characterized by

$$\mathcal{H}^{\hat{x}(t), \hat{u}(t)}(t, \hat{x}(t), \hat{u}(t)) = \inf_v \mathcal{H}^{\hat{x}(t), \hat{u}(t)}(t, \hat{x}(t), v), \text{ for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s.}$$

There is also third case: if the state is given by (10) but the action space U is convex, it is possible to derive the maximum principle in a local form. This is accomplished by using a convex perturbation of the control instead of a spike variation, see Bensoussan (1982). The necessary condition for optimality is then the following.

$$\frac{d}{dv}H(t, \hat{x}(t), \hat{u}(t), \hat{p}(t), \hat{q}(t))(v - \hat{u}(t)) \geq 0, \text{ for all } 0 \leq t \leq T, \mathbb{P}\text{-a.s.}$$

Both methods – the spike and the convex perturbations – are used in this thesis.

Relaxed control

An important part of this thesis is the notion of relaxed control, which we will briefly explain next.

Since the maximum principle describes a necessary condition for optimality we know that if there exists an optimal control, there also exists at least one solution to the Hamiltonian system in the maximum principle. However, the control problem may fail to have an optimal solution even for quite simple situations. At least if by a solution we mean a process $u(t)$ taking values in U . The idea is then to compactify the space of controls by extending the definition of controls to include the space of probability measures on U . The set of relaxed controls $\mu_t(du)dt$, where μ is a probability measure, is the closure under weak* topology of the measures $\delta_{u(t)}(du)dt$ corresponding to usual, or strict, controls.

Here is a simple example from deterministic control. Consider the cost functional

$$J(u) = \int_0^T \left(x^2(t) + (1 - u^2(t))^2 \right) dt,$$

to be minimized over the set of controls $u : [0, T] \mapsto U = [-1, 1]$. The state of the system is given by

$$\begin{cases} dx(t) &= u(t)dt, \\ x(0) &= 0. \end{cases}$$

Now, consider the following sequence of controls

$$u^{(n)}(t) = (-1)^k \text{ if } t \in \left[\frac{k}{n}, \frac{k+1}{n} \right), \text{ } 0 \leq k \leq n-1.$$

Then we have $|x^{(n)}(t)| \leq 1/n$ and $|J(u^{(n)})| \leq T/n^2$, which implies that $\inf_u J(u) = 0$. The limit of $u^{(n)}$ is however not in the space of strict controls. Instead the sequence $\delta_{u^{(n)}(t)}(du)dt$ converges weakly to $1/2(\delta_{-1} + \delta_1)(du)dt$. Thus, there does not exist an optimal strict control in this case but only a relaxed one. But since we can construct a sequence of strict controls such that the cost functional is arbitrarily

close to its infimum, it is clear that there does exist an optimal solution, albeit in a wider sense.

This notion of relaxed control is introduced for deterministic optimal control problems in Young (1969), and generalized to the stochastic control problem (9) in Fleming (1977). A maximum principle for this relaxed control problem is established in Mezerdi and Bahlali (2002). As for the case (10), i.e. with controlled diffusion coefficient, the existence of an optimal relaxed control is proved in El Karoui et al. (1987), and a maximum principle for this problem is established in Bahlali et al. (2006). In these papers, control is defined as the solution to an appropriate martingale problem, and the relaxation is performed by integrating the generator with respect to a probability measure on U . In Ma and Yong (1995) another type of relaxation of the problem (10) is studied. The relaxed control problem is then defined by integrating the coefficients in the SDE with respect to a probability measure. It should be noted that this is not equivalent to the relaxation in El Karoui et al. (1987). Roughly, the difference is that in the former one integrates the diffusion coefficient $\sigma(\cdot, \cdot, u)$, and in the latter $\sigma(\cdot, \cdot, u)\sigma(\cdot, \cdot, u)^*$.

In the second and third paper of this thesis we study relaxed control problems of the form suggested in Ma and Yong (1995), where the more direct form of relaxation is used to formulate an optimal bond portfolio problem. In the fourth paper we obtain a maximum principle for a general relaxed singular control problem, where the relaxation of the problem is constructed as in El Karoui et al. (1987).

Summary of the papers

Paper I: A maximum principle for SDEs of mean-field type

We consider stochastic control problems where the state process is governed by an SDE of the so-called mean-field type. That is, the coefficients of the SDE depend on the marginal law of state as well as the state and the control. More specifically, the SDE is defined as

$$\begin{cases} dx(t) &= b(t, x(t), \mathbb{E}\psi(x(t)), u(t)) dt + \sigma(t, x(t), \mathbb{E}\phi(x(t)), u(t)) dB(t), \\ x(0) &= x_0, \end{cases}$$

for some functions b , σ , ψ and ϕ , and a Brownian motion $B(t)$. This mean-field SDE is obtained as the mean square limit of an interacting particle system of the form

$$\begin{aligned} dx^{i,n}(t) &= b\left(t, x^{i,n}(t), \frac{1}{n} \sum_{j=1}^n \psi(x^{j,n}(t)), u(t)\right) dt \\ &\quad + \sigma\left(t, x^{i,n}(t), \frac{1}{n} \sum_{j=1}^n \phi(x^{j,n}(t)), u(t)\right) dB^i(t), \end{aligned}$$

when $n \rightarrow \infty$. Here $\{B^i(t)\}_{i=1}^n$ are independent Brownian motions. That is, each “particle” is weakly interacting with the others, and they all have the same law. The classical example is the McKean-Vlasov model, see e.g. Sznitman (1989) and the references therein.

The cost functional is of the form

$$J(u) = \mathbb{E} \left(\int_0^T h(t, x(t), \mathbb{E}\varphi(x(t)), u(t)) dt + g(x(T), \mathbb{E}\chi(x(T))) \right),$$

for given functions h , g , φ and χ . This cost functional is also of mean-field type, as the functions h and g depend on the law of the state process. This leads to a so-called time inconsistent control problem – that is, the Bellman optimality principle does not hold – since one cannot apply the law of iterated expectations on the cost functional.

We derive necessary and sufficient conditions for optimality of this control problem in the form of a stochastic maximum principle. It turns out that the adjoint equation is a mean-field backward SDE. We apply the methods in Bensoussan (1982) to obtain the necessary conditions, i.e. we assume that the action space is convex, which allows us to make a convex perturbation of the optimal control and obtain a maximum principle in a local form.

We illustrate the result by applying it to the mean-variance optimization problem. That is, the continuous version of a Markowitz investment problem where one constructs a portfolio by investing in a risk free bank account and a risky asset. The objective is to maximize the expected terminal wealth while minimizing its variance. Since this cost function involves the variance – which is quadratic in the expected value – the problem is time inconsistent as explained above.

Paper II: A maximum principle for relaxed stochastic control of linear SDEs with application to bond portfolio optimization

Consider an investor with initial capital x_0 who invests in bonds, i.e. financial contracts that are bought today and pay a fixed amount at some future time, called the maturity time. If the interest rate is stochastic, so are the prices of the bonds since the interest rate of outstanding bonds can be higher or lower than its current market value.

Assume that there exists a market of bonds with time to maturity in some set $U \subset \mathbb{R}_+$. There is one major difference between such a bond market and a stock market. In the bond market there is naturally a continuum of assets, while in the standard model of a stock market there is only a finite number of assets. Thus, an appropriate definition of a portfolio strategy should include such portfolios that may contain a continuum of assets. Here we define a portfolio strategy as a measure valued process $\rho_t(du)$, reflecting the “number” of bonds in the portfolio at time t , which have time to maturity in the interval $[u, u + du]$. In the Heath-Jarrow-Morton-Musiela framework, the portfolio value $x(t)$ can be derived as

$$dx(t) = \int_U p_t(u) (r_t^0 - v_t(u)\Theta_t) \rho_t(du)dt + \int_U p_t(u)v_t(u)\rho_t(du)dB(t),$$

where $p_t(u)$ is the price of the bond with time to maturity equal to u , $v_t(u)$ is its volatility, r_t^0 the short rate and Θ_t is the so-called market price of risk. If the investor only takes long positions (positive portfolio weights) we can introduce the relative portfolio μ_t as a process taking values in the space of probability measures, by

$$\rho_t(du) = \frac{x(t)}{p_t(u)}\mu_t(du).$$

The relative portfolio, which is the proportion of the portfolio value invested in different bonds, has the following dynamics.

$$dx(t) = x(t) \int_U (r_t^0 - v_t(u)\Theta_t) \mu_t(du)dt + x(t) \int_U v_t(u)\mu_t(du)dB(t).$$

We consider this as a controlled SDE with action space U and relaxed control μ . Given a cost functional, we derive existence of an optimal relaxed control as well as necessary conditions for optimality, when the state equation is given by this type of linear SDEs with random and unbounded coefficients.

In the example above motivating the use of relaxed controls, we found an approximating strict control sequence which converges to the optimal relaxed control. This is actually true in general and the result is known as the chattering lemma. That is, we can always find a sequence of strict controls which are near optimal. We derive a maximum principle for these near optimal controls. Passing to the limit and using the chattering lemma as well as some stability properties of the state and adjoint processes, we obtain our main result: if $\hat{\mu}$ is an optimal relaxed control, it fulfills the usual maximization condition, however with the Hamiltonian extended to a functional of probability measures on U .

This paper originally appeared in my licentiate thesis.

Paper III: A mixed relaxed singular maximum principle for linear SDEs with random coefficients

We study an extension of the previous problem by adding a singular control component. The controlled process is the following two-dimensional SDE.

$$\begin{cases} dx(t) = \int_U b^x(t, x(t), u) \mu_t(du)dt + \int_U \sigma^x(t, x(t), u) \mu_t(du)dB(t) + G_t^x d\xi(t), \\ dy(t) = b^y(t, y(t)) dt + \sigma^y(t, y(t)) dB(t) + G_t^y d\xi(t). \end{cases}$$

The SDE for $x(t)$ is of the type studied in paper II, i.e. linear with random coefficients, with an added singular component (the SDE for $y(t)$ is linear with constant coefficients). The singular control component is the process ξ , which is increasing and càglàd (french acronym for “left continuous with right limits”). I.e., as opposed to the control u , its influence on the state is not continuous in time. We refer to u as the absolutely continuous part of the control.

As for the optimal portfolio application, the absolutely continuous control corresponds to the continuous rebalancing of a portfolio, which becomes too expensive when one takes into account the transaction costs. The singular control $\xi(t)$ is in this context the total transactions up to time t , cf. Shreve and Soner (1994). We can interpret $x(t)$ and $y(t)$ as the value of a bond portfolio and a portfolio of stocks respectively, and with proportional transaction costs incurred whenever money is moved between the stock and the bonds.

The main result of the paper is a stochastic maximum principle for this type of controlled system. It consists of a maximum condition as in the case of only an

absolutely continuous control, plus a second condition characterizing the singular part of the optimal control. The method used is different from that in the previous paper in that by performing the perturbation directly on the relaxed control, it allows us to use a convex perturbation. This ultimately gives us the maximum principle of the first order (c.f. (12)). This, in fact, shows that the maximum principle in paper II also can be strengthened. We also establish existence of an optimal control which is derived using a similar scheme as in paper II.

Paper IV: The relaxed general maximum principle for singular optimal control of diffusions

In this paper we study a general stochastic control problem with both continuous and singular control components. The state equation is given by

$$dx(t) = b(t, x(t), u(t)) dt + \sigma(t, x(t), u(t)) dB(t) + G_t d\xi(t),$$

where the control is the vector (u, ξ) .

The question of existence of an optimal control is basically the same as in the non-singular case. Indeed, by relaxing the absolutely continuous control an existence result is obtained in Haussmann and Suo (1995). A control is defined as a solution to the martingale problem corresponding to the state equation. The relaxation is performed by integrating the generator with respect to a probability measure. This however leads to problems of interpretation; what kind of process induces a measure which is a solution to this relaxed martingale problem? It turns out that the Itô integral should be replaced by a stochastic integral with respect to a continuous orthogonal martingale measure, i.e. a process which for fixed t is a measure on sets $A \subset U$, and for a fixed set $A \subset U$ is a continuous martingale, see e.g. El Karoui and Méléard (1990) .

By the chattering lemma one can construct a sequence of strict controls approximating the relaxed optimal control. Therefore, using methods from the derivation of the maximum principle for strict singular control in Mezerdi and Bahlali (2002), we are able to derive the corresponding maximum principle for near optimal controls. Passing to the limit gives the relaxed maximum principle. This is a generalization of a result in Bahlali et al. (2007), where the diffusion coefficient σ is independent of the control.

This paper appeared in my licentiate thesis under the title “The relaxed stochastic maximum principle in singular optimal control of diffusions with controlled diffusion coefficient”.

Bibliography

- Bahlali, S., Djehiche, B. and Mezerdi, B., 2007. The relaxed stochastic maximum principle in singular optimal control of diffusions. *SIAM J. Control Optim.* 46(2), 427–444.
- Bahlali, S., Djehiche, B. and Mezerdi, B., 2006. Approximation and optimality necessary conditions in relaxed stochastic control problems. *J. Appl. Math. Stoch. Anal.*, Article ID 72762, 1–23.
- Bahlali, S. and Mezerdi, B., 2005. A general stochastic maximum principle for singular control problems. *Electron. J. Probab.* 10, Paper no. 30, 988–1004.
- Bellman, R., 1957. *Dynamic Programming*. Princeton Univ. Press.
- Bensoussan, A., 1982. *Lectures on stochastic control, Nonlinear Filtering and Stochastic Control* (S.K. Mitter, A. Moro, eds.). *Lecture Notes in Math.* 972, Springer-Verlag.
- Bismut, J.M., 1978. An introductory approach to duality in optimal stochastic control. *SIAM Rev.* 20, 62–78.
- El Karoui, N. and Méléard, S., 1990. Martingale measures and stochastic calculus. *Probab. Theory Rel.* 84, 83–101.
- El Karoui, N., Hùu Nguyen, D. and Jeanblanc-Picqué, M., 1987. Compactification methods in the control of degenerate diffusions: existence of an optimal control. *Stochastics* 20, 169–221.
- Fleming, W.H. and Soner, H.M., 2006. *Controlled Markov Processes and Viscosity Solutions*, Second Edition. Springer-Verlag.
- Fleming, W.H., 1977. *Generalized solutions in optimal stochastic control. Differential Games and Control Theory II* (E. Roxin, P.T. Liu, R. Sternberg, eds.), Marcel Dekker.
- Glad, T. and Ljung, L., 1997. *Reglerteori*. Studentlitteratur.
- Haussmann, U.G. and Suo, W., 1995. Singular optimal stochastic controls I: Existence. *SIAM J. Control Optim.* 33(3), 916–936.
- Haussmann, U.G., 1986. *A stochastic maximum principle for optimal control of diffusions*. Pitman Research Notes in Mathematics Series 151, Longman Scientific & Technical, Harlow.

- Kushner, H.J. and Schweppe, F.C., 1964. A maximum principle for stochastic control systems. *J. Math. Anal. Appl.* 8, 287–302.
- Ma, J. and Yong, J., 1995. Solvability of forward-backward SDEs and the nodal set of Hamilton-Jacobi-Bellman equations. *Chin. Ann. Math. Ser. B* 16, 279–298.
- Ma, J. and Yong, J., 1999. Forward-backward stochastic differential equations and their applications. *Lecture Notes in Math.* 1702, Springer-Verlag.
- Mezerdi, B. and Bahlali, S., 2002. Necessary conditions for optimality in relaxed stochastic control problems. *Stoch. Stoch. Rep.* 73(3), 201–218.
- Peng, S., 1990. A general stochastic maximum principle for optimal control problems. *SIAM J. Control Optim.* 28(4), 966–979.
- Pontryagin, L.S., Boltyanskii, V.G., Gamkrelidze, R.V. and Mishchenko, E.F., 1962. *The Mathematical Theory of Optimal Processes.* Interscience.
- Shreve, S.E. and Soner, H.M., 1994. Optimal investment and consumption with transaction costs. *Ann. Appl. Probab.* 4(3), 609–692.
- Sznitman, A.S., 1989. Topics in propagation of chaos. *École de Probabilités de Saint Flour, XIX-1989, Lecture Notes in Math.* 1464, Springer-Verlag, Berlin, 165–251.
- Sontag, E.D., 1990. *Mathematical Control Theory.* Springer-Verlag.
- Yong, J. and Zhou, X.Y., 1999. *Stochastic Controls: Hamiltonian Systems and HJB Equations.* Springer-Verlag.
- Young, L.C., 1969. *Lectures on the calculus of variations and optimal control theory.* W.B. Saunders Co.