[SDE V] Forward–Backward SDEs and the Stochastic Maximum Principle

Anqiao Ouyang

June 6, 2026

1 Introduction

In [SDE IV] we developed BSDEs and the nonlinear Feynman–Kac representation. The Markovian setting there had a strict structure: a forward diffusion X whose coefficients (b, σ) depend only on (t, X_t), coupled to a backward equation for (Y, Z) whose driver f depends on (t, X_t, Y_t, Z_t). Crucially, the *forward* coefficients did not feel Y or Z: the BSDE rode on top of an autonomous diffusion.

That asymmetry was not a coincidence. It is exactly what makes the Picard-by-contraction proof of Pardoux–Peng possible: X is fixed once and for all, and only (Y, Z) is being iterated. Many natural problems break this asymmetry.

Motivating example: LQ control with feedback. Consider the linear–quadratic problem on $[0, T]$

$$dX_t = (AX_t + B\alpha_t)\,dt + \sigma\,dB_t, \qquad X_0 = x,$$

with cost

$$J(\alpha) = \mathbb{E}\left[\int_0^T \big(X_t^\top Q X_t + \alpha_t^\top R \alpha_t\big)dt + X_T^\top G X_T\right].$$

With the cost convention above, the Pontryagin maximum principle (§5) gives

$$\alpha_t^* = -\tfrac{1}{2} R^{-1} B^\top Y_t,$$

where Y is the adjoint process. Substituting the feedback into the state equation gives

$$dX_t = \big(AX_t - \tfrac{1}{2} B R^{-1} B^\top Y_t\big)dt + \sigma\,dB_t.$$

The forward equation now depends on Y. The pair (X, Y) is genuinely coupled: X pushes information forward, Y pushes information backward, and *both depend on the other.* This is a **forward–backward SDE (FBSDE)**.

What changes. The decoupled case of [SDE IV] could be solved in two passes: solve X forward, then solve (Y, Z) backward. A coupled FBSDE has no such order; a fixed-point iteration on the whole triple (X, Y, Z) may not contract because forward errors propagate into the backward equation and *vice versa.* Three distinct techniques have emerged to handle this:

- **Monotonicity (Hu–Peng, Peng–Wu, Yong).** An algebraic structural condition on the coefficients that restores enough contraction in the right inner product.

- **Small-horizon (Antonelli).** On a short enough interval $[0,T]$, the contraction works as in the decoupled case.
- **The four-step scheme (Ma–Protter–Yong).** When the associated nonlinear PDE has a classical solution, construct the FBSDE solution from the PDE.

Each technique illuminates a different facet of the same object. In this article we will:

- Define coupled FBSDEs precisely and identify the obstacle to naive iteration (§2).
- Prove existence and uniqueness under the Hu–Peng monotonicity condition (§3).
- Develop the four-step scheme and connect it to the nonlinear Feynman–Kac formula from [SDE IV] (§4).
- Establish the stochastic maximum principle in its general form and check duality with HJB (§5).
- Apply the machinery to LQ control, constrained utility maximization, and the Deep FBSDE method (§6).
- Close with reflected, mean-field, and second-order extensions (§7).

Standing notation. As in [SDE IV], we work on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0\le t\le T}, \mathbb{P})$ carrying a d-dimensional Brownian motion B, with $(\mathcal{F}_t)$ the augmented Brownian filtration and $T > 0$ fixed. State X lives in $\mathbb{R}^n$, adjoint Y in $\mathbb{R}^m$, integrand Z in $\mathbb{R}^{m\times d}$. The shorthand $\mathcal{S}^2, \mathcal{H}^2$ for the solution spaces is the same as in [SDE IV §2]. We will reuse K as a generic Lipschitz / growth constant.

2 Coupled FBSDEs and the obstacle to naive iteration

2.1 Definition

Definition 2.1 (FBSDE)**.** Given measurable coefficients

$$b : [0,T]\times\mathbb{R}^n\times\mathbb{R}^m\times\mathbb{R}^{m\times d}\to\mathbb{R}^n, \quad \sigma : [0,T]\times\mathbb{R}^n\times\mathbb{R}^m\times\mathbb{R}^{m\times d}\to\mathbb{R}^{n\times d},$$

$$f : [0,T]\times\mathbb{R}^n\times\mathbb{R}^m\times\mathbb{R}^{m\times d}\to\mathbb{R}^m, \quad g : \mathbb{R}^n\to\mathbb{R}^m,$$

and an initial point $x\in\mathbb{R}^n$, the (coupled) *forward–backward SDE* for (X,Y,Z) is

$$\begin{cases} dX_t = b(t,X_t,Y_t,Z_t)\,dt + \sigma(t,X_t,Y_t,Z_t)\,dB_t, & X_0 = x, \\ -dY_t = f(t,X_t,Y_t,Z_t)\,dt - Z_t\,dB_t, & Y_T = g(X_T), \end{cases} \tag{1}$$

on $t\in[0,T]$. A *solution* is a triple (X,Y,Z) with $X\in\mathcal{S}^2(\mathbb{R}^n)$, $Y\in\mathcal{S}^2(\mathbb{R}^m)$, $Z\in\mathcal{H}^2(\mathbb{R}^{m\times d})$ satisfying (1) $\mathbb{P}$-a.s.

The system reduces to the decoupled BSDE setting of [SDE IV] when b,σ do not depend on (Y,Z). The genuinely new feature is the dependence of b and σ on Y and Z.

2.2 Why naive iteration fails

A first instinct is to iterate. Pick $(Y^{(0)}, Z^{(0)})$, solve the forward SDE

$$dX_t^{(k)} = b(t, X_t^{(k)}, Y_t^{(k)}, Z_t^{(k)})\, dt + \sigma(t, X_t^{(k)}, Y_t^{(k)}, Z_t^{(k)})\, dB_t,$$

then solve the backward equation with terminal $g(X_T^{(k)})$ to obtain $(Y^{(k+1)}, Z^{(k+1)})$, and repeat. The map $(Y^{(k)}, Z^{(k)}) \mapsto (Y^{(k+1)}, Z^{(k+1)})$ has the right type, so in principle Banach gives a fixed point if it is contractive.

The trouble is that the forward step amplifies errors in Y: a perturbation of Y by δ becomes a perturbation of X_T of order roughly $L_b \cdot T \cdot \delta$ at the worst, where L_b is the Lipschitz constant of b in Y. The backward step in turn passes the X_T perturbation back into Y at $t < T$. Composing the two, the contraction constant ends up scaling like

$$C \cdot T \cdot (L_b + L_\sigma + L_f)(1 + L_g),$$

which is below 1 only when T is sufficiently small.

Theorem 2.2 (Antonelli, 1993)**.** *Under uniform Lipschitz hypotheses on b, σ, f, g, there exists $T_0 > 0$ depending only on the Lipschitz constants such that for every $T \leq T_0$ the FBSDE* (1) *admits a unique solution.*

The proof is the contraction argument outlined above, with the weighted-norm trick from [SDE IV §3] absorbing the bookkeeping. We omit it; the short-horizon restriction is essential in general, as the next example shows.

Remark. Examples of Antonelli show that for T above some threshold, naive Picard iteration can fail to converge or the FBSDE may have multiple solutions. The classical (linear, harmonic-oscillator-type) example is

$$dX_t = Y_t\, dt, \qquad -dY_t = X_t\, dt - Z_t\, dB_t, \qquad X_0 = 0,\ Y_T = 0. \tag{2}$$

Its deterministic skeleton has solutions $X_t = b \sin t$, $Y_t = b \cos t$. The terminal condition $Y_T = 0$ is satisfied non-trivially iff $T \in \{\pi/2, 3\pi/2, 5\pi/2, \ldots\}$. So for $T < \pi/2$ the solution is unique (and trivial); at $T = \pi/2$ uniqueness fails and a one-parameter family of solutions appears; for $\pi/2 < T < 3\pi/2$ uniqueness returns. The same picture repeats at every odd multiple of $\pi/2$. The phenomenon is a genuine eigenvalue effect of the coupled boundary-value structure (see Figure 1); length matters.

2.3 The decoupling field

For Markovian FBSDEs (coefficients depending on time and the processes pointwise, not path-wise), a key heuristic is the **decoupling field**: a function $u : [0, T] \times \mathbb{R}^n \to \mathbb{R}^m$ such that, along any solution,

$$Y_t = u(t, X_t), \qquad Z_t = (\sigma^\top \partial_x u)(t, X_t, u(t, X_t), \ldots).$$

Plugging this into the forward equation removes Y and Z explicitly:

$$dX_t = b\big(t, X_t, u(t, X_t), (\sigma^\top \partial_x u)(t, X_t)\big) dt + \sigma\big(t, X_t, u(t, X_t), \ldots\big) dB_t,$$

and the FBSDE collapses to a standard SDE for X alone. If we can find u, we have decoupled the problem. The function u itself satisfies a quasilinear PDE (see §4). This is the heart of the four-step scheme.

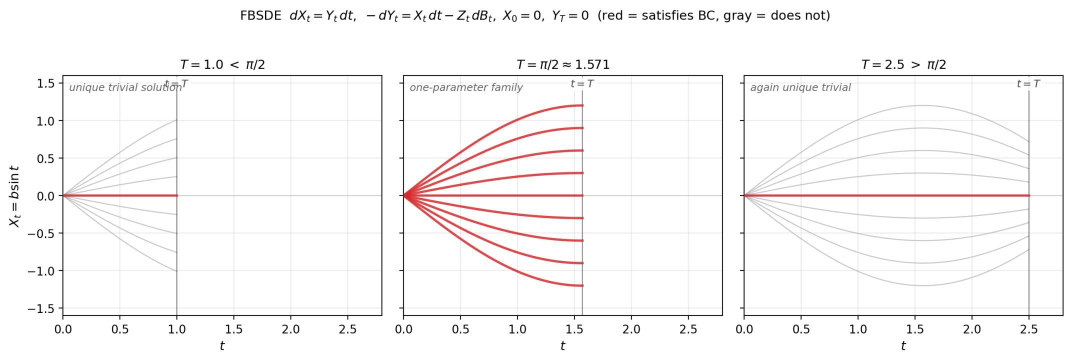

Figure 1: Antonelli's linear example (2). For the deterministic family $X_t = b\sin t$, $Y_t = b\cos t$, the initial condition $X_0 = 0$ is automatic and the terminal condition $Y_T = 0$ becomes $b\cos T = 0$. Red curves satisfy the terminal condition and gray curves do not. Left: for $T = 1 < \pi/2$, only $b = 0$ is admissible. Center: for $T = \pi/2$, every b is admissible, so uniqueness fails. Right: for $T = 2.5 \in (\pi/2, 3\pi/2)$, only $b = 0$ is again admissible. The singularity is therefore a boundary-value eigenvalue phenomenon rather than a monotone loss of well-posedness after a fixed time.

3 Monotonicity: the Hu–Peng method

To go beyond short horizons, we need an algebraic condition on the coefficients that prevents the forward and backward errors from amplifying one another. Hu and Peng (1995) identified one standard condition of this type. To avoid rectangular-matrix bookkeeping, we state the clean square version: $X, Y \in \mathbb{R}^n$, $Z \in \mathbb{R}^{n\times d}$, and $G = I$. The general G-monotonicity condition is similar, but uses $G\Delta x$, $G^\top \Delta y$, and $G^\top \Delta z$ in the estimates.

For $\theta = (x, y, z)$, define the signed coefficient vector

$$A(t,\theta) := (-f(t,\theta),\ b(t,\theta),\ \sigma(t,\theta)).$$

The signs are chosen so that Itô's formula for $\langle \Delta X_t, \Delta Y_t\rangle$ produces exactly the inner product below.

Definition 3.1 (Monotonicity)**.** The coefficients (b, σ, f) are *monotone* with constant $\beta > 0$ if, for all t and all θ_1, θ_2,

$$\big\langle A(t,\theta_1) - A(t,\theta_2),\ \theta_1 - \theta_2 \big\rangle \le -\beta\, |\theta_1 - \theta_2|^2.$$

The terminal map is compatible with this monotonicity if

$$\langle g(x_1) - g(x_2),\ x_1 - x_2\rangle \ge 0 \qquad \text{for all } x_1, x_2 \in \mathbb{R}^n.$$

Theorem 3.2 (Hu–Peng, simplified form)**.** *Assume the coefficients are progressively measurable, uniformly Lipschitz in (x, y, z), have linear growth, and satisfy Definition 3.1. Then the FBSDE* (1) *has a unique solution $(X, Y, Z) \in \mathcal{S}^2 \times \mathcal{S}^2 \times \mathcal{H}^2$ for any finite horizon $T > 0$.*

Sketch. The technique is the *continuation method*. Embed (1) in a one-parameter family $(\Pi_\lambda)_{\lambda\in[0,1]}$, with Π_0 chosen to be a simple solvable monotone FBSDE and Π_1 the target system. Let Λ be the set of λ for which Π_λ is solvable. One proves:

(a) $0 \in \Lambda$.

(b) If $\lambda \in \Lambda$, then $\lambda + \varepsilon \in \Lambda$ for some $\varepsilon > 0$ depending only on the monotonicity and Lipschitz constants.

The uniform step size is the point of the monotonicity assumption. For two candidate triples, apply Itô's formula to $\langle \Delta X_t, \Delta Y_t \rangle$ and use

$$\langle \Delta b, \Delta Y \rangle - \langle \Delta f, \Delta X \rangle + \langle \Delta \sigma, \Delta Z \rangle = \langle \Delta A, \Delta \theta \rangle \leq -\beta |\Delta \theta|^2.$$

The terminal condition contributes a nonnegative boundary term. These estimates make the perturbation map a contraction with a step size independent of λ, so the continuation reaches $\lambda = 1$. $\square$

Remark. The continuation method is the same machinery used in implicit function theorems on Banach spaces. In the fully rectangular case $n \neq m$, one inserts a full-rank matrix G to compare the forward and backward components in compatible dimensions.

Remark (Other sufficient conditions).

- Peng–Wu (1999): monotonicity-type conditions can be adapted to stochastic Hamiltonian systems over arbitrary horizons.
- Delarue (2002): when the diffusion is non-degenerate and does not depend on (Y, Z), PDE methods give another route to global solvability.
- Ma–Yong (1999): the decoupling-field perspective gives a systematic way to organize the continuation, four-step, and PDE viewpoints.

4 The Four-Step Scheme

The four-step scheme is the cleanest constructive proof of FBSDE solvability when the coefficients are sufficiently regular and a certain PDE has a classical solution. It is named for the four stages of the construction (Ma–Protter–Yong, 1994).

For a transparent statement, assume here that the diffusion is Markovian and independent of the backward variables:

$$\sigma = \sigma(t, x).$$

The forward drift may still depend on (Y, Z). With this restriction, the candidate decoupling field $u : [0, T] \times \mathbb{R}^n \to \mathbb{R}^m$ satisfies a quasilinear parabolic system.

4.1 The four steps

Step 1 — Guess the decoupling field. Suppose $u \in C^{1,2}([0, T] \times \mathbb{R}^n; \mathbb{R}^m)$ exists with $Y_t = u(t, X_t)$. Set

$$Z_t = D_x u(t, X_t)\, \sigma(t, X_t),$$

where $D_x u$ is the $m \times n$ Jacobian. Itô's formula gives

$$dY_t = \Big(\partial_t u + \tfrac{1}{2} \operatorname{Tr}\big[\sigma\sigma^\top D^2_{xx} u\big] + D_x u\, b(t, X_t, u, Z_t)\Big) dt + D_x u(t, X_t) \sigma(t, X_t) dB_t.$$

Comparing this with $dY_t = -f(t, X_t, Y_t, Z_t)dt + Z_t dB_t$, the field u must solve

$$\partial_t u(t,x) + \tfrac{1}{2} \operatorname{Tr}\big[\sigma\sigma^\top(t,x) D^2_{xx} u(t,x)\big] + D_x u(t,x) b\big(t, x, u(t,x), D_x u(t,x)\sigma(t,x)\big) + f\big(t, x, u(t,x), D_x u(t,x)\sigma(t,x)\big) = 0, \tag{3}$$

with terminal condition $u(T, x) = g(x)$.

Step 2 — Solve the PDE. Under smoothness, boundedness, and uniform parabolicity assumptions, classical parabolic PDE theory gives a classical solution u to (3). This is the analytically difficult step.

Step 3 — Solve the now-decoupled forward SDE. With u in hand, define

$$\tilde{b}(t,x) := b\big(t,x,u(t,x),D_x u(t,x)\sigma(t,x)\big),$$

and solve

$$dX_t = \tilde{b}(t,X_t)dt + \sigma(t,X_t)dB_t, \qquad X_0 = x.$$

This is a standard SDE.

Step 4 — Read off (Y,Z). Set

$$Y_t := u(t,X_t), \qquad Z_t := D_x u(t,X_t)\sigma(t,X_t).$$

By Step 1, (X,Y,Z) solves the original FBSDE.

Theorem 4.1 (Ma–Protter–Yong, simplified form)**.** *Assume b,σ,f,g are sufficiently smooth and Lipschitz, $\sigma\sigma^\top$ is uniformly elliptic, and the PDE* (3) *admits a classical solution $u \in C^{1,2}$ of at most polynomial growth. Then the FBSDE* (1)*, with $\sigma = \sigma(t,x)$, has a solution given by Steps 3–4. Uniqueness holds in the class generated by this classical decoupling field, and standard Lipschitz estimates give uniqueness in the corresponding square-integrable solution class.*

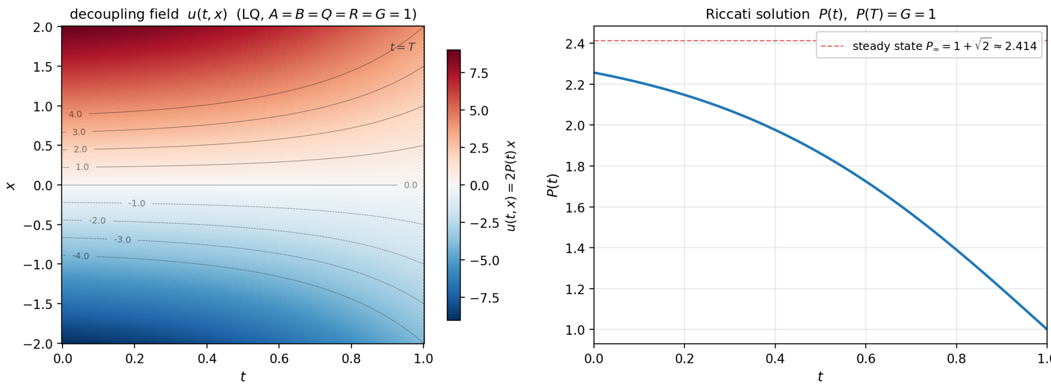

Figure 2: Decoupling field for the scalar LQ problem with $A = B = Q = R = G = 1$ and $T = 1$. The left panel plots $u(t,x) = 2P(t)x$ on $[0,1]\times[-2,2]$; this is the value of the adjoint variable Y_t when the state is $X_t = x$. The right panel plots the Riccati solution $P(t)$ with terminal condition $P(T) = 1$. Since $\dot{P}+2P-P^2+1 = 0$, one has $P(t) = 1+\sqrt{2}\tanh(\sqrt{2}(1-t))$, and longer backward horizons approach the positive algebraic Riccati root $1+\sqrt{2}$.

4.2 Relation to nonlinear Feynman–Kac

[SDE IV Thm. 5.1] said: a classical PDE solution u to the semilinear PDE produces the BSDE solution by setting $Y_s = u(s,X_s)$ and $Z_s = D_x u(s,X_s)\,\sigma(s,X_s)$. The four-step scheme is the natural *coupled* extension: a classical PDE solution to the quasilinear PDE (3) produces the FBSDE triple through the same read-off rule, with the added twist that the PDE itself contains u and $D_x u\,\sigma$ inside its coefficients.

Remark (When the PDE has no classical solution)*.* The four-step scheme requires Step 2 to succeed in the classical sense, which restricts to uniformly parabolic, smooth-coefficient regimes. Delarue (2002) showed how to extend it to degenerate cases via a quasi-classical regularization, and the monotonicity method of §3 covers many cases where the PDE has only viscosity solutions.

5 The Stochastic Maximum Principle

We now connect FBSDEs to optimal control. To keep the statement first-order and avoid the second-order adjoint terms that appear when the control enters the diffusion, we use the following standard setting:

- State $X \in \mathbb{R}^n$ with dynamics
$$dX_t = b(t, X_t, \alpha_t)dt + \sigma(t, X_t)dB_t, \qquad X_0 = x.$$
- The control α_t takes values in a convex set $A \subset \mathbb{R}^k$ and is progressively measurable.
- The cost is
$$J(\alpha) = \mathbb{E}\left[\int_0^T \ell(t, X_t, \alpha_t)dt + h(X_T)\right],$$
and is to be minimized.

Definition 5.1 (Hamiltonian)**.** The Hamiltonian is
$$H(t, x, y, z, a) := \langle b(t, x, a), y\rangle + \operatorname{tr}\big(\sigma^\top(t, x)z\big) + \ell(t, x, a),$$
where $a \in A$, $y \in \mathbb{R}^n$, and $z \in \mathbb{R}^{n\times d}$.

5.1 The principle

Theorem 5.2 (Stochastic maximum principle, first-order form)**.** *Suppose α^* is optimal and X^* is the corresponding state. Assume b, σ, ℓ, h are differentiable in x, and b, ℓ are differentiable in a. Then there exist adapted processes $(Y, Z) \in \mathcal{S}^2 \times \mathcal{H}^2$ solving*
$$-dY_t = \partial_x H(t, X_t^*, Y_t, Z_t, \alpha_t^*)dt - Z_t dB_t, \qquad Y_T = \partial_x h(X_T^*), \tag{4}$$
such that, for $dt \otimes d\mathbb{P}$-a.e. (t, ω),
$$\big\langle \partial_a H(t, X_t^*, Y_t, Z_t, \alpha_t^*),\, a - \alpha_t^*\big\rangle \geq 0 \qquad \text{for all } a \in A. \tag{5}$$
If $a \mapsto H(t, X_t^*, Y_t, Z_t, a)$ is convex, then (5) is equivalent to
$$\alpha_t^* \in \arg\min_{a\in A} H(t, X_t^*, Y_t, Z_t, a). \tag{6}$$

The pair (X^*, Y, Z) together with the feedback condition is therefore a coupled forward–backward system: the state determines the adjoint equation, while the adjoint determines the optimal feedback.

Theorem 5.3 (Sufficient condition, convex case)**.** *Assume h is convex and $(x, a) \mapsto H(t, x, y, z, a)$ is convex for each (t, y, z). If an admissible quadruple (X^*, Y, Z, α^*) satisfies (4) and the pointwise minimization condition (6), then α^* is optimal.*

Sketch. For the necessary condition, take a convex perturbation $\alpha^\varepsilon = \alpha^* + \varepsilon(a - \alpha^*)$ and compute the first variation δX. The adjoint BSDE is chosen so that the integration-by-parts identity for $\langle Y_t, \delta X_t\rangle$ cancels the state variation terms in the derivative of J. Since α^* minimizes J, the remaining first-order term is nonnegative for every admissible direction, which gives (5). For the sufficient condition, convexity gives
$$J(\alpha) - J(\alpha^*) \geq \mathbb{E}\int_0^T \big[H(t, X_t^*, Y_t, Z_t, \alpha_t) - H(t, X_t^*, Y_t, Z_t, \alpha_t^*)\big]dt \geq 0.$$
□

5.2 Duality with HJB

[SDE IV §6.1] gave the HJB approach: the value function $V(t,x) := \inf_\alpha \mathbb{E}[\int_t^T \ell ds + h(X_T)|X_t = x]$ satisfies

$$\partial_t V + \inf_{a\in A} \tilde{H}(t,x,\partial_x V, \partial_{xx} V, a) = 0, \qquad V(T,x) = h(x),$$

where $\tilde{H}$ is the *dynamic-programming Hamiltonian*. The relation to Definition 5.1 is:

$$Y_t = \partial_x V(t, X_t^*), \qquad Z_t = D_{xx}^2 V(t, X_t^*)\,\sigma(t, X_t^*). \tag{7}$$

The adjoint variable Y is the gradient of the value function along the optimal trajectory; Z is the second-order correction. The two approaches are dual: HJB lives on the PDE side, the maximum principle on the FBSDE side, and (7) bridges them. When $V \in C^{1,2}$, either approach gives the same answer; when V is only viscosity, the maximum principle is often easier to apply because it does not require V's regularity.

5.3 LQ control, worked example

Return to the LQ problem from the introduction:

$$dX_t = (AX_t + B\alpha_t)dt + \sigma dB_t, \quad J = \mathbb{E}\Big[\int_0^T (X^\top QX + \alpha^\top R\alpha)dt + X_T^\top G X_T\Big],$$

with Q, R, G symmetric positive (semi)definite and $R \succ 0$. The Hamiltonian is

$$H(t,x,y,z,a) = \langle Ax + Ba, y\rangle + \mathrm{tr}(\sigma^\top z) + x^\top Qx + a^\top Ra.$$

First-order condition $\partial_a H = 0$ gives

$$\alpha_t^* = -\tfrac{1}{2} R^{-1} B^\top Y_t.$$

The adjoint BSDE is

$$-dY_t = (A^\top Y_t + 2QX_t^*)dt - Z_t dB_t, \qquad Y_T = 2GX_T^*.$$

Combined with the controlled state SDE, the system is a linear FBSDE. Make the ansatz $Y_t = 2P(t)X_t^*$ for a deterministic $P(t)$, plug into the adjoint equation, and the function P must solve the (matrix) *Riccati equation*

$$\dot{P} + PA + A^\top P - PBR^{-1}B^\top P + Q = 0, \qquad P(T) = G.$$

This is a closed-form classical result: existence and uniqueness of P on $[0,T]$ follows from standard ODE theory under mild conditions on the matrices. The optimal feedback is then

$$\alpha_t^* = -R^{-1}B^\top P(t) X_t^*,$$

i.e. linear feedback with a time-varying gain.

Remark. Setting up the LQ problem via Pontryagin and via HJB gives the same Riccati equation. HJB writes $V(t,x) = x^\top P(t)x + q(t)$ and substitutes; Pontryagin writes $Y_t = 2P(t)X_t^*$ and substitutes. Both routes hit the same nonlinear matrix ODE because LQ enjoys quadratic value functions and linear adjoints — this is the simplest case where (7) can be verified by hand.

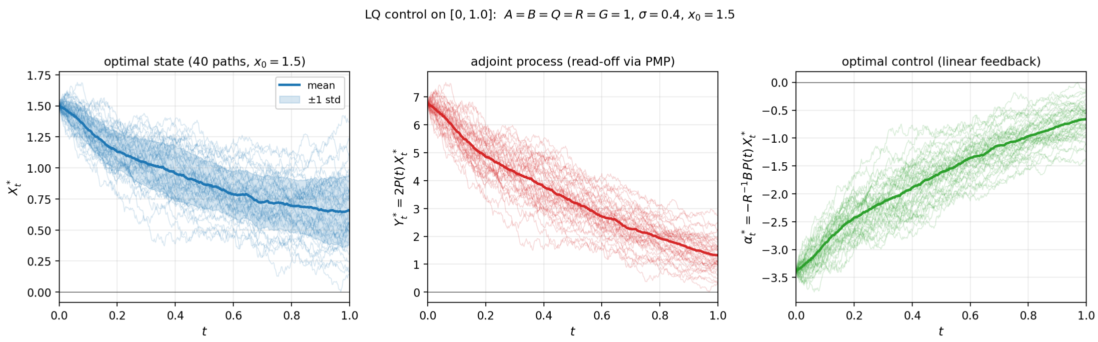

Figure 3: Closed-loop scalar LQ control with $A = B = Q = R = G = 1$, $\sigma = 0.4$, $T = 1$, and $x_0 = 1.5$, simulated by Euler–Maruyama. The left panel shows 40 sample paths of X_t^* and the mean ± 1 standard-deviation envelope. The center panel shows the adjoint process $Y_t^* = 2P(t)X_t^*$ obtained from the maximum principle. The right panel shows the feedback control $\alpha_t^* = -R^{-1}B^\top P(t)X_t^*$, which is strongly negative near $t = 0$ and relaxes as the state approaches the origin. The same Riccati gain $P(t)$ is the one plotted in Figure 2.

6 Applications

6.1 Constrained utility maximization

Consider an investor with terminal wealth W_T generated by a self-financing portfolio,

$$dW_t = \pi_t \cdot \mu\, dt + \pi_t \cdot \sigma\, dB_t, \qquad W_0 = w,$$

where π_t is the dollar amount in risky assets. The investor maximizes $\mathbb{E}[U(W_T)]$ for a concave utility U, subject to the portfolio constraint $\pi_t \in K$ for a closed convex set or cone K.

The unconstrained Merton problem is usually solved directly by HJB. With convex constraints, the Cvitanić–Karatzas duality method introduces a dual process taking values in the polar constraint set. The primal wealth equation and the dual martingale/adjoint equation form a forward–backward system. The exact sign convention depends on whether the dual variable is placed in the drift of the state-price deflator or in the support function of K, but the structural point is invariant: the forward wealth process and the backward marginal-utility process determine each other through complementary slackness. Under the regularity and convexity assumptions of the duality theorem, this FBSDE formulation is equivalent to the constrained optimum and reduces to the Merton solution when $K = \mathbb{R}^d$.

6.2 Deep FBSDE

The Deep BSDE method of [SDE IV §6.2] extends naturally to the coupled setting (Han–Long, 2020; see also Huré–Pham–Warin, 2020, for related backward schemes). Discretize $0 = t_0 < t_1 < \cdots < t_M = T$. Parameterize:

- the initial value $Y_0 = \theta_0 \in \mathbb{R}^m$;
- the integrand $Z_k \approx z^\theta(t_k, X_{t_k})$ by a per-step neural network.

Simulate $X_{t_{k+1}} = X_{t_k} + b\,\Delta t + \sigma\,\Delta B_k$ *using the current estimates* of Y_{t_k} and $Z_{t_k} = z^\theta(t_k, X_{t_k})$ in b, σ. Propagate Y by Euler:

$$Y_{t_{k+1}} = Y_{t_k} - f(t_k, X_{t_k}, Y_{t_k}, Z_{t_k})\,\Delta t + Z_{t_k} \cdot \Delta B_k.$$

Minimize the terminal mismatch

$$\mathcal{J}(\theta) := \mathbb{E}\big[|g(X_{t_M}) - Y_{t_M}|^2\big]$$

by SGD. The key difference from Deep BSDE: X depends on (Y, Z), so the forward chain itself contains learnable parameters. The variational characterization is the same: $\mathcal{J}(\theta) = 0$ iff (X, Y, Z) is the unique FBSDE solution.

6.3 A note on coupled diffusion

When σ depends on (Y, Z), the clean version of the four-step scheme above no longer applies directly. The relation $Z = D_x u\, \sigma(t, x, u, Z)$ becomes an implicit algebraic equation for Z, and the associated PDE may become fully nonlinear or even ill-posed without additional structure. In robust-control and model-uncertainty problems, the fully nonlinear PDE side is often represented probabilistically by second-order BSDEs (2BSDEs) in the sense of Soner–Touzi–Zhang.

7 Outlook

Reflected FBSDEs. Adding a lower obstacle L_t to the backward equation, as in [SDE IV §7], gives reflected FBSDEs. These arise in optimal stopping games where both players control their state. The associated PDE is a variational inequality with obstacle.

Mean-field FBSDEs (McKean–Vlasov). Letting b, σ, f, g depend on the marginal law $\mathcal{L}(X_t), \mathcal{L}(Y_t)$ leads to mean-field FBSDEs. These are the natural state-equation form of mean-field games (Lasry–Lions, Caines–Huang–Malhamé) and mean-field control. The associated PDE, the *master equation*, lives on $[0, T] \times \mathbb{R}^n \times \mathcal{P}_2(\mathbb{R}^n)$, and its existence-uniqueness theory is currently very active (Cardaliaguet–Delarue–Lasry–Lions 2019).

Second-order BSDEs (2BSDE). As mentioned in §6.3, when the diffusion is controlled, the dual is a 2BSDE. These solve fully nonlinear PDEs probabilistically. The Soner–Touzi–Zhang framework now subsumes both BSDE and 2BSDE under a single theory based on nonlinear expectations.

Rough BSDEs and rough FBSDEs. If the driver is rougher than the semimartingale setting, Itô calculus must be replaced by rough-path or pathwise stochastic analysis. This is a more delicate and still-developing direction, so it is better treated separately from the Brownian theory above.

Acknowledgments. As always, comments and corrections are welcome.

References

[1] Pardoux, E.; Peng, S. *Adapted solution of a backward stochastic differential equation.* Systems & Control Letters 14 (1990), 55–61.

[2] Antonelli, F. *Backward–forward stochastic differential equations.* Annals of Applied Probability 3 (1993), 777–793.

[3] Ma, J.; Protter, P.; Yong, J. *Solving forward–backward stochastic differential equations explicitly: a four step scheme.* Probability Theory and Related Fields 98 (1994), 339–359.

[4] Hu, Y.; Peng, S. *Solution of forward–backward stochastic differential equations.* Probability Theory and Related Fields 103 (1995), 273–283.

[5] Peng, S.; Wu, Z. *Fully coupled forward–backward stochastic differential equations and applications to optimal control.* SIAM Journal on Control and Optimization 37 (1999), 825–843.

[6] Ma, J.; Yong, J. *Forward–Backward Stochastic Differential Equations and Their Applications.* Lecture Notes in Mathematics 1702, Springer, 1999.

[7] Yong, J.; Zhou, X. Y. *Stochastic Controls: Hamiltonian Systems and HJB Equations.* Springer, 1999.

[8] Delarue, F. *On the existence and uniqueness of solutions to FBSDEs in a non-degenerate case.* Stochastic Processes and their Applications 99 (2002), 209–286.

[9] Cvitanić, J.; Karatzas, I. *Convex duality in constrained portfolio optimization.* Annals of Applied Probability 2 (1992), 767–818.

[10] Soner, H. M.; Touzi, N.; Zhang, J. *Wellposedness of second order backward SDEs.* Probability Theory and Related Fields 153 (2012), 149–190.

[11] Han, J.; Long, J. *Convergence of the deep BSDE method for coupled FBSDEs.* Probability, Uncertainty and Quantitative Risk 5, article 5 (2020).

[12] Huré, C.; Pham, H.; Warin, X. *Deep backward schemes for high-dimensional nonlinear PDEs.* Mathematics of Computation 89 (2020), 1547–1579.

[13] Cardaliaguet, P.; Delarue, F.; Lasry, J.-M.; Lions, P.-L. *The Master Equation and the Convergence Problem in Mean Field Games.* Princeton University Press, 2019.

[14] Ladyzhenskaya, O. A.; Solonnikov, V. A.; Ural'tseva, N. N. *Linear and Quasilinear Equations of Parabolic Type.* Translations of Mathematical Monographs 23, American Mathematical Society, 1968.

[15] Lieberman, G. M. *Second Order Parabolic Differential Equations.* World Scientific, 1996.