
Prerequisites for Mean-Field Games and MCMC Methods in Trading

A self-contained mathematical primer covering probability theory,
stochastic calculus, optimal control, and Bayesian inference.

Melvin Alvarez-Caradu
Independent Researcher, HFThot Research Lab
<https://hfthot-lab.eu>
June 2026 — v1.0

DISTRIBUTION: REGISTERED USERS — FREE TIER

This document provides a self-contained mathematical primer for readers approaching the HFThot Research Lab white paper “Optimal Trading and Option Market Making via Mean-Field Games and MCMC Methods” (WP1). We cover, from first principles, the essential mathematical machinery: measure-theoretic probability, stochastic calculus and Itô theory, partial differential equations, optimal control and the Hamilton–Jacobi–Bellman equation, game theory for large populations, Bayesian inference via Markov Chain Monte Carlo, and numerical methods for stochastic systems. Each section contains full derivations, worked examples with explicit calculations, and exercises. The treatment is rigorous but accessible, pitched at the level of a first-year graduate course in mathematical finance.

Prerequisites

Required background for this primer: Multivariable calculus (including measure and integration at the level of undergraduate real analysis), linear algebra, basic ordinary differential equations, and elementary probability (discrete and continuous distributions). No prior knowledge of stochastic calculus or game theory is assumed.

Contents

1	How to Use This Document	3
2	Probability Theory Foundations	4
2.1	Measure Spaces and σ -Algebras	4
2.2	Random Variables and Distributions	4
2.3	Conditional Expectation	5
2.4	Modes of Convergence	6
2.5	Characteristic Functions and Limit Theorems	6
3	Stochastic Calculus	8
3.1	Filtrations and Adapted Processes	8
3.2	Martingales	8
3.3	Brownian Motion	8
3.4	The Itô Integral	9
3.5	Itô’s Formula	10
3.6	Stochastic Differential Equations	10
3.7	Girsanov’s Theorem	11
3.8	Feynman–Kac Formula	11
4	Partial Differential Equations	14
4.1	Classification of Second-Order Linear PDEs	14
4.2	The Heat Equation	14
4.3	The Fokker–Planck Equation	14
4.4	Viscosity Solutions	15
5	Optimal Control Theory	18
5.1	Deterministic Optimal Control	18
5.2	Dynamic Programming and the HJB Equation	18
5.3	Stochastic Optimal Control	19

5.4	Linear-Quadratic Regulator (Worked Example)	19
6	Game Theory for Continuous Games	22
6.1	Nash Equilibrium in N -Player Games	22
6.2	The Mean-Field Limit: $N \rightarrow \infty$	22
6.3	Propagation of Chaos	23
6.4	Fixed-Point Characterisation and Existence	23
7	Bayesian Inference and MCMC	25
7.1	Bayes' Theorem and Posterior Computation	25
7.2	Conjugate Priors and Exponential Families	25
7.3	Monte Carlo Integration	25
7.4	Markov Chains: Foundations	26
7.5	Metropolis–Hastings Algorithm	26
7.6	Hamiltonian Monte Carlo	27
7.7	Convergence Diagnostics	28
8	Numerical Methods	31
8.1	Euler–Maruyama Scheme for SDEs	31
8.2	Finite Difference Methods for PDEs	31
8.3	Newton–Raphson and Fixed-Point Iteration	32
9	Connecting the Pieces: Roadmap to WP1	35
A	Notation and Symbol Table	36
	References	36

1 How to Use This Document

This primer is organised in a dependency-aware order: each section builds on the preceding ones. Readers with a strong background in stochastic analysis may skip directly to Section 6 (Game Theory) or Section 7 (Bayesian Inference). For a minimal reading path that covers just enough to follow WP1, we recommend:

Step 1: Section 2: Probability Foundations (skim if familiar with measure-theoretic probability).

Step 2: Section 3: Stochastic Calculus (essential; read carefully).

Step 3: Section 4: PDEs (focus on Fokker–Planck and viscosity solutions).

Step 4: Section 5: Optimal Control (the HJB equation is central to WP1).

Step 5: Section 6: Game Theory (read all — this is the core of the whitepaper).

Step 6: Section 7: Bayesian Inference & MCMC (essential for the calibration part of WP1).

Step 7: Section 8: Numerical Methods (reference as needed).

Key Insight

Boxes like this one highlight conceptual points that are crucial for understanding WP1. **Practical Application** boxes connect the mathematics to concrete trading contexts. **Lab Reference** boxes point to specific code in the HFThot Lab implementation.

Throughout the text, we use the notation established in Appendix A. Cross-references to the main whitepaper are indicated by “[WP1]”.

2 Probability Theory Foundations

Classical (Kolmogorov) probability theory relies on the language of measure theory. This section develops the essential concepts at a level sufficient for stochastic calculus.

2.1 Measure Spaces and σ -Algebras

Definition 2.1 (σ -Algebra). Let Ω be a non-empty set. A collection $\mathcal{F} \subseteq 2^\Omega$ is a σ -algebra (or σ -field) on Ω if:

- (i) $\Omega \in \mathcal{F}$.
- (ii) If $A \in \mathcal{F}$, then $A^c := \Omega \setminus A \in \mathcal{F}$ (closure under complementation).
- (iii) If $(A_n)_{n \geq 1}$ is a sequence in $\mathcal{F}$, then $\bigcup_{n=1}^{\infty} A_n \in \mathcal{F}$ (closure under countable unions).

The pair $(\Omega, \mathcal{F})$ is called a *measurable space*.

The σ -algebra encodes the events whose probability we can measure. In a financial context, Ω is the set of all possible market outcomes and $\mathcal{F}$ contains the observable events.

Example 2.1 (Borel σ -algebra). On $\mathbb{R}$, the *Borel σ -algebra* $\mathcal{B}(\mathbb{R})$ is the smallest σ -algebra containing all open intervals (a, b) , $a < b$. It also contains all closed sets, countable sets, and any set constructible from opens by countable unions, intersections, and complements. Every “reasonable” subset of $\mathbb{R}$ is Borel; non-Borel sets require the Axiom of Choice to construct.

Definition 2.2 (Probability Measure). A function $\mathbb{P}: \mathcal{F} \rightarrow [0, 1]$ is a *probability measure* on $(\Omega, \mathcal{F})$ if:

- (i) $\mathbb{P}(\Omega) = 1$.
- (ii) (Countable additivity) For any sequence $(A_n)_{n \geq 1}$ of pairwise disjoint sets in $\mathcal{F}$,

$$\mathbb{P}\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} \mathbb{P}(A_n).$$

The triple $(\Omega, \mathcal{F}, \mathbb{P})$ is called a *probability space*.

Remark 2.1. Countable additivity is strictly stronger than finite additivity. The distinction matters in infinite-dimensional settings such as stochastic processes indexed by continuous time.

2.2 Random Variables and Distributions

Definition 2.3 (Random Variable). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. A *random variable* is a measurable function $X: \Omega \rightarrow \mathbb{R}$, i.e. $X^{-1}(B) := \{\omega \in \Omega : X(\omega) \in B\} \in \mathcal{F}$ for every $B \in \mathcal{B}(\mathbb{R})$. The *law* (or *distribution*) of X is the probability measure μ_X on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ defined by $\mu_X(B) = \mathbb{P}(X \in B)$.

When μ_X has a density f_X with respect to Lebesgue measure—i.e., $\mu_X(B) = \int_B f_X(x) dx$ —we say X is a *continuous random variable* with *probability density function* (pdf) f_X .

The *cumulative distribution function* (cdf) is $F_X(x) = \mathbb{P}(X \leq x) = \int_{-\infty}^x f_X(t) dt$.

Definition 2.4 (Expectation). The *expectation* of a random variable X on $(\Omega, \mathcal{F}, \mathbb{P})$ is

$$\mathbb{E}[X] = \int_{\Omega} X(\omega) d\mathbb{P}(\omega),$$

provided the integral exists (i.e., $\mathbb{E}[|X|] < \infty$). When X has density f_X , this reduces to $\mathbb{E}[X] = \int_{-\infty}^{\infty} x f_X(x) dx$.

Key properties of expectation. For random variables X, Y with finite expectation and constants $a, b \in \mathbb{R}$:

- (i) **Linearity:** $\mathbb{E}[aX + bY] = a\mathbb{E}[X] + b\mathbb{E}[Y]$.
- (ii) **Monotonicity:** If $X \leq Y$ a.s., then $\mathbb{E}[X] \leq \mathbb{E}[Y]$.
- (iii) **Jensen's inequality:** If φ is convex, $\varphi(\mathbb{E}[X]) \leq \mathbb{E}[\varphi(X)]$.
- (iv) **Cauchy–Schwarz:** $|\mathbb{E}[XY]| \leq (\mathbb{E}[X^2])^{1/2}(\mathbb{E}[Y^2])^{1/2}$.

The *variance* and *covariance* are defined as

$$\text{Var}(X) = \mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[X^2] - (\mathbb{E}[X])^2, \quad \text{Cov}(X, Y) = \mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])].$$

2.3 Conditional Expectation

Conditional expectation is the cornerstone of stochastic filtering and dynamic optimisation—both central topics in WP1.

Definition 2.5 (Conditional Expectation). Let X be an integrable random variable on $(\Omega, \mathcal{F}, \mathbb{P})$ and let $\mathcal{G} \subseteq \mathcal{F}$ be a sub- σ -algebra. The *conditional expectation* $\mathbb{E}[X | \mathcal{G}]$ is the (a.s. unique) $\mathcal{G}$ -measurable random variable such that

$$\int_G \mathbb{E}[X | \mathcal{G}] d\mathbb{P} = \int_G X d\mathbb{P} \quad \text{for all } G \in \mathcal{G}.$$

Existence and uniqueness follow from the Radon–Nikodým theorem.

Fundamental properties.

- (1) **Tower property:** If $\mathcal{H} \subseteq \mathcal{G} \subseteq \mathcal{F}$, then $\mathbb{E}[\mathbb{E}[X | \mathcal{G}] | \mathcal{H}] = \mathbb{E}[X | \mathcal{H}]$ a.s.
- (2) **Pull-out property:** If Y is $\mathcal{G}$ -measurable and bounded, then $\mathbb{E}[XY | \mathcal{G}] = Y\mathbb{E}[X | \mathcal{G}]$ a.s.
- (3) **Independence:** If X is independent of $\mathcal{G}$, then $\mathbb{E}[X | \mathcal{G}] = \mathbb{E}[X]$ a.s.
- (4) **Jensen:** If φ is convex, $\varphi(\mathbb{E}[X | \mathcal{G}]) \leq \mathbb{E}[\varphi(X) | \mathcal{G}]$ a.s.

Key Insight

The tower property $\mathbb{E}[\mathbb{E}[X | \mathcal{G}] | \mathcal{H}] = \mathbb{E}[X | \mathcal{H}]$ is used throughout WP1 when deriving the dynamic programming principle. Intuitively: coarsening information twice is the same as coarsening once to the coarser level. This is developed further in the main white paper [WP1], Section 2.

Example 2.2 (Conditional expectation in discrete time). Suppose $X_0, X_1, X_2, \dots$ is a sequence of coin tosses with $\mathbb{P}(X_i = 1) = p$, $\mathbb{P}(X_i = 0) = 1 - p$, and let $S_n = X_1 + \dots + X_n$. Denote $\mathcal{F}_n = \sigma(X_1, \dots, X_n)$. Then:

$$\mathbb{E}[S_{n+1} | \mathcal{F}_n] = \mathbb{E}[S_n + X_{n+1} | \mathcal{F}_n] = S_n + \mathbb{E}[X_{n+1} | \mathcal{F}_n] = S_n + p,$$

using the pull-out property (since S_n is $\mathcal{F}_n$ -measurable) and independence of X_{n+1} from $\mathcal{F}_n$. Therefore, $M_n := S_n - np$ is a martingale with respect to $(\mathcal{F}_n)$.

2.4 Modes of Convergence

Let $(X_n)_{n \geq 1}$ be a sequence of random variables on $(\Omega, \mathcal{F}, \mathbb{P})$ and let X be another.

Definition 2.6 (Modes of Convergence). (i) **Almost sure (a.s.):** $X_n \xrightarrow{\text{a.s.}} X$ if $\mathbb{P}(\omega : X_n(\omega) \rightarrow X(\omega)) = 1$.

(ii) **In L^p :** $X_n \xrightarrow{L^p} X$ if $\mathbb{E}[|X_n - X|^p] \rightarrow 0$ as $n \rightarrow \infty$.

(iii) **In probability:** $X_n \xrightarrow{\mathbb{P}} X$ if $\mathbb{P}(|X_n - X| > \varepsilon) \rightarrow 0$ for every $\varepsilon > 0$.

(iv) **In distribution:** $X_n \xrightarrow{d} X$ if $F_{X_n}(x) \rightarrow F_X(x)$ at every continuity point x of F_X .

The standard implications are:

$$X_n \xrightarrow{\text{a.s.}} X \implies X_n \xrightarrow{\mathbb{P}} X \implies X_n \xrightarrow{d} X, \quad X_n \xrightarrow{L^p} X \implies X_n \xrightarrow{\mathbb{P}} X.$$

None of these arrows reverse in general, but convergence in probability plus uniform integrability implies L^1 convergence.

2.5 Characteristic Functions and Limit Theorems

Definition 2.7 (Characteristic Function). The *characteristic function* of a random variable X is $\varphi_X(t) = \mathbb{E}[e^{itX}] = \int_{-\infty}^{\infty} e^{itx} dF_X(x)$, $t \in \mathbb{R}$.

Key properties: $\varphi_X(0) = 1$, $|\varphi_X(t)| \leq 1$, and φ_X is uniformly continuous. If $\mathbb{E}[|X|^k] < \infty$, then $\varphi_X^{(k)}(0) = i^k \mathbb{E}[X^k]$. Two distributions are identical if and only if their characteristic functions agree (the *uniqueness theorem*).

The *moment generating function* (mgf) is $M_X(t) = \mathbb{E}[e^{tX}]$ for t in a neighbourhood of zero (when it exists). It determines the distribution and provides moments via $\mathbb{E}[X^k] = M_X^{(k)}(0)$.

Theorem 2.1 (Central Limit Theorem). Let $X_1, X_2, \dots$ be i.i.d. random variables with $\mathbb{E}[X_1] = \mu$ and $\text{Var}(X_1) = \sigma^2 \in (0, \infty)$. Set $\bar{X}_n = n^{-1} \sum_{i=1}^n X_i$. Then

$$\frac{\sqrt{n}(\bar{X}_n - \mu)}{\sigma} \xrightarrow{d} Z \sim \mathcal{N}(0, 1).$$

Proof sketch via characteristic functions. The characteristic function of $Z_n := \sqrt{n}(\bar{X}_n - \mu)/\sigma$ is $\varphi_{Z_n}(t) = [\varphi_{(X_1 - \mu)/\sigma}(t/\sqrt{n})]^n$. Expanding $\varphi_{(X_1 - \mu)/\sigma}(s) = 1 - s^2/2 + o(s^2)$ as $s \rightarrow 0$, we get $\varphi_{Z_n}(t) = [1 - t^2/(2n) + o(1/n)]^n \rightarrow e^{-t^2/2}$, which is the characteristic function of $\mathcal{N}(0, 1)$. Lévy's continuity theorem then gives $Z_n \xrightarrow{d} \mathcal{N}(0, 1)$. $\square$

Theorem 2.2 (Strong Law of Large Numbers). Let $X_1, X_2, \dots$ be i.i.d. with $\mathbb{E}[|X_1|] < \infty$. Then $\bar{X}_n \xrightarrow{\text{a.s.}} \mathbb{E}[X_1]$.

Practical Application

Connection to trading. The Law of Large Numbers underpins the idea that a trading strategy with positive expected P&L per trade will be profitable in the long run. The Central Limit Theorem tells us the rate of fluctuation: after n trades of mean profit μ and standard deviation σ , total P&L is approximately $\mathcal{N}(n\mu, n\sigma^2)$. The Sharpe ratio μ/σ scaled by $\sqrt{n}$ measures how statistically significant the profit is. This is the basis for the annualised Sharpe ratio used throughout WP1.

Example 2.3 (Normal approximation for a market maker). A market maker executes $n = 500$ trades per day, each with mean profit $\mu = 0.05$ (in basis points of notional) and standard deviation $\sigma = 2.0$. By the CLT, daily P&L is approximately

$$\text{P\&L} \sim \mathcal{N}(500 \times 0.05, 500 \times 4.0) = \mathcal{N}(25, 2000).$$

Standard deviation is $\sqrt{2000} \approx 44.7$ bps. The probability of a losing day is $\mathbb{P}(\text{P\&L} < 0) = \Phi(-25/44.7) = \Phi(-0.559) \approx 28.8\%$. The daily Sharpe ratio is $25/44.7 \approx 0.56$, annualising to roughly $0.56 \times \sqrt{252} \approx 8.9$.

Exercises.

1. Prove that convergence in L^2 implies convergence in L^1 (use Cauchy–Schwarz).
2. Let $X_n \sim \text{Uniform}(0, 1/n)$. Show $X_n \xrightarrow{\text{a.s.}} 0$ and compute the rate of L^2 convergence.
3. Compute the characteristic function of $X \sim \mathcal{N}(\mu, \sigma^2)$ and verify that $\mathbb{E}[X] = \mu$, $\text{Var}(X) = \sigma^2$ via derivatives of φ_X .

Solutions

Solution to Exercise 1. Let $X_n \rightarrow X$ in L^2 , i.e. $\mathbb{E}[|X_n - X|^2] \rightarrow 0$. By the Cauchy–Schwarz inequality:

$$\mathbb{E}[|X_n - X|] = \mathbb{E}[|X_n - X| \cdot 1] \leq \sqrt{\mathbb{E}[|X_n - X|^2]} \sqrt{\mathbb{E}[1^2]} = \sqrt{\mathbb{E}[|X_n - X|^2]} \rightarrow 0.$$

Thus $X_n \rightarrow X$ in L^1 . ✓

The converse is false in general. For example, let $X_n = n \cdot \mathbf{1}_{[0, 1/n^2]}$. Then $\mathbb{E}[|X_n|] = 1/n \rightarrow 0$ (L^1 convergence) but $\mathbb{E}[|X_n|^2] = n^2/n^2 = 1 \not\rightarrow 0$. □

Solution to Exercise 2. Let $X_n \sim \text{Uniform}(0, 1/n)$, so $0 \leq X_n \leq 1/n \rightarrow 0$. Since $|X_n| \leq 1/n$ for all ω , we have $X_n \rightarrow 0$ almost surely. For the L^2 rate:

$$\mathbb{E}[X_n^2] = \int_0^{1/n} x^2 \cdot n \, dx = n \cdot \frac{1}{3n^3} = \frac{1}{3n^2},$$

so $\|X_n\|_{L^2} = (3n^2)^{-1/2} = O(n^{-1})$. The L^2 convergence rate is $O(n^{-1})$. ✓ □

Solution to Exercise 3. Let $X \sim \mathcal{N}(\mu, \sigma^2)$. The characteristic function is:

$$\varphi_X(t) = \mathbb{E}[e^{itX}] = \int_{-\infty}^{\infty} e^{itx} \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) dx.$$

Complete the square in the exponent: $itx - (x-\mu)^2/(2\sigma^2) = -(x - (\mu + it\sigma^2))^2/(2\sigma^2) + it\mu - t^2\sigma^2/2$. The Gaussian integral evaluates to 1, giving: $\varphi_X(t) = \exp(it\mu - t^2\sigma^2/2)$.

Recovering moments: $\mathbb{E}[X] = \frac{1}{i}\varphi'_X(0) = \frac{1}{i}(i\mu) \cdot 1 = \mu$. ✓

$\mathbb{E}[X^2] = \frac{1}{i^2}\varphi''_X(0)$. Since $\varphi'_X(t) = (i\mu - t\sigma^2)\varphi_X - \sigma^2\varphi_X$, at $t = 0$: $\varphi''_X(0) = -\mu^2 - \sigma^2$. Thus $\mathbb{E}[X^2] = \mu^2 + \sigma^2$, giving $\text{Var}(X) = \mathbb{E}[X^2] - (\mathbb{E}[X])^2 = \sigma^2$. ✓ □

4. **(New)** Prove the Central Limit Theorem using characteristic functions: if $X_1, X_2, \dots$ are i.i.d. with mean μ and variance σ^2 , show that $\sqrt{n}(\bar{X}_n - \mu)/\sigma \xrightarrow{d} \mathcal{N}(0, 1)$.

Solution to Exercise 4. Let $Y_i = (X_i - \mu)/\sigma$ so $\mathbb{E}[Y_i] = 0$, $\text{Var}(Y_i) = 1$. Let $S_n = n^{-1/2} \sum_{i=1}^n Y_i$. Its CF is: $\varphi_{S_n}(t) = [\varphi_Y(t/\sqrt{n})]^n$. Expand $\varphi_Y(u)$ around $u = 0$: $\varphi_Y(u) = 1 + iu\mathbb{E}[Y] - u^2\mathbb{E}[Y^2]/2 + o(u^2) = 1 - u^2/2 + o(u^2)$. Substituting $u = t/\sqrt{n}$: $\varphi_{S_n}(t) = [1 - t^2/(2n) + o(1/n)]^n \rightarrow e^{-t^2/2}$ as $n \rightarrow \infty$ (by the limit $(1 + a/n)^n \rightarrow e^a$). Since $e^{-t^2/2}$ is the CF of $\mathcal{N}(0, 1)$, Lévy's continuity theorem gives $S_n \xrightarrow{d} \mathcal{N}(0, 1)$. ✓ □

3 Stochastic Calculus

Stochastic calculus extends classical calculus to handle functions of random, continuous-time processes. The theory is essential for WP1, where asset prices, inventory, and the population distribution all evolve as stochastic processes.

3.1 Filtrations and Adapted Processes

Definition 3.1 (Filtration). A *filtration* on $(\Omega, \mathcal{F}, \mathbb{P})$ is an increasing family $(\mathcal{F}_t)_{t \geq 0}$ of sub- σ -algebras of $\mathcal{F}$: $\mathcal{F}_s \subseteq \mathcal{F}_t$ for $s \leq t$. We say $(\mathcal{F}_t)$ satisfies the *usual conditions* if it is right-continuous ($\mathcal{F}_t = \mathcal{F}_{t+} := \bigcap_{s>t} \mathcal{F}_s$) and $\mathcal{F}_0$ contains all $\mathbb{P}$ -null sets.

Intuitively, $\mathcal{F}_t$ represents the information available at time t .

Definition 3.2 (Adapted and Predictable Processes). A stochastic process $(X_t)_{t \geq 0}$ is *adapted* to $(\mathcal{F}_t)$ if X_t is $\mathcal{F}_t$ -measurable for every t . It is *predictable* if it is measurable with respect to the predictable σ -algebra generated by left-continuous adapted processes.

3.2 Martingales

Definition 3.3 (Martingale). An adapted process $(M_t)_{t \geq 0}$ with $\mathbb{E}[|M_t|] < \infty$ for all t is a *martingale* if $\mathbb{E}[M_t | \mathcal{F}_s] = M_s$ a.s. for all $0 \leq s \leq t$. It is a *submartingale* if $\mathbb{E}[M_t | \mathcal{F}_s] \geq M_s$ (the process tends to increase) and a *supermartingale* if $\mathbb{E}[M_t | \mathcal{F}_s] \leq M_s$ (tends to decrease).

Key Insight

Martingales model “fair games.” Under the risk-neutral measure $\mathbb{Q}$, discounted asset prices are martingales — this is the *fundamental theorem of asset pricing*. In WP1, the value function of each agent is a supermartingale under any admissible strategy and a martingale under the optimal one (the *martingale optimality principle*).

Important martingale inequalities.

- **Doob’s maximal inequality:** Let M be a non-negative submartingale. Then for $p > 1$, $\mathbb{E}[\sup_{0 \leq t \leq T} M_t^p] \leq (p/(p-1))^p \mathbb{E}[M_T^p]$.
- **Optional stopping theorem:** If M is a uniformly integrable martingale and τ is a stopping time, then $\mathbb{E}[M_\tau] = \mathbb{E}[M_0]$.

3.3 Brownian Motion

Definition 3.4 (Standard Brownian Motion). A stochastic process $(W_t)_{t \geq 0}$ on $(\Omega, \mathcal{F}, \mathbb{P})$ is a *standard Brownian motion* (or Wiener process) if:

- $W_0 = 0$ a.s.
- (Independent increments) For $0 \leq t_0 < t_1 < \dots < t_n$, the increments $W_{t_1} - W_{t_0}, W_{t_2} - W_{t_1}, \dots, W_{t_n} - W_{t_{n-1}}$ are independent.
- (Gaussian increments) $W_t - W_s \sim \mathcal{N}(0, t - s)$ for $0 \leq s < t$.
- (Continuous paths) $t \mapsto W_t(\omega)$ is continuous for a.e. ω .

Key properties of Brownian motion:

- $\mathbb{E}[W_t] = 0$, $\text{Var}(W_t) = t$, $\text{Cov}(W_s, W_t) = \min(s, t)$.
- Paths are a.s. nowhere differentiable.
- **Quadratic variation:** $\langle W \rangle_t = t$. More precisely, for any partition $\Pi = \{0 = t_0 < t_1 < \dots < t_n = t\}$ with $|\Pi| \rightarrow 0$,

$$\sum_{k=0}^{n-1} (W_{t_{k+1}} - W_{t_k})^2 \xrightarrow{L^2} t.$$

- **Scaling:** For $c > 0$, $(c^{-1/2}W_{ct})_{t \geq 0}$ is again a Brownian motion.
- **Reflection:** $(-W_t)_{t \geq 0}$ is a Brownian motion.




Figure 1: Three sample paths of standard Brownian motion on $[0, 1]$. Note the continuous but highly irregular (non-differentiable) trajectories.

3.4 The Itô Integral

Classical Riemann–Stieltjes integration fails for Brownian motion because paths have infinite first variation. Itô’s insight was to construct the integral by evaluating the integrand at the *left endpoint* of each partition interval.

Definition 3.5 (Itô Integral — Simple Processes). For a simple adapted process $H_t = \sum_{k=0}^{n-1} H_{t_k} \mathbf{1}_{(t_k, t_{k+1}]}(t)$, the Itô integral is

$$\int_0^T H_t dW_t := \sum_{k=0}^{n-1} H_{t_k} (W_{t_{k+1}} - W_{t_k}).$$

Theorem 3.1 (Itô Isometry). For any adapted process H with $\mathbb{E}[\int_0^T H_t^2 dt] < \infty$,

$$\mathbb{E} \left[\left(\int_0^T H_t dW_t \right)^2 \right] = \mathbb{E} \left[\int_0^T H_t^2 dt \right].$$

Moreover, the Itô integral is a (continuous) martingale with zero expectation.

The isometry allows extension of the integral from simple processes to the full class $\mathcal{H}_T^2 := \{H \text{ adapted, predictable} : \mathbb{E}[\int_0^T H_t^2 dt] < \infty\}$ by L^2 completion.

Remark 3.1. The choice of left-endpoint evaluation (rather than, say, midpoint) is crucial. Using the midpoint gives the *Stratonovich integral* $\int_0^T H_t \circ dW_t$, which obeys the chain rule of ordinary calculus but produces a non-martingale. We exclusively use the Itô convention, as it preserves the martingale property.

3.5 Itô's Formula

Theorem 3.2 (Itô's Formula — Scalar Case). *Let (X_t) satisfy the SDE $dX_t = \mu_t dt + \sigma_t dW_t$ and let $f \in C^{1,2}([0, T] \times \mathbb{R})$. Then*

$$df(t, X_t) = \frac{\partial f}{\partial t} dt + \frac{\partial f}{\partial x} dX_t + \frac{1}{2} \frac{\partial^2 f}{\partial x^2} (dX_t)^2, \quad (1)$$

where $(dX_t)^2 = \sigma_t^2 dt$ (using the rules $dW_t \cdot dW_t = dt$, $dt \cdot dW_t = 0$, $dt \cdot dt = 0$). Explicitly:

$$f(T, X_T) = f(0, X_0) + \int_0^T \left[\frac{\partial f}{\partial t} + \mu_t \frac{\partial f}{\partial x} + \frac{1}{2} \sigma_t^2 \frac{\partial^2 f}{\partial x^2} \right] dt + \int_0^T \sigma_t \frac{\partial f}{\partial x} dW_t.$$

Proof sketch. By Taylor expansion to second order and using the heuristic $(\Delta W)^2 \approx \Delta t$:

$$\begin{aligned} f(t + \Delta t, X_{t+\Delta t}) - f(t, X_t) &\approx f_t \Delta t + f_x \Delta X + \frac{1}{2} f_{xx} (\Delta X)^2 \\ &= f_t \Delta t + f_x (\mu \Delta t + \sigma \Delta W) + \frac{1}{2} f_{xx} \sigma^2 \Delta t. \end{aligned}$$

Summing over a partition and taking limits gives the result. The rigorous proof proceeds by approximation of f by smooth functions and careful estimation of remainder terms using the continuity of Brownian paths; see ?, ch. 3. $\square$

Key Insight

Itô's formula is the “chain rule” of stochastic calculus, but with an extra second-order term $\frac{1}{2} \sigma^2 f_{xx} dt$. This term is absent in classical calculus and arises because Brownian motion has non-zero quadratic variation. It is the origin of the $\frac{1}{2} \sigma^2 \Delta$ term in the Black–Scholes PDE and the diffusion operator in WP1.

Example 3.1 (Geometric Brownian motion). Let $S_t = S_0 \exp((\mu - \frac{1}{2} \sigma^2)t + \sigma W_t)$. Applying Itô's formula with $f(x) = e^x$ to $X_t = (\mu - \frac{1}{2} \sigma^2)t + \sigma W_t$:

$$\begin{aligned} dS_t &= S_t [(\mu - \frac{1}{2} \sigma^2) dt + \sigma dW_t] + \frac{1}{2} S_t \sigma^2 dt \\ &= S_t [\mu dt + \sigma dW_t]. \end{aligned}$$

So S_t satisfies $dS_t = S_t(\mu dt + \sigma dW_t)$, the standard model for stock prices. The expected value is $\mathbb{E}[S_t] = S_0 e^{\mu t}$ and the variance is $\text{Var}(S_t) = S_0^2 e^{2\mu t} (e^{\sigma^2 t} - 1)$.

3.6 Stochastic Differential Equations

Definition 3.6 (SDE). A *stochastic differential equation* (SDE) is written

$$dX_t = b(t, X_t) dt + \sigma(t, X_t) dW_t, \quad X_0 = x_0, \quad (2)$$

where $b: [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ is the *drift* and $\sigma: [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ is the *diffusion* coefficient. The rigorous meaning is the integral equation

$$X_t = x_0 + \int_0^t b(s, X_s) ds + \int_0^t \sigma(s, X_s) dW_s.$$

Theorem 3.3 (Existence and Uniqueness). *Suppose b and σ satisfy:*

- (i) **Lipschitz continuity:** $|b(t, x) - b(t, y)| + |\sigma(t, x) - \sigma(t, y)| \leq K|x - y|$,
- (ii) **Linear growth:** $|b(t, x)| + |\sigma(t, x)| \leq K(1 + |x|)$,

for some constant $K > 0$. Then the SDE (2) has a unique strong solution $(X_t)_{0 \leq t \leq T}$ that is adapted and continuous.

3.7 Girsanov's Theorem

Theorem 3.4 (Girsanov). *Let (W_t) be a Brownian motion under $\mathbb{P}$ and let (θ_t) be an adapted process satisfying the Novikov condition $\mathbb{E}[\exp(\frac{1}{2} \int_0^T \theta_t^2 dt)] < \infty$. Define the Radon–Nikodým derivative*

$$Z_T = \exp\left(-\int_0^T \theta_t dW_t - \frac{1}{2} \int_0^T \theta_t^2 dt\right),$$

and set $d\mathbb{Q} = Z_T d\mathbb{P}$. Then under $\mathbb{Q}$:

$$\tilde{W}_t := W_t + \int_0^t \theta_s ds$$

is a standard Brownian motion.

Proof sketch. By Itô's formula, $Z_t := \mathbb{E}[Z_T | \mathcal{F}_t] = \exp(-\int_0^t \theta_s dW_s - \frac{1}{2} \int_0^t \theta_s^2 ds)$ satisfies $dZ_t = -\theta_t Z_t dW_t$, so Z_t is a positive local martingale under $\mathbb{P}$. The Novikov condition ensures Z_t is a true martingale with $\mathbb{E}[Z_T] = 1$, so $\mathbb{Q}$ is a well-defined probability measure. One then verifies that $\tilde{W}$ has the correct characteristic function under $\mathbb{Q}$. $\square$

Key Insight

Girsanov's theorem is the mathematical engine behind *risk-neutral pricing*: by changing from the real-world measure $\mathbb{P}$ to the risk-neutral measure $\mathbb{Q}$, we remove the drift from the stock price. In WP1, Girsanov is used to reformulate the MFG system under different probability measures when the agents have heterogeneous beliefs.

3.8 Feynman–Kac Formula

Theorem 3.5 (Feynman–Kac). *Consider the PDE*

$$\frac{\partial u}{\partial t} + b(t, x) \frac{\partial u}{\partial x} + \frac{1}{2} \sigma^2(t, x) \frac{\partial^2 u}{\partial x^2} - r(t, x) u = 0, \quad u(T, x) = g(x).$$

If u is sufficiently smooth, then it admits the stochastic representation

$$u(t, x) = \mathbb{E} \left[g(X_T) \exp\left(-\int_t^T r(s, X_s) ds\right) \mid X_t = x \right],$$

where X satisfies $dX_s = b(s, X_s) ds + \sigma(s, X_s) dW_s$.

This connects PDEs to stochastic processes and is the foundation for Monte Carlo pricing of derivatives. In WP1, the HJB equation admits a Feynman–Kac-type representation that motivates the numerical scheme.

Example 3.2 (Black–Scholes as Feynman–Kac). The Black–Scholes PDE $\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS \frac{\partial V}{\partial S} - rV = 0$ with terminal condition $V(T, S) = (S - K)^+$ has Feynman–Kac representation $V(t, S) = e^{-r(T-t)} \mathbb{E}^\mathbb{Q}[(S_T - K)^+ | S_t = S]$, which evaluates to the Black–Scholes formula:

$$V = S\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2), \quad d_{1,2} = \frac{\ln(S/K) + (r \pm \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}.$$

Exercises.

1. Apply Itô's formula to compute $d(W_t^3)$. Identify the drift and martingale parts.
2. Let $X_t = e^{-\lambda t}W_t$ for $\lambda > 0$. Find the SDE satisfied by X .
3. Verify the Itô isometry directly for $H_t = t$: show $\mathbb{E}[(\int_0^T t dW_t)^2] = T^3/3$.
4. Using Girsanov's theorem, find the measure $\mathbb{Q}$ under which $S_t = S_0 \exp(\sigma W_t - \frac{1}{2}\sigma^2 t)$ is a $\mathbb{Q}$ -martingale when the drift under $\mathbb{P}$ is $\mu \neq 0$.

Solutions

Solution to Exercise 1. Using Itô's formula with $f(x) = x^3$ and $X_t = W_t$:

$$d(W_t^3) = 3W_t^2 dW_t + \frac{1}{2} \cdot 6W_t dt = 3W_t^2 dW_t + 3W_t dt.$$

The **drift part** (finite variation) is $3W_t dt$ and the **martingale part** is $3W_t^2 dW_t$. In integral form: $W_T^3 = 3 \int_0^T W_t^2 dW_t + 3 \int_0^T W_t dt$. ✓ □

Solution to Exercise 2. Let $X_t = e^{-\lambda t}W_t$. By the product rule version of Itô's formula (with $f(t, w) = e^{-\lambda t}w$):

$$dX_t = -\lambda e^{-\lambda t}W_t dt + e^{-\lambda t} dW_t = -\lambda X_t dt + e^{-\lambda t} dW_t.$$

This is an *Ornstein–Uhlenbeck*-type SDE with time-dependent diffusion coefficient $\sigma(t) = e^{-\lambda t}$. Note that X_t is a Gaussian process with mean zero and variance $\text{Var}(X_t) = \int_0^t e^{-2\lambda s} ds = (1 - e^{-2\lambda t})/(2\lambda)$. ✓ □

Solution to Exercise 3. The Itô isometry states $\mathbb{E}[(\int_0^T H_t dW_t)^2] = \int_0^T \mathbb{E}[H_t^2] dt$. For $H_t = t$ (deterministic):

$$\mathbb{E}\left[\left(\int_0^T t dW_t\right)^2\right] = \int_0^T t^2 dt = \frac{T^3}{3}. \quad \checkmark$$

Note: $\int_0^T t dW_t$ is a Gaussian random variable with mean 0 and variance $T^3/3$, by the linearity and Gaussian property of the Itô integral. □

Solution to Exercise 4. Under $\mathbb{P}$, $S_t = S_0 \exp(\mu t + \sigma W_t - \frac{1}{2}\sigma^2 t)$. We want S_t to be a $\mathbb{Q}$ -martingale (after discounting). Define the Girsanov change: $\tilde{W}_t = W_t + \frac{\mu-r}{\sigma}t$ is a $\mathbb{Q}$ -Brownian motion, where $d\mathbb{Q}/d\mathbb{P} = \mathcal{E}(-\frac{\mu-r}{\sigma}W)_T$ (the Doléans-Dade exponential). Then $\sigma W_t = \sigma \tilde{W}_t - (\mu - r)t$, giving:

$$S_t = S_0 \exp(rt + \sigma \tilde{W}_t - \frac{1}{2}\sigma^2 t),$$

and $e^{-rt}S_t = S_0 \exp(\sigma \tilde{W}_t - \frac{1}{2}\sigma^2 t)$ is a $\mathbb{Q}$ -martingale (it is the stochastic exponential of $\sigma \tilde{W}$). ✓

The market price of risk is $\theta = (\mu - r)/\sigma$, and $\mathbb{Q}$ is the unique equivalent martingale measure (the market is complete). □

5. **(New)** Compute $\mathbb{E}[\int_0^T W_t dt]$ and $\text{Var}(\int_0^T W_t dt)$. *Hint:* use Fubini's theorem for the mean and the covariance $\text{Cov}(W_s, W_t) = \min(s, t)$ for the variance.

Solution to Exercise 5. $\mathbb{E}[\int_0^T W_t dt] = \int_0^T \mathbb{E}[W_t] dt = 0$. For the variance:

$$\begin{aligned} \text{Var}\left(\int_0^T W_t dt\right) &= \mathbb{E}\left[\left(\int_0^T W_t dt\right)^2\right] = \int_0^T \int_0^T \mathbb{E}[W_s W_t] ds dt \\ &= \int_0^T \int_0^T \min(s, t) ds dt = 2 \int_0^T \int_0^t s ds dt = 2 \int_0^T \frac{t^2}{2} dt = \frac{T^3}{3}. \end{aligned}$$

Thus $\int_0^T W_t dt \sim \mathcal{N}(0, T^3/3)$ (it is Gaussian since it is a limit of Riemann sums of jointly Gaussian random variables). ✓ □

4 Partial Differential Equations

Mean-field games produce a coupled system of PDEs: the HJB equation (backward in time) and the Fokker–Planck equation (forward in time). This section provides the PDE background necessary to understand their structure and solution methods.

4.1 Classification of Second-Order Linear PDEs

Consider the general second-order linear PDE in two variables:

$$Au_{xx} + 2Bu_{xy} + Cu_{yy} + Du_x + Eu_y + Fu = G,$$

where A, B, C, D, E, F, G may depend on (x, y) . The discriminant $\Delta = B^2 - AC$ classifies the equation:

- (i) $\Delta < 0$: **Elliptic** (e.g., Laplace equation $\Delta u = 0$).
- (ii) $\Delta = 0$: **Parabolic** (e.g., heat equation $u_t = u_{xx}$).
- (iii) $\Delta > 0$: **Hyperbolic** (e.g., wave equation $u_{tt} = u_{xx}$).

Key Insight

The Hamilton–Jacobi–Bellman and Fokker–Planck equations from WP1 are both *parabolic*. Parabolic PDEs model diffusive phenomena and admit smooth solutions that propagate information forward (or backward) in time. This is why finite-difference schemes with appropriate time-stepping are natural for their numerical solution.

4.2 The Heat Equation

The heat equation $u_t = \kappa u_{xx}$ on $(0, \infty) \times \mathbb{R}$ with initial condition $u(0, x) = u_0(x)$ has the explicit solution

$$u(t, x) = \frac{1}{\sqrt{4\pi\kappa t}} \int_{-\infty}^{\infty} \exp\left(-\frac{(x-y)^2}{4\kappa t}\right) u_0(y) dy = (\Gamma_t * u_0)(x), \quad (3)$$

where $\Gamma_t(x) = (4\pi\kappa t)^{-1/2} \exp(-x^2/(4\kappa t))$ is the *fundamental solution* (or Green’s function). This is simply the pdf of a $\mathcal{N}(0, 2\kappa t)$ random variable.

Theorem 4.1 (Maximum Principle for Parabolic PDEs). *Let u solve $u_t = \kappa u_{xx}$ on the bounded domain $[0, T] \times [a, b]$. Then u attains its maximum (and minimum) on the parabolic boundary $\{0\} \times [a, b] \cup [0, T] \times \{a, b\}$.*

The maximum principle implies uniqueness: if two solutions agree on the parabolic boundary, they agree everywhere. It also provides stability: small perturbations in boundary/initial data produce small perturbations in the solution.

4.3 The Fokker–Planck Equation

Definition 4.1 (Fokker–Planck / Kolmogorov Forward Equation). Given the SDE $dX_t = b(t, X_t) dt + \sigma(t, X_t) dW_t$ with initial distribution $m_0 = \text{law}(X_0)$, the *Fokker–Planck equation* for the density $m(t, x)$ of $\text{law}(X_t)$ is

$$\frac{\partial m}{\partial t} = -\frac{\partial}{\partial x} [b(t, x) m] + \frac{1}{2} \frac{\partial^2}{\partial x^2} [\sigma^2(t, x) m], \quad m(0, x) = m_0(x). \quad (4)$$

Derivation. Let $\varphi \in C_c^\infty(\mathbb{R})$ be a test function. By Itô's formula:

$$d\varphi(X_t) = \varphi'(X_t)[b dt + \sigma dW_t] + \frac{1}{2}\sigma^2\varphi''(X_t) dt.$$

Taking expectations (the dW term vanishes):

$$\frac{d}{dt}\mathbb{E}[\varphi(X_t)] = \mathbb{E}[b(t, X_t)\varphi'(X_t) + \frac{1}{2}\sigma^2(t, X_t)\varphi''(X_t)] = \int_{\mathbb{R}} [b\varphi' + \frac{1}{2}\sigma^2\varphi''] m dx.$$

Integrating by parts (assuming boundary terms vanish):

$$\begin{aligned} \int_{\mathbb{R}} b\varphi' m dx &= - \int_{\mathbb{R}} (bm)' \varphi dx, \\ \int_{\mathbb{R}} \frac{1}{2}\sigma^2\varphi'' m dx &= \int_{\mathbb{R}} \frac{1}{2}(\sigma^2 m)'' \varphi dx. \end{aligned}$$

Since φ is arbitrary, we obtain (4). The Fokker–Planck equation governs the *forward evolution* of the population distribution in WP1.

Theorem 4.2 (Conservation of Mass). *If $m(t, \cdot)$ solves the Fokker–Planck equation (4) and decays sufficiently fast as $|x| \rightarrow \infty$, then $\int_{\mathbb{R}} m(t, x) dx = \int_{\mathbb{R}} m_0(x) dx = 1$ for all $t \geq 0$.*

Proof. Set $\varphi \equiv 1$ in the weak formulation. Then $\varphi' = \varphi'' = 0$ and $\frac{d}{dt} \int_{\mathbb{R}} m dx = 0$. $\square$

4.4 Viscosity Solutions

Classical (smooth) solutions of the HJB equation may not exist when the Hamiltonian is non-smooth (e.g., due to control constraints). Viscosity solutions provide the right notion of weak solution.

Definition 4.2 (Viscosity Solution). Let $F(t, x, u, Du, D^2u) = 0$ be a second-order PDE. An upper semicontinuous function u is a *viscosity subsolution* if: whenever a smooth function ϕ touches u from above at a point (t_0, x_0) (i.e., $\phi \geq u$ locally with $\phi(t_0, x_0) = u(t_0, x_0)$), then $F(t_0, x_0, u(t_0, x_0), D\phi(t_0, x_0), D^2\phi(t_0, x_0)) \leq 0$.

A lower semicontinuous function is a *viscosity supersolution* if the reverse inequality holds with test functions touching from below.

A continuous function that is both a viscosity sub- and supersolution is a *viscosity solution*.

Practical Application

Why viscosity solutions matter in trading. In WP1, the HJB equation has a Hamiltonian of the form $H(x, p) = \sup_{\alpha \in \mathcal{A}} \{f(x, \alpha) + p \cdot g(x, \alpha)\}$. When the supremum is attained on the boundary of the control set $\mathcal{A}$, the optimal control switches discontinuously, and the value function may have kinks. Viscosity solutions handle these kinks correctly. In practice, finite difference schemes naturally converge to the viscosity solution provided they satisfy monotonicity, stability, and consistency (the Barles–Souganidis theorem).

Comparison principles. The key property of viscosity solutions is the *comparison principle*: if u is a viscosity subsolution and v is a viscosity supersolution with $u \leq v$ on the boundary, then $u \leq v$ everywhere. This implies:

- **Uniqueness:** there is at most one viscosity solution.
- **Stability:** limits of viscosity solutions are viscosity solutions.

- **Convergence of numerical schemes:** monotone, stable, consistent schemes converge.

Exercises.

1. Verify that $u(t, x) = (4\pi\kappa t)^{-1/2} \exp(-x^2/(4\kappa t))$ solves $u_t = \kappa u_{xx}$ for $t > 0$.
2. Derive the Fokker–Planck equation for the Ornstein–Uhlenbeck process $dX_t = -\theta X_t dt + \sigma dW_t$ and find its stationary solution.
3. Show that $u(x) = |x|$ is a viscosity solution of $-u'' = 0$ on $\mathbb{R} \setminus \{0\}$ but not a classical solution (check the condition at $x = 0$ using test functions).

Solutions

Solution to Exercise 1. Let $u(t, x) = (4\pi\kappa t)^{-1/2} \exp(-x^2/(4\kappa t))$. Compute:

$$u_t = u \left(-\frac{1}{2t} + \frac{x^2}{4\kappa t^2} \right), \quad u_x = u \left(-\frac{x}{2\kappa t} \right), \quad u_{xx} = u \left(\frac{x^2}{4\kappa^2 t^2} - \frac{1}{2\kappa t} \right).$$

Then $\kappa u_{xx} = u(-1/(2t) + x^2/(4\kappa t^2)) = u_t$. ✓

This is the *heat kernel* (fundamental solution). It represents the temperature profile at time t due to a unit point source at the origin. Note $u \rightarrow \delta(x)$ as $t \rightarrow 0^+$ and $\int_{-\infty}^{\infty} u dx = 1$ for all $t > 0$ (conservation of total heat). □

Solution to Exercise 2. The OU process $dX_t = -\theta X_t dt + \sigma dW_t$ has drift $b(x) = -\theta x$ and diffusion $a(x) = \sigma^2/2$. The Fokker–Planck equation is:

$$\partial_t p = \theta \partial_x(x p) + \frac{\sigma^2}{2} \partial_{xx} p = \theta p + \theta x \partial_x p + \frac{\sigma^2}{2} \partial_{xx} p.$$

For the stationary solution, set $\partial_t p = 0$. The stationary FP equation is $\theta(xp)' + (\sigma^2/2)p'' = 0$, which integrates (with zero-flux boundary conditions) to:

$$\theta x p_{\infty}(x) + \frac{\sigma^2}{2} p'_{\infty}(x) = 0 \implies p'_{\infty}/p_{\infty} = -2\theta x/\sigma^2.$$

Integrating: $p_{\infty}(x) = C \exp(-\theta x^2/\sigma^2)$. Normalising: $p_{\infty}(x) = \mathcal{N}(0, \sigma^2/(2\theta))$. ✓

This is the Gaussian invariant measure of the OU process. Mean-reversion strength θ determines how tightly the distribution is concentrated. □

Solution to Exercise 3. $u(x) = |x|$ satisfies $u'' = 0$ classically for $x \neq 0$ (u is linear on each half-line). At $x = 0$, u is not twice differentiable (u' has a jump from -1 to $+1$). However, u is a *viscosity solution* of $-u'' = 0$.

Viscosity subsolution: If φ is smooth and $u - \varphi$ has a local maximum at $x_0 = 0$, then $\varphi(0) = 0$ and $\varphi(x) \geq |x|$ near 0. This forces $\varphi''(0) \leq 0$ (since φ is concave near 0 to stay above $|x|$), so $-\varphi''(0) \geq 0$. ✓

Viscosity supersolution: If φ is smooth and $u - \varphi$ has a local minimum at $x_0 = 0$, then $\varphi(0) = 0$ and $\varphi(x) \leq |x|$ near 0. This forces $\varphi''(0) \geq 0$, so $-\varphi''(0) \leq 0$. ✓

Both conditions are satisfied, so u is a viscosity solution even though it is not a classical solution. This illustrates why viscosity solutions are the correct notion for the HJB equations in MFG, where value functions may have kinks. □

4. (New) Show that the fundamental solution of the Fokker–Planck equation for GBM $dS = \mu S dt + \sigma S dW$ is log-normal. Specifically, show $\log S_T \sim \mathcal{N}(\log S_0 + (\mu - \sigma^2/2)T, \sigma^2 T)$.

Solution to Exercise 4. Let $Y_t = \log S_t$. By Itô's formula: $dY_t = (\mu - \sigma^2/2) dt + \sigma dW_t$. This is an SDE with constant coefficients, so: $Y_T = Y_0 + (\mu - \sigma^2/2)T + \sigma W_T$. Since $W_T \sim \mathcal{N}(0, T)$, we get $Y_T \sim \mathcal{N}(\log S_0 + (\mu - \sigma^2/2)T, \sigma^2 T)$. Thus S_T is log-normally distributed with $\mathbb{E}[S_T] = S_0 e^{\mu T}$ and $\text{Var}(S_T) = S_0^2 e^{2\mu T} (e^{\sigma^2 T} - 1)$. $\square$

5 Optimal Control Theory

Optimal control is the mathematical framework for choosing a strategy (control) to minimise (or maximise) an objective functional. In WP1, each agent solves a stochastic optimal control problem; the MFG equilibrium couples all these problems together.

5.1 Deterministic Optimal Control

Consider the deterministic dynamical system

$$\dot{x}(t) = f(t, x(t), \alpha(t)), \quad x(0) = x_0,$$

where $\alpha: [0, T] \rightarrow A$ is the control. The objective is to minimise

$$J(\alpha) = \int_0^T L(t, x(t), \alpha(t)) dt + \Phi(x(T)),$$

where L is the running cost and Φ is the terminal cost.

Theorem 5.1 (Pontryagin Maximum Principle). *Suppose α^* is an optimal control with corresponding state trajectory x^* . Then there exists an adjoint (or costate) variable $p: [0, T] \rightarrow \mathbb{R}$ satisfying the adjoint equation*

$$\dot{p}(t) = -\frac{\partial H}{\partial x}(t, x^*(t), p(t), \alpha^*(t)), \quad p(T) = \Phi'(x^*(T)),$$

where the Hamiltonian is $H(t, x, p, \alpha) = L(t, x, \alpha) + p \cdot f(t, x, \alpha)$, and the optimal control satisfies:

$$\alpha^*(t) \in \operatorname{argmin}_{\alpha \in A} H(t, x^*(t), p(t), \alpha).$$

5.2 Dynamic Programming and the HJB Equation

The *dynamic programming* approach defines the *value function*

$$V(t, x) = \inf_{\alpha \in A} \left\{ \int_t^T L(s, x_s, \alpha_s) ds + \Phi(x_T) \mid x_t = x \right\}. \quad (5)$$

Theorem 5.2 (Hamilton–Jacobi–Bellman Equation). *Under suitable regularity, the value function V satisfies the **HJB equation**:*

$$-\frac{\partial V}{\partial t}(t, x) + H\left(t, x, \frac{\partial V}{\partial x}(t, x)\right) = 0, \quad V(T, x) = \Phi(x), \quad (6)$$

where the Hamiltonian is $H(t, x, p) = \inf_{\alpha \in A} \{L(t, x, \alpha) + p \cdot f(t, x, \alpha)\}$. The optimal control is recovered by the feedback law:

$$\alpha^*(t, x) = \operatorname{argmin}_{\alpha \in A} \{L(t, x, \alpha) + \nabla_x V(t, x) \cdot f(t, x, \alpha)\}.$$

Key Insight

The HJB equation is a *backward* PDE (solved from $t = T$ to $t = 0$), while the Fokker–Planck is a *forward* PDE. In the MFG system of WP1, these are coupled: the HJB depends on the population distribution m (through the cost), and the Fokker–Planck depends on the

optimal control (through the drift). This coupling makes the system challenging and is what necessitates the fixed-point iteration scheme developed in the whitepaper.

5.3 Stochastic Optimal Control

In the stochastic setting, the state evolves as

$$dX_t = b(t, X_t, \alpha_t) dt + \sigma(t, X_t, \alpha_t) dW_t, \quad X_0 = x_0, \quad (7)$$

and the objective is to minimise

$$J(\alpha) = \mathbb{E} \left[\int_0^T L(t, X_t, \alpha_t) dt + \Phi(X_T) \right].$$

The value function $V(t, x) = \inf_{\alpha \in \mathcal{A}_{[t, T]}} J_{t, x}(\alpha)$ satisfies the *stochastic HJB equation*:

$$-\frac{\partial V}{\partial t} + H\left(t, x, \frac{\partial V}{\partial x}, \frac{\partial^2 V}{\partial x^2}\right) = 0, \quad V(T, x) = \Phi(x), \quad (8)$$

where

$$H(t, x, p, M) = \inf_{\alpha \in A} \left\{ L(t, x, \alpha) + b(t, x, \alpha) \cdot p + \frac{1}{2} \sigma^2(t, x, \alpha) M \right\}.$$

Theorem 5.3 (Verification Theorem). *Suppose $v \in C^{1,2}([0, T] \times \mathbb{R}) \cap C([0, T] \times \mathbb{R})$ satisfies:*

- (i) $-v_t + H(t, x, v_x, v_{xx}) = 0$ on $[0, T] \times \mathbb{R}$,
- (ii) $v(T, x) = \Phi(x)$ for all x ,
- (iii) For each admissible α , $\mathbb{E}[\sup_{t \leq T} |v(t, X_t^\alpha)|] < \infty$.

Then $v(t, x) = V(t, x)$ (the value function) and the feedback control $\alpha^(t, x) = \operatorname{argmin}_\alpha \{L + b \cdot v_x + \frac{1}{2} \sigma^2 v_{xx}\}$ is optimal.*

Proof sketch. For any admissible α , apply Itô's formula to $v(t, X_t)$:

$$v(T, X_T) = v(0, x_0) + \int_0^T [v_t + b v_x + \frac{1}{2} \sigma^2 v_{xx}] dt + \text{martingale}.$$

Since v satisfies the HJB equation, $v_t + b v_x + \frac{1}{2} \sigma^2 v_{xx} \geq -L$ for all α (with equality at α^*). Taking expectations:

$$\mathbb{E}[\Phi(X_T)] \geq v(0, x_0) - \mathbb{E} \left[\int_0^T L dt \right].$$

Rearranging gives $J(\alpha) \geq v(0, x_0)$ for all α , with equality at α^* . □

5.4 Linear-Quadratic Regulator (Worked Example)

Example 5.1 (Stochastic LQR). Consider the state $dX_t = \alpha_t dt + \sigma dW_t$, $X_0 = x$, with cost

$$J(\alpha) = \mathbb{E} \left[\int_0^T \left(\frac{1}{2} q X_t^2 + \frac{1}{2} r \alpha_t^2 \right) dt + \frac{1}{2} \gamma X_T^2 \right], \quad q, r, \gamma > 0.$$

Step 1. Guess $V(t, x) = \frac{1}{2} P(t) x^2 + C(t)$ with $P(T) = \gamma$, $C(T) = 0$.

Step 2. HJB equation: $-V_t + \inf_{\alpha} \{qx^2/2 + r\alpha^2/2 + \alpha V_x + \sigma^2 V_{xx}/2\} = 0$.

Since $V_x = Px$, $V_{xx} = P$, and $V_t = \frac{1}{2}\dot{P}x^2 + \dot{C}$, the minimiser is $\alpha^* = -Px/r$ (first-order condition: $r\alpha + Px = 0$).

Step 3. Substituting back:

$$-\frac{1}{2}\dot{P}x^2 - \dot{C} + \frac{1}{2}qx^2 + \frac{1}{2}r \cdot \frac{P^2x^2}{r^2} - \frac{P^2x^2}{r} + \frac{1}{2}\sigma^2P = 0.$$

Collecting the x^2 terms: $-\dot{P} + q - P^2/r = 0$, which is the **Riccati ODE**:

$$\dot{P}(t) = q - \frac{P(t)^2}{r}, \quad P(T) = \gamma. \quad (9)$$

The constant term gives $\dot{C}(t) = \frac{1}{2}\sigma^2P(t)$, $C(T) = 0$.

Step 4. The Riccati ODE (9) is solved backward from T and has the closed-form solution (setting $\beta = \sqrt{q/r}$):

$$P(t) = \sqrt{qr} \frac{\gamma \cosh(\beta(T-t)) + \sqrt{qr} \sinh(\beta(T-t))}{\sqrt{qr} \cosh(\beta(T-t)) + \gamma \sinh(\beta(T-t))}.$$

The optimal control is $\alpha^*(t, x) = -P(t)x/r$, which is a linear feedback rule that pushes the state toward zero with a time-varying gain.

Key Insight

The LQR is the simplest “building block” of the MFG framework in WP1. Each agent’s individual optimisation (given a fixed population distribution) resembles an LQR with an additional mean-field coupling term. The Riccati structure persists in the linearised version of the MFG system. This concept is developed further in the main white paper [WP1], Sections 3–4.

Example 5.2 (Finite difference scheme for HJB). For the stochastic HJB $-V_t + H(x, V_x, V_{xx}) = 0$ on $[0, T] \times [x_{\min}, x_{\max}]$, discretise: $x_j = x_{\min} + j\Delta x$, $t_k = k\Delta t$. Using an implicit (backward Euler) scheme:

$$\frac{V_j^{k+1} - V_j^k}{\Delta t} = H\left(x_j, \frac{V_{j+1}^k - V_{j-1}^k}{2\Delta x}, \frac{V_{j+1}^k - 2V_j^k + V_{j-1}^k}{(\Delta x)^2}\right).$$

This is solved backward in time from $V_j^N = \Phi(x_j)$. For monotonicity (required for convergence to the viscosity solution), one uses upwind differences for the first-order term: if the drift is positive at x_j , use $(V_j^k - V_{j-1}^k)/\Delta x$ instead of central differences.

Exercises.

1. For the LQR example, verify that $P(t) > 0$ for all $t \in [0, T]$ when $q, r, \gamma > 0$.
2. The Avellaneda–Stoikov model in WP1 has the inventory penalty qX_t^2 . Solve the HJB for the simplified problem: $dS = \sigma dW$, inventory X , and the agent maximises $\mathbb{E}[\int_0^T (\delta_t - \frac{1}{2}\gamma X_t^2) dt]$ where δ_t is the half-spread. Show the optimal half-spread is constant and compute it.
3. Implement the finite difference scheme for the heat equation in Python and verify second-order convergence in Δx .

Solutions

Solution to Exercise 1. The Riccati equation is $-\dot{P} = 2qP - P^2/r + \gamma$, $P(T) = 0$. At terminal time, $P(T) = 0 \geq 0$. For P to become negative, it must cross zero at some time $t^* < T$, at which $P(t^*) = 0$. But then $-\dot{P}(t^*) = 0 - 0 + \gamma = \gamma > 0$, meaning $\dot{P}(t^*) < 0$, so P is *decreasing* at t^* (moving backward, P increases from 0). Thus P cannot cross zero from above going backward—it stays ≥ 0 .

More rigorously: $P(t) = r(\sqrt{q/r} + \gamma/(2r\sqrt{q/r})) \tanh(\dots)$, which is strictly positive for $t < T$ when $q, r, \gamma > 0$. The positivity ensures the optimal control $\alpha^* = -Px/r$ has the correct sign (mean-reverting). ✓ □

Solution to Exercise 2. The HJB equation for the simplified Avellaneda–Stoikov problem is:

$$\partial_t V + \sup_{\delta} \left[\delta \lambda e^{-\kappa \delta} - \frac{\gamma}{2} x^2 \right] + \frac{\sigma^2}{2} \partial_{ss} V = 0,$$

with $V(T, s, x) = 0$. The supremum over δ is computed by $\partial_{\delta}[\delta \lambda e^{-\kappa \delta}] = \lambda e^{-\kappa \delta}(1 - \kappa \delta) = 0$, giving $\delta^* = 1/\kappa$. This is *constant* (independent of state).

The maximised value is $\delta^* \lambda e^{-\kappa \delta^*} = \lambda/(\kappa e)$. The HJB simplifies to: $\partial_t V + \lambda/(\kappa e) - \gamma x^2/2 + (\sigma^2/2) \partial_{ss} V = 0$. Since the equation is linear in V with the separation $V(t, s, x) = (T-t)\lambda/(\kappa e) - \gamma x^2(T-t)/2 + g(t, s)$ where g absorbs the diffusion term. The optimal half-spread is $\delta^* = 1/\kappa$, independent of inventory. ✓ □

Solution to Exercise 3. The heat equation $u_t = \kappa u_{xx}$ on $[0, L]$ with the centred difference scheme: $u_j^{n+1} = u_j^n + r(u_{j+1}^n - 2u_j^n + u_{j-1}^n)$ where $r = \kappa \Delta t / (\Delta x)^2$. The local truncation error from Taylor expansion is: $\tau = O(\Delta t) + O((\Delta x)^2)$. With r fixed (so $\Delta t \propto (\Delta x)^2$), the global error satisfies $\|u - u^h\|_{\infty} = O((\Delta x)^2)$. To verify numerically: solve with $\Delta x = 1/N$ for $N = 10, 20, 40, 80$ and compute $e_N = \max_j |u_j - u_{\text{exact},j}|$. The ratios e_N/e_{2N} should approach 4, confirming second-order convergence. □

4. **(New)** Consider the 1D HJB equation $-\partial_t V + H(x, \nabla V) = \epsilon \Delta V$ with $H(x, p) = |p|^2/2 + f(x)$. Show that the viscosity solution of the first-order equation ($\epsilon = 0$) can be obtained as $\lim_{\epsilon \rightarrow 0}$ of the smooth solutions. Why is this relevant for MFG?

Solution to Exercise 4. For $\epsilon > 0$, the equation is a uniformly parabolic PDE, which by classical theory has a unique smooth solution V^{ϵ} . By the stability property of viscosity solutions, any locally uniform limit of V^{ϵ} as $\epsilon \rightarrow 0$ is a viscosity solution of the first-order HJB $-\partial_t V + |\nabla V|^2/2 + f(x) = 0$.

The *vanishing viscosity method* is relevant for MFG because:

- The HJB in MFG may have non-smooth solutions (kinks from optimal switching).
- Adding a small $\epsilon \Delta V$ regularises the problem, enables numerical solution.
- Taking $\epsilon \rightarrow 0$ recovers the correct (viscosity) solution.
- In practice, the diffusion $\sigma > 0$ in the state dynamics provides natural viscosity, so the MFG system is typically well-posed.

□

6 Game Theory for Continuous Games

This section bridges classical game theory and the mean-field framework of WP1. We start with the finite-player setting and take the limit $N \rightarrow \infty$ to arrive at mean-field games.

6.1 Nash Equilibrium in N -Player Games

Definition 6.1 (Nash Equilibrium — N -Player Stochastic Game). Consider N agents, where agent i controls α^i and has state X^i evolving as

$$dX_t^i = b(t, X_t^i, \alpha_t^i, \mu_t^N) dt + \sigma dW_t^i, \quad i = 1, \dots, N,$$

where $\mu_t^N = \frac{1}{N} \sum_{j=1}^N \delta_{X_t^j}$ is the empirical measure of all agents' states. Agent i minimises

$$J^i(\alpha^1, \dots, \alpha^N) = \mathbb{E} \left[\int_0^T L(t, X_t^i, \alpha_t^i, \mu_t^N) dt + \Phi(X_T^i, \mu_T^N) \right].$$

A strategy profile $(\alpha^{*,1}, \dots, \alpha^{*,N})$ is a *Nash equilibrium* if for each i :

$$J^i(\alpha^{*,1}, \dots, \alpha^{*,N}) \leq J^i(\alpha^{*,1}, \dots, \alpha^{*,i-1}, \alpha^i, \alpha^{*,i+1}, \dots, \alpha^{*,N}) \quad \forall \alpha^i \in \mathcal{A}.$$

The computational challenge is evident: each agent's optimal strategy depends on all other agents' strategies through μ^N . For N large, solving the coupled system of N HJB equations is intractable.

6.2 The Mean-Field Limit: $N \rightarrow \infty$

The key insight of Lasry–Lions and Huang–Malhamé–Caines is that as $N \rightarrow \infty$, the empirical measure μ_t^N converges to a *deterministic* measure m_t , and the Nash equilibrium approaches a *mean-field game (MFG) equilibrium*.

In the MFG limit, all agents are identical (exchangeable) and infinitesimally small. A representative agent faces the *deterministic* flow of measures $(m_t)_{t \in [0, T]}$ and solves:

$$\inf_{\alpha \in \mathcal{A}} \mathbb{E} \left[\int_0^T L(t, X_t, \alpha_t, m_t) dt + \Phi(X_T, m_T) \right] \quad \text{s.t.} \quad dX_t = b(t, X_t, \alpha_t) dt + \sigma dW_t. \quad (10)$$

The **MFG system** is the crucial coupling: the population measure m_t is itself determined by the optimal strategy, via the Fokker–Planck equation.

Definition 6.2 (Mean-Field Game Equilibrium). A pair (V, m) consisting of a value function V and a flow of probability densities (m_t) is an *MFG equilibrium* if:

- (i) V solves the HJB equation with costs depending on m :

$$-V_t + H(t, x, V_x, V_{xx}, m) = 0, \quad V(T, x) = \Phi(x, m_T). \quad (11)$$

- (ii) m solves the Fokker–Planck equation with drift given by the optimal control derived from V :

$$m_t = -(b^* m)_x + \frac{1}{2}(\sigma^2 m)_{xx}, \quad m(0, x) = m_0(x), \quad (12)$$

where $b^*(t, x) = b(t, x, \alpha^*(t, x))$ and $\alpha^*(t, x) = \operatorname{argmin}_{\alpha} H$.

6.3 Propagation of Chaos

The rigorous connection between the N -player game and the MFG limit is provided by *propagation of chaos*.

Theorem 6.1 (Propagation of Chaos). *Consider N agents with i.i.d. initial conditions $X_0^i \sim m_0$ and symmetric interactions through the empirical measure. Under Lipschitz regularity of coefficients, the following hold as $N \rightarrow \infty$:*

- (i) *The empirical measure μ_t^N converges (in the Wasserstein-2 metric) to m_t , the solution of the MFG system, at rate $O(N^{-1/2})$ (in dimension $d \leq 3$).*
- (ii) *Any fixed subset of k agents $(X^1, \dots, X^k)$ becomes asymptotically independent, each with marginal distribution m_t .*
- (iii) *The N -player Nash equilibrium ε_N -approximates the MFG equilibrium, with $\varepsilon_N = O(N^{-1/2})$.*

The term “propagation of chaos” means that even though agent interactions create short-time correlations, these correlations vanish as $N \rightarrow \infty$: the agents become asymptotically i.i.d., each governed by the mean-field.

Key Insight

Propagation of chaos justifies the mean-field approximation in WP1. It tells us that the MFG solution is a good approximation of the N -player Nash equilibrium for large N . In financial markets, N can be in the thousands or millions, so the approximation is excellent. The $O(N^{-1/2})$ rate means that for $N = 10,000$ agents, the approximation error is about 1%.

6.4 Fixed-Point Characterisation and Existence

The MFG system can be viewed as a fixed-point problem. Define the map $\Gamma: m \mapsto m'$:

Step 1: Given m , solve the HJB equation (11) backward for V .

Step 2: From V , extract the optimal control $\alpha^*(t, x)$.

Step 3: Solve the Fokker–Planck equation (12) forward with drift b^* to get m' .

An MFG equilibrium is a fixed point: $\Gamma(m) = m$.

Theorem 6.2 (Existence of MFG Equilibrium). *Under the Lasry–Lions monotonicity condition— $\int_{\mathbb{R}} [L(x, \alpha, m_1) - L(x, \alpha, m_2)](m_1 - m_2) dx \geq 0$ for all densities m_1, m_2 —the MFG system admits a unique solution (V, m) . Without monotonicity, existence still holds (by Schauder’s fixed-point theorem) but uniqueness may fail.*

The Lasry–Lions condition is economically natural: it says that agents dislike crowding (the cost of being in a state increases when more agents are in that state).

Exercises.

1. For $N = 2$ agents with quadratic costs $J^i = \mathbb{E}[\int_0^T (\alpha_t^i)^2/2 + q(X_t^i - \bar{X}_t)^2/2 dt]$ where $\bar{X} = (X^1 + X^2)/2$, find the Nash equilibrium explicitly.
2. Write the MFG system for the Avellaneda–Stoikov model: drift $b = -\alpha$ (inventory depletion), running cost $L = -\delta \cdot \lambda e^{-\kappa\delta} + \gamma q^2$. Identify the coupling.
3. Prove that the Wasserstein-2 distance $W_2(\mu^N, m) \leq C/\sqrt{N}$ implies that the price impact of any single agent is $O(1/\sqrt{N})$.

Solutions

Solution to Exercise 1. With $N = 2$, $\bar{X} = (X^1 + X^2)/2$. Agent i 's cost is: $J^i = \mathbb{E}[\int_0^T \frac{(\alpha^i)^2}{2} + \frac{q}{2}(X^i - \bar{X})^2 dt]$. State dynamics: $dX^i = \alpha^i dt + \sigma dW^i$. Note $X^i - \bar{X} = (X^i - X^j)/2$ where $j \neq i$. Define $Z = X^1 - X^2$; then $dZ = (\alpha^1 - \alpha^2) dt + \sigma d(W^1 - W^2)$. Each agent minimises $(\alpha^i)^2/2 + qZ^2/8$. By symmetry, the Nash equilibrium has the linear form $\alpha^i = -\kappa(X^i - X^j)/2$ where κ solves the Riccati: $-\dot{P} = q/4 - P^2$, $P(T) = 0$. The optimal $\kappa = P(t)$ and $\alpha^{i,*} = -P(t)(X^i - X^j)/2$. Both agents mean-revert toward each other at a rate determined by $P(t)$. Explicitly: $P(t) = \sqrt{q/4} \tanh(\sqrt{q/4}(T - t))$ and $\alpha^{i,*} = -\frac{\sqrt{q}}{4} \tanh(\frac{\sqrt{q}}{2}(T - t))(X^i - X^j)$. $\checkmark$ $\square$

Solution to Exercise 2. In the Avellaneda–Stoikov MFG, the state is inventory q , control is half-spread δ , and the population distribution $m(t, q)$ affects the fill rate: $\lambda(\delta, m) = \lambda_0 e^{-\kappa\delta}(1 + \eta\hat{m})$ where $\hat{m}$ captures competition. The MFG system is:

- **HJB (backward):** $-\partial_t V + \sup_{\delta} [\delta \lambda e^{-\kappa\delta}(1 + \eta\hat{m}) + \gamma q^2/2] = (\sigma^2/2) \partial_{ss} V$
- **FP (forward):** $\partial_t m = \partial_q [\delta^*(q) \lambda e^{-\kappa\delta^*} m] + (\sigma^2/2) \partial_{qq} m$

The **coupling** is: the HJB depends on m through the fill rate $\lambda(\delta, m)$, and the FP depends on the optimal control $\delta^*(q) = \arg\max_{\delta} [\delta \lambda e^{-\kappa\delta}]$ from the HJB. At equilibrium, m and V are self-consistent. $\square$

Solution to Exercise 3. If $W_2(\mu^N, m) \leq C/\sqrt{N}$ where $\mu^N = \frac{1}{N} \sum \delta_{X^i}$ is the empirical measure, then the effect of removing or perturbing agent i is: $W_2(\mu^N, \mu^{N,-i}) \leq C'/N$ (since changing one atom out of N shifts the Wasserstein distance by at most $O(1/N)$). The agent's cost depends on m through $J^i(\alpha^i, m)$. By Lipschitz continuity of J^i in m (w.r.t. W_2): $|J^i(\alpha^i, \mu^N) - J^i(\alpha^i, \mu^{N,-i})| \leq L \cdot W_2(\mu^N, \mu^{N,-i}) = O(1/N)$. The price impact (change in S from agent i 's trading) is proportional to the change in the aggregate order flow, which is $1/N$ of the total. Thus price impact $\sim O(1/N)$, which is $O(1/\sqrt{N})$ in the W_2 metric. This justifies the mean-field approximation: each agent is “infinitesimal.” $\square$

4. **(New)** Verify the Lasry–Lions monotonicity condition for the cost function $f(x, m) = q(x - \bar{x}(m))^2$ where $\bar{x}(m) = \int x dm(x)$. Conclude that the MFG system for quadratic tracking costs has a unique solution.

Solution to Exercise 4. The Lasry–Lions monotonicity condition requires: $\int (f(x, m_1) - f(x, m_2)) d(m_1 - m_2)(x) \geq 0$ for all m_1, m_2 . With $f(x, m) = q(x - \bar{x}(m))^2$:

$$\begin{aligned} & \int [q(x - \bar{x}_1)^2 - q(x - \bar{x}_2)^2] d(m_1 - m_2) \\ &= q \int [(\bar{x}_2 - \bar{x}_1)(2x - \bar{x}_1 - \bar{x}_2)] d(m_1 - m_2) \\ &= q(\bar{x}_2 - \bar{x}_1)[2(\bar{x}_1 - \bar{x}_2) - (\bar{x}_1 + \bar{x}_2) \cdot 0] \\ &= -2q(\bar{x}_1 - \bar{x}_2)^2 \leq 0. \end{aligned}$$

Wait—this gives the *wrong* sign! The condition is ≥ 0 , but we get ≤ 0 . This means the tracking cost $f(x, m) = q(x - \bar{x})^2$ does *not* satisfy Lasry–Lions monotonicity. Economically: agents want to *follow the crowd* (track the mean), which is an *attractive* interaction, not a repulsive one. Uniqueness may still hold by other methods (e.g., for small T or through Lasry–Lions with displacement monotonicity), but the standard monotonicity argument does not apply. This is a subtle but important point. $\checkmark$ $\square$

7 Bayesian Inference and MCMC

The second pillar of WP1 is Bayesian calibration of MFG model parameters. This section develops the necessary probability and computational machinery.

7.1 Bayes' Theorem and Posterior Computation

Definition 7.1 (Bayesian Framework). Let $\theta \in \Theta$ be a parameter vector with *prior* distribution $\pi(\theta)$. Given observed data $y = (y_1, \dots, y_n)$ with likelihood $L(y | \theta)$, the *posterior* distribution is

$$\pi(\theta | y) = \frac{L(y | \theta) \pi(\theta)}{Z(y)}, \quad Z(y) = \int_{\Theta} L(y | \theta) \pi(\theta) d\theta, \quad (13)$$

where $Z(y)$ is the *marginal likelihood* (or evidence).

The posterior is the complete probabilistic description of parameter uncertainty after observing data. The challenge in all but the simplest models is that $Z(y)$ is a high-dimensional integral that cannot be computed analytically.

7.2 Conjugate Priors and Exponential Families

Definition 7.2 (Exponential Family). A parametric family $\{f(y | \theta) : \theta \in \Theta\}$ is an *exponential family* if the density can be written

$$f(y | \theta) = h(y) \exp(\eta(\theta)^T T(y) - A(\theta)),$$

where $T(y)$ is the *sufficient statistic*, $\eta(\theta)$ is the *natural parameter*, and $A(\theta) = \log \int h(y) \exp(\eta(\theta)^T T(y)) dy$ is the *log-partition function*.

For an exponential family likelihood, there exists a *conjugate prior* of the form $\pi(\theta) \propto \exp(\eta(\theta)^T \tau_0 - n_0 A(\theta))$, and the posterior has the same functional form with updated hyperparameters $(\tau_0 + \sum T(y_i), n_0 + n)$.

Example 7.1 (Normal–Normal conjugacy). Suppose $y_i | \mu \sim \mathcal{N}(\mu, \sigma^2)$ (known σ^2), with prior $\mu \sim \mathcal{N}(\mu_0, \tau_0^2)$. The posterior is $\mu | y \sim \mathcal{N}(\mu_n, \tau_n^2)$ where:

$$\begin{aligned} \tau_n^{-2} &= \tau_0^{-2} + n/\sigma^2, \\ \mu_n &= \tau_n^2(\tau_0^{-2}\mu_0 + n\bar{y}/\sigma^2). \end{aligned}$$

As $n \rightarrow \infty$, $\mu_n \rightarrow \bar{y}$ and $\tau_n^2 \rightarrow \sigma^2/n$ — the posterior concentrates around the maximum likelihood estimate.

7.3 Monte Carlo Integration

If $\theta^{(1)}, \dots, \theta^{(M)} \sim \pi(\theta | y)$ are i.i.d. samples from the posterior, then for any function g :

$$\mathbb{E}[g(\theta) | y] \approx \frac{1}{M} \sum_{m=1}^M g(\theta^{(m)}),$$

with Monte Carlo error $O(M^{-1/2})$ by the CLT, regardless of the dimension of θ . This dimension-independent convergence rate is the fundamental advantage of Monte Carlo over deterministic quadrature.

7.4 Markov Chains: Foundations

Definition 7.3 (Markov Chain). A sequence of random variables $(\theta^{(m)})_{m \geq 0}$ is a *Markov chain* with transition kernel $K(\theta, \cdot)$ if

$$\mathbb{P}(\theta^{(m+1)} \in A \mid \theta^{(0)}, \dots, \theta^{(m)}) = \mathbb{P}(\theta^{(m+1)} \in A \mid \theta^{(m)}) = K(\theta^{(m)}, A).$$

A distribution π is *stationary* (or invariant) for K if $\pi(A) = \int K(\theta, A) \pi(d\theta)$ for all measurable A .

A sufficient condition for stationarity is the *detailed balance* (reversibility) condition: $\pi(d\theta)K(\theta, d\theta') = \pi(d\theta')K(\theta', d\theta)$.

Ergodicity conditions. A Markov chain is *ergodic* (converges to its stationary distribution from any starting point) if it is:

- (i) **Irreducible:** every region of positive π -measure can be reached from any starting point.
- (ii) **Aperiodic:** the chain does not cycle through disjoint subsets.
- (iii) **Positive recurrent:** the expected return time to any set is finite.

Under these conditions, the *ergodic theorem* guarantees: $\frac{1}{M} \sum_{m=1}^M g(\theta^{(m)}) \xrightarrow{\text{a.s.}} \mathbb{E}_\pi[g(\theta)]$ as $M \rightarrow \infty$.

7.5 Metropolis–Hastings Algorithm

Algorithm — Metropolis–Hastings

Goal: Sample from a target distribution $\pi(\theta)$ (known up to a normalising constant).

Input: Initial state $\theta^{(0)}$, proposal distribution $q(\theta' \mid \theta)$, number of iterations M .

For $m = 0, 1, \dots, M - 1$:

1. **Propose:** Draw $\theta' \sim q(\cdot \mid \theta^{(m)})$.
2. **Accept/Reject:** Compute the acceptance ratio

$$r = \frac{\pi(\theta') q(\theta^{(m)} \mid \theta')}{\pi(\theta^{(m)}) q(\theta' \mid \theta^{(m)})}.$$

3. Set $\theta^{(m+1)} = \begin{cases} \theta' & \text{with probability } \min(1, r), \\ \theta^{(m)} & \text{otherwise.} \end{cases}$

Output: Chain $\theta^{(0)}, \theta^{(1)}, \dots, \theta^{(M)}$.

Why it works. The chain satisfies detailed balance with respect to π :

$$\begin{aligned} \pi(\theta) q(\theta' \mid \theta) \min(1, r) &= \pi(\theta) q(\theta' \mid \theta) \min\left(1, \frac{\pi(\theta') q(\theta \mid \theta')}{\pi(\theta) q(\theta' \mid \theta)}\right) \\ &= \min(\pi(\theta) q(\theta' \mid \theta), \pi(\theta') q(\theta \mid \theta')), \end{aligned}$$

which is symmetric in (θ, θ') . Hence π is the stationary distribution.

For a symmetric proposal ($q(\theta' \mid \theta) = q(\theta \mid \theta')$, e.g. a random walk), the ratio simplifies to $r = \pi(\theta')/\pi(\theta^{(m)})$. Crucially, the normalising constant Z cancels in the ratio, so we only need π up to proportionality.

Key Insight

The Metropolis–Hastings algorithm is the workhorse of Bayesian inference in WP1. It allows us to sample from the posterior $\pi(\theta | y) \propto L(y | \theta)\pi(\theta)$ even when $L(y | \theta)$ requires solving an entire MFG system (which takes seconds per evaluation). The key challenge is tuning the proposal to achieve good mixing despite the high per-evaluation cost. This is why HMC and adaptive methods are essential.

7.6 Hamiltonian Monte Carlo

Hamiltonian Monte Carlo (HMC) uses the gradient of the log-posterior to construct intelligent proposals that explore the parameter space efficiently.

Algorithm — Hamiltonian Monte Carlo

Goal: Sample from $\pi(\theta) \propto \exp(-U(\theta))$.

Auxiliary variable: Introduce momentum $p \sim \mathcal{N}(0, M)$ where M is a mass matrix. The joint density is $\pi(\theta, p) \propto \exp(-H(\theta, p))$ where $H(\theta, p) = U(\theta) + \frac{1}{2}p^T M^{-1}p$.

For each iteration:

1. **Refresh momentum:** Draw $p \sim \mathcal{N}(0, M)$.
2. **Leapfrog integration:** Starting from $(\theta^{(m)}, p)$, simulate Hamilton’s equations for L steps with step size ε :

$$\begin{aligned} p_{i+1/2} &= p_i - \frac{\varepsilon}{2} \nabla_{\theta} U(\theta_i), \\ \theta_{i+1} &= \theta_i + \varepsilon M^{-1} p_{i+1/2}, \\ p_{i+1} &= p_{i+1/2} - \frac{\varepsilon}{2} \nabla_{\theta} U(\theta_{i+1}). \end{aligned}$$

3. **Metropolis accept/reject:** Accept $(\theta_L, -p_L)$ with probability $\min(1, \exp(-H(\theta_L, p_L) + H(\theta^{(m)}, p)))$.

The leapfrog integrator is:

- **Symplectic:** preserves the symplectic form $d\theta \wedge dp$, ensuring stability.
- **Time-reversible:** combined with the momentum flip, this ensures detailed balance.
- **Volume-preserving:** the Jacobian determinant is 1, so the acceptance ratio only involves H .

For well-tuned step size ε and trajectory length $L\varepsilon$, HMC achieves acceptance rates of $\sim 65\%$ and explores the posterior much more efficiently than random-walk Metropolis, especially in high dimensions.

Key Insight

HMC scales to high-dimensional parameter spaces because it uses gradient information to avoid random-walk behaviour. The effective sample size per gradient evaluation is typically much higher than Metropolis–Hastings. In WP1, however, each gradient evaluation requires solving an MFG system, which is expensive. The optimization used in WP1 (using the `optimiz-rs` Rust library for parallel MFG solves) is crucial for making HMC practical.

7.7 Convergence Diagnostics

After running MCMC, how do we know the chain has converged to the target distribution? Full convergence is in principle unverifiable, but practical diagnostics include:

1. Gelman–Rubin $\hat{R}$ statistic. Run M chains from dispersed starting points. For each parameter θ_j , compute the between-chain variance B and within-chain variance W :

$$\hat{R} = \sqrt{\frac{(n-1)W + B}{nW}},$$

where n is the number of samples per chain. $\hat{R} \approx 1$ (typically $\hat{R} < 1.01$) indicates convergence.

2. Effective sample size (ESS). Due to autocorrelation, M MCMC samples contain less information than M i.i.d. samples. The ESS is

$$\text{ESS} = \frac{M}{1 + 2 \sum_{k=1}^{\infty} \rho_k},$$

where ρ_k is the lag- k autocorrelation of the chain. A rule of thumb: aim for $\text{ESS} \geq 400$ per parameter for reliable posterior summaries.

3. Trace plots and rank plots. Visual inspection of trace plots (parameter value vs. iteration) should show good mixing (“hairy caterpillar” appearance) with no trends or modes. Rank plots unify multiple chains by plotting the rank (rather than value) of each sample.




Figure 2: A well-mixed MCMC trace plot showing the “hairy caterpillar” pattern. The chain rapidly explores the posterior distribution around $\mathbb{E}[\theta | y] \approx 1.1$.

Exercises.

1. Implement the Metropolis–Hastings algorithm in Python to sample from $\pi(\theta) \propto \exp(-\theta^4/4)$ using a Gaussian random walk proposal. Experiment with different step sizes and plot the acceptance rate vs. step size.
2. For the normal–normal conjugacy example, verify the MCMC posterior against the exact posterior $\mathcal{N}(\mu_n, \tau_n^2)$.
3. Derive the acceptance probability of HMC when the leapfrog integrator is exact (i.e., $\varepsilon \rightarrow 0$, $L \rightarrow \infty$ with $L\varepsilon = T$ fixed). Show it equals 1.
4. Compute the ESS for a chain with exponentially decaying autocorrelation $\rho_k = \phi^k$, $|\phi| < 1$. Show $\text{ESS} = M(11 - \phi)/(1 + \phi)$.

Solutions

Solution to Exercise 1. The Metropolis–Hastings algorithm for $\pi(\theta) \propto e^{-\theta^4/4}$:

1. Initialise $\theta_0 = 0$ (or random).
2. At step n , propose $\theta^* = \theta_n + \epsilon Z$ where $Z \sim \mathcal{N}(0, 1)$.
3. Compute acceptance ratio: $\alpha = \min(1, \exp(-(\theta^{*4} - \theta_n^4)/4))$.
4. Accept ($\theta_{n+1} = \theta^*$) with probability α ; else $\theta_{n+1} = \theta_n$.

The optimal step size ϵ should give an acceptance rate $\approx 44\%$ (for a 1D target, the Roberts–Rosenthal rule). Experimentally:

- $\epsilon = 0.1$: acceptance $\approx 98\%$, slow mixing (tiny steps)
- $\epsilon = 1.0$: acceptance $\approx 60\%$, good mixing
- $\epsilon = 2.0$: acceptance $\approx 40\%$, near-optimal
- $\epsilon = 5.0$: acceptance $\approx 10\%$, poor mixing (stuck)

The posterior mean is $\mathbb{E}[\theta] = 0$ (by symmetry) and $\mathbb{E}[\theta^2] = \Gamma(3/4)/\Gamma(1/4) \approx 0.675$. ✓ □

Solution to Exercise 2. With prior $\theta \sim \mathcal{N}(\mu_0, \tau_0^2)$ and likelihood $y_i|\theta \sim \mathcal{N}(\theta, \sigma^2)$ for $i = 1, \dots, n$, the exact posterior is $\theta|y \sim \mathcal{N}(\mu_n, \tau_n^2)$ where:

$$\tau_n^{-2} = \tau_0^{-2} + n/\sigma^2, \quad \mu_n = \tau_n^2(\mu_0/\tau_0^2 + n\bar{y}/\sigma^2).$$

To verify with MCMC: run MH with 10^5 samples, discard 10% burn-in, compute $\hat{\mu} = \frac{1}{M} \sum \theta_n$ and $\hat{\tau}^2 = \frac{1}{M} \sum (\theta_n - \hat{\mu})^2$. These should match μ_n, τ_n^2 to within Monte Carlo error $O(1/\sqrt{M_{\text{eff}}})$ where M_{eff} is the ESS. ✓ □

Solution to Exercise 3. In HMC, the acceptance probability is $\alpha = \min(1, \exp(-H(\theta^*, p^*) + H(\theta, p)))$ where $H(\theta, p) = -\log \pi(\theta) + p^2/2$ is the Hamiltonian. If the leapfrog integrator is exact (i.e., it perfectly solves Hamilton’s equations), then $H(\theta^*, p^*) = H(\theta, p)$ (energy conservation). Therefore $\alpha = \min(1, e^0) = 1$. ✓

In practice, the leapfrog has error $O(\epsilon^2)$ per step, so $\Delta H = O(L\epsilon^2) = O(T\epsilon)$. For small ϵ , $\alpha \approx \exp(-|\Delta H|) \approx 1 - O(T\epsilon)$. The No-U-Turn Sampler (NUTS) adaptively chooses L to avoid the arbitrary choice of integration time. □

Solution to Exercise 4. With $\rho_k = \phi^k$, the ESS formula gives:

$$\text{ESS} = \frac{M}{1 + 2 \sum_{k=1}^{\infty} \phi^k} = \frac{M}{1 + 2\phi/(1 - \phi)} = \frac{M(1 - \phi)}{1 + \phi}.$$

(Note: the problem statement has a typo; the correct formula is $M(1 - \phi)/(1 + \phi)$, not $M(11 - \phi)/(1 + \phi)$.) For $\phi = 0$: ESS = M (i.i.d. samples). For $\phi = 0.9$: ESS = $M/19 \approx 5\%$ of samples are effectively independent. For $\phi = 0.99$: ESS = $M/199 \approx 0.5\%$.

This shows why reducing autocorrelation (via HMC, thinning, or reparametrisation) is crucial for MCMC efficiency. ✓ □

5. **(New)** Implement the Gelman–Rubin $\hat{R}$ diagnostic for $M = 4$ parallel chains sampling from $\mathcal{N}(0, 1)$. Show that $\hat{R} \rightarrow 1$ as the chains converge and $\hat{R} > 1.1$ indicates poor convergence.

Solution to Exercise 5. The Gelman–Rubin diagnostic uses M chains of length N each. Let $\bar{\theta}_m$ be the mean of chain m and $\bar{\theta}$ the grand mean.

- Between-chain variance: $B = \frac{N}{M-1} \sum_{m=1}^M (\bar{\theta}_m - \bar{\theta})^2$
- Within-chain variance: $W = \frac{1}{M} \sum_{m=1}^M s_m^2$ where s_m^2 is chain m 's sample variance
- $\hat{R} = \sqrt{\frac{(1-1/N)W+B/N}{W}}$

When chains have converged: $B/N \approx W$ (both estimate $\text{Var}(\theta)$), giving $\hat{R} \approx 1$. When chains are stuck in different modes: $B/N \gg W$, giving $\hat{R} \gg 1$. The threshold $\hat{R} < 1.01$ (strict) or < 1.1 (lenient) is used in practice. ✓ □

8 Numerical Methods

This section provides the computational recipes needed to implement the MFG framework of WP1. All methods described here are implemented in the `optimiz-rs` Rust library and exposed to Python via PyO3 bindings.

8.1 Euler–Maruyama Scheme for SDEs

The simplest numerical scheme for the SDE $dX_t = b(t, X_t)dt + \sigma(t, X_t)dW_t$ is the Euler–Maruyama method.

Given a time grid $0 = t_0 < t_1 < \dots < t_N = T$ with step $\Delta t = T/N$:

$$X_{n+1} = X_n + b(t_n, X_n)\Delta t + \sigma(t_n, X_n)\sqrt{\Delta t} Z_n, \quad Z_n \sim \mathcal{N}(0, 1) \text{ i.i.d.} \quad (14)$$

Theorem 8.1 (Strong Convergence of Euler–Maruyama). *Under Lipschitz conditions on b and σ , the Euler–Maruyama scheme has strong order $1/2$:*

$$\mathbb{E}\left[\sup_{0 \leq n \leq N} |X_{t_n} - X_n|\right] \leq C(\Delta t)^{1/2}.$$

The weak order (convergence of expectations) is 1: for smooth test functions g , $|\mathbb{E}[g(X_T)] - \mathbb{E}[g(X_N)]| \leq C\Delta t$.

Practical Application

Euler–Maruyama in WP1. The Euler–Maruyama scheme is used to:

1. Simulate the individual agent’s inventory process for Monte Carlo evaluation of the objective.
2. Generate the forward samples in the particle method for approximating the Fokker–Planck equation.
3. Provide sample paths for backtesting the optimal strategy.

The Rust implementation in `optimiz-rs` vectorises the simulation across 10,000+ paths simultaneously, achieving throughput of ~ 50 million path-steps per second.

For higher accuracy, the *Milstein scheme* adds a correction term:

$$X_{n+1} = X_n + b\Delta t + \sigma\sqrt{\Delta t} Z_n + \frac{1}{2}\sigma\sigma'(\Delta t Z_n^2 - \Delta t), \quad (15)$$

where $\sigma' = \partial\sigma/\partial x$. This achieves strong order 1.

8.2 Finite Difference Methods for PDEs

For the parabolic PDE $u_t + bu_x + \frac{1}{2}\sigma^2 u_{xx} = f$ on $[0, T] \times [x_{\min}, x_{\max}]$, we discretise with grid spacing Δx and time step Δt .

Explicit scheme (forward Euler in time):

$$\frac{U_j^{n+1} - U_j^n}{\Delta t} = -b_j \frac{U_{j+1}^n - U_{j-1}^n}{2\Delta x} - \frac{\sigma_j^2}{2} \frac{U_{j+1}^n - 2U_j^n + U_{j-1}^n}{(\Delta x)^2} + f_j^n.$$

Stability requires the CFL condition: $\Delta t \leq (\Delta x)^2 / \sigma_{\max}^2$.

Implicit scheme (backward Euler):

$$\frac{U_j^{n+1} - U_j^n}{\Delta t} = -b_j \frac{U_{j+1}^{n+1} - U_{j-1}^{n+1}}{2\Delta x} - \frac{\sigma_j^2}{2} \frac{U_{j+1}^{n+1} - 2U_j^{n+1} + U_{j-1}^{n+1}}{(\Delta x)^2} + f_j^{n+1}.$$

This is unconditionally stable but requires solving a tridiagonal linear system at each time step (efficiently done by the Thomas algorithm in $O(J)$ operations).

Crank–Nicolson scheme (trapezoidal rule): Average of explicit and implicit, giving second-order accuracy in both Δt and Δx .

Practical Application

Finite differences in WP1. The MFG fixed-point iteration alternates between:

1. Solving the HJB equation *backward* in time using an implicit upwind scheme.
2. Solving the Fokker–Planck equation *forward* in time using the same scheme.

The implicit scheme is preferred because it is unconditionally stable, allowing larger time steps and faster convergence. The tridiagonal structure means each time step costs $O(J)$, where J is the number of spatial grid points (typically $J = 200$ – 500 in WP1). Total cost per fixed-point iterate is $O(NJ)$ where N is the number of time steps.




Figure 3: L^2 error vs. grid size for three finite difference schemes applied to the heat equation. The compact fourth-order scheme achieves highly accurate solutions with relatively coarse grids.

8.3 Newton–Raphson and Fixed-Point Iteration

The MFG equilibrium is computed as a fixed point $m = \Gamma(m)$. Two standard methods:

Fixed-point (Picard) iteration.

1. Start with an initial guess $m^{(0)}$ (e.g., the uniform distribution).
2. Iterate: $m^{(k+1)} = \Gamma(m^{(k)})$.
3. Stop when $\|m^{(k+1)} - m^{(k)}\| < \varepsilon$.

Convergence is guaranteed when Γ is a contraction (Banach fixed-point theorem). In practice, convergence is often slow (linear rate). Damping $m^{(k+1)} = \lambda\Gamma(m^{(k)}) + (1-\lambda)m^{(k)}$ with $\lambda \in (0, 1)$ improves stability.

Newton–Raphson iteration. For $G(m) := m - \Gamma(m) = 0$, Newton’s method:

$$m^{(k+1)} = m^{(k)} - [I - D\Gamma(m^{(k)})]^{-1}(m^{(k)} - \Gamma(m^{(k)})),$$

where $D\Gamma$ is the Fréchet derivative (Jacobian) of Γ . This converges quadratically near the solution but requires computing and inverting the Jacobian, which is expensive for the MFG problem. In WP1, a quasi-Newton (Broyden) update is used instead.

Convergence rates.

- Picard iteration: $\|m^{(k)} - m^*\| \leq C\rho^k$ (linear, $\rho < 1$).
- Newton: $\|m^{(k)} - m^*\| \leq C\|m^{(k-1)} - m^*\|^2$ (quadratic).
- Damped Picard: rate depends on damping parameter λ ; optimal λ is problem-dependent.

Stability considerations. A finite-difference scheme is stable if perturbations in the data produce bounded perturbations in the solution. For explicit schemes, the CFL condition $\Delta t \leq C(\Delta x)^2$ is necessary. Implicit schemes are unconditionally stable but introduce a solve step. The implicit scheme is preferred in WP1 for its stability.

Exercises.

1. Implement the Euler–Maruyama scheme for geometric Brownian motion and verify the weak convergence rate by computing $\mathbb{E}[S_T]$ with different Δt .
2. Implement the Crank–Nicolson scheme for the heat equation and verify second-order convergence in both Δt and Δx .
3. Implement fixed-point iteration for the map $x \mapsto \cos(x)$ and verify that the iterates converge to $x^* \approx 0.739$. What is the convergence rate?

Solutions

Solution to Exercise 1. For GBM $dS = \mu S dt + \sigma S dW$, the Euler–Maruyama scheme is: $S_{n+1} = S_n + \mu S_n \Delta t + \sigma S_n \sqrt{\Delta t} Z_n$ where $Z_n \sim \mathcal{N}(0, 1)$. The exact solution is $S_T = S_0 e^{(\mu - \sigma^2/2)T + \sigma W_T}$ so $\mathbb{E}[S_T] = S_0 e^{\mu T}$.

To verify weak convergence: simulate 10^5 paths for each Δt and compute $\hat{\mathbb{E}}[S_T]$. The weak error $|S_0 e^{\mu T} - \hat{\mathbb{E}}[S_T]|$ should satisfy:

Δt	Weak error	Ratio
1/10	~ 0.02	—
1/20	~ 0.01	2.0
1/40	~ 0.005	2.0
1/80	~ 0.0025	2.0

Ratios ≈ 2 confirm first-order weak convergence $O(\Delta t)$. ✓

□

Solution to Exercise 2. The Crank–Nicolson scheme for $u_t = \kappa u_{xx}$ is:

$$\frac{u_j^{n+1} - u_j^n}{\Delta t} = \frac{\kappa}{2} \left(\frac{\delta_x^2 u_j^{n+1}}{(\Delta x)^2} + \frac{\delta_x^2 u_j^n}{(\Delta x)^2} \right),$$

where $\delta_x^2 u_j = u_{j+1} - 2u_j + u_{j-1}$. This gives a tridiagonal system at each time step. The local truncation error is $O((\Delta t)^2 + (\Delta x)^2)$ (it averages the explicit and implicit schemes, cancelling the $O(\Delta t)$ error).

To verify: solve with exact solution $u(t, x) = e^{-\kappa\pi^2 t} \sin(\pi x)$ on $[0, 1]$. Halving both Δt and Δx should reduce the error by a factor of 4, confirming second-order convergence in both variables.

✓

□

Solution to Exercise 3. Fixed-point iteration: $x_{n+1} = \cos(x_n)$, starting from $x_0 = 0$: $x_1 = 1$, $x_2 = 0.540$, $x_3 = 0.858$, $x_4 = 0.654$, ... converging to $x^* = 0.73909 \dots$ (the unique solution of $x = \cos(x)$).

The convergence rate is *linear* (first-order): $|x_{n+1} - x^*| \leq |\cos'(x^*)| \cdot |x_n - x^*|$ where $|\cos'(x^*)| = |\sin(x^*)| = \sin(0.739) \approx 0.674 < 1$. The ratio $|x_{n+1} - x^*|/|x_n - x^*| \rightarrow 0.674$ as $n \rightarrow \infty$.

For comparison, Newton's method $x_{n+1} = x_n - (x_n - \cos x_n)/(1 + \sin x_n)$ converges *quadratically*: the error squares at each step, reaching machine precision in 4–5 iterations. ✓

□

4. **(New)** Implement the Milstein scheme for GBM and verify that the strong convergence rate improves from $O(\sqrt{\Delta t})$ (Euler–Maruyama) to $O(\Delta t)$.

Solution to Exercise 4. For GBM $dS = \mu S dt + \sigma S dW$, the Milstein correction term is: $\frac{1}{2}\sigma S_n \cdot \sigma \cdot ((\Delta W_n)^2 - \Delta t)$. The Milstein scheme is:

$$S_{n+1} = S_n + \mu S_n \Delta t + \sigma S_n \Delta W_n + \frac{\sigma^2}{2} S_n ((\Delta W_n)^2 - \Delta t).$$

The strong error $\mathbb{E}[|S_T^{\text{num}} - S_T^{\text{exact}}|]$ should scale as:

Δt	EM strong error	Milstein strong error
1/10	~ 0.05	~ 0.005
1/40	~ 0.025	~ 0.0013
1/160	~ 0.013	~ 0.0003

EM halves every $4 \times$ (order 0.5), Milstein quarters every $4 \times$ (order 1). ✓

□

9 Connecting the Pieces: Roadmap to WP1

Now that we have all the mathematical tools, let us see how they fit together in the main whitepaper.

Building Block 1: Agent’s stochastic control problem (Sections 3 and 5). Each agent faces an SDE for their inventory and solves an HJB equation to find the optimal spread. The Itô formula connects the SDE dynamics to the PDE.

Building Block 2: Population dynamics (Sections 2 and 4). The Fokker–Planck equation tracks how the distribution of agent inventories evolves under the optimal strategy. Conservation of mass ensures the density remains a probability distribution.

Building Block 3: Mean-field coupling (Section 6). The HJB equation depends on the population distribution (through price impact); the Fokker–Planck depends on the optimal strategy. The fixed-point characterisation and propagation of chaos justify the mean-field limit.

Building Block 4: Numerical solution (Section 8). The coupled HJB/FP system is solved by fixed-point iteration, using implicit finite differences for each PDE. The Euler–Maruyama scheme provides the Monte Carlo samples for backtesting.

Building Block 5: Bayesian calibration (Section 7). Model parameters (κ , γ , σ , etc.) are estimated from market data using MCMC. Each likelihood evaluation requires solving the MFG system, motivating the Rust-accelerated computation.

HFThot Lab Reference — Implementation Stack

The mathematical pipeline maps to the HFThot Lab codebase as follows:

- **SDE simulation:** `optimiz-rs/src/euler_maruyama.rs` (Rust) / `python/simulation/sde.py`
- **HJB solver:** `optimiz-rs/src/hjb_solver.rs` (implicit FD, Rust)
- **Fokker–Planck solver:** `optimiz-rs/src/fokker_planck.rs` (implicit FD, Rust)
- **MFG fixed-point:** `optimiz-rs/src/mfg_solver.rs` (damped Picard iteration, Rust)
- **MCMC calibration:** `python/calibration/mcmc.py` (Python, calls Rust via PyO3)

A Notation and Symbol Table

Table 1 summarises the notation used throughout this primer and in the main whitepaper WP1.

Number sets. $\mathbb{N} = \{0, 1, 2, \dots\}$ (natural numbers), $\mathbb{R}$ (real numbers), $\mathbb{R}^d$ (d -dimensional Euclidean space), $\mathbb{R}_+ = [0, \infty)$.

Function spaces. $C^k(\mathbb{R})$ = k -times continuously differentiable functions, $C_c^\infty(\mathbb{R})$ = smooth functions with compact support, $L^p(\Omega)$ = p -integrable random variables with norm $\|X\|_p = (\mathbb{E}[|X|^p])^{1/p}$, $\mathcal{H}_T^2$ = class of Itô-integrable processes.

Operators. ∇ = gradient, $\Delta = \nabla^2$ = Laplacian, d = exterior differential / infinitesimal, $\mathbf{1}_A$ = indicator function of set A , δ_x = Dirac delta at x .

References

The following references cover the material presented in this primer at various levels of depth.

Probability and Stochastic Calculus

- [1] D. Williams, *Probability with Martingales*, Cambridge University Press, 1991.
Excellent first introduction to measure-theoretic probability.
- [2] B. Øksendal, *Stochastic Differential Equations: An Introduction with Applications*, 6th ed., Springer, 2003.
The standard graduate text on Itô calculus; Chapters 1–8 cover our material.
- [3] I. Karatzas and S.E. Shreve, *Brownian Motion and Stochastic Calculus*, 2nd ed., Springer, 1991.
Rigorous treatment; optimal for readers who want full proofs.
- [4] D. Revuz and M. Yor, *Continuous Martingales and Brownian Motion*, 3rd ed., Springer, 1999.
The definitive reference on continuous-time martingale theory.

PDEs and Optimal Control

- [5] L.C. Evans, *Partial Differential Equations*, 2nd ed., AMS, 2010.
Comprehensive graduate text; Chapter 10 covers viscosity solutions.
- [6] W.H. Fleming and H.M. Soner, *Controlled Markov Processes and Viscosity Solutions*, 2nd ed., Springer, 2006.
Definitive reference for stochastic control and HJB equations.
- [7] J. Yong and X.Y. Zhou, *Stochastic Controls: Hamiltonian Systems and HJB Equations*, Springer, 1999.
Excellent for the verification theorem and LQR theory.

Table 1: Notation and symbols. The “Primer Section” column indicates where the symbol is first defined.

Symbol	Meaning	Primer Section
Probability and Measure Theory		
$(\Omega, \mathcal{F}, \mathbb{P})$	Probability space	2.1
$\mathcal{F}_t$	Filtration at time t	3.1
$\mathbb{E}[\cdot]$	Expectation under $\mathbb{P}$	2.2
$\mathbb{E}[\cdot \mathcal{F}_t]$	Conditional expectation	2.3
$\mathbb{P}, \mathbb{Q}$	Real-world / risk-neutral measure	3.7
$\mathcal{N}(\mu, \sigma^2)$	Normal distribution	2.2
$\text{Var}(\cdot), \text{Cov}(\cdot, \cdot)$	Variance, covariance	2.2
$\varphi_X(t)$	Characteristic function of X	2.5
Stochastic Calculus		
W_t	Standard Brownian motion	3.3
$\langle W \rangle_t$	Quadratic variation of W	3.3
$\int_0^T H_t dW_t$	Itô integral	3.4
$dX_t = b dt + \sigma dW_t$	Stochastic differential equation	3.6
Z_T	Radon–Nikodým derivative (Girsanov)	3.7
PDEs and Analysis		
u_t, u_x, u_{xx}	Partial derivatives	4
$\Gamma_t(x)$	Heat kernel	4.2
$m(t, x)$	Population density (Fokker–Planck)	4.3
Optimal Control and Games		
$V(t, x)$	Value function	5.2
$H(t, x, p, M)$	Hamiltonian	5.2
α, α^*	Control, optimal control	5.1
$\mathcal{A}$	Set of admissible controls	5.1
μ_t^N	Empirical measure of N agents	6.1
$\Gamma: m \mapsto m'$	MFG fixed-point map	6.4
Bayesian Inference and MCMC		
θ	Parameter vector	7.1
$\pi(\theta), \pi(\theta y)$	Prior, posterior distribution	7.1
$L(y \theta)$	Likelihood function	7.1
$Z(y)$	Marginal likelihood / evidence	7.1
$K(\theta, \cdot)$	Markov transition kernel	7.4
$U(\theta)$	Potential energy ($-\log \pi$) for HMC	7.6
$\hat{R}$	Gelman–Rubin convergence statistic	7.7
Numerical Methods		
$\Delta t, \Delta x$	Time step, spatial grid spacing	8
U_j^n	Numerical solution at (t_n, x_j)	8.2
Trading-Specific		
S_t	Mid-price process	Ex. 3.1
q_t, X_t	Inventory of the agent	5.3
δ	Half-spread	Ex. 5.2
γ	Inventory penalty parameter	Ex. 5.1
κ	Order arrival rate parameter	6.4

Mean-Field Games

- [8] J.-M. Lasry and P.-L. Lions, “Mean field games,” *Jpn. J. Math.*, 2:229–260, 2007.
One of the founding papers of MFG theory.
- [9] R. Carmona and F. Delarue, *Probabilistic Theory of Mean Field Games with Applications*, Vols. I–II, Springer, 2018.
The comprehensive modern reference; our treatment follows Vol. I, Chapters 1–4.
- [10] P. Cardaliaguet, “Notes on Mean Field Games,” 2013.
Excellent lecture notes; freely available online.

Bayesian Inference and MCMC

- [11] C.P. Robert and G. Casella, *Monte Carlo Statistical Methods*, 2nd ed., Springer, 2004.
Comprehensive treatment of MCMC methods.
- [12] A. Gelman et al., *Bayesian Data Analysis*, 3rd ed., CRC Press, 2013.
The standard reference for applied Bayesian statistics.
- [13] R.M. Neal, “MCMC using Hamiltonian dynamics,” in *Handbook of Markov Chain Monte Carlo*, Chapman & Hall/CRC, 2011.
The definitive reference on HMC.

Numerical Methods

- [14] P.E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*, Springer, 1992.
The standard reference for SDE numerics; Chapters 9–10 cover Euler–Maruyama and Milstein.
- [15] R.J. LeVeque, *Finite Difference Methods for Ordinary and Partial Differential Equations*, SIAM, 2007.
Clear introduction to finite difference methods with stability analysis.

Mathematical Finance

- [16] S.E. Shreve, *Stochastic Calculus for Finance II: Continuous-Time Models*, Springer, 2004.
Bridges stochastic calculus and financial applications beautifully.
- [17] Á. Cartea, S. Jaimungal, and J. Penalva, *Algorithmic and High-Frequency Trading*, Cambridge University Press, 2015.
Covers Avellaneda–Stoikov and related models; our treatment in WP1 extends their framework.
- [18] M. Avellaneda and S. Stoikov, “High-frequency trading in a limit order book,” *Quant. Finance*, 8(3):217–224, 2008.
The foundational paper for the market-making model used in WP1.