Convex Duality Method for Constrained Quadratic Risk Minimization via FBSDEs

Harry Zheng
Imperial College
(joint work with Ashley Davey)

Workshop on Stochastic Control in Finance
National University of Singapore
22-26 July 2019

Outlines

- Constrained LQ problems and FBSDEs for primal and dual problems
- Examples to constrained quadratic risk minimization problems
- Deep reinforcement learning to solve FBSDEs

Constrained Linear Quadratic Problem

The state $X^u \in \mathbb{R}^d$ has linear dynamics controlled by some process $u \in \mathcal{A}$, where $\mathcal{A}$ is the set of progressively measurable admissible processes that are constrained to some closed convex set $K \in \mathbb{R}^m$.

$$dX^u(t) = (A(t)X^u(t)+B(t)u(t))dt+(C(t)X^u(t)+D(t)u(t))dW(t) \quad X_0 = x_0,$$

where W is a 1 dimensional Brownian motion. We wish to minimise a quadratic of the system, that is, a cost function J of the form

$$J(u) = E\left[\int_0^T f(t, X^u(t), u(t))dt + g(X^u(T))\right],$$

where

$$\begin{aligned} f(t,x,u) &= x^\top Q(t)x + 2x^\top S(t)u + u^\top R(t)u + \xi(t)^\top x + \psi(t)^\top u \\ g(x) &= x^\top M x + \eta^\top x. \end{aligned}$$

We assume that the matrices

$$\begin{pmatrix} Q(t) & S(t) \\ S(t)^\top & R(t) \end{pmatrix}$$

and M are positive definite, so J is a convex function.

Dual Controlled State Process

Assume D has full rank, that is, $D(t)D(t)^\top(t)$ invertible for each t. Define $\mathcal{H}^2(0,T;\mathbb{R}^n)$ to be the set of progressively measurable $\mathbb{R}^n$ valued processes x satisfying $E[\int_0^T ||x(t)||^2 dt] < \infty$. Let $y_0 \in \mathbb{R}^d$ and $\alpha, \beta \in \mathcal{H}^2(0,T;\mathbb{R}^n)$. The dual state process $Y^{y_0,\alpha,\beta}$ is given by

$$dY^{y_0,\alpha,\beta}(t) = (\alpha(t)+\tilde{A}(t)Y^{y_0,\alpha,\beta}(t)+\tilde{B}(t)\beta(t))dt+(\tilde{C}(t)Y^{y_0,\alpha,\beta}(t)+\tilde{D}(t)\beta(t))dW(t),$$

where $Y(0) = y_0$ and

$$\begin{aligned}
\tilde{A}(t) &= \left(B(t)D^\top(t)\left(D(t)D(t)^\top\right)^{-1}C(t) - A(t)\right)^\top \\
\tilde{B}(t) &= -\left(D^\top(t)\left(D(t)D(t)^\top\right)^{-1}C(t)\right)^\top \\
\tilde{C}(t) &= -\left(B(t)D^\top(t)\left(D(t)D(t)^\top\right)^{-1}\right)^\top \\
\tilde{D}(t) &= \left(D^\top(t)\left(D(t)D(t)^\top\right)^{-1}\right)^\top
\end{aligned}$$

and (y_0, α, β) are admissible dual controls. Then

$$X^u(t)^\top Y^{y_0,\alpha,\beta}(t) - \int_0^t (X^u(s)^\top \alpha(s) + u(s)^\top \beta(s))ds$$

is a super martingale

Dual Problem

Define the dual functions

$$\tilde{f}(t,\alpha,\beta) = \sup_{x,u\in K}\{-f(t,x,u) + x^\top\alpha + u^\top\beta\}$$
$$\tilde{g}(y) = \sup_{x}\{-g(x) - x^\top y\}$$
$$= (y+\eta)^\top\tilde{M}(y+\eta),$$

where the matrix $\tilde{M}$ is given by

$$\tilde{M} = [(M+M^\top)^{-1} - (M+M^\top)^{-1}M(M+M^\top)^{-1}].$$

In particular, if M is symmetric then $\tilde{M} = \frac{1}{2}M^{-1}$. These dual functions allow us to form the following dual relation

$$\inf_u J(u) \geqslant -\inf_{y_0,\alpha,\beta}\left\{x_0^\top y_0 + E\left[\int_0^T \tilde{f}(t,\alpha(t),\beta(t))dt + \tilde{g}(Y^{y_0,\alpha,\beta}(T))\right]\right\}.$$

Primal Optimality Condition

Define $\mathcal{S}^2(0,T;\mathbb{R}^n)$ to be the set of progressively measurable $\mathbb{R}^n$ valued processes x such that $E[\sup_{0\leqslant t\leqslant T}||x(t)||^2] < \infty$.

Theorem 1. Let $\hat{u} \in \mathcal{A}$ be admissible. Then $\hat{u}$ is optimal for the primal problem if and only if the solution $X^{\hat{u}}$ and the adjoint process $(\hat{p}_1, \hat{q}_1) \in \mathcal{S}^2(0,T;\mathbb{R}^n)$ of the FBSDE

$$\begin{aligned}
dX^{\hat{u}}(t) &= (A(t)X^{\hat{u}}(t) + B(t)\hat{u}(t))dt + (C(t)X^{\hat{u}}(t) + D(t)\hat{u}(t))dW(t) \\
X(0) &= x_0 \\
d\hat{p}_1(t) &= \Big(-\left(A(t)^\top \hat{p}_1(t) + C(t)^\top \hat{q}_1(t)\right) + 2Q(t)X^{\hat{u}}(t) + 2S(t)^\top \hat{u}(t) + \xi(t)\Big)\, dt \\
&\quad + \hat{q}_1(t)dW(t) \\
\hat{p}_1(T) &= -(M + M^\top)X^{\hat{u}}(t) - \eta
\end{aligned} \tag{1}$$

satisfy the condition

$$[\hat{u} - u]^\top \left[B(t)^\top \hat{p}_1(t) + D(t)^\top \hat{q}_1(t) + 2S_t X^{\hat{u}}(t) + R(t)\hat{u}(t) + \psi(t)\right] \geqslant 0 \tag{2}$$

almost surely, for a.e $t \in [0,T]$, and $u \in K$.

Dual Optimality Condition

The associated adjoint equation is the following linear BSDE in unknown processes $p_2, q_2 \in \mathcal{S}^2(0,T;\mathbb{R}^d)$.

Theorem 2. Let $(\hat{y}_0, \hat{\alpha}, \hat{\beta})$ be admissible dual controls. Then $(\hat{y}_0, \hat{\alpha}, \hat{\beta})$ is a solution of the dual problem if and only if the solution $Y^{\hat{y}_0,\hat{\alpha},\hat{\beta}}$ and the adjoint process $(\hat{p}_2, \hat{q}_2)$ of the FBSDE

$$\begin{aligned}
dY^{\hat{y}_0,\hat{\alpha},\hat{\beta}}(t) &= (\hat{\alpha}(t) + \tilde{A}(t)Y^{\hat{y}_0,\hat{\alpha},\hat{\beta}}(t) + \tilde{B}(t)\beta(t))dt \\
&\qquad + (\tilde{C}(t)Y^{\hat{y}_0,\hat{\alpha},\hat{\beta}}(t) + \tilde{D}(t)\hat{\beta}(t))dW(t) \\
Y^{\hat{y}_0,\hat{\alpha},\hat{\beta}}(0) &= y_0 \\
d\hat{p}_2(t) &= -\left[\tilde{A}(t)\hat{p}_2(t) + \tilde{C}(t)\hat{q}_2(t)\right] dt + \hat{q}_2(t)dW(t) \\
d\hat{p}_2(T) &= -M^{-1}(Y^{\hat{y}_0,\hat{\alpha},\hat{\beta}}(T) + \eta)
\end{aligned}$$

satisfy the conditions

$$\begin{gathered}
\hat{p}_2(0) = x_0 \\
\tilde{B}(t)^\top \hat{p}_2(t) + \tilde{D}(t)^\top \hat{q}_2(t) \in K \\
\big(\hat{p}_2(t), \tilde{B}(t)^\top \hat{p}_2(t) + \tilde{D}(t)^\top \hat{q}_2(t)\big) \in \partial \tilde{f}(\hat{\alpha}(t), \hat{\beta}(t))
\end{gathered} \tag{3}$$

Dual to Primal Relations

Theorem 3. Suppose $(\hat{y}_0, \hat{\alpha}, \hat{\beta})$ is an optimal dual control. Let $Y^{\hat{y}_0,\hat{\alpha},\hat{\beta}}, \hat{p}_2, \hat{q}_2$ be the corresponding FBSDE solutions. Define

$$\hat{u}(t) = \tilde{B}(t)^\top \hat{p}_2(t) + \tilde{D}(t)^\top q_2(t).$$

Then $\hat{u} \in K$ is an optimal control for the primal problem. Furthermore, letting $X^{\hat{u}}, \hat{p}_1, \hat{q}_1$ be the corresponding FBSDE solutions, we have the following relations.

$$\begin{aligned} X^{\hat{u}}(t) &= \hat{p}_2(t) \\ \hat{p}_1(t) &= Y^{\hat{y}_0,\hat{\alpha},\hat{\beta}}(t) \\ \hat{q}_1(t) &= \tilde{D}(t)\hat{\beta}(t) + \tilde{C}(t)Y^{\hat{y}_0,\hat{\alpha},\hat{\beta}}(t). \end{aligned}$$

Primal to Dual Relations

Theorem 4. Suppose $\hat{u} \in \mathcal{A}$ is an optimal primal control. Let $X^{\hat{u}}, \hat{p}_1, \hat{q}_1$ be the corresponding FBSDE solutions. Define

$$\begin{aligned}
\hat{y}_0 &= \hat{p}_1(0) \\
\hat{\alpha}(t) &= 2Q(t)X^{\hat{u}}(t) + 2S(t)^\top \hat{u}(t) \\
\hat{\beta}(t) &= D(t)^\top q_1(t) - B(t)^\top \hat{p}_1(t).
\end{aligned}$$

Then $(\hat{y}_0, \hat{\alpha}, \hat{\beta})$ is an optimal control for the dual problem. Furthermore, letting $Y^{\hat{y}_0, \hat{\alpha}, \hat{\beta}}, \hat{p}_2, \hat{q}_2$ be the corresponding FBSDE solutions, we have the following relations.

$$\begin{aligned}
Y^{\hat{y}_0, \hat{\alpha}, \hat{\beta}}(t) &= \hat{p}_1(t) \\
\hat{p}_2(t) &= X^{\hat{u}}(t) \\
\hat{q}_2(t) &= C(t)X^{\hat{u}}(t) + D(t)\hat{u}(t).
\end{aligned}$$

Special Case

Assume that

$$d = 1,\ m = 1,\ A(t) = r(t),\ B(t) = \mu(t) - r(t),\ C(t) = 0,\ D(t) = \sigma(t)$$

and K is a closed convex cone. Then $X(t)$ is the wealth process satisfying the SDE

$$\begin{cases} dX^{\pi}(t) = \{r(t)X^{\pi}(t) + \pi^{\intercal}(t)\sigma(t)\theta(t)\}dt + \pi^{\intercal}(t)\sigma(t)dW(t), \\ X^{\pi}(0) = x_0, \end{cases} \tag{4}$$

where $\theta(t) \triangleq \sigma^{-1}(t)\left[b(t) - r(t)\mathbf{1}\right]$ is the market price of risk at time t and $\pi(t) \in K$.

Assume there is no running cost. Then quadratic risk minimization problem is defined by

$$\begin{cases} \textit{Minimize } J(\pi(\cdot)) = E\left[\dfrac{1}{2}aX(T)^2\right], \\ \textit{Subject to } \ (X(\cdot), \pi(\cdot)) \ \textit{ is admissible.} \end{cases} \tag{5}$$

Since there is no running cost, the dual control $\alpha(t) = 0$ for all t. The dual state process satisfies the SDE

$$\begin{cases} dY(t) = -r(t)Y(t)dt + \left[\sigma^{-1}(t)\beta(t) - \theta(t)Y(t)\right]^{\intercal} dW(t) \\ Y(0) = y, \end{cases} \tag{6}$$

where $\beta(t) \in K^0 = \{v : v^T \pi \leqslant 0, \forall \pi \in K\}$, the negative polar cone of K.

The dual problem is given by

$$Minimize \ x_0 y + E\left[\frac{Y(T)^2}{2a}\right] \tag{7}$$

over $(y, \beta) \in \mathbb{R} \times \mathcal{H}^2(0, t; \mathbb{R}^N)$.

Unconstrained Case

Assume $K = \mathbb{R}^N$. Primal FBSDE for primal optimal solution is given by

$$\begin{aligned} dX^\pi(t) &= \left[r(t)X^\pi(t) + \pi(t)^\top \sigma(t)\theta(t)\right] dt + \pi(t)^\top \sigma(t) dW(t) \\ X^\pi(0) &= x_0 \\ dp_1(t) &= -r(t)p_1(t)dt + q_1(t)^\top dW(t) \\ p_1(T) &= -aX^\pi(t). \end{aligned}$$

The necessary and sufficient optimality condition for optimal control $\hat{\pi}$ is

$$p_1(t)\sigma(t)\theta(t) + \sigma(t)q_1(t) = 0 \tag{8}$$

almost surely, for a.e $t \in [0, T]$. It is difficult to solve the above FBSDE. Note that (8 implies

$$dp_1(t) = -r(t)p_1(t)dt - p_1(t)\theta(t)^\top dW(t).$$

We now use the dual method to solve it. Since $K = \mathbb{R}^N$, then $K^0 = \{0\}$, so $\beta(t) = 0$ for $\forall t \in [0, T]$ a.e. The dual problem is

$$\min_y \left\{ x_0 y + E\left[\frac{Y(T)^2}{2a}\right] \right\}.$$

Dual FBSDE for dual optimal solution is given by

$$\begin{aligned}
dY(t) &= -r(t)Y(t)dt + \left[-\theta(t)Y(t)\right]^\top dW(t) \\
Y(0) &= y \\
dp_2(t) &= \left[r(t)p_2(t) + q_2(t)\theta(t)\right] dt + q_2(t)^\top dW(t) \\
p_2(T) &= -\frac{Y(T)}{a}
\end{aligned}$$

The necessary and sufficient optimality conditions for optimal control $\hat{y}$ is

$$p_2(0) = x_0$$

almost surely. We have

$$Y(t) = y\Gamma(t),$$

where Γ satisfies the linear SDE

$$d\Gamma(t) = \Gamma(t)[-r(t)dt - \theta^\intercal(t)dW(t)],\ \Gamma(0) = 1.$$

The dual objective function is a quadratic function of y:

$$x_0 y + y^2 E\left[\frac{\Gamma(T)^2}{2a}\right].$$

The minimum point $\hat{y}$ is given by

$$\hat{y} = -\frac{x_0}{E\left[\frac{\Gamma(T)^2}{a}\right]}.$$

Optimal dual process is $\hat{Y}(t) = \hat{y}\Gamma(t)$. Let $(\hat{p}_2, \hat{q}_2)$ be adjoint process associated with optimal control $(\hat{y}, 0)$, satisfying BSDE

$$d\hat{p}_2(t) = [r(t)\hat{p}_2(t) + \theta^\intercal(t)\hat{q}_2(t)]dt + \hat{q}_2^\intercal(t)dW(t),\ \hat{p}_2(T) = -\frac{\hat{Y}(T)}{a},$$

Hence we obtain that

$$\hat{p}_2(t) = \Gamma(t)^{-1}E\left[\Gamma(T)\hat{p}_2(T)\middle|\mathcal{F}_t\right] = -\hat{y}\Gamma(t)^{-1}E\left[\frac{1}{a}\Gamma(T)^2\middle|\mathcal{F}_t\right],$$

which shows that $\hat{p}_2(t) \neq 0$ for $t \in [0, T]$ a.e. if and only if $\hat{y} \neq 0$.

Define a process $P(t) \triangleq -\frac{\hat{Y}(t)}{\hat{p}_2(t)}$, $\forall t \in [0, T]$. Applying Ito's formula, we obtain

$$dP(t) = -2r(t)P(t)dt + \frac{1}{\hat{p}_2(t)^2} P(t)\hat{q}_2(t)^T \hat{q}_2(t)dt - P(t)\left(\theta(t) + \frac{1}{\hat{p}_2(t)}\hat{q}_2(t)\right)^T dW(t).$$

Define a process

$$\Lambda(t) \triangleq \frac{\hat{q}_2(t)\hat{Y}(t)}{\hat{p}_2(t)^2} + \frac{\theta(t)\hat{Y}(t)}{\hat{p}_2(t)}.$$

Substituting Λ into the above equation and rearranging, we have

$$dP(t) = \left[-2r(t)P(t) + 2\theta^{\intercal}(t)\Lambda(t) + \theta^T(t)\theta(t)P(t) + \frac{\Lambda^{\intercal}(t)\Lambda(t)}{P(t)}\right] dt + \Lambda^{\intercal}(t)dW(t),$$

which is the stochastic Riccati equation (SRE) introduced in Lim-Zhou (2002). Using the dual approach, we obtain an explicit representation of the unique solution to the SRE.

Dual-Primal Relations

Let $(\hat{y}, 0, 0)$ be optimal dual controls. Then the optimal primal control and processes are given by

$$\begin{aligned}
\pi(t) &= \sigma^{-1}(t)^{\top} q_2(t) \\
X^{\pi}(t) &= p_2(t) \\
p_1(t) &= Y^{(\hat{y},0,0)}(t) \\
q_1(t) &= -\theta(t) Y^{(\hat{y},\hat{\alpha},\hat{\beta})}(t)
\end{aligned}$$

Conversely, suppose $\pi(t)$ is an optimal primal control. Then the optimal dual controls and processes are given by

$$\begin{aligned}
y &= p_1(0) \\
\alpha(t) &= 0 \\
\beta(t) &= 0 \\
Y^{(y,\alpha,\beta)}(t) &= p_1(t) \\
p_2(t) &= X^{\hat{\pi}}(t) \\
q_2(t) &= \sigma(t)^{\top} \hat{\pi}(t)
\end{aligned}$$

Constrained Case

Assume K is a closed convex cone. Heunis-Labbe (2007) states that there exists optimal control $(\hat{y}, \hat{\beta})$ to (7) with assiciated optimal state process $\hat{Y}$. Define $\hat{\gamma}$ by $\hat{\beta}(t) = \hat{\gamma}(t)\hat{Y}(t)$ for $\forall t \in [0, T]$, we obtain that $\hat{Y}$ follows the SDE

$$\begin{cases} d\hat{Y}(t) = -r(t)\hat{Y}(t)dt + \left[\sigma^{-1}(t)\hat{\gamma}(t) - \theta(t)\right]^{\intercal} \hat{Y}(t)dW(t) \\ \hat{Y}(0) = \hat{y}. \end{cases}$$

Hence, we have

$$\hat{Y}(t) = \hat{y}\hat{H}(t),$$

where

$$\hat{H}(t) \triangleq \exp\left(\int_0^t \left\{ -r(s) - \frac{1}{2}\left[\sigma^{-1}(s)\hat{\gamma}(s) - \theta(s)\right]^{\intercal} \left[\sigma^{-1}(s)\hat{\gamma}(s) - \theta(s)\right] \right\} ds \right.$$
$$\left. + \left[\sigma^{-1}(s)\hat{\gamma}(s) - \theta(s)\right]^{\intercal} dW(s) \right).$$

By the dual FBSDE, we obtain

$$\hat{p}_2(0) = E\left[\Gamma(T)\hat{p}_2(T)\right] = E\left[-\Gamma(T)\frac{\hat{Y}(T)}{a}\right] = -\hat{y}E\left[\Gamma(T)\frac{\hat{H}(T)}{a}\right] = x_0,$$

which implies

$$\hat{y} = -\frac{x_0}{E\left[\dfrac{\Gamma(T)\hat{H}(T)}{a}\right]}.$$

Moreover, we have

$$\hat{p}_2(t) = \Gamma(t)^{-1}E\left[-\Gamma(T)\frac{\hat{Y}(T)}{a}\middle|\mathcal{F}_t\right] = -\hat{y}\Gamma(t)^{-1}E\left[\Gamma(T)\frac{\hat{H}(T)}{a}\middle|\mathcal{F}_t\right],$$

which shows that $\hat{p}_2(t) \neq 0$ $\mathbb{P}$-a.e. for $\forall t \in [0,T]$.

If $x_0 > 0$ then $\hat{Y}(t) < 0$ and $\hat{p}_2(t) > 0$ for $\forall t \in [0,T]$, $\mathbb{P}$-a.e. Define

$$P_+(t) \triangleq -\frac{\hat{Y}(t)}{\hat{p}_2(t)} = -\frac{\hat{p}_1(t)}{\hat{X}(t)},\ \forall t \in [0,T].$$

Applying Ito's formula, we have

$$dP_+(t) = \left\{-2r(t)P_+(t) - \hat{\xi}_+^\intercal(t)\left[\sigma(t)\theta(t)P_+(t) + \sigma(t)\Lambda_+(t)\right]\right\}dt + \Lambda_+(t)dW(t),$$

where

$$\Lambda_+(t) \triangleq -\frac{\hat{q}_1(t)}{\hat{X}(t)} - \frac{P_+(t)\pi^\intercal(t)\sigma(t)}{\hat{X}(t)},\ \hat{\xi}_+(t) \triangleq \frac{\hat{\pi}(t)}{\hat{X}(t)}.$$

Define the following functions:

$$\begin{aligned} H_+(t, v, P, \Lambda) &\triangleq v^\intercal P\sigma(t)\sigma^\intercal(t)v + 2v^\intercal \left[\sigma(t)\theta(t)P + \sigma(t)\Lambda\right], \\ H_+^*(t, P, \Lambda) &\triangleq \inf_{v\in K} H_+(t, v, P, \Lambda). \end{aligned}$$

After some technical discussions involving Clarke (1990) nonsmooth optimization, we can show that P_+ is the solution to the following nonlinear BSDE

$$\begin{cases} dP_+(t) = -\left[2r(t)P_+(t) + H_+^*(t, P_+(t), \Lambda_+(t))\right] dt + \Lambda_+^\intercal(t)dW(t), \\ P_+(T) = a, \\ P_+(t) > 0. \ \forall t \in [0, T]. \end{cases} \tag{9}$$

If $x_0 < 0$, then we can define $P_-(t)$ that satisfies a similar nonlinear BSDE to that of (9). Using the dual approach, we have obtained an explicit representation of the unique solution to the ESREs, introduced in Hu-Zhou (2005), in terms of optimal dual state and adjoint processes. Optimal solution to the primal problem is given by

$$\begin{cases} \hat{\pi}^\intercal(t) = [\sigma^\intercal]^{-1}(t)\hat{q}_2(t), \\ \hat{X}(t) = \hat{p}_2(t) = -\hat{Y}(t)\left[\dfrac{1_{\{x_0>0\}}}{P_+(t)} + \dfrac{1_{\{x_0<0\}}}{P_-(t)}\right]. \end{cases}$$

Deep Reinforcement Learning

Suppose we want to approximate a function $\varphi\colon \mathbb{R}^p \to \mathbb{R}^q$ for some $p, q \in \mathbb{N}$. We use the following algorithm to construct an approximate function.

We take several elements of $\mathbb{R}^p$, say a batch of size $k \in \mathbb{N}$. For our purposes, k will be how many Brownian paths we generate. There is number of 'layers' $L \in \mathbb{N}$. Suppose we have a matrix $X \in \mathbb{R}^{k\times p}$ as our input. Set $N_0 = X$, then for $i = 1, \ldots, L$, perform the update

$$N_i = h(N_{i-1}M_i + v_i)$$

where $M_1 \in \mathbb{R}^{p\times\ell}$, $M_j \in \mathbb{R}^{\ell\times\ell}$ for $j > 1$, $v_i \in \mathbb{R}^{k\times\ell}$ and $h\colon \mathbb{R}^{k\times\ell} \to \mathbb{R}^{k\times\ell}$ applies the non-linear function $x \mapsto \max(x, 0)$ element-wise. Finally, we output

$$N_\theta(X) = N_L M_{L+1} + v_{L+1}$$

where $M_{L+1} \in \mathbb{R}^{\ell\times q}$ and $v_{L+1} \in \mathbb{R}^{k\times q}$. The parameter vector θ representing our network is M_i and v_i matrices, and is optimised at the training step. We do not know a priori the function φ, so we train the weights against some other performance metric. In our case, this metric is the Hamiltonian of the control problem, or the squared discrepancy of the terminal condition. This is what is known as reinforcement learning.

Solving Dual FBSDE via Deep Learning

Assume K is the whole space. Recall dual FBSDE for optimal dual solution is given by

$$
\begin{aligned}
dY(t) &= [-r(t)Y(t)]\,dt + [-\theta(t)Y(t)]^\top\,dW(t)\\
Y(0) &= y\\
dp_2(t) &= [r(t)p_2(t) + q_2(t)\theta(t)]\,dt + q_2(t)^\top dW(t)\\
p_2(T) &= -\frac{Y(T)}{a}\\
p_2(0) &= x_0.
\end{aligned}
$$

To solve this problem via machine learning we search for a control

$$q_2 = q_2(t, Y(t))$$

which is function determined by neural networks and depending on a finite number of parameters which are chosen to ensure optimality condition hold, and BSDEs is solved. We simulate all processes in the forward direction, set $p_2(0) = x_0$ and choose the controls such that

$$y, q_2 \in \arg\min_{y,q_2} E\left[\left|p_2(T) + \frac{Y(T)}{a}\right|^2\right].$$

Solving Primal FBSDE via Deep Learning

Recall primal FBSDE for optimal primal solution is given by

$$
\begin{aligned}
dX^\pi(t) &= \left[r(t)X^\pi(t) + \pi(t)^\top\sigma(t)\theta(t)\right]dt + \pi(t)^\top\sigma(t)dW(t)\\
X^\pi(0) &= x_0\\
dp_1(t) &= -r(t)p_1(t)dt - p_1(t)\theta(t)^\top dW(t)\\
p_1(T) &= -aX^\pi(t).
\end{aligned}
$$

To solve this problem via machine learning we search for a control

$$\pi = \pi(t, X^\pi(t))$$

which is function determined by neural networks and depending on a finite number of parameters which are chosen to ensure optimality condition hold, and BSDEs is solved. We simulate all processes in the forward direction, and choose the controls such that

$$p_1(0), \pi \in \arg\min_{p_1(0),\pi} E\left[\left|p_1(T) + aX^\pi(t)\right|^2\right].$$

Each if these minimisations is performed using the ADAM algorithm, which is a variant of stochastic gradient descent.

Explicit Dual Solution for Unconstrained Problem

The state process is simply a geometric Brownian motion, with solution

$$Y(t) = y \exp\left(\int_0^t \left(-r(s) - \frac{1}{2}|\theta(s)|^2\right) ds - \int_0^t \theta(s)^\top dW(s)\right).$$

Under some integrability condition on theta, there exists a measure Q such that the process

$$W_Q(t) := \int_0^t \theta(s)ds + W(t)$$

is a Q- Brownian motion. p_2 has the following condition expectation representation:

$$p_2(t) = e^{-\int_t^T r(s)ds} E^Q\left[p_2(T)|\mathcal{F}_t\right] = -\frac{e^{-\int_t^T r(s)ds}}{a} E^Q\left[Y(T)|\mathcal{F}_t\right]$$
$$= -\frac{e^{-\int_t^T r(s)ds}}{a} y \exp\left(\int_0^T \left(-r(s) + \frac{1}{2}|\theta(s)|^2\right) ds + \int_t^T \frac{1}{2}|\theta(s)|^2 ds - \int_0^t |\theta(s)|^2 ds - \int_0^t \theta(s)^\top dW(s)\right)$$

Finally, the optimality condition $p_2(0) = x_0$ requires us to take

$$y = -a x_0 e^{\int_0^T r(s)ds} \exp\left(\int_0^T \left(r(s) - |\theta(s)|^2\right) ds\right)$$

Numerical Tests

We take the processes r, σ, b to be constant, and set $T = 0.5$. We simulate 3 paths of each process with 50 time steps, and compare the processes which the optimality relations say must be equal. The exact analytical path-wise solutions are also given in red.

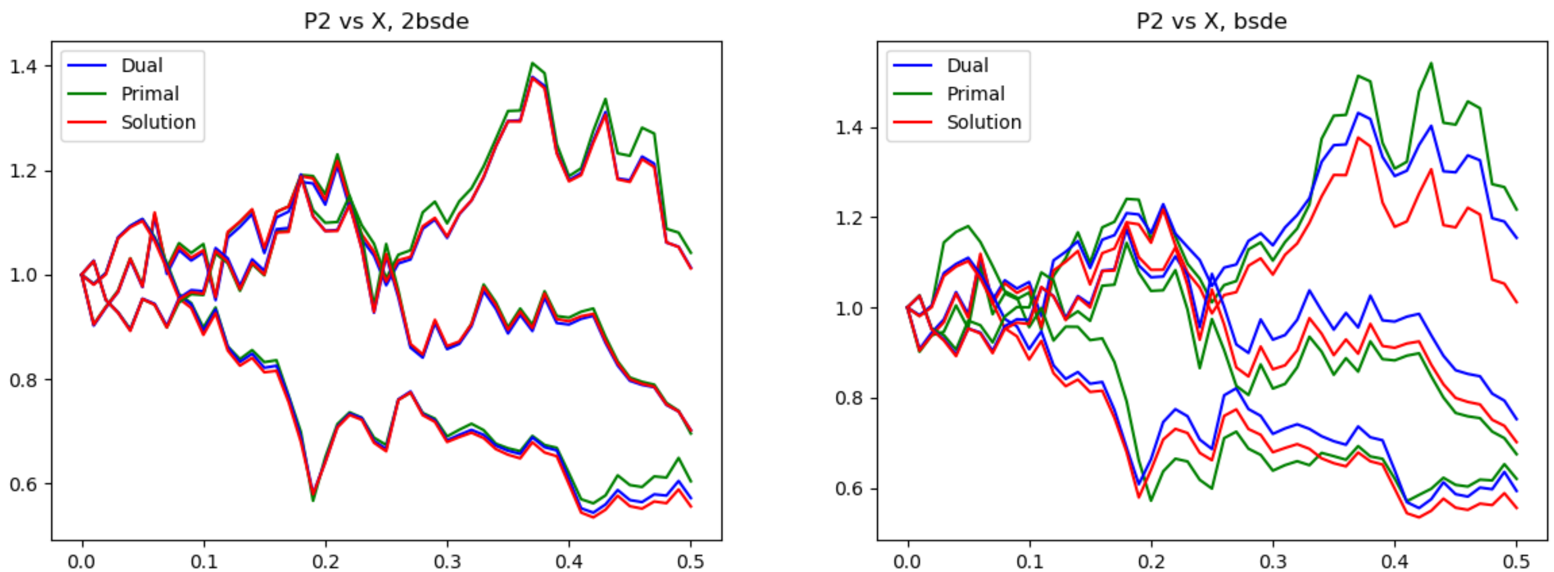

Figure 1: Comparison of the primal state process X and the dual adjoint process p_2.

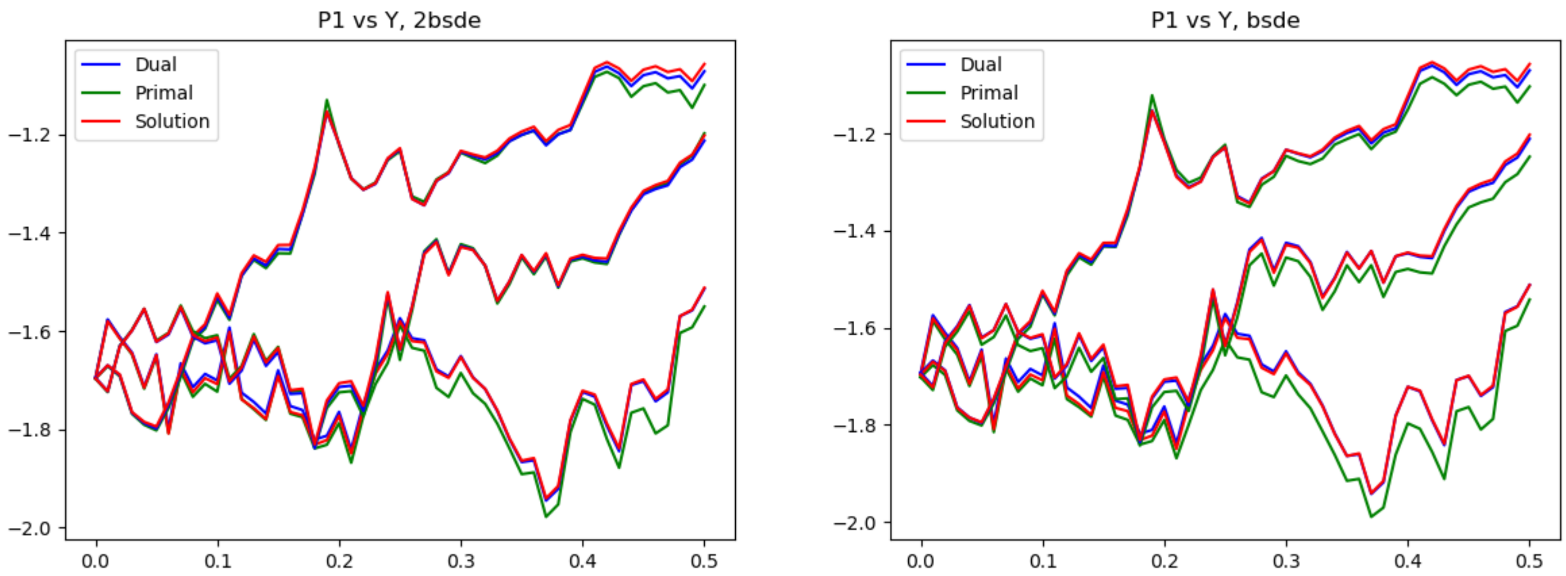

Figure 2: Comparison of the dual state process Y and the dual adjoint process p_1.

Conclusions

- We study a constrained LQ problem with random market parameters.
- We characterise optimal conditions for both the primal and dual problem in terms of FBSDEs plus additional conditions, and their relations.
- We apply the results to solve both unconstrained and cone-constrained quadratic risk minimization problems. Solutions to SREs can be recovered from optimal solutions to dual problem.
- We suggest to use deep reinforcement learning to solve FBSDEs and show a numerical example for a simple case.