C. R. Acad. Sci. Paris, Ser. I 352 (2014) 707–713

Contents lists available at ScienceDirect

C. R. Acad. Sci. Paris, Ser. I

www.sciencedirect.com

Optimal control/Calculus of variations

Dynamic programming for mean-field type control

Programmation dynamique pour les problèmes de contrôle à champs moyen

Mathieu Laurière, Olivier Pironneau

LJLL, Université Pierre-et-Marie-Curie (Paris-6), 4, place Jussieu, 75005 Paris, France

A R T I C L E   I N F O

Article history:
Received 13 April 2014
Accepted 17 July 2014
Available online 6 September 2014

Presented by Alain Bensoussan

A B S T R A C T

For mean-field type control problems, stochastic dynamic programming requires adaptation. We propose to reformulate the problem as a distributed control problem by assuming that the PDF ρ of the stochastic process exists. Then we show that Bellman's principle applies to the dynamic programming value function $V(\tau,\rho_\tau)$, where the dependency on ρ_τ is functional as in P.-L. Lions' analysis of mean-field games (2007) [10]. We derive HJB equations and apply them to two examples, a portfolio optimization and a systemic risk model.

R É S U M É

Pour les problèmes de contrôle stochastique à champs moyen, la programmation dynamique ne s'applique pas sans adaptation ; mais si l'on reformule le problème avec l'équation de Fokker–Planck, on peut le faire en utilisant une fonctionnelle valeur $\{\tau,\rho_\tau(\cdot)\}\to V(\tau,\rho_\tau)$ comme dans l'analyse des problèmes de jeux à champs moyen par P.-L. Lions (2007) [10]. Les résultats sont appliqués à un problème d'optimisation de portefeuille et à un problème de risque systémique.

1. Introduction

Stochastic control is an old topic [5,12,14,15], which benefits from a renewed interest in economy and finance due to mean-field games [6,7,9,13]. They lead, among other things, to stochastic control problems which involve statistics of the Markov process like means and variance. Optimality conditions for these are derived either by stochastic calculus of variation [1] or by stochastic dynamic programming in the quadratic case [2,3], but not in the general case for the fundamental reason that Bellman's principle does not apply in its original form on the stochastic trajectories of say X_t if those depend upon statistics of X_t like its mean value. As noticed earlier in [10] and in [4],[1] there seems to be no such restriction if one works with the probability measure of X_t and uses the Fokker–Planck equation.

E-mail addresses: Mathieu.Lauriere@UPMC.fr (M. Laurière), Olivier.Pironneau@UPMC.fr (O. Pironneau).

[1] This preprint came to our knowledge after the submission of this note.

http://dx.doi.org/10.1016/j.crma.2014.07.008
1631-073X/

In this note we apply the dynamic programming argument to the value functional $V(\tau, \rho_\tau(\cdot))$, where ρ_τ is the PDF of X_τ. Of course, this is at the cost of several regularity assumptions; in particular, it requires the existence of PDF at all times.

Once the problem is reformulated with the Fokker–Planck equation, it becomes a somewhat standard exercise to find the optimality necessary conditions by a calculus of variations. So the note begins likewise. Then a similar result is obtained by using dynamic programming, and the connection with the previous approach and with stochastic dynamic programming is established, with the advantage that sufficient conditions for optimality are obtained. Finally, we apply the method to two mean-field type control problems stated in [1] and [6].

2. The problem

Let $d, s, r \in \mathbf{N}^+$. Consider a stochastic differential equation

$$dX_t = u(X_t, t)dt + \sigma\big(X_t, t, u(X_t, t)\big)dW_t \tag{2.1}$$

where $T > 0$, $u : \mathbf{R}^d \times (0, T) \to \mathbf{R}^d$, $\sigma : \mathbf{R}^d \times (0, T) \times \mathbf{R}^d \to \mathbf{R}^{d\times d}$ and W_t is a d-vector of independent Brownian motions. We make the usual assumptions for X_t to exist once X_0 is known [14].

Let $\tilde H : \mathbf{R}^d \times (0, T) \times \mathbf{R}^d \times \mathbf{R}^r \to \mathbf{R}$, $\tilde h : \mathbf{R}^d \times (0, T) \times \mathbf{R}^d \to \mathbf{R}^r$, $G : \mathbf{R}^d \times \mathbf{R}^s \to \mathbf{R}$, $g : \mathbf{R}^d \to \mathbf{R}^s$. Assume also that ρ_0 is positive with unit measure on $\mathbf{R}^d$.

Let $\mathcal{V}_d \subset \mathbf{R}^d$, $\mathcal{U}_d = \{u \in (L^\infty(\mathbf{R}^d \times \mathbf{R}))^d \ : \ u(x, t) \in \mathcal{V}_d \ \forall x, t\}$ and consider the problem

$$\begin{aligned}&\min_{u\in\mathcal{U}_d} J := \int_0^T \mathbf{E}\Big[\tilde H\big(X_t, t, u(X_t, t), \mathbf{E}\big[\tilde h\big(X_t, t, u(X_t, t)\big)\big]\big)\Big]dt + \mathbf{E}\big[G\big(X_T, \mathbf{E}\big[g(X_T)\big]\big)\big]\\ &\text{subject to (2.1) and such that } \rho_0 \text{ is the PDF of } X_0\end{aligned} \tag{2.2}$$

Andersson et al. [1] analyzed this problem using stochastic calculus of variations, claiming rightly that dynamic programming is not possible unless $\tilde h = 0$, $g = 0$. Yet denoting $Q = \mathbf{R}^d \times (0, T)$ and $\mu_{ij} = \frac{1}{2}\sum_k \sigma_{ik}\sigma_{jk}$, with sufficient regularity, namely if X_t has a PDF ρ_t (for weaker hypotheses see [11]), the problem is equivalent to

$$\begin{aligned}&\min_{u\in\mathcal{U}_d} J = \int_Q H\big(x, t, u(x, t), \rho_t(x), \chi(t)\big)\rho_t(x)dxdt + \int_{\mathbf{R}^d} G(x, \xi)\rho_{|T}dx\\ &\quad \text{where } \chi(t) = \int_{\mathbf{R}^d} h\big(x, t, u(x, t), \rho_t(x)\big)\rho_t(x)dx,\ \xi = \int_{\mathbf{R}^d} g(x)\rho_T(x)dx \text{ and } \rho_t\\ &\text{s.t.}\quad \partial_t\rho + \nabla\cdot(u\rho) - \nabla\cdot\nabla\cdot(\mu\rho) = 0, \quad \rho|_0 = \rho_0(x), \quad x \in \mathbf{R}^d\end{aligned} \tag{2.3}$$

where $\tilde H = H$, $\tilde h = h$ if these are not functions of $\rho_t(x)$.

Hypothesis 1. Assume that all data are continuously differentiable with respect to u and ρ and have additional regularity so that the solution to the Fokker–Planck equation is unique and uniformly continuously differentiable with respect to u and μ.

3. Calculus of variations

Proposition 1. *Let $A : B = trace(A^{\mathrm{T}}B)$. A control u is optimal for (2.3) only if*

$$\int_{\mathbf{R}^d}\bigg(H'_u + h'_u\int_{\mathbf{R}^d} H'_\chi\rho dx + \nabla\rho^* - \mu'_u : \big(\nabla\nabla\rho^*\big)\bigg)(v - u)\rho dx \geqslant 0 \quad \forall t, \forall v \in \mathcal{U}_d \tag{3.4}$$

$$\textit{with } \partial_t\rho^* + u\nabla\rho^* + \mu : \nabla\nabla\rho^* = -\bigg[H'_\rho\rho + H + \big(h'_\rho\rho + h\big)\int_{\mathbf{R}^d} H'_\chi\rho dx\bigg],\ \rho^*_{|T} = g\int_{\mathbf{R}^d} G'_\xi\rho_{|T}dx + G \tag{3.5}$$

Proof. See http://hal.archives-ouvertes.fr:hal-01018361. □

4. Dynamic programming

For notational clarity, let us consider the more general case where H, G are functionals of $\rho_t(\cdot)$. For any $\tau \in [0, T]$ and any $\rho_\tau \geqslant 0$ with unit measure on $\mathbf{R}^d$, let:

$$V(\tau;\rho_\tau)=\min_{u\in\mathcal{U}_d} J(\tau;\rho_\tau,u):=\int_\tau^{\mathrm{T}}\int_{\mathbf{R}^d} H\big(x,t,u(x,t);\rho_t\big)\rho_t(x)\mathrm{d}x\mathrm{d}t+\int_{\mathbf{R}^d} G(x;\rho_{|T})\rho_{|T}\mathrm{d}x$$

subject to (2.3), i.e. such that ρ_t is the PDF of X_t given by (2.1) starting with ρ_τ at time τ (4.6)

Note that the second parameter in V is a function of x, yet it is not $V(\tau,\rho_\tau(x))$ but $V(\tau;\rho_\tau(\cdot))$. We now prove the following version of *Bellman's principle of optimality*:

Proposition 2. *If the problem is regular, then for any $\tau\in[0,T]$ and any positive ρ_τ with unit measure on $\mathbf{R}^d$, we have:*

$$V(\tau;\rho_\tau)=\min_{u\in\mathcal{U}_d}\int_\tau^{\tau+\delta\tau}\int_{\mathbf{R}^d} H\big(x,t,u(x,t);\rho_t\big)\rho_t(x)\mathrm{d}x\mathrm{d}t+V(\tau+\delta\tau;\rho_{\tau+\delta\tau})$$

subject to ρ_t given by (2.3) on $[\tau,\tau+\delta\tau]$ initialized by ρ_τ at time τ (4.7)

Proof. See http://hal.archives-ouvertes.fr:hal-01018361. □

Proposition 3 *(HJB minimum principle). There exists $\{x,\tau,\rho_\tau(\cdot)\}\to V'\in\mathbf{R}$ such that:*

$$0=\min_{v\in\mathcal{V}_d}\int_{\mathbf{R}^d}\Big(H\big(x,\tau,v(x);\rho_\tau\big)+H'_\rho\big(x,\tau,v(x);\rho_\tau\big)\cdot\rho_\tau$$
$$+\partial_\tau V'+\mu\big(x,\tau,v(x)\big):\nabla_x\nabla_x V'+v(x)\cdot\nabla_x V'\Big)\rho_\tau\mathrm{d}x \tag{4.8}$$

where $H'_\rho\cdot\nu=\lim_{\lambda\to0}[H(x,\tau;\rho+\lambda\nu)-H(x,\tau,\rho)]/\lambda$.

Proof. A first-order approximation of the time derivative in the Fokker–Planck equation gives

$$\delta\rho_\tau:=\rho_{\tau+\delta\tau}-\rho_\tau=\delta\tau\big[\nabla\cdot\nabla\cdot(\mu_\tau\rho_\tau)-\nabla\cdot(u_\tau\rho)\big]+o(\delta\tau) \tag{4.9}$$

When everything is differentiable and smooth,

$$V(\tau+\delta\tau;\rho_{\tau+\delta\tau})=V(\tau;\rho_\tau)+\partial_\tau V(\tau;\rho_\tau)\delta\tau+V'_\rho(\tau;\rho_\tau)\cdot\delta\rho_\tau+o(\delta\tau) \tag{4.10}$$

Using (4.10) and the mean value theorem for the time integral, (4.7) yields

$$V(\tau;\rho_\tau)=\min_{u\in\mathcal{U}_d}\Big\{\delta\tau\int_{\mathbf{R}^d}H\rho_\tau\mathrm{d}x+V(\tau;\rho_\tau)+\partial_\tau V(\tau;\rho_\tau)\delta\tau+V'_\rho(\tau;\rho_\tau)\cdot\delta\rho_\tau+o(\delta\tau)\Big\}$$

The terms $V(\tau;\rho_\tau)$ cancel; divided by $\delta\tau$ and combined with (4.9) and letting $\delta\tau\to0$, (4.11) gives

$$0=\min_{u\in\mathcal{U}_d}\Big\{\int_{\mathbf{R}^d}H\rho_\tau\mathrm{d}x+\partial_\tau V(\tau;\rho_\tau)+V'_\rho(\tau;\rho_\tau)\cdot\big[\nabla\cdot\nabla\cdot(\mu_\tau\rho_\tau)-\nabla\cdot(u_\tau\rho)\big]\Big\} \tag{4.11}$$

To finalize the proof, we need to relate V to V'_ρ and to its Riesz representative V':

Proposition 4. *For any $\tau\in[0,T]$ and any initial PDF ρ_τ, let $\hat{u}$ and $\hat{\rho}$ denote respectively the optimal control and the corresponding solution of (2.3). Then:*

$$\int_{\mathbf{R}^d}V'(\tau;\rho_\tau)\rho_\tau\mathrm{d}x=V'_\rho(\tau;\rho_\tau)\cdot\rho_\tau=V(\tau;\rho_\tau)+\int_\tau^{\mathrm{T}}\int_{\mathbf{R}^d}\big(H'_\rho\big(x,t,\hat{u}(x,t);\hat{\rho}_t\big)\cdot\hat{\rho}_t\big)\hat{\rho}_t(x)\mathrm{d}x\mathrm{d}t$$
$$+\int_{\mathbf{R}^d}\big(G'_\rho(x;\hat{\rho}_T)\cdot\hat{\rho}_T\big)\hat{\rho}_T(x)\mathrm{d}x \tag{4.12}$$

Proof. Notice that the Fokker–Planck equation implies $\rho_t=\mathbf{G}(t-\tau)*\rho_\tau$ where $\mathbf{G}$ is a semi-group operator. Let $(\hat{u}_t)_{t\in[0,T]}$ be the optimal control and $(\hat{\rho}_t)_{t\in[0,T]}$ the corresponding solution of (2.3). Then:

$$V(\tau;\hat\rho_\tau)=\int_\tau^T\int_{\mathbf{R}^d} H\big(x,t,\hat u(x,t);\mathbf{G}(t-\tau)*\hat\rho_\tau\big)\mathbf{G}(t-\tau)*\hat\rho_\tau\,\mathrm{d}x\mathrm{d}t+\int_{\mathbf{R}^d} G(x;\hat\rho_{|T})\hat\rho_{|T}\,\mathrm{d}x$$

This can be differentiated with respect to ρ by computing $\lim_{\lambda\to 0}\frac{1}{\lambda}[V(\tau;\hat\rho_\tau+\lambda\nu)-V(\tau;\hat\rho_\tau)]$, for a given function $x\to\nu(x)$:

$$\begin{aligned}V'_\rho(\tau;\hat\rho_\tau)\cdot\nu=&\int_\tau^T\int_{\mathbf{R}^d} H\big(x,t,\hat u(x,t);\mathbf{G}(t-\tau)*\hat\rho_\tau\big)\mathbf{G}(t-\tau)*\nu\,\mathrm{d}x\mathrm{d}t+\int_{\mathbf{R}^d} G(x;\hat\rho_{|T})\mathbf{G}(T-\tau)*\nu\,\mathrm{d}x\\&+\int_\tau^T\int_{\mathbf{R}^d}\big(H'_\rho\big(x,t,\hat u(x,t);\mathbf{G}(t-\tau)*\hat\rho_\tau\big)\cdot\big[\mathbf{G}(t-\tau)*\nu\big]\big)\mathbf{G}(t-\tau)*\hat\rho_\tau\,\mathrm{d}x\mathrm{d}t\\&+\int_{\mathbf{R}^d}\big(G'_\rho\big(x;\mathbf{G}(T-\tau)*\hat\rho_\tau\big)\cdot\big[\mathbf{G}(T-\tau)*\nu\big]\big)\mathbf{G}(T-\tau)*\hat\rho_\tau\,\mathrm{d}x\end{aligned}$$

Taking $\nu=\hat\rho_\tau$ leads to (4.12). □

End of proof of Proposition 3 Differentiating (4.12) w.r.t. τ leads to

$$\partial_\tau V(\tau;\rho_\tau)=\big(\partial_\tau V'_\rho(\tau;\rho_\tau)\big)\cdot\rho_\tau+\int_{\mathbf{R}^d}\big(H'_\rho\big(x,\tau,\hat u_\tau(x);\rho_\tau\big)\cdot\rho_\tau\big)\rho_\tau(x)\mathrm{d}x$$

where $\hat u_\tau$ is the optimal control at time τ. Now, let us go back to (4.11), which we rewrite:

$$\begin{aligned}0=\min_{u_\tau}\Big\{&\int_{\mathbf{R}^d}\big(H\big(x,\tau,u_\tau(x);\rho_\tau\big)+H'_\rho\big(x,\tau,u_\tau(x);\rho_\tau\big)\cdot\rho_\tau\big)\rho_\tau(x)\mathrm{d}x\\&+\big(\partial_\tau V'_\rho(\tau;\rho_\tau)\big)\cdot\rho_\tau+V'_\rho(\tau;\rho_\tau)\cdot\big[\nabla\cdot\nabla\cdot(\mu_\tau\rho_\tau)-\nabla\cdot(u_\tau\rho_\tau)\big]\Big\}\end{aligned}\tag{4.13}$$

By integrating by parts the last term, Proposition 3 is proved. □

Remark 1. Notice that (4.12) and (4.6) imply:

$$\int_{\mathbf{R}^d} V'_{|T}\hat\rho_T\,\mathrm{d}x=V(T,\hat\rho_T)=\int_{\mathbf{R}^d}\Big(G+g\int_{\mathbf{R}^d}\partial_\xi G\hat\rho_T\,\mathrm{d}x\Big)\hat\rho_T\,\mathrm{d}x,\qquad \hat\xi=\int_{\mathbf{R}^d} g(x,\hat\rho_T)\hat\rho_T\,\mathrm{d}x\tag{4.14}$$

Remark 2. By taking $\rho_\tau=\delta(x-x_0)$ the usual HJB principle is found if $h=g=0$.

Proposition 5 *(Hamilton–Jacobi–Bellman equation). When $\mathcal{V}_d=\mathbf{R}^d$, at the optimal solution $\hat u$*

$$\nabla_u H+\nabla_u H'\cdot\hat\rho_\tau+\nabla_x V'+\partial_u\mu:\nabla_x\nabla_x V'=0\tag{4.15}$$

$$\int_{\mathbf{R}^d}\big(H+H'\cdot\hat\rho_\tau+\partial_\tau V'+\hat\mu:\nabla_x\nabla_x V'+\hat u\cdot\nabla_x V'\big)\hat\rho_\tau\,\mathrm{d}x=0\tag{4.16}$$

Remark 3. When $H=H(x,t,u(x,t),\rho_t(x),\chi(t))$ with $\chi(t)=\int_{\mathbf{R}^d} h(x,t,u(x,t),\rho_t(x))\rho_t(x)\mathrm{d}x$,

$$H'_\rho\big(x,\tau,u(x,\tau);\rho_\tau\big)\cdot\rho_\tau=\rho_\tau\partial_\rho H+\Big(\int_{\mathbf{R}^d}\partial_\chi H\rho_\tau\,\mathrm{d}x\Big)(h+\rho_\tau\partial_\rho h)\tag{4.17}$$

Then for the optimal $\hat u$ and $\hat\rho$ (4.16) yields

$$\partial_\tau V'+\hat\mu:\nabla_x\nabla_x V'+\hat u\cdot\nabla_x V'=-\Big[H+\hat\rho\partial_\rho H+\Big(\int_{\mathbf{R}^d}\partial_\chi H\hat\rho\,\mathrm{d}x\Big)(h+\hat\rho\partial_\rho h)\Big]\tag{4.18}$$

The link with Section 3 is established: (3.5) and (4.18) coincide with $V'=\rho^*$.

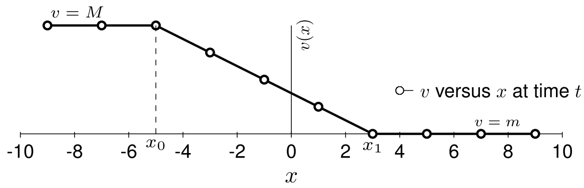

Fig. 1. The control is $v=M$ when $x<x_0$, affine when $x_0<v<x_1$, and $v=m$ when $x>x_1$ with $x_0=-\frac{M\sigma^2}{b}+(\frac{1}{2\gamma}+\mathbf{E}_T x)\,\mathrm{e}^{-a(T-t)}$, $x_1=-\frac{m\sigma^2}{b}+(\frac{1}{2\gamma}+\mathbf{E}_T x)\,\mathrm{e}^{-a(T-t)}$.

5. Portfolio optimization

Following [1], a portfolio of value x_t made of a risky asset and a riskless one is optimally managed at t if the quantity v invested at t in the risky asset minimizes, with $\rho_{|0}$ given,

$$J=\frac{1}{2}\int_{\Omega}(\gamma x^2-x)\rho_{|T}\,\mathrm{d}x-\frac{\gamma}{2}\Big[\int_{\Omega}x\rho_{|T}\,\mathrm{d}x\Big]^2,\qquad \partial_t\rho+\partial_x\big[(ax+bv)\rho\big]-\partial_{xx}\Big[\frac{\sigma^2v^2}{2}\rho\Big]=0 \tag{5.19}$$

where a is the interest rate b is a minus the drift of the risky asset and σ is its volatility. We assume that v is a feedback function $x,t\to v(x,t)$, there are bounds on v, at each time $m\leqslant v\leqslant M$. Thus $d=1$ and

$$H=0,\qquad h=0,\qquad G=\frac{1}{2}\Big(\gamma x^2-x-\gamma x\int_{\Omega}x\rho_{|T}\,\mathrm{d}x\Big),\qquad u=ax+bv,\qquad \mu=\frac{\sigma^2v^2}{2} \tag{5.20}$$

The problem deviates slightly from framework (2.3) but the methodology is the same and gives:

$$\delta J=\int_{Q}\big[(b\partial_x\rho^*+\sigma^2v\partial_{xx}\rho^*)\rho\delta v\big]\quad\text{with }\rho^*(\pm\infty)=0,\quad\text{and}$$

$$\partial_t\rho^*+(ax+bv)\partial_x\rho^*+\frac{\sigma^2v^2}{2}\partial_{xx}\rho^*=0,\qquad \rho^*_{|T}=\frac{1}{2}(\gamma x^2-x)-\gamma x\int_{\mathbf{R}}[x\rho_{|T}] \tag{5.21}$$

5.1. Polynomial solution

Assume $\rho^*=qx^2+rx+s$ and $v=Ax+B\in(m,M)$. Then the adjoint equation gives solvable ODEs for $q(t)$, $r(t)$ and $s(t)$. Because of the constraints, the general solution has three regimes, as shown in Fig. 1.

Proof. See http://hal.archives-ouvertes.fr:hal-01018361. □

Remark 4. The advantage here compared with [1] is that we do not need to guess the shape of the control nor of the adjoint state, once it is assumed polynomial. The analysis also handles constraints.

6. Numerical solution of a systemic risk problem

In [6] it is shown that the rare-event probability that the state of a system of N banks, depending on the mean situation of all, transits from a stable situation ρ_0 to a critical one ρ_T at time T is given finding the minimum in g of J with

$$J(g)=\frac{1}{2\sigma^2}\int_{Q}g^2\rho\ :\ \partial_t\rho+\partial_x\big(b(x,g)\rho\big)-\frac{\sigma^2}{2}\partial_{xx}\rho=0,\qquad \rho(x,0)=\rho_0(x) \tag{6.22}$$

subject to $b=-hx^3+(h-\theta)x-\theta\int_{\mathbf{R}}x\rho-g$, $\rho(x,T)=\rho_T(x)$ where $h,\theta\in\mathbf{R}$ are given. With $\kappa=h-\theta$, this is also

$$\min_u J=\frac{1}{2\sigma^2}\int_{Q}\Big(hx^3-\kappa x-\chi+\frac{u}{\rho}\Big)^2\rho\ :\ \partial_t\rho-\frac{\sigma^2}{2}\partial_{xx}\rho=-\partial_x u,\quad \rho_{|0},\rho_{|T}\text{ given},\quad \chi=\theta\int_{\mathbf{R}}x\rho$$

Now we notice that $\rho=\tilde\rho+\frac{t}{T}(\rho_T-\tilde\rho_{|T})$ satisfies the conditions at 0 and T and the PDE with $u=\tilde u-\frac{1}{T}\int^x(\rho_T-\tilde\rho_{|T})\mathrm{d}x+\frac{t\sigma^2}{2T}\partial_x(\rho_T-\tilde\rho_{|T})$ provided that $\partial_t\tilde\rho-\frac{\sigma^2}{2}\partial_{xx}\tilde\rho=-\partial_x\tilde u$. This means that the problem is in the form analyzed above with state variable $\{\tilde\rho,\rho\}$ and control $\tilde u$; naturally the adjoint state has also two components: $\{\rho^*,\tilde\rho^*\}$.

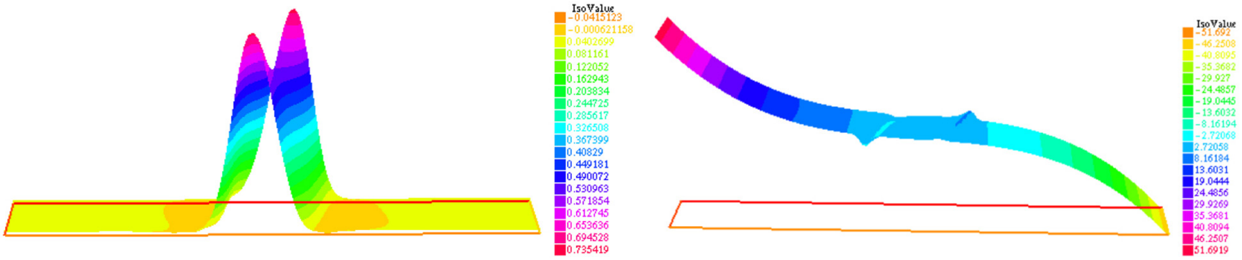

Fig. 2. Left: $x, t \to \rho(x, t)$; $x \in (-15, 15)$ is horizontal, $t \in (0, 2)$ is from front to back with origin $x = -15, t = 0$ on the lower left corner. Right: $x, t \to g(x, t)$. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

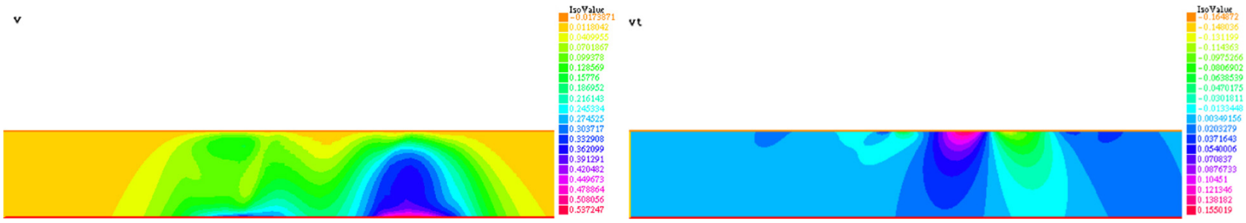

Fig. 3. ρ^* (left) and $\tilde{\rho}^*$ (right) versus x (horizontal) and t (vertical). (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

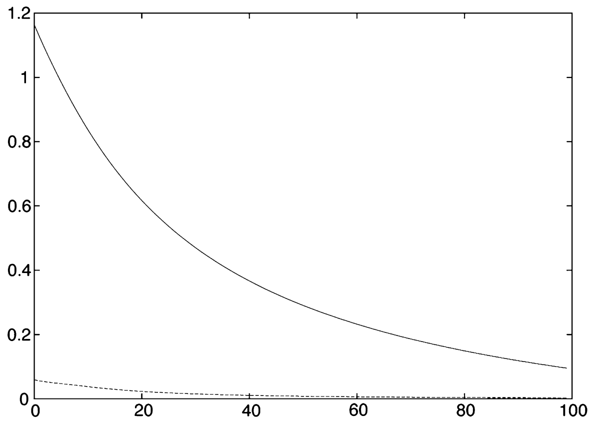

Fig. 4. Iteration history: values of J (top curve) and $\|\text{grad}_u J\|^2$ versus iteration count.

Based on the variation of J with respect to u we have used 100 iterations of a gradient method with fixed step size, $\omega = 0.3$. The parameters of the problem are $T = 2$, $h = 0.1$, $\theta = 1$, $\sigma^2 = \frac{2}{3}\theta$ and ρ_0, ρ_T are as in [6]. The numerical method for the PDEs is a centered space–time finite element method of degree 1, implemented with freefem++ [8], on a mesh of 150 points and 40 time steps. Results are displayed in Figs. 2–4.

References

[1] D. Andersson, B. Derviche, A maximum principle for SDEs of mean-field type, Appl. Math. Optim. 63 (2011) 341–356.
[2] A. Bensoussan, J. Frehse, Control and Nash games with mean field effect, Chin. Ann. Math., Ser. B 34B (2) (2013) 161–192.
[3] A. Bensoussan, J. Frehse, S.C.P. Yam, Mean-Field Games and Mean-Field Type Control, Springer Briefs in Mathematics, 2014.
[4] A. Bensoussan, J. Frehse, S.C.P. Yam, The master equation in mean-field theory, Asymptot. Anal. (2014), in press.
[5] W.H. Fleming, H.M. Soner, Controlled Markov Process and Viscosity Solutions, 2nd ed., Springer, 2006.
[6] J. Garnier, G. Papanicolaou, Tzu-Wei Yang, Large deviations for a mean field model of systemic risk, SIAM J. Financ. Math. 4 (1) (2013) 151–184.
[7] O. Guéant, M. Lasry, P.-L. Lions, Mean field games and applications, in: Paris–Princeton Lectures on Mathematical Finance, in: Lecture Notes in Mathematics, Springer, 2011.
[8] F. Hecht, New development in freefem++, J. Numer. Math. 20 (3–4) (2012) 251–265.
[9] J.-M. Lasry, P.-L. Lions, Mean-field games, Jpn. J. Math. 2 (2007) 229–260.
[10] P.-L. Lions, Mean-field games, Cours au Collège de France (2007–2008), http://www.college-de-france.fr/site/pierre-louis-lions/course-2007-2008_1.htm.

[11] A. Neufeld, M. Nutz, Nonlinear Lévy processes and their characteristics, arXiv:1401.7253v1.
[12] B. Øksendal, A. Sulem, Applied Stochastic Control of Jump Diffusions, Springer, 2005.
[13] Min Shen, G. Turinici, Liquidity generated by heterogeneous beliefs and costly estimations, Netw. Heterog. Media 7 (2) (2012) 349–361.
[14] N. Touzi, Optimal Stochastic Control, Stochastic Target Problems and Backard SDE, Field Inst. Monogr., vol. 29, Springer, 2013.
[15] Jiongmin Yong, Xun Yu Zhou, Stochastic Control, Applications of Mathematics Series, vol. 43, Springer, 1999.