Robust deep learning algorithms for coupled FBSDEs

Kristoffer Andersson, Adam Andersson and Kees Oosterlee[1]

[1]Utrecht University

COST Action Stochastica, Cork, Ireland, April 2026

Agenda

- Coupled BSDEs, stochastic control
- Deep BSDE algorithm issues
- Proposed alternative algorithm for coupled BSDEs
- Numerical examples

Stochastic control $\rightarrow$ FBSDE

Stochastic control problem:

$$\mathrm{d}X_t^u = \bar{b}(t, X_t^u, u_t)\mathrm{d}t + \sigma(t, X_t^u)\mathrm{d}W_t,$$

$$\mathrm{J}^u(t,x) = \mathbb{E}\Big[\int_t^T \bar{f}(s, X_s^u, u_s)\mathrm{d}s + g(X_T^u)\Big]$$

V(t,x)=$\inf_u J^u(t,x)$

Connection via Itô: $Y_t = V(t, X_t), \qquad Z_t = \nabla_x V(t, X_t)^\top \sigma(t, X_t)$

$$\Rightarrow \begin{cases} X_t = x_0 + \int_0^t b(s, X_s, Z_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\mathrm{d}W_s, \\ Y_t = g(X_T) + \int_t^T f(s, X_s, Z_s)\mathrm{d}s - \int_t^T Z_s \mathrm{d}W_s \end{cases}$$

- optimal control $\longleftrightarrow$ FBSDE solution (X, Y, Z)
- Initial value $Y_0 = V(0, x_0)$

Reformulation of the FBSDE to a stochastic control problem

The Deep BSDE method[1] uses the following reformulation of a FBSDE

$$\begin{cases} \inf_{y_0,(z_t)_{t\in[0,T]}} \mathbb{E}|Y_T - g(X_T)|^2, \quad \text{subject to} \\ X_t = x_0 + \int_0^t b(s, X_s, z_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\mathrm{d}W_s, \\ Y_t = y_0 - \int_0^t f(s, X_s, z_s)\mathrm{d}s + \int_0^t z_s \mathrm{d}W_s, \end{cases} \tag{1}$$

compared to

$$\begin{cases} X_t = x_0 + \int_0^t b(s, X_s, z_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\mathrm{d}W_s, \\ Y_t = g(X_T) + \int_t^T f(s, X_s, z_s)\mathrm{d}s - \int_t^T z_s \mathrm{d}W_s, \end{cases} \tag{2}$$

Motivation:

- A solution to (2) solves (1);
- By wellposedness of the FBSDE, the solution is unique.

[1] A. Jentzen et. al. (2018): 8505-8510.

After time discretization, we solve: optimize over y_0 and ζ

Fix y_0, optimize controls ζ (NN approximation of $\bar{z}_k$)

$$\mathrm{MSE}(y_0) = \inf_{\zeta} \mathbb{E}\left[\left| Y_N^{h,y_0,\zeta} - g(X_N^{h,y_0,\zeta}) \right|^2 \right]$$

If method works: $\arg\min_{y_0} \mathrm{MSE}(y_0) \approx Y_0$

What do we observe?

- LQ problems: clear minimum ✓
- Strong coupling: flat / misleading landscape $\times$

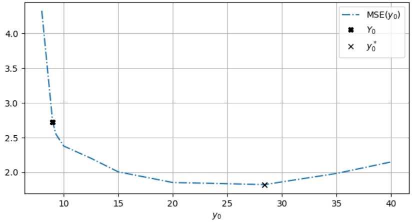

Result Deep FBSDE method

Although the initial condition is significantly incorrect, the terminal condition is well approximated, i.e., $g(X_N^{h,y_0,\zeta}) - Y_N^{h,y_0,\zeta}$ is small both on the pathwise level and in the mean-squared sense.

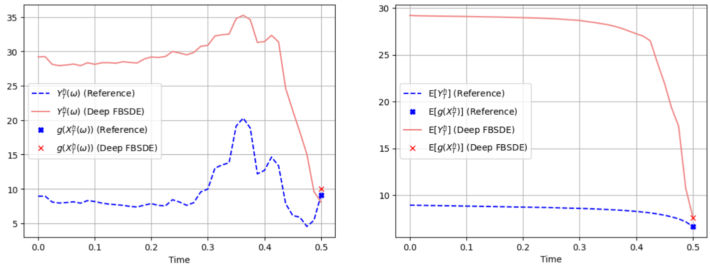

Figure: Left: One representative path. Right: A sample mean of size 2^{12}.

$\Rightarrow$ One can construct low-dimensional (e.g., $d = 2$) decoupled BSDEs with smooth solutions to the PDE, where the deep FBSDE method fails.

Deep BSDE: what does convergence theory say?

Available convergence results

- Han & Long (2020): *a posteriori* error estimates for the deep BSDE method applied to coupled FBSDEs
- Generalized convergence analysis: Négyesi, Huang & CWO (2025)

Structural assumptions behind convergence

- Lipschitz continuity of b, σ, f, g in (x, y, z)
- Stability / well-posedness of the FBSDE
- Moderate coupling:
  - typically: forward equation depends on (X, Y), not strongly on Z
- Existence of contraction-type constants $(\bar{A}, \bar{B} < 1)$

If these conditions hold, then:

$$\text{small terminal loss } \Rightarrow \text{small error in } (Y, Z)$$

$\Rightarrow$ Convergence is guaranteed *within this regime*

References: Han & Long (2020); Négyesi et al. (2024/2025)

Why do we still observe non-convergence?

Violation of structural assumptions

- Strong coupling (e.g. stochastic control): $b = b(t, X, Z)$
- Driver f highly nonlinear or only locally Lipschitz
- Gradients $Z_t = \nabla_x V \, \sigma$ can be large or unstable

Breakdown of contraction conditions

- Constants in the convergence analysis $(\bar{A}, \bar{B})$ are not < 1
- $\Rightarrow$ no theoretical guarantee of convergence

Optimization pathology

- Loss only enforces: $Y_T \approx g(X_T)$
- But does *not uniquely identify* the correct trajectory (Y_t, Z_t)

Discretization vs. optimization

- $\Rightarrow$ No convergence as $N \to \infty$ in practice

General Problem Setting Recap

- Let $(W_t)_{t\in[0,T]}$ be a k-D BM on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P})$;
- $b\colon [0,T]\times\mathbb{R}^d\times\mathbb{R}\times\mathbb{R}^d \to \mathbb{R}^d$, $\sigma\colon [0,T]\times\mathbb{R}^d \to \mathbb{R}^{d\times k}$, $g\colon \mathbb{R}^d\to\mathbb{R}$ and $f\colon [0,T]\times\mathbb{R}^d\times\mathbb{R}\times\mathbb{R}^d \to \mathbb{R}$.
- The general form of the coupled FBSDE system, $t\in[0,T]$:

$$\begin{aligned}
X_t &= x_0 + \int_0^t b(s, X_s, Y_s, Z_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\mathrm{d}W_s,\\
Y_t &= g(X_T) + \int_t^T f(s, X_s, Y_s, Z_s)\mathrm{d}s - \int_t^T \langle Z_s, \sigma(s, X_s)\mathrm{d}W_s\rangle.
\end{aligned}$$

- From here on, we write $Z_t = D_x v(t, X_t)$ so the martingale term is $\langle Z_s, \sigma(s, X_s)\mathrm{d}W_s\rangle$. Earlier Z_t absorbed σ.

Associated PDE

Associated to our coupled FBSDE system is a semilinear parabolic PDE:

$$\frac{\partial v}{\partial t}(t,x) + \mathcal{L}v(t,x) + f\big(t,x,v(t,x),\mathrm{D}_x v(t,x)\big) = 0, \quad (t,x) \in [0,T) \times \mathbb{R}^d,$$
$$v(T,x) = g(x), \quad x \in \mathbb{R}^d, \text{ with:}$$

$$\mathcal{L}\rho(t,x) = b\big(t,x,\rho(t,x),\mathrm{D}_x\rho(t,x)\big)^\top \mathrm{D}_x\rho(t,x) + \frac{1}{2}\mathrm{Tr}\big(a(t,x)\mathrm{D}_x^2\rho(t,x)\big)$$
$$= \sum_{i=1}^{d} b_i\big(t,x,\rho(t,x),\mathrm{D}_x\rho(t,x)\big)\partial_i\rho(t,x) + \frac{1}{2}\sum_{i,j=1}^{d} a_{ij}(t,x)\partial_{i,j}\rho(t,x),$$

with $a = \sigma\sigma^\top$, $\mathrm{Tr}(\cdot)$ the trace of a matrix. The connection is established by the nonlinear Feynman–Kac formula:

$$Y_t = v(t,X_t), \text{ and } Z_t = \mathrm{D}_x v(t,X_t).$$

Deep multi-FBSDE method

For sufficiently regular $\psi\colon [0,T] \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \to \mathbb{R}^d$, it holds, for $Y_t^\psi = v(t, X_t^\psi)$ and $Z_t^\psi = \mathrm{D}_x v(t, X_t^\psi)$ that

$$\begin{aligned} X_t^\psi =& x_0 + \int_0^t \big(b(s, X_s^\psi, Y_s^\psi, Z_s^\psi) - \psi(s, X_s^\psi, Y_s^\psi, Z_s^\psi)\big)\mathrm{d}s + \int_0^t \sigma(s, X_s^\psi)\mathrm{d}W_s \\ Y_t^\psi =& g(X_T^\psi) + \int_t^T \big(f(s, X_s^\psi, Y_s^\psi, Z_s^\psi) + \langle \psi(s, X_s^\psi, Y_s^\psi, Z_s^\psi), Z_s^\psi \rangle\big)\mathrm{d}s \\ & - \int_t^T \langle Z_s^\psi, \sigma(s, X_s^\psi)\mathrm{d}W_s \rangle, \quad t \in [0,T]. \end{aligned} \tag{3}$$

This defines a family of FBSDEs, which are all equivalent to the PDE.

For all ψ we have $Y_0^\psi = v(0, X_0^\psi) = v(0, x_0)$.

Deep multi-FBSDE method

Introducing the shorthand notation

$$
\begin{aligned}
b^{\psi}(t,x,y,z) &:= b(t,x,y,z) - \psi(t,x,y,z)\\
f^{\psi}(t,x,y,z) &:= f(t,x,y,z) + \langle z, \psi(t,x,y,z)\rangle,
\end{aligned}
$$

we have for any ψ the variational formulation

$$
\begin{cases}
\underset{y_0,\zeta}{\text{minimize}}\ \mathbb{E}\left|Y_T^{y_0,\zeta,\psi} - g(X_T^{y_0,\zeta,\psi})\right|^2, \quad \text{where}\\
X_t^{y_0,\zeta,\psi} = x_0 + \int_0^t b^{\psi}(s, X_s^{y_0,\zeta,\psi}, Y_s^{y_0,\zeta,\psi}, Z_s^{y_0,\zeta,\psi})\mathrm{d}s + \int_0^t \sigma(s, X_s^{y_0,\zeta,\psi})\mathrm{d}W_s\\
Y_t^{y_0,\zeta,\psi} = y_0 - \int_0^t f^{\psi}(s, X_s^{y_0,\zeta,\psi}, Y_s^{y_0,\zeta,\psi}, Z_s^{y_0,\zeta,\psi})\mathrm{d}s\\
\qquad\qquad + \int_0^t \left\langle Z_s^{y_0,\zeta,\psi}, \sigma(s, X_s^{y_0,\zeta,\psi})\mathrm{d}W_s\right\rangle,\\
Z_t^{y_0,\zeta,\psi} = \zeta(t, X_t^{y_0,\zeta,\psi}), \quad t \in [0,T].
\end{cases}
\tag{4}
$$

The deep multi-BSDE scheme, Phase I

Let ζ represent $\zeta_0, \ldots, \zeta_{N-1}$.

Phase I - Approximation of Y_0:

$$
\begin{cases}
\underset{y_0,\zeta_0,\ldots,\zeta_{N-1}}{\text{minimize}} \ \sum_{i=1}^{K} \mathbb{E}\left| Y_N^{h,y_0,\zeta,\psi_i} - g(X_N^{h,y_0,\zeta,\psi_i})\right|^2, \quad \text{where for } \psi \in \{\psi_1, \ldots, \psi_K\} \\
X_{n+1}^{h,y_0,\zeta,\psi} = X_n^{h,y_0,\zeta,\psi} + b^{\psi}\big(t_n, X_n^{h,y_0,\zeta,\psi}, Y_n^{h,y_0,\zeta,\psi}, Z_n^{h,y_0,\zeta,\psi}\big)h \\
\qquad\qquad\quad + \sigma(t_n, X_n^{h,y_0,\zeta,\psi})\Delta W_n, \\
Y_{n+1}^{h,y_0,\zeta,\psi} = Y_n^{h,y_0,\zeta,\psi} - f^{\psi}\big(t_n, X_n^{h,y_0,\zeta,\psi}, Y_n^{h,y_0,\zeta,\psi}, Z_n^{h,y_0,\zeta,\psi}\big)h \\
\qquad\qquad\quad + \big\langle Z_n^{h,y_0,\zeta,\psi}, \sigma\big(t_n, X_n^{h,y_0,\zeta,\psi}\big)\Delta W_n\big\rangle, \\
Z_n^{h,y_0,\zeta,\psi} = \zeta_n(X_n^{h,y_0,\zeta,\psi}), \ X_0^{h,y_0,\zeta,\psi} = x_0, \ Y_0^{h,y_0,\zeta,\psi} = y_0, \\
\text{for } n \in \{0, 1, \ldots, N-1\}.
\end{cases} \tag{5}
$$

The deep multi-BSDE scheme, Phase II

Phase II - Approximation of the full FBSDE with known initial value:

$$\begin{cases} \underset{\zeta_0,\ldots,\zeta_{N-1}}{\operatorname{minimize}} \, \mathbb{E}\left|Y_N^{h,\zeta} - g(X_N^{h,\zeta})\right|^2, \quad \text{where} \\ X_{n+1}^{h,\zeta} = X_n^{h,\zeta} + b\big(t_n, X_n^{h,\zeta}, Y_n^{h,\zeta}, Z_n^{h,\zeta}\big)h + \sigma(t_n, X_n^{h,\zeta})\Delta W_n, \\ Y_{n+1}^{h,\zeta} = Y_n^{h,\zeta} - f\big(t_n, X_n^{h,\zeta}, Y_n^{h,\zeta}, Z_n^{h,\zeta}\big)h + \big\langle Z_n^{h,\zeta}, \sigma\big(t_n, X_n^{h,\zeta}\big)\Delta W_n \big\rangle, \\ Z_n^{h,\zeta} = \zeta_n\big(X_n^{h,\zeta}\big), \;\; Y_0^{h,\zeta} = y_0^*, \quad \text{for } n \in \{0, 1, \ldots, N-1\}. \end{cases} \tag{6}$$

Choosing ψ_i to be constant in time (or Hölder–$1/2$ continuous) and Lipschitz continuous in the spatial variables leads to stable optimization; it preserves the qualitative regularity of b^ψ and f^ψ.

More on the choice of ψ

If Phase I yields consistent Y_0 across different ψ-families, this indicates a reliable approximation of Y_0.

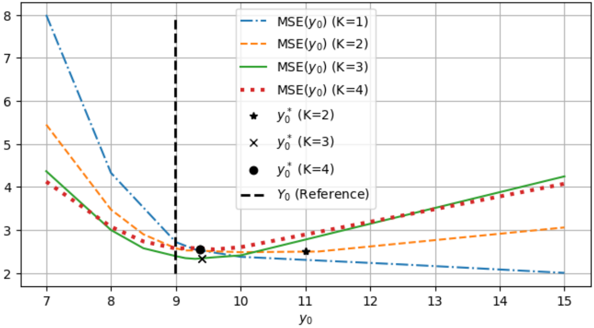

MSE curves for $K = 1, \ldots, 4$, with $\psi_1 = 0$, $(\psi_1, \psi_2) = (0, b_1)$, $(\psi_1, \psi_2, \psi_3) = (0, b_1, -b_1)$ and $(\psi_1, \psi_2, \psi_3, \psi_4) = (0, b_1, -b_1, -0.5b_1)$, with $b(x, y, z) = b_1(x) + b_2(z)$. For visualization, we have scaled curves.

Empirical convergence for Y_0 approximations

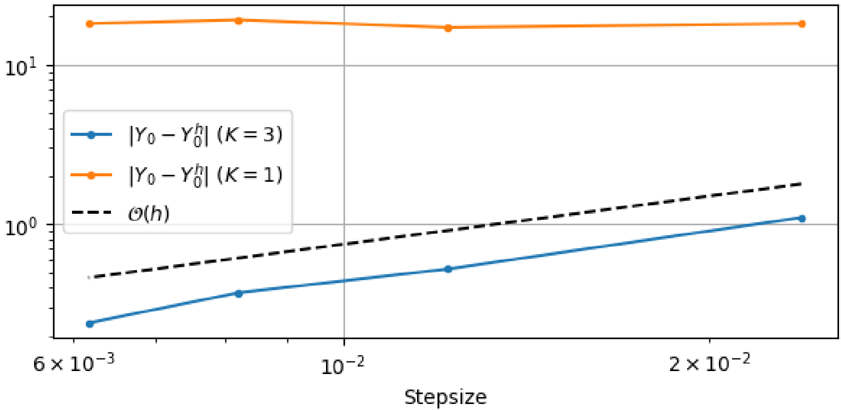

The deep multi-FBSDE method Phase I. $K = 1$ coincides with the original deep FBSDE method.

$\Rightarrow$ We use standard fully connected neural networks (ReLU); the observed behaviour is not driven by architectural choices.

Cole–Hopf transformation

Idea: remove quadratic gradient nonlinearity

$$\partial_t v + \alpha \Delta v + 2\beta \|\nabla v\|^2 - \gamma v = 0, \qquad v(T, x) = g(x)$$

Transform:

$$\rho(t,x) = \exp\Big(- \frac{2\beta}{\alpha} v(t,x) \Big) \quad \Longrightarrow \quad v = -\frac{\alpha}{2\beta} \log \rho$$

Result: $\partial_t \rho + \alpha \Delta \rho - \gamma \rho \log \rho = 0$

Example (controlled Brownian motion):

$$\partial_t v + \frac{\sigma^2}{2} \Delta v - \frac{1}{2r} |\nabla v|^2 = 0 \quad \Rightarrow \quad \partial_t \rho + \frac{\sigma^2}{2} \Delta \rho = 0$$

$$Y_t = -r\sigma^2 \log \mathbb{E}\Big[e^{-g(X_T)/(r\sigma^2)} \mid X_t \Big]$$

- nonlinear PDE $\rightarrow$ (semi-)linear PDE
- enables probabilistic (Feynman–Kac) representations
- analytical benchmarks and intuition

Controlled Brownian motion and Cole–Hopf

Stochastic control problem:

$$\mathrm{d}X_t = u_t \mathrm{d}t + \sigma \mathrm{d}W_t, \qquad J(u; t, x) = \mathbb{E}\Big[\int_t^T \tfrac{r}{2} \|u_s\|^2 \mathrm{d}s + g(X_T) \,\Big|\, X_t = x \Big]$$

$$v(t, x) = \inf_u J(u; t, x), \qquad u^* = -\tfrac{1}{r} \nabla v$$

Associated FBSDE:

$$X_t = x_0 - \int_0^t \tfrac{1}{r} Z_s \,\mathrm{d}s + \sigma W_t$$

$$Y_t = g(X_T) + \int_t^T \tfrac{1}{2r} \|Z_s\|^2 \mathrm{d}s - \int_t^T \langle Z_s, \sigma \mathrm{d}W_s \rangle$$

Cole–Hopf $\Rightarrow$ explicit representation:

$$Y_t = -r\sigma^2 \log \mathbb{E}\Big[\exp\Big(-\tfrac{1}{r\sigma^2} g(X_t + \sigma\sqrt{T-t}\,\xi) \Big) \,\Big|\, X_t \Big]$$

Deep multi-FBSDE Method Specific Settings

We opt for $K = 2$, setting $\psi_1(t,x,y,z) = 0$ and $\psi_2(t,x,y,z) = \frac{r_1^2}{2r_2} z$. So, $(X^{\psi_1}, Y^{\psi_1}, Z^{\psi_1})$ is the solution satisfying the following decoupled equation:

$$
\begin{aligned}
X_t^{\psi_2} &= x_0 + \sigma W_t, \\
Y_t^{\psi_2} &= g(X_T^{\psi_2}) - \int_t^T \frac{1}{2r} \|Z_s^{\psi_2}\|^2 \mathrm{d}s - \int_t^T \left\langle Z_s^{\psi_2}, \sigma \mathrm{d}W_s \right\rangle, \quad t \in [0, T].
\end{aligned}
$$

$N = 40$, $T = 0.5$, $d = \ell = k = 2$, $r = 1$, $\sigma = 0.25$; $x = (x_1, x_2)^\top$, $x_0 = (-0.1, 0.1)^\top$, $g(x) = -|x_1 - x_2|$.

$\Rightarrow$ Control the components of Brownian motion such that they diverge as much as possible by time $t = T$, while minimizing the use of control force.

Loss and approximation of Y_0 during training

The loss approaches zero for both methods, but only the deep multi-FBSDE method converges to the correct Y_0.

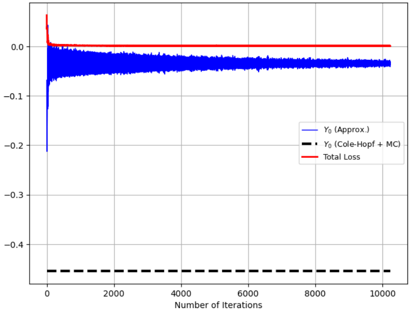

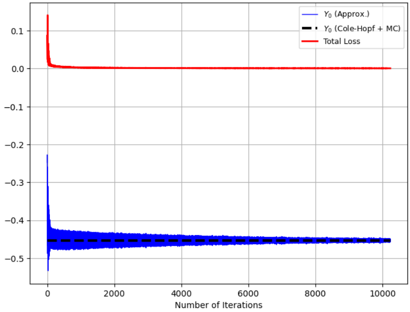

Loss and Y_0 for the deep FBSDE method (left) and the deep multi-FBSDE method (right). "Number of iterations" refers to gradient updates.

The deep FBSDE method fails to converge to the true solution. To gain some insight, the sample mean and a representative path of X and Y:

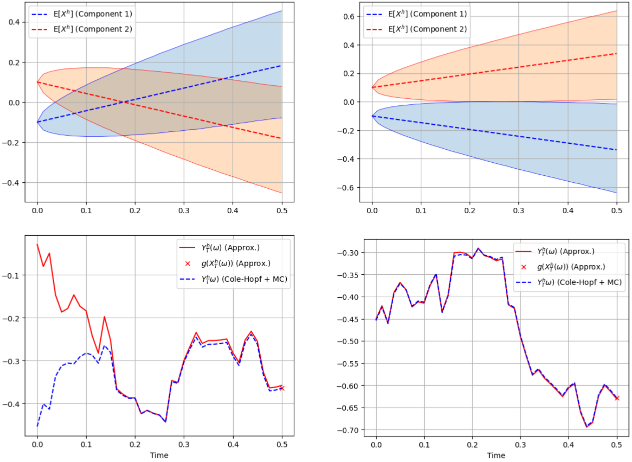

Empirical mean of X and 95th and 5th percentiles for the deep FBSDE method (left) and deep multi-FBSDE method (right), plus a representative trajectory of the approximate Y, with semi-analytic counterparts.

Observations

- The optimal strategy pushes X^1 to lower values and X^2 to higher values, since $X_0^1 < X_0^2$ and the task is to maximize $|X_T^1 - X_T^2|$ with low control cost.
- The deep FBSDE method causes X^1 and X^2 to cross each other, and naturally uses a more expensive control for the task, reading the initially high values of the Y-process.
- A similar MSE-analysis before shows that the problem with the deep FBSDE solution is due to a bad global minimum.

Parabolic PDE with quadratic gradient nonlinearity

Let $\alpha, \beta \in \mathbb{R}^+$, $\gamma \in \mathbb{R}$, and $x_0 \in \mathbb{R}^d$. With $g\colon \mathbb{R}^d \to \mathbb{R}$ we define

$$\begin{aligned} \frac{\partial v}{\partial t} + \alpha \Delta v + 2\beta \|\mathrm{D}_x v\|^2 - \gamma v =& 0, \quad (t,x) \in [0,T) \times \mathbb{R}^d, \\ v(T,x) =& g(x), \quad x \in \mathbb{R}^d. \end{aligned} \tag{7}$$

Since $\beta > 0$, the PDE does not have an associate stochastic control problem, and hence does not represent an HJB equation.

We consider the equivalent FBSDE:

$$\begin{aligned} X_t =& x_0 + \int_0^t \beta Z_t \mathrm{d}s + \sqrt{2\alpha} W_t, \\ Y_t =& g(X_T) + \int_t^T \left(\beta \|Z_s\|^2 - \gamma Y_s\right) \mathrm{d}s - \sqrt{2\alpha} \int_t^T \langle Z_s, \mathrm{d}W_s \rangle, \quad t \in [0,T]. \end{aligned} \tag{8}$$

Cole-Hopf transformation

Defining $\rho(t,x) = \log\left(\frac{\alpha}{2\beta} v(t,x)\right)$ leads directly to the simplified PDE:

$$\begin{aligned} \frac{\partial \rho}{\partial t}(t,x) + \alpha \Delta \rho - \gamma \rho \log \rho =& 0, \quad (t,x) \in [0,T) \times \mathbb{R}^d, \\ \rho(T,x) =& \exp\left(\frac{2\beta}{\alpha} g(x)\right), \quad x \in \mathbb{R}^d. \end{aligned} \tag{9}$$

Back transformation yields $\rho(t,x) = \exp\left(\frac{2\beta}{\alpha} v(t,x)\right)$.
In the special case $\gamma = 0$, we have for $t \in [0, T]$

$$Y_t = -\frac{\alpha}{2\beta} \log\left(E\left[\mathrm{e}^{-\frac{2\beta}{\alpha} g(X_t + \sqrt{2\alpha(T-t)}\,\xi)} \Big| X_t\right]\right), \tag{10}$$

where ξ is a standard d-dimensional normally distributed random vector.

PDE experiment: robustness in β

Setup (deep multi-FBSDE):

- $K = 2$, $\psi_1 = 0$, $\psi_2(t, x, y, z) = -\frac{\beta}{2} z$, $\quad N = 40$
- Terminal condition: $g(x) = -|x_1 - x_2|$, $\quad x_0 = (-0.1, 0.1)^\top$

Experiment: $\alpha = 0.0315$, $\quad \gamma = 1$, $\quad \beta \in [0.00125, 5]$
Reference solution: finite difference PDE on $[-2, 2]^2$

- Deep FBSDE fails for $\beta > 0.25$
- Deep multi-FBSDE remains accurate across all β

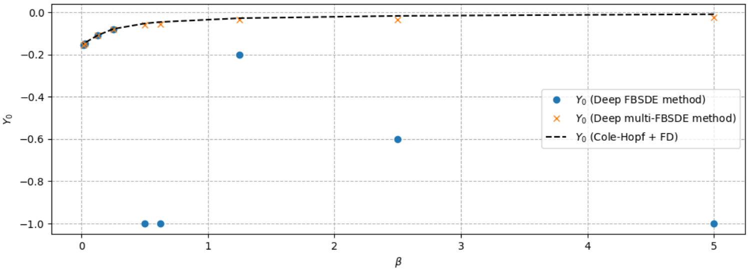

Summary and conclusion

- Empirically, the direct deep BSDE method may fail for strongly coupled FBSDEs, even in simple benchmark problems;
- The optimization landscape can be misleading: a small terminal loss does not necessarily imply an accurate approximation of the full solution, in particular of Y_0;
- The deep multi-FBSDE approach improves robustness by replacing one formulation by a family of equivalent formulations;
- Numerical experiments indicate significantly improved approximation of Y_0 and much better stability across parameter regimes.

References

⇒ K. Andersson, A. Andersson, *C.W. Oosterlee*, Convergence of a robust deep FBSDE method for stochastic control. *SIAM J. Sci. Comp., 45(1), (2023)*.

- K. Andersson, A. Andersson, *C.W. Oosterlee,* The deep multi-FBSDE method: a robust deep learning method for coupled FBSDEs. *J. Sci. Comput.* 106:77 (2026).
- B. Negyesi, Z. Huang, *C.W. Oosterlee*, Generalized convergence of the deep BSDE method: a step towards fully-coupled FBSDEs and applications in stochastic control. ArXiv:2403.18552 (2025).
- A. Jentzen et. al. Solving high-dimensional partial differential equations using deep learning. *Procs. National Academy of Sciences* 115.34 (2018).
- Han, J and Long J. Convergence of the deep BSDE method for coupled FBSDEs. *Probability, Uncertainty and Quantitative Risk* 5.1 (2020).
- C. Reisinger, W. Stockinger, Y. Zhang, A posteriori error estimates for fully coupled McKean-Vlasov forward-backward SDEs. *IMA J. Num. Analysis*, 44(4), (2024).