

Foundations and Strategies in Stochastic Control: Dynamic Optimization in Financial Engineering

Fatim Majumder

March 25, 2025

Martingales

Objective: Continuous-Time Dynamic Portfolio Problem

Goal: Maximize expected utility over investment horizon $[0, T]$.

$$\text{Maximize } E \left[\int_0^T e^{-\rho s} u(c(s)) ds + e^{-\rho T} b(W(T)) \right],$$

subject to $dW(s) = rW(s) ds + \theta(s)(\mu - r) ds + \sigma\theta(s) dZ(s) - c(s) ds$

$$W(s) \geq 0, \quad \forall s \in [0, T].$$

Key Variables and Parameters:

- $\theta(s)$: Portfolio allocation in risky asset at time s .
- $c(s)$: Consumption rate at each instant s .
- $u(\cdot)$: Utility function representing consumption preferences.
- $b(W(T))$: Utility derived from terminal wealth.
- ρ : Rate of time preference, r : Risk-free interest rate.
- μ : Expected return on the risky asset, σ^2 : Variance of return.
- $dZ(s)$: Increment of a Wiener process, representing stochastic market movements.

Martingale in Optimal Strategy: The Value Process $\mathcal{M}(t)$

Martingale Dynamics in Financial Optimization:

- A process $\mathcal{M}(t)$ is a **martingale** if, for any $s < t$, its expected value conditioned on the information up to s equals its current value:

$$E[\mathcal{M}(t)|\mathcal{F}_s] = \mathcal{M}(s),$$

indicating no expected change in value in an optimally controlled financial model.

Role of the Law of Iterated Expectations:

- The law of iterated expectations states that for any integrable random variable X and times $s < t$:

$$E[E[X|\mathcal{F}_t]|\mathcal{F}_s] = E[X|\mathcal{F}_s].$$

- In the context of $\mathcal{M}(t)$, this ensures that the expected value of the process remains consistent across different time intervals, a fundamental aspect in stochastic control theory.

Application in Optimization Problem

Role of $\mathcal{M}(t)$ in Expected Utility Optimization:

- The value process $\mathcal{M}(t)$ is integral to understanding the dynamics of expected utility within our optimization problem.
- It represents the cumulative utility from consumption up to time t and the expected utility of future wealth, formulated as:

$$\mathcal{M}(t) = \int_0^t e^{-\rho\tau} u(c(\tau)) d\tau + e^{-\rho t} V(W(t), t),$$

where $V(W(t), t)$ denotes the value function, which quantifies the maximum expected utility from time t onwards, given the wealth level $W(t)$.

Implications for Dynamic Portfolio Management:

- This formulation allows for the analysis of consumption and investment decisions over time in a stochastic financial environment.
- The insights derived from $\mathcal{M}(t)$ are crucial for developing optimal strategies that adapt to varying market conditions.

Itô's Lemma

Key Concepts in Stochastic Dynamic Programming

The Itô Process in Financial Modeling:

- The Itô process is a fundamental stochastic model in financial engineering, capturing the nuanced dynamics of financial variables:

$$dX(t) = \mu(t)X(t) dt + \sigma(t)X(t) dW(t),$$

where $X(t)$ represents a financial variable at time t , $\mu(t)$ is the drift component, $\sigma(t)$ denotes volatility, and $dW(t)$ is the increment of a Wiener process.

Drift and Volatility:

- **Drift** $\mu(t)$: This term signifies the expected rate of change or the trend in the asset price over an infinitesimally small time interval, playing a key role in forecasting asset price movements.
- **Volatility** $\sigma(t)$: Captures the degree of uncertainty or risk associated with the asset's returns, essential for understanding market fluctuations and informing risk management strategies.

Itô's Lemma in Stochastic Calculus

Itô's Lemma:

- Itô's Lemma is a cornerstone in stochastic calculus, crucial for the analysis and modeling of dynamic systems in financial engineering.
- It provides a framework for differentiating functions of stochastic processes, uncovering the subtle and complex changes in financial models over time.

Mathematical Formulation and Implications:

- For a function $f(X(t), t)$ dependent on a stochastic process $X(t)$, Itô's Lemma delineates:

$$df(X(t), t) = \left(\frac{\partial f}{\partial t} + \mu(t) \frac{\partial f}{\partial X} + \frac{1}{2} \sigma(t)^2 \frac{\partial^2 f}{\partial X^2} \right) dt + \sigma(t) \frac{\partial f}{\partial X} dW(t).$$

- In this formulation, $\mu(t)$ symbolizes the drift, $\sigma(t)$ the volatility of the process $X(t)$, and $dW(t)$ represents the increment of the Wiener process at time t .

Multidimensional Itô's Lemma in Stochastic Control

Itô's Lemma for Multidimensional Processes:

- The extension of Itô's Lemma to multidimensional stochastic processes is pivotal for sophisticated financial modeling, offering a more comprehensive analysis of dynamic systems influenced by multiple variables.

Multidimensional Context:

- Consider a function $H(\mathbf{X}(t), t)$, where $\mathbf{X}(t)$ represents a d -dimensional stochastic vector in \mathbb{R}^d over time $[0, T]$.
- With continuous partial derivatives H_t , $H_{\mathbf{X}}$, and $H_{\mathbf{X}\mathbf{X}}$, the dynamics of $Y(t) = H(\mathbf{X}(t), t)$ are articulated as:

$$dY(t) = H_t dt + H_{\mathbf{X}} d\mathbf{X}(t) + \frac{1}{2} \text{tr}(\mathbf{G}(t)\mathbf{G}(t)' H_{\mathbf{X}\mathbf{X}}) dt,$$

where $\mathbf{G}(t)$ is the volatility matrix corresponding to $\mathbf{X}(t)$.

- This lemma is integral for analyzing complex financial instruments and risk models that involve interactions across multiple variables.

The Bellman Equation

Transition from Value Process $\mathcal{M}(t)$ to Bellman Equation

Martingale and Supermartingale Concepts:

- In an optimally controlled strategy, the value process $\mathcal{M}(t)$ functions as a **martingale**, where its drift $a(t) = 0$ indicates a stable expected value over time.
- Conversely, under suboptimal strategies, $\mathcal{M}(t)$ takes on **supermartingale** characteristics, with drift $a(t) \leq 0$, suggesting a potential decline in expected value.

Criterion for Achieving Optimality in Strategy:

- The criterion for an optimal strategy is to nullify the drift of $\mathcal{M}(t)$:

$$\max_{\theta(t), c(t)} \text{Drift}(\mathcal{M}(t)) = 0,$$

signifying the absence of expected value change and reflecting the core principle of strategic asset allocation.

Applying Itô's Lemma to $\mathcal{M}(t)$: Wealth Dynamics

- **Itô's Lemma** is utilized to dissect the differential $d\mathcal{M}(t)$, which reflects the dynamic interplay between consumption decisions, investment strategies, and market fluctuations in shaping wealth.
- The formulation provides insights into the cumulative effects of consumption rate $c(t)$, portfolio allocation $\theta(t)$, and market parameters (μ, r, σ) on wealth evolution $W(t)$ over time.

$$\begin{aligned}d\mathcal{M}(t) = & e^{-\rho t}(u(c(t)) + \frac{\partial V}{\partial t}(W(t), t) - \rho V(W(t), t)) \\ & + \frac{\partial V}{\partial W}(W(t), t)(rW(t) + \theta(t)(\mu - r) - c(t)) dt \\ & + \frac{1}{2}\theta(t)^2\sigma^2\frac{\partial^2 V}{\partial W^2}(W(t), t) dt,\end{aligned}$$

Deriving the Bellman Equation for Dynamic Optimization

Fundamental Role of the Bellman Equation:

- Pivotal in dynamic optimization, necessitating the drift of the value process $\mathcal{M}(t)$ to be zero in an optimized strategy.

$$\begin{aligned} \max_{\theta(t), c(t)} & \left(u(c(t)) + \frac{\partial V}{\partial t}(W(t), t) - \rho V(W(t), t) \right. \\ & + \frac{\partial V}{\partial W}(W(t), t)(rW(t) + \theta(t)(\mu - r) - c(t)) \\ & \left. + \frac{\theta(t)^2 \sigma^2}{2} \frac{\partial^2 V}{\partial W^2}(W(t), t) \right) = 0, \end{aligned}$$

Characterizing Optimal Strategies in Financial Models:

- Aids in identifying the optimal consumption rate $c(t)$ and portfolio allocation $\theta(t)$ for maximizing expected utility.
- Establishes a recursive relationship for the value function $V(W(t), t)$, critical in strategic planning under stochastic conditions, incorporating risk, return, and temporal preferences.

Optimal Strategies

Derivation of Optimal Strategies from First-Order Conditions:

- The optimal consumption rate $c(t)$ and portfolio allocation $\theta(t)$ are determined from the first-order conditions of the Bellman equation, essential for maximizing expected utility.

Mathematical Formulation:

- Marginal Utility of Consumption:

$$\frac{\partial u}{\partial c}(c(t)) = \frac{\partial V}{\partial W}(W(t), t),$$

- Balancing Risk and Return in Portfolio Choice:

$$(\mu - r) \frac{\partial V}{\partial W}(W(t), t) + \frac{\sigma^2 \theta(t)}{2} \frac{\partial^2 V}{\partial W^2}(W(t), t) = 0.$$

Optimal Consumption and Portfolio Allocation

Determining Optimal Consumption $c^*(t)$:

- The optimal consumption rate $c^*(t)$ is derived by inverting the marginal utility function, linking the utility gained from consumption to the wealth level:

$$c^*(t) = I \left(\frac{\partial V}{\partial W}(W(t), t) \right),$$

where I is the inverse of the marginal utility function, and $\frac{\partial V}{\partial W}$ represents the marginal value of wealth.

Calculating Optimal Portfolio Weight $\theta^*(t)$:

- Optimal portfolio weight $\theta^*(t)$ balances risk-return trade-off, accounting for market dynamics and individual risk preferences:

$$\theta^*(t) = -\frac{\mu - r}{\sigma^2} \frac{\frac{\partial V}{\partial W}(W(t), t)}{\frac{\partial^2 V}{\partial W^2}(W(t), t)},$$

where $\mu - r$ reflects the excess return over the risk-free rate, and σ^2 represents the variance of the risky asset.

Bellman Equation with Optimized Values

- The integration of optimal consumption $c^*(t)$ and portfolio allocation $\theta^*(t)$ into the Bellman equation crystallizes the framework for dynamic optimization in financial settings.
- This approach provides a nuanced understanding of the interconnections among consumption choices, investment strategies, and wealth evolution.

$$\begin{aligned} & u \left(I \left(\frac{\partial V}{\partial W}(W(t), t) \right) \right) + \frac{\partial V}{\partial t}(W(t), t) - \rho V(W(t), t) \\ & + \left(rW(t) - I \left(\frac{\partial V}{\partial W}(W(t), t) \right) \right) \frac{\partial V}{\partial W}(W(t), t) \\ & - \frac{(\mu - r)^2}{2\sigma^2} \left(\frac{\frac{\partial V}{\partial W}(W(t), t)}{\frac{\partial^2 V}{\partial W^2}(W(t), t)} \right)^2 = 0, \end{aligned}$$

where u denotes the utility function, I is the inverse of the marginal utility, and V is the value function.

Algorithm

Algorithm for Dynamic Portfolio Optimization

Algorithm 1 Dynamic Portfolio Optimization

1: **Model Asset Dynamics (Implement Itô's Process):**

$$dX(t) = \mu(t)X(t) dt + \sigma(t)X(t) dW(t).$$

2: **Formulate Bellman Equation:** Construct

$$V(W(t), t) = \max_{\theta(t), c(t)} E \left[\int_t^T e^{-\rho s} u(c(s)) ds + e^{-\rho T} b(W(T)) | \mathcal{F}_t \right].$$







3: **Apply Itô's Lemma:** Decompose the dynamics of the value function and other relevant financial variables.

4: **Derive Optimal Strategies:** Calculate optimal consumption $c^*(t)$ and portfolio weight $\theta^*(t)$ using the first-order conditions.

5: **Numerical Solution:** Solve the resulting partial differential equation (PDE) using numerical methods. =0

References

References

-  K. Itô, *On Stochastic Differential Equations*, Memoirs of the American Mathematical Society, 1951.
-  F. Black and M. Scholes, *The Pricing of Options and Corporate Liabilities*, Journal of Political Economy, 1973.
-  R. C. Merton, *Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case*, Review of Economics and Statistics, 1969.
-  R. Bellman, *Dynamic Programming*, Princeton University Press, 1957.
-  W. H. Fleming and R. W. Rishel, *Deterministic and Stochastic Optimal Control*, Springer Verlag, 1975.
-  J. C. Hull, *Options, Futures, and Other Derivatives*, Prentice Hall, 8th Edition, 2012.