[SDE IV] Backward SDEs and the Nonlinear Feynman–Kac Formula

Anqiao Ouyang, Xuecheng Liu

May 27, 2026

1 Introduction

The previous three articles in this series built the forward picture of stochastic dynamics. In [SDE I] we constructed Brownian motion and the surrounding measure-theoretic infrastructure. In [SDE II] we developed the Itô integral and the chain rule of stochastic calculus. In [SDE III] we proved existence and uniqueness of solutions to forward SDEs and crossed over to PDEs via the Feynman–Kac formula: for the linear parabolic problem

$$\partial_t u + \mathcal{L}u = 0, \qquad u(T,x) = g(x),$$

where $\mathcal{L}$ is the infinitesimal generator of X, the solution admits the probabilistic representation

$$u(t,x) = \mathbb{E}\big[g(X_T) \,\big|\, X_t = x\big].$$

A natural question is what happens if we add a nonlinear source term to the PDE:

$$\partial_t u + \mathcal{L}u + f\big(t,x,u,\sigma^\top \nabla u\big) = 0, \qquad u(T,x) = g(x).$$

The expectation in linear Feynman–Kac is taken at the fixed time T, and it does not see u or ∇u inside the integrand. To bring those in, we need a representation that propagates information *backward in time*, conditioning on the path of X all along $[t,T]$. The object that does this is a **backward stochastic differential equation (BSDE)**.

A BSDE is, schematically, an SDE run from a terminal condition $Y_T = \xi$ back to time 0:

$$-dY_t = f(t,Y_t,Z_t)\,dt - Z_t\,dB_t, \qquad Y_T = \xi.$$

The unknown is the pair (Y_t, Z_t), where Y is a real-valued process and Z is the integrand against Brownian motion. As we will see in §2, Z is forced upon us by adaptedness: it is not an extra data choice but a consequence of the martingale representation theorem. When $\xi = g(X_T)$ and f depends on (t, X_t, Y_t, Z_t) in a Markovian way, the BSDE solves precisely the semilinear PDE above, with the identification

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

This is the **nonlinear Feynman–Kac formula** of Pardoux and Peng.

The theory was initiated by Pardoux and Peng in 1990 and now sits at the heart of stochastic control, mathematical finance, and — more recently — deep-learning-based PDE solvers. In this article we will:

- Define BSDEs precisely and explain why a pair (Y, Z) is the right unknown (§2).

- Prove the Pardoux–Peng existence and uniqueness theorem under Lipschitz conditions (§3).
- Establish the comparison theorem and observe how it reflects the no-arbitrage structure of pricing (§4).
- State and prove the nonlinear Feynman–Kac formula for Markovian BSDEs, side by side with the linear version from [SDE III] (§5).
- Apply the machinery to stochastic optimal control (HJB) and outline the Deep BSDE numerical method (§6).
- Close with an outlook on forward–backward coupling, mean-field BSDEs, and rough paths (§7).

Standing notation. We work on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0\le t\le T}, \mathbb{P})$ carrying a d-dimensional Brownian motion $B = (B_t)_{0\le t\le T}$, with $(\mathcal{F}_t)$ the augmented filtration of B and a fixed horizon $T > 0$. Vectors are columns; $\sigma^\top$ is matrix transpose; $|\cdot|$ denotes Euclidean norm. "Adapted" means $(\mathcal{F}_t)$-adapted; "predictable" means $(\mathcal{F}_t)$-predictable. We will reuse K as a generic constant whose value may change from line to line.

2 Definition of BSDEs and the role of Z

2.1 Why a pair?

Suppose we try to define a backward SDE in the naive way: pick a terminal condition $\xi \in L^2(\mathcal{F}_T)$ and a driver $f(t, y)$, and look for an $(\mathcal{F}_t)$-adapted process Y satisfying

$$Y_T = \xi, \qquad -dY_t = f(t, Y_t)\, dt.$$

This is just an ODE run backward, pathwise. The trouble is adaptedness: $Y_T = \xi$ is $\mathcal{F}_T$-measurable, so Y_t defined by an ODE backward from T inherits all the information in ξ, which violates $\mathcal{F}_t$-measurability at any earlier time $t < T$. In other words, knowing Y_t would let us forecast ξ.

The remedy is to allow a martingale correction. By the martingale representation theorem (e.g. [5], Ch. V), every square-integrable Brownian martingale M with $M_0 = \mathbb{E}[\xi]$ admits a unique predictable integrand $Z \in \mathcal{H}^2$ such that

$$M_t = M_0 + \int_0^t Z_s\, dB_s.$$

Choosing $M_t := \mathbb{E}[\xi \mid \mathcal{F}_t]$ produces the Z that "carries" the news about ξ revealed up to time t. The BSDE is the integral identity that combines the deterministic driver f with this Brownian correction $Z\, dB$.

Definition 2.1 (BSDE solution)**.** Let $\xi \in L^2(\mathcal{F}_T; \mathbb{R})$ and $f : [0, T] \times \Omega \times \mathbb{R} \times \mathbb{R}^d \to \mathbb{R}$ be progressively measurable. A pair of processes (Y, Z) — with Y continuous, real-valued, $(\mathcal{F}_t)$-adapted, and Z $(\mathcal{F}_t)$-predictable, $\mathbb{R}^d$-valued — is called a *solution* to the BSDE (ξ, f) if

$$Y_t = \xi + \int_t^T f(s, Y_s, Z_s)\, ds - \int_t^T Z_s\, dB_s, \qquad t \in [0, T], \tag{1}$$

$\mathbb{P}$-a.s., and additionally

$$\mathbb{E}\Big[\sup_{0\le t\le T} |Y_t|^2\Big] < \infty, \qquad \mathbb{E}\int_0^T |Z_s|^2\, ds < \infty.$$

We give the solution spaces names:

$$\mathcal{S}^2 := \Big\{ Y \text{ adapted, continuous} : \|Y\|_{\mathcal{S}^2}^2 := \mathbb{E}\big[\sup_t |Y_t|^2\big] < \infty \Big\},$$

$$\mathcal{H}^2 := \Big\{ Z \text{ predictable} : \|Z\|_{\mathcal{H}^2}^2 := \mathbb{E}\int_0^T |Z_s|^2\, ds < \infty \Big\}.$$

Both are Banach (in fact Hilbert), and the equation (1) is to be read in $\mathcal{S}^2 \times \mathcal{H}^2$.

Remark. Differentiating (1) gives the differential form

$$-dY_t = f(t, Y_t, Z_t)\, dt - Z_t\, dB_t, \qquad Y_T = \xi.$$

The minus sign in front of dY is bookkeeping: it records that time is being run backward from T. Some authors absorb it into the driver by writing $\widetilde{f} = -f$.

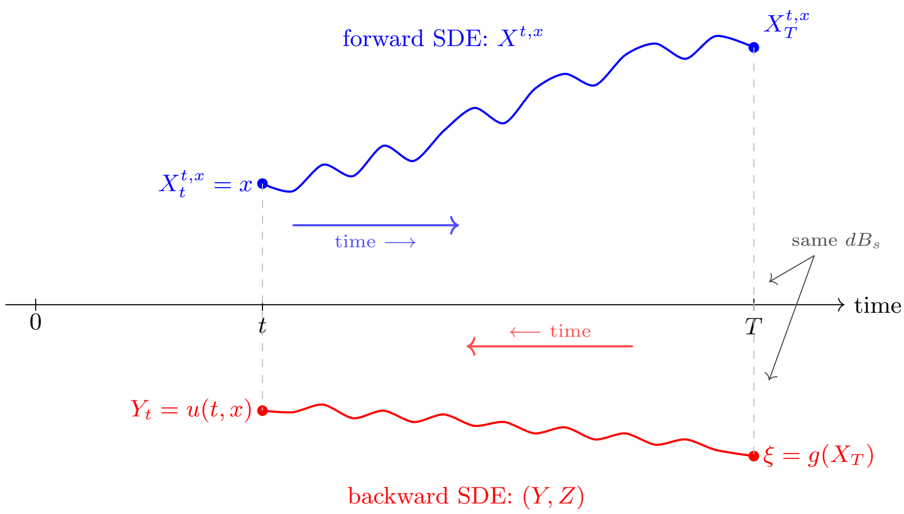

Figure 1: The forward–backward picture underlying the nonlinear Feynman–Kac formula. The diffusion $X^{t,x}$ is launched from x at time t and runs *forward* to a random terminal value X_T; the terminal payoff $\xi = g(X_T)$ is then transported *backward* by the BSDE to produce $Y_t = u(t,x)$. Both processes are driven by the same Brownian increments dB_s on $[t,T]$: the forward equation uses them through the diffusion coefficient σ, the backward equation through the integrand Z. Adaptedness forces Y_t to be a function of the history up to t only; this is what makes Z unavoidable.

2.2 Two warm-up examples

Example 2.2 ($f \equiv 0$: linear Feynman–Kac in disguise)**.** The BSDE $(\xi, 0)$ reduces to $Y_t = \xi - \int_t^T Z_s\, dB_s$. Taking conditional expectations,

$$Y_t = \mathbb{E}[\xi \mid \mathcal{F}_t],$$

and Z is the integrand from the martingale representation of $M_t = \mathbb{E}[\xi \mid \mathcal{F}_t]$. The BSDE with zero driver is exactly the linear Feynman–Kac formula. Specializing to $\xi = g(X_T)$ recovers [SDE III, Thm. 4.1].

Example 2.3 (Linear driver and explicit solution)**.** For $f(t,y,z) = \alpha_t\, y + \beta_t \cdot z + \gamma_t$ with bounded adapted coefficients, define the discount and change-of-measure factors

$$\Gamma_t := \exp\Big(\int_0^t \alpha_s\, ds \Big), \qquad \frac{d\mathbb{P}^\beta}{d\mathbb{P}}\Big|_{\mathcal{F}_T} := \exp\Big(\int_0^T \beta_s\, dB_s - \tfrac{1}{2}\int_0^T |\beta_s|^2 ds \Big).$$

Then

$$\Gamma_t\, Y_t = \mathbb{E}^{\mathbb{P}^\beta}\Big[\Gamma_T\, \xi + \int_t^T \Gamma_s\, \gamma_s\, ds \,\Big|\, \mathcal{F}_t\Big].$$

This formula is the BSDE analogue of variation of parameters; we will use it inside the proof of the comparison theorem.

3 Existence and Uniqueness

The cornerstone result of the linear theory is the following.

Theorem 3.1 (Pardoux–Peng, 1990)**.** *Assume*

(i) $\xi \in L^2(\mathcal{F}_T)$*;*

(ii) $(t,\omega) \mapsto f(t,\omega,0,0)$ *belongs to* $\mathcal{H}^2$*;*

(iii) f *is uniformly Lipschitz in* (y,z)*: there exists* $K \geq 0$ *with*

$$|f(t,\omega,y,z) - f(t,\omega,y',z')| \leq K\big(|y-y'| + |z-z'|\big)$$

$dt \otimes d\mathbb{P}$*-a.e.*

Then the BSDE (1) *admits a unique solution* $(Y,Z) \in \mathcal{S}^2 \times \mathcal{H}^2$*.*

The proof rests on a contraction argument in the spirit of Picard iteration for forward SDEs ([SDE III, §2]), with one new ingredient: the contraction happens in a *weighted* norm whose exponential weight absorbs the Lipschitz constant.

3.1 An a priori estimate

Lemma 3.2 (A priori bound)**.** *If* $(Y,Z) \in \mathcal{S}^2 \times \mathcal{H}^2$ *solves* (1)*, then for every* $\beta > 2K + 2K^2$*,*

$$\mathbb{E}\int_0^T e^{\beta s}\big(|Y_s|^2 + |Z_s|^2\big)\, ds \leq C\left(\mathbb{E}\big[e^{\beta T}|\xi|^2\big] + \mathbb{E}\int_0^T e^{\beta s}|f(s,0,0)|^2\, ds\right),$$

for a constant C *depending only on* K *and* β*.*

Sketch. Apply Itô's formula to $\varphi_t := e^{\beta t}|Y_t|^2$:

$$d\varphi_t = \beta e^{\beta t}|Y_t|^2\, dt + 2e^{\beta t} Y_t\, dY_t + e^{\beta t}|Z_t|^2\, dt.$$

Substituting $dY_t = -f(t,Y_t,Z_t)\, dt + Z_t\, dB_t$ and integrating from t to T,

$$\begin{aligned} e^{\beta t}|Y_t|^2 + \int_t^T e^{\beta s}\big(\beta|Y_s|^2 + |Z_s|^2\big)\, ds = e^{\beta T}|\xi|^2 &+ 2\int_t^T e^{\beta s}\, Y_s\, f(s,Y_s,Z_s)\, ds \\ &- 2\int_t^T e^{\beta s}\, Y_s\, Z_s\, dB_s. \end{aligned}$$

Take expectations (the Itô integral has zero mean given local boundedness, which one bootstraps from $(Y,Z) \in \mathcal{S}^2 \times \mathcal{H}^2$). Decompose

$$Y_s\, f(s,Y_s,Z_s) = Y_s\, f(s,0,0) + Y_s\big(f(s,Y_s,Z_s) - f(s,0,0)\big),$$

use Lipschitz on the second piece and the elementary $2|y||z| \le \varepsilon|y|^2 + \varepsilon^{-1}|z|^2$, then choose ε small and β large enough to absorb $\beta|Y|^2$ and $|Z|^2$ on the left. The constant $\beta > 2K + 2K^2$ does the job. $\square$

A consequence (apply the lemma with $\xi = 0$ and $f \equiv 0$, to two solutions): *uniqueness.* Two solutions (Y, Z), (Y', Z') to the same BSDE have difference $(Y - Y', Z - Z')$ solving a BSDE with zero data; the estimate forces both to vanish.

3.2 Existence by contraction

Define a map $\Phi : \mathcal{S}^2 \times \mathcal{H}^2 \to \mathcal{S}^2 \times \mathcal{H}^2$ as follows. Given (y, z), set $(Y, Z) := \Phi(y, z)$ to be the solution of the BSDE with *frozen* driver $f^{y,z}(t, \omega) := f(t, \omega, y_t, z_t)$ and terminal condition ξ:

$$Y_t = \xi + \int_t^T f(s, y_s, z_s)\, ds - \int_t^T Z_s\, dB_s.$$

For this linear sub-problem (no Y- or Z-dependence on the right) the solution can be written explicitly: set

$$N_t := \mathbb{E}\left[\xi + \int_0^T f(s, y_s, z_s)\, ds \,\middle|\, \mathcal{F}_t\right].$$

Then N is a square-integrable martingale; by Brownian representation, $N_t = N_0 + \int_0^t Z_s\, dB_s$ for a unique predictable $Z \in \mathcal{H}^2$. Put

$$Y_t := N_t - \int_0^t f(s, y_s, z_s)\, ds,$$

and verify directly that (Y, Z) satisfies the sub-problem and lies in $\mathcal{S}^2 \times \mathcal{H}^2$.

The next lemma turns the a priori bound into a quantitative Lipschitz estimate on Φ.

Lemma 3.3 (Difference estimate)**.** *For $(y, z), (y', z') \in \mathcal{S}^2 \times \mathcal{H}^2$, let $(Y, Z) := \Phi(y, z)$, $(Y', Z') := \Phi(y', z')$. For every $\beta \ge 2$,*

$$\|(Y - Y', Z - Z')\|_\beta^2 \;\le\; \frac{4K^2}{\beta}\, \|(y - y', z - z')\|_\beta^2,$$

where $\|(U, V)\|_\beta^2 := \mathbb{E}\int_0^T e^{\beta s}\big(|U_s|^2 + |V_s|^2\big)\, ds$ and K is the Lipschitz constant of f in (y, z).

Proof. Write $\Delta y := y - y'$, $\Delta z := z - z'$, $\Delta Y := Y - Y'$, $\Delta Z := Z - Z'$, and $\Delta f_s := f(s, y_s, z_s) - f(s, y'_s, z'_s)$. The two frozen BSDEs are identical at time T ($\Delta Y_T = 0$) and yield

$$d\Delta Y_t = -\Delta f_t\, dt + \Delta Z_t\, dB_t.$$

Apply Itô's formula to $e^{\beta s}|\Delta Y_s|^2$ on $[0, T]$, take expectations, and use $\Delta Y_T = 0$:

$$\mathbb{E}\int_0^T e^{\beta s}\big(\beta|\Delta Y_s|^2 + |\Delta Z_s|^2\big)\, ds \;=\; 2\,\mathbb{E}\int_0^T e^{\beta s}\, \Delta Y_s\, \Delta f_s\, ds.$$

By Lipschitz, $|\Delta f_s| \le K(|\Delta y_s| + |\Delta z_s|)$. Applying $2ab \le \frac{\beta}{4}a^2 + \frac{4}{\beta}b^2$ to each cross term,

$$2\,\Delta Y_s\, \Delta f_s \;\le\; \tfrac{\beta}{2}|\Delta Y_s|^2 + \tfrac{4K^2}{\beta}\big(|\Delta y_s|^2 + |\Delta z_s|^2\big).$$

Absorb the $\frac{\beta}{2}|\Delta Y|^2$ on the left and use $\beta \ge 2$ to bound the $|\Delta Z|^2$ coefficient below by 1, giving

$$\mathbb{E}\int_0^T e^{\beta s}\big(|\Delta Y_s|^2 + |\Delta Z_s|^2\big)\, ds \;\le\; \frac{4K^2}{\beta}\,\mathbb{E}\int_0^T e^{\beta s}\big(|\Delta y_s|^2 + |\Delta z_s|^2\big)\, ds.$$

□

Choosing any $\beta > 4K^2$ (e.g. $\beta = 8K^2 + 2$) makes Φ a strict contraction in $\|\cdot\|_\beta$ with ratio $\sqrt{4K^2/\beta} < 1$. Banach's fixed-point theorem on the complete metric space $(\mathcal{S}^2 \times \mathcal{H}^2, \|\cdot\|_\beta)$ produces a unique fixed point, which by construction solves the BSDE (1). Combined with the uniqueness corollary, this completes the proof of Theorem 3.1. □

Remark (Why a weighted norm?). The unweighted L^2-norm fails because the contraction constant from the same calculation reads $4K^2T$, which is ≥ 1 once T is moderately large. The exponential weight $e^{\beta s}$ rescales time so that one can spend the full Lipschitz budget locally without the global horizon T entering. The same trick (under the name "Bielecki norm") is the standard cure for Picard-type contractions on long intervals — it appeared in [SDE III, §2] for forward SDEs and is even more crucial here, because the BSDE cannot be solved by short-step concatenation: shrinking T does not help, since the terminal datum lives at T.

Remark (What if f is only monotone in y?). Lipschitz can be replaced by "monotonicity in y plus polynomial growth in y, Lipschitz in z" (Pardoux 1999). The proof becomes more delicate but follows the same outline. Quadratic growth in z (Kobylanski 2000) requires bounded ξ and entirely different techniques (exponential transforms).

4 The Comparison Theorem

A defining feature of BSDEs — and one with no counterpart in the forward theory — is the following monotonicity in the data.

Theorem 4.1 (Comparison). *Let (Y^1, Z^1) and (Y^2, Z^2) solve the BSDEs (ξ^1, f^1) and (ξ^2, f^2) respectively, with both drivers Lipschitz. Assume*

$$\xi^1 \leq \xi^2 \ \ a.s., \qquad f^1(t, Y_t^1, Z_t^1) \leq f^2(t, Y_t^1, Z_t^1) \quad dt \otimes d\mathbb{P}\text{-}a.e.$$

Then $Y_t^1 \leq Y_t^2$ a.s. for every $t \in [0, T]$.

Sketch. Set $\Delta Y := Y^2 - Y^1$ and $\Delta Z := Z^2 - Z^1$. Subtracting the two BSDEs,

$$-d\Delta Y_t = \left[f^2(t, Y^2, Z^2) - f^1(t, Y^1, Z^1)\right] dt - \Delta Z_t \, dB_t.$$

Linearize the driver difference:

$$f^2(t, Y_t^2, Z_t^2) - f^1(t, Y_t^1, Z_t^1) = \alpha_t \Delta Y_t + \beta_t \cdot \Delta Z_t + \gamma_t,$$

where α, β are bounded (by Lipschitz) and $\gamma_t := f^2(t, Y_t^1, Z_t^1) - f^1(t, Y_t^1, Z_t^1) \geq 0$ by hypothesis. So ΔY solves a *linear* BSDE with positive source. By a change of measure (Girsanov, [SDE III, §5]) under

$$\left.\frac{d\mathbb{P}^\beta}{d\mathbb{P}}\right|_{\mathcal{F}_T} = \exp\Big(\int_0^T \beta_s \, dB_s - \tfrac{1}{2}\int_0^T |\beta_s|^2 \, ds\Big),$$

the term $\beta \cdot \Delta Z$ is absorbed into the Brownian motion. Applying the linear-driver formula from §2.2,

$$\Delta Y_t = \mathbb{E}^{\mathbb{P}^\beta}\left[e^{\int_t^T \alpha_s \, ds} \Delta\xi + \int_t^T e^{\int_t^s \alpha_r \, dr} \gamma_s \, ds \,\middle|\, \mathcal{F}_t \right].$$

The integrand is a sum of two nonnegative quantities ($\Delta\xi \geq 0$ and $\gamma \geq 0$), so $\Delta Y_t \geq 0$. □

Remark (Finance interpretation). In an option-pricing model, Y_t represents the price at time t of a contract paying ξ at maturity T, and f encodes the cost of hedging (interest rates, dividends, portfolio constraints). The comparison theorem says: *a more valuable payoff or a cheaper hedge yields a higher price.* This is the dynamic-programming form of no-arbitrage; it is what powers, for instance, the upper-bound pricing under model uncertainty.

Corollary 4.2 (Strict comparison). *Under the hypotheses of Theorem 4.1, if additionally*

$$\mathbb{P}(\xi^1 < \xi^2) > 0 \qquad \text{or} \qquad f^1 < f^2 \text{ at } (Y^1, Z^1) \text{ on a set of positive } dt \otimes d\mathbb{P}\text{-measure,}$$

then $Y_0^1 < Y_0^2$.

Proof. In the representation derived in the proof of Theorem 4.1,

$$\Delta Y_0 = \mathbb{E}^{\mathbb{P}^\beta}\left[e^{\int_0^T \alpha_s\, ds} \Delta\xi + \int_0^T e^{\int_0^s \alpha_r\, dr} \gamma_s\, ds \right],$$

the integrand is nonnegative and, under one of the two hypotheses, strictly positive on a set of positive $\mathbb{P}^\beta$-measure (which coincides with positive $\mathbb{P}$-measure by absolute continuity of $\mathbb{P}^\beta$). Hence $\Delta Y_0 > 0$. □

4.1 g-expectations and dynamic risk measures

The comparison theorem is what makes the following construction well-defined. Fix a Lipschitz driver $g : [0,T] \times \mathbb{R} \times \mathbb{R}^d \to \mathbb{R}$ satisfying $g(t, y, 0) = 0$, and for $\xi \in L^2(\mathcal{F}_T)$ let (Y^ξ, Z^ξ) be the unique solution of the BSDE (ξ, g). Define the *g-expectation* and its conditional version by

$$\mathcal{E}_g[\xi] := Y_0^\xi, \qquad \mathcal{E}_g[\xi \mid \mathcal{F}_t] := Y_t^\xi.$$

The condition $g(\cdot, \cdot, 0) = 0$ ensures that $\mathcal{E}_g[c] = c$ for constants. The operator $\mathcal{E}_g$ is in general nonlinear in ξ (linearity holds iff g is linear in (y, z)), but it retains the structural features of an expectation:

(M) *Monotonicity:* $\xi^1 \le \xi^2 \Rightarrow \mathcal{E}_g[\xi^1 \mid \mathcal{F}_t] \le \mathcal{E}_g[\xi^2 \mid \mathcal{F}_t]$ (Theorem 4.1).

(T) *Time-consistency:* $\mathcal{E}_g[\mathcal{E}_g[\xi \mid \mathcal{F}_s] \mid \mathcal{F}_t] = \mathcal{E}_g[\xi \mid \mathcal{F}_t]$ for $t \le s$ (flow property of the BSDE).

(L) *Locality / zero-one law:* $\mathbf{1}_A \mathcal{E}_g[\xi \mid \mathcal{F}_t] = \mathbf{1}_A \mathcal{E}_g[\mathbf{1}_A \xi \mid \mathcal{F}_t]$ for $A \in \mathcal{F}_t$.

When g is positively homogeneous and subadditive in (y, z), $\rho_t(\xi) := \mathcal{E}_g[-\xi \mid \mathcal{F}_t]$ is a *dynamic coherent risk measure* in the sense of Artzner–Delbaen–Eber–Heath; convex g gives a dynamic convex risk measure. Conversely, a theorem of Coquet–Hu–Mémin–Peng characterizes which dynamically consistent nonlinear expectations arise this way: precisely those dominated by a g_μ-expectation for the linear driver $g_\mu(z) = \mu|z|$. Nonlinearity in g encodes hedging frictions (bid–ask spreads, collateral, capital charges) and Knightian uncertainty about the law of X.

5 The Nonlinear Feynman–Kac Formula

We now specialize to the Markovian setting that makes contact with PDEs.

5.1 Setting

Fix $(t, x) \in [0, T) \times \mathbb{R}^n$. Consider the forward SDE

$$dX_s^{t,x} = b(s, X_s^{t,x})\, ds + \sigma(s, X_s^{t,x})\, dB_s, \qquad X_t^{t,x} = x, \ s \in [t, T], \tag{2}$$

with $b : [0, T] \times \mathbb{R}^n \to \mathbb{R}^n$ and $\sigma : [0, T] \times \mathbb{R}^n \to \mathbb{R}^{n\times d}$ satisfying the standing Lipschitz and linear-growth hypotheses from [SDE III, §2]. Coupled to it, consider the BSDE

$$Y_s^{t,x} = g(X_T^{t,x}) + \int_s^T f\big(r, X_r^{t,x}, Y_r^{t,x}, Z_r^{t,x}\big)\, dr - \int_s^T Z_r^{t,x}\, dB_r, \tag{3}$$

with $g : \mathbb{R}^n \to \mathbb{R}$ continuous of polynomial growth and $f : [0, T] \times \mathbb{R}^n \times \mathbb{R} \times \mathbb{R}^d \to \mathbb{R}$ continuous and uniformly Lipschitz in (y, z). By Theorem 3.1, the BSDE (3) has a unique solution.

Define the candidate PDE solution

$$u(t, x) := Y_t^{t,x}.$$

Note that $Y_t^{t,x}$ is $\mathcal{F}_t$-measurable, and because $X_t^{t,x} = x$ deterministically, the value $Y_t^{t,x}$ is actually deterministic in (t, x) — so $u(t, x)$ is a number, not a random variable.

Recall the generator of X:

$$\mathcal{L}\varphi(t, x) := b(t, x) \cdot \nabla\varphi(t, x) + \tfrac{1}{2} \operatorname{tr}\!\big(\sigma\sigma^\top(t, x)\, \nabla^2\varphi(t, x)\big).$$

Theorem 5.1 (Nonlinear Feynman–Kac)**.**

(a) *(Verification) If a function $v \in C^{1,2}([0, T] \times \mathbb{R}^n)$ of polynomial growth solves the semilinear parabolic problem*

$$\partial_t v(t, x) + \mathcal{L}v(t, x) + f\big(t, x, v(t, x), \sigma^\top(t, x)\, \nabla v(t, x)\big) = 0, \tag{4}$$

with terminal condition $v(T, x) = g(x)$, then for every (t, x),

$$Y_s^{t,x} = v(s, X_s^{t,x}), \qquad Z_s^{t,x} = \sigma^\top(s, X_s^{t,x})\, \nabla v(s, X_s^{t,x}),$$

is the unique solution to the BSDE (3)*. In particular $u(t, x) = v(t, x)$.*

(b) *(Representation) Conversely, if the candidate $u(t, x) := Y_t^{t,x}$ defined from the BSDE is in $C^{1,2}$, then u solves* (4) *with terminal condition g.*

Verification half. Suppose $v \in C^{1,2}$ solves (4). Apply Itô's formula to $v(s, X_s^{t,x})$:

$$dv(s, X_s) = \big(\partial_s v + \mathcal{L}v\big)(s, X_s)\, ds + \nabla v(s, X_s)^\top \sigma(s, X_s)\, dB_s.$$

By the PDE (4), the drift becomes $-f(s, X_s, v, \sigma^\top\nabla v)$, so

$$-dv(s, X_s) = f\big(s, X_s, v(s, X_s), \sigma^\top(s, X_s)\nabla v(s, X_s)\big)\, ds - \sigma^\top(s, X_s)\nabla v(s, X_s)\, dB_s,$$

with terminal value $v(T, X_T) = g(X_T)$. This is exactly the BSDE (3) with $Y_s := v(s, X_s)$ and $Z_s := \sigma^\top(s, X_s)\nabla v(s, X_s)$. By the uniqueness half of Theorem 3.1, this is the only solution. □

The converse (representation) half is more subtle: showing $u \in C^{1,2}$ requires either Malliavin smoothness of the flow $x \mapsto X_T^{t,x}$, or a regularization-by-uniformly-elliptic-noise argument. We refer to [2] and [4] for two complementary approaches.

Remark (Why $\sigma^\top\nabla u$ and not ∇u?). The forward SDE drives the noise through σ. Itô's formula on $v(s, X_s)$ produces $(\nabla v)^\top \sigma\, dB$, so the integrand against dB — which is precisely what Z is — equals $\sigma^\top\nabla v$. When $\sigma = I_d$ (driftless Brownian noise on $\mathbb{R}^d$), this collapses to the familiar $Z = \nabla u$.

Remark (The Malliavin viewpoint). If ξ and the data are Malliavin-differentiable in the sense of Nualart, then so is the solution (Y, Z), and one has the identification

$$Z_t \;=\; D_t Y_t \qquad dt \otimes d\mathbb{P}\text{-a.e.}, \tag{5}$$

where D_t is the Malliavin derivative (Pardoux–Peng 1992; El Karoui–Peng–Quenez 1997). The pair $(D_sY_t, D_sZ_t)_{s\le t}$ itself solves a linear BSDE obtained by differentiating (1) formally, which feeds back the regularity of the data into regularity of the solution. In the Markovian setting the chain rule and the tangent-flow formula $D_sX_t = \mathbf{1}_{s\le t}\, \nabla_x X_t^{s,X_s}\, \sigma(s, X_s)$ collapse (5) to

$$Z_t = \sigma^\top(t, X_t)\, \nabla u(t, X_t),$$

matching part (a) of Theorem 5.1 *without* assuming $u \in C^{1,2}$ a priori. The Malliavin route is in fact how the converse half is proved: one establishes Sobolev (rather than classical) regularity of u on the support of the law of X, then upgrades to $C^{1,2}$ via parabolic Schauder estimates when $\sigma\sigma^\top$ is uniformly elliptic.

5.2 Linear vs. nonlinear: a side-by-side comparison

Linear Feynman–Kac ([SDE III])	**Nonlinear Feynman–Kac (this article)**
PDE: $\partial_t u + \mathcal{L}u = 0$, $u(T, x) = g(x)$	PDE: $\partial_t u + \mathcal{L}u + f(t, x, u, \sigma^\top \nabla u) = 0$, $u(T, x) = g(x)$
Probabilistic side: forward SDE only ($X^{t,x}$)	Probabilistic side: forward SDE *plus* BSDE for (Y, Z)
Representation: $u(t, x) = \mathbb{E}[g(X_T^{t,x})]$	Representation: $u(t, x) = Y_t^{t,x}$, the initial value of the BSDE
Gradient ∇u: not represented; need Bismut–Elworthy or Malliavin	Gradient: $Z_t^{t,x} = \sigma^\top(t, x)\, \nabla u(t, x)$ — read off directly from the BSDE
PDE drift has no u-dependence	f may depend on u and ∇u in any (Lipschitz) nonlinear way

The right column reduces to the left when $f \equiv 0$, which is the warm-up example from §2.2.

5.3 A worked example: a Fisher–KPP-type equation

Consider on $\mathbb{R}$ the semilinear equation

$$\partial_t u + \tfrac{1}{2}\partial_{xx} u + u(1 - u) = 0, \qquad u(T, x) = \mathbf{1}_{x>0},$$

with $b \equiv 0$, $\sigma \equiv 1$. The associated forward process is Brownian motion $X_s = x + B_s - B_t$, and the BSDE is

$$Y_s = \mathbf{1}_{x+B_T-B_t>0} + \int_s^T Y_r(1 - Y_r)\, dr - \int_s^T Z_r\, dB_r.$$

Theorem 5.1 identifies $u(t, x) = Y_t^{t,x}$. Note that $u \in [0, 1]$ pathwise — one can check that the driver $y \mapsto y(1-y)$ preserves the interval $[0, 1]$ via the comparison theorem applied to the constant solutions $y \equiv 0$ and $y \equiv 1$. This is exactly the maximum principle for the Fisher–KPP equation, recovered probabilistically.

The BSDE route also gives a Monte Carlo algorithm for u that never touches a grid: simulate X forward, then march (Y, Z) backward by a discretization. This is the entry point to the Deep BSDE method of §6.2.

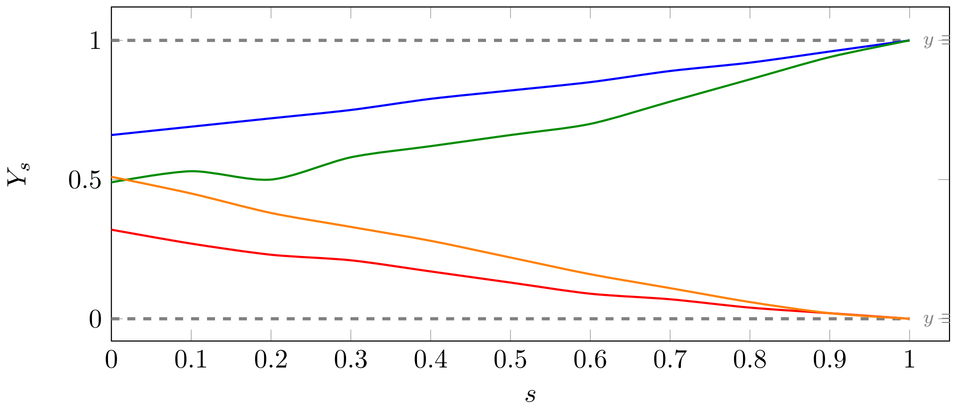

Figure 2: Four simulated Y-trajectories for the Fisher–KPP BSDE on $[0, T] = [0, 1]$ with $x = 0$, generated by Euler–Maruyama discretization of the forward Brownian path together with an explicit backward sweep for (Y, Z). All trajectories remain in $[0, 1]$: the constants $y \equiv 0$ and $y \equiv 1$ are themselves BSDE solutions for the boundary terminal data $\xi = 0$ and $\xi = 1$, and the comparison theorem (Theorem 4.1) sandwiches every solution between them. The terminal values realize the indicator $\mathbf{1}_{B_T > 0}$ along each Brownian path; the initial values Y_0 are Monte Carlo estimates of $u(0, 0) \approx \frac{1}{2}$.

6 Applications

6.1 Stochastic optimal control and HJB

Consider a controlled diffusion

$$dX_s^\alpha = b(s, X_s^\alpha, \alpha_s)\, ds + \sigma(s, X_s^\alpha)\, dB_s,$$

where the control α takes values in a compact action set $A \subset \mathbb{R}^k$ and is required to be adapted. With running cost ℓ and terminal cost g, define the value function

$$V(t, x) := \inf_\alpha \mathbb{E}\left[\int_t^T \ell(s, X_s^\alpha, \alpha_s)\, ds + g(X_T^\alpha) \,\middle|\, X_t^\alpha = x\right].$$

By the dynamic programming principle, V formally satisfies the Hamilton–Jacobi–Bellman equation

$$\partial_t V(t, x) + \inf_{a \in A} \Big\{ b(t, x, a) \cdot \nabla V + \tfrac{1}{2} \operatorname{tr}(\sigma\sigma^\top \nabla^2 V) + \ell(t, x, a) \Big\} = 0, \quad V(T, x) = g(x).$$

In the case where σ does not depend on a (so-called *drift control*), the HJB equation is of the semilinear form (4) with

$$f(t, x, y, z) = \inf_{a \in A} \Big\{ b(t, x, a) \cdot \sigma^{-\top}(t, x)\, z + \ell(t, x, a) \Big\}$$

(assuming σ is invertible; if not, work with $\sigma^\top \nabla V = Z$ as the primary object and a generalized Hamiltonian). The BSDE associated with this driver therefore gives a probabilistic representation of V that bypasses the regularity demands of viscosity-solution theory, and produces an optimal feedback control $\alpha^*(s, X_s) \in \arg\min_a \{\cdots\}$ as a measurable selector.

When σ *does* depend on a, the BSDE picture must be replaced by a *second-order* BSDE (2BSDE, Soner–Touzi–Zhang); this is one of the threads of §7.

6.2 Deep BSDE for high-dimensional PDEs

The PDE (4) in high spatial dimension n is out of reach for grid-based solvers: storage is $\Theta(N^n)$ for N points per axis. The Deep BSDE method of Han, Jentzen, and E ([3]) recasts solving the PDE as a single global stochastic optimization. Discretize time as $0 = t_0 < t_1 < \cdots < t_M = T$ with $\Delta t = T/M$. Parameterize:

- the initial value by a scalar $\theta_0 \in \mathbb{R}$ (the candidate for $u(0, x_0)$);
- the spatial gradient $z(t, x) := \sigma^\top(t, x)\,\nabla u(t, x)$ by a neural network $z^\theta(t, x)$ with parameters θ.

Simulate the forward Euler–Maruyama discretization of X starting from x_0, and propagate

$$Y^\theta_{t_{k+1}} = Y^\theta_{t_k} - f\big(t_k, X_{t_k}, Y^\theta_{t_k}, z^\theta(t_k, X_{t_k})\big)\,\Delta t + z^\theta(t_k, X_{t_k}) \cdot \Delta B_k,$$

with $Y^\theta_0 = \theta_0$. Train $\theta = (\theta_0, \theta_{\text{net}})$ by stochastic gradient descent on the terminal-condition mismatch loss

$$\mathcal{J}(\theta) := \mathbb{E}\Big[\big|g(X_T) - Y^\theta_T\big|^2\Big].$$

At a minimizer, $\theta_0 \approx u(0, x_0)$ and $z^\theta(\cdot, \cdot) \approx \sigma^\top \nabla u$.

The mathematical content is the *variational characterization*: the function $u(0, x_0)$ is the unique value such that, when $Y_0 = u(0, x_0)$ and Z is the BSDE solution, the residual $g(X_T) - Y_T$ vanishes a.s. Replacing "a.s." by "in L^2" gives the loss, and replacing "BSDE solution Z" by "best-NN approximation z^θ" gives the algorithm. The connection between L^2-loss and a.s. equality is precisely the BSDE uniqueness lemma in disguise.

In practice, Deep BSDE handles $n = 100$ on a laptop in minutes, well beyond the reach of finite differences. Reported accuracy is 10^{-3} to 10^{-4} relative error on benchmark problems (HJB, Black–Scholes with default risk, Allen–Cahn).

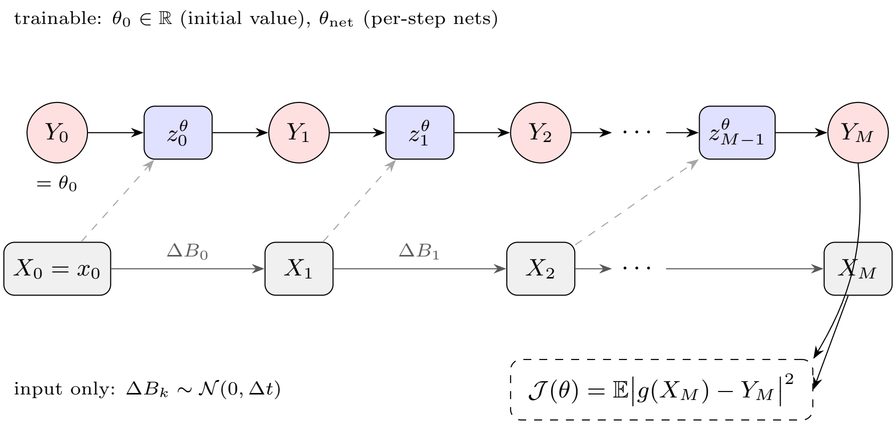

Figure 3: Architecture of the Deep BSDE method. The forward chain (gray) is the Euler–Maruyama discretization of X driven by sampled Brownian increments ΔB_k. The backward chain (red/blue) propagates Y^θ *forward in time* from a trainable scalar θ_0 using $Y^\theta_{k+1} = Y^\theta_k - f(t_k, X_k, Y^\theta_k, z^\theta_k)\Delta t + z^\theta_k \cdot \Delta B_k$, where each z^θ_k is a small neural network mapping $X_k \mapsto \sigma^\top(t_k, X_k)\nabla u(t_k, X_k)$. Training by SGD on the terminal mismatch loss $\mathcal{J}(\theta)$ drives $\theta_0 \to u(0, x_0)$. The variational characterization that makes this work is, in essence, the uniqueness half of Theorem 3.1.

7 Outlook

We close with four threads that the next article in this series could follow.

Forward–backward coupling (FBSDEs). If the forward coefficients depend on Y or Z, we obtain a *coupled* forward–backward SDE:

$$\begin{aligned} dX_s &= b(s,X_s,Y_s,Z_s)\,ds + \sigma(s,X_s,Y_s,Z_s)\,dB_s, \\ -dY_s &= f(s,X_s,Y_s,Z_s)\,ds - Z_s\,dB_s, \qquad Y_T = g(X_T). \end{aligned}$$

This is the natural language of fully coupled stochastic optimal control via Pontryagin's maximum principle. Existence requires either a monotonicity condition (Ma–Yong), a small-time hypothesis (Antonelli), or the four-step scheme (Ma–Protter–Yong) when the PDE side has a classical solution.

Mean-field BSDEs. Letting the driver depend on the marginal law $\mathcal{L}(Y_s)$ leads to McKean–Vlasov BSDEs, the natural language of mean-field games. The corresponding PDE lives on $[0,T]\times\mathbb{R}^n\times\mathcal{P}_2(\mathbb{R}^n)$ — the so-called master equation — and links back to the Wasserstein calculus of Lions.

Reflected BSDEs and American options. A reflected BSDE (El Karoui–Kapoudjian–Pardoux–Peng–Quenez 1997) adds an obstacle process $L=(L_t)$ to the equation: a solution is a triple (Y,Z,K) where K is a continuous nondecreasing process with $K_0=0$, $Y_t\ge L_t$ a.s., and

$$Y_t = \xi + \int_t^T f(s,Y_s,Z_s)\,ds + K_T - K_t - \int_t^T Z_s\,dB_s,$$

together with the Skorokhod condition $\int_0^T (Y_t - L_t)\,dK_t = 0$. The process K pushes Y upward minimally, only when Y touches the obstacle. The associated PDE is the obstacle problem

$$\min\bigl\{-\partial_t u - \mathcal{L}u - f(t,x,u,\sigma^\top\nabla u),\ u - L\bigr\} = 0, \qquad u(T,x) = g(x),$$

and the BSDE representation gives $Y_t = \operatorname{esssup}_{\tau\in[t,T]}\mathbb{E}[\int_t^\tau f\,ds + L_\tau\mathbf{1}_{\tau<T} + \xi\mathbf{1}_{\tau=T}\mid\mathcal{F}_t]$, i.e. the value of an optimal stopping problem with running cost f. The textbook application is American-option pricing: $L_t=(K-X_t)^+$ for a put with strike K, $f\equiv -ry$ for discounting; the early-exercise boundary is the contact set $\{Y=L\}$. Double-obstacle (Dynkin game) and second-order variants follow the same template.

Rough paths and rough BSDEs. If the driving noise is rougher than Brownian motion (e.g. fractional Brownian motion with Hurst parameter $H<1/2$), the Itô calculus underlying BSDEs breaks down. Lyons' rough path theory and Hairer's regularity structures supply the replacement; rough BSDEs are an active area, with the Diehl–Friz–Gassiat formulation as the standard reference.

Acknowledgments. We thank our reader for following the series. As always, comments, corrections, and counter-examples are welcome.

References

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

[2] El Karoui, N., Peng, S., Quenez, M. C. *Backward stochastic differential equations in finance.* Math. Finance 7 (1997), 1–71.

[3] Han, J., Jentzen, A., E, W. *Solving high-dimensional partial differential equations using deep learning.* PNAS 115 (2018), 8505–8510.

[4] Pardoux, É., Răşcanu, A. *Stochastic Differential Equations, Backward SDEs, Partial Differential Equations.* Springer, 2014.

[5] Revuz, D., Yor, M. *Continuous Martingales and Brownian Motion*, 3rd ed., Springer, 1999.