Stochastic Optimal Control with Finance Applications

Tomas Björk,
Department of Finance,
Stockholm School of Economics,

KTH, February, 2010

Contents

- Dynamic programming.
- Investment theory.
- The martingale approach.
- Markov processes
- Filtering theory.

1. Dynamic Programming

- The basic idea.
- Deriving the HJB equation.
- The verification theorem.
- The linear quadratic regulator.

Problem Formulation

$$\max_{u} \quad E\left[\int_0^T F(t, X_t, u_t)dt + \Phi(X_T)\right]$$

subject to

$$\begin{aligned} dX_t &= \mu\left(t, X_t, u_t\right)dt + \sigma\left(t, X_t, u_t\right)dW_t \\ X_0 &= x_0, \\ u_t &\in U(t, X_t), \quad \forall t. \end{aligned}$$

We will only consider **feedback control laws**, i.e. controls of the form

$$u_t = \mathbf{u}(t, X_t)$$

Terminology:

$$\begin{aligned} X &= \text{state variable} \\ u &= \text{control variable} \\ U &= \text{control constraint} \end{aligned}$$

Note: No state space constraints.

Main idea

- Embedd the problem above in a family of problems indexed by starting point in time and space.
- Tie all these problems together by a PDE–the Hamilton Jacobi Bellman equation.
- The control problem is reduced to the problem of solving the deterministic HJB equation.

Some notation

- For any fixed vector $u \in R^k$, the functions μ^u, σ^u and C^u are defined by

$$
\begin{aligned}
\mu^u(t,x) &= \mu(t,x,u), \\
\sigma^u(t,x) &= \sigma(t,x,u), \\
C^u(t,x) &= \sigma(t,x,u)\sigma(t,x,u)'.
\end{aligned}
$$

- For any control law $\mathbf{u}$, the functions $\mu^{\mathbf{u}}$, $\sigma^{\mathbf{u}}$, $C^{\mathbf{u}}(t,x)$ and $F^{\mathbf{u}}(t,x)$ are defined by

$$
\begin{aligned}
\mu^{\mathbf{u}}(t,x) &= \mu(t,x,\mathbf{u}(t,x)), \\
\sigma^{\mathbf{u}}(t,x) &= \sigma(t,x,\mathbf{u}(t,x)), \\
C^{\mathbf{u}}(t,x) &= \sigma(t,x,\mathbf{u}(t,x))\sigma(t,x,\mathbf{u}(t,x))', \\
F^{\mathbf{u}}(t,x) &= F(t,x,\mathbf{u}(t,x)).
\end{aligned}
$$

More notation

- For any fixed vector $u \in R^k$, the partial differential operator $\mathcal{A}^u$ is defined by

$$\mathcal{A}^u = \sum_{i=1}^{n} \mu_i^u(t,x)\frac{\partial}{\partial x_i} + \frac{1}{2}\sum_{i,j=1}^{n} C_{ij}^u(t,x)\frac{\partial^2}{\partial x_i \partial x_j}.$$

- For any control law $\mathbf{u}$, the partial differential operator $\mathcal{A}^{\mathbf{u}}$ is defined by

$$\mathcal{A}^{\mathbf{u}} = \sum_{i=1}^{n} \mu_i^{\mathbf{u}}(t,x)\frac{\partial}{\partial x_i} + \frac{1}{2}\sum_{i,j=1}^{n} C_{ij}^{\mathbf{u}}(t,x)\frac{\partial^2}{\partial x_i \partial x_j}.$$

- For any control law $\mathbf{u}$, the process $X^{\mathbf{u}}$ is the solution of the SDE

$$dX_t^{\mathbf{u}} = \mu\left(t, X_t^{\mathbf{u}}, \mathbf{u}_t\right)dt + \sigma\left(t, X_t^{\mathbf{u}}, \mathbf{u}_t\right)dW_t,$$

where

$$\mathbf{u}_t = \mathbf{u}(t, X_t^{\mathbf{u}})$$

Embedding the problem

For every fixed (t, x) the control problem $\mathcal{P}(t, x)$ is defined as the problem to maximize

$$E_{t,x}\left[\int_t^T F(s, X_s^{\mathbf{u}}, u_s)ds + \Phi\left(X_T^{\mathbf{u}}\right)\right],$$

given the dynamics

$$\begin{aligned} dX_s^{\mathbf{u}} &= \mu\left(s, X_s^{\mathbf{u}}, \mathbf{u}_s\right)ds + \sigma\left(s, X_s^{\mathbf{u}}, \mathbf{u}_s\right)dW_s, \\ X_t &= x, \end{aligned}$$

and the constraints

$$\mathbf{u}(s, y) \in U, \quad \forall (s, y) \in [t, T] \times R^n.$$

The original problem was $\mathcal{P}(0, x_0)$.

The optimal value function

- The **value function**

$$\mathcal{J} : R_+ \times R^n \times \mathcal{U} \to R$$

is defined by

$$\mathcal{J}(t, x, \mathbf{u}) = E\left[\int_t^T F(s, X_s^{\mathbf{u}}, \mathbf{u}_s) ds + \Phi\left(X_T^{\mathbf{u}}\right)\right]$$

given the dynamics above.

- The **optimal value function**

$$V : R_+ \times R^n \to R$$

is defined by

$$V(t, x) = \sup_{\mathbf{u} \in \mathcal{U}} \mathcal{J}(t, x, \mathbf{u}).$$

- We want to derive a PDE for V.

Assumptions

We assume:

- There exists an optimal control law $\hat{\mathbf{u}}$.
- The optimal value function V is regular in the sense that $V \in C^{1,2}$.
- A number of limiting procedures in the following arguments can be justified.

The Bellman Optimality Principle

Dynamic programming relies heavily on the following basic result.

Proposition: If $\hat{\mathbf{u}}$ is optimal on the time interval $[t, T]$ then it is also optimal on every subinterval $[s, T]$ with $t \leq s \leq T$.

Proof: Iterated expectations.

Basic strategy

To derive the PDE do as follows:

- Fix $(t, x) \in (0, T) \times R^n$.
- Choose a real number h (interpreted as a "small" time increment).
- Choose an arbitrary control law $\mathbf{u}$.

Now define the control law $\mathbf{u}^\star$ by

$$\mathbf{u}^\star(s, y) = \begin{cases} \mathbf{u}(s, y), & (s, y) \in [t, t+h] \times R^n \\ \hat{\mathbf{u}}(s, y), & (s, y) \in (t+h, T] \times R^n. \end{cases}$$

In other words, if we use $\mathbf{u}^\star$ then we use the arbitrary control $\mathbf{u}$ during the time interval $[t, t+h]$, and then we switch to the optimal control law during the rest of the time period.

Basic idea

The whole idea of DynP boils down to the following procedure.

- Given the point (t, x) above, we consider the following two strategies over the time interval $[t, T]$:

  I: Use the optimal law $\hat{\mathbf{u}}$.
  II: Use the control law $\mathbf{u}^{\star}$ defined above.

- Compute the expected utilities obtained by the respective strategies.

- Using the obvious fact that Strategy I is least as good as Strategy II, and letting h tend to zero, we obtain our fundamental PDE.

Strategy values

Expected utility for strategy I:

$$\mathcal{J}(t, x, \hat{\mathbf{u}}) = V(t, x)$$

Expected utility for strategy II:

- The expected utility for $[t, t+h)$ is given by

$$E_{t,x}\left[\int_t^{t+h} F\left(s, X_s^{\mathbf{u}}, \mathbf{u}_s\right) ds\right].$$

- Conditional expected utility over $[t+h, T]$, given (t, x):

$$E_{t,x}\left[V(t+h, X_{t+h}^{\mathbf{u}})\right].$$

- Total expected utility for Strategy II is

$$E_{t,x}\left[\int_t^{t+h} F\left(s, X_s^{\mathbf{u}}, \mathbf{u}_s\right) ds + V(t+h, X_{t+h}^{\mathbf{u}})\right].$$

Comparing strategies

We have trivially

$$V(t,x) \geq E_{t,x}\left[\int_t^{t+h} F\left(s, X_s^{\mathbf{u}}, \mathbf{u}_s\right) ds + V(t+h, X_{t+h}^{\mathbf{u}})\right].$$

Remark
We have equality above if and only if the control law $\mathbf{u}$ is an optimal law $\hat{\mathbf{u}}$.

Now use Itô to obtain

$$V(t+h, X_{t+h}^{\mathbf{u}}) = V(t,x)$$

$$+ \int_t^{t+h} \left\{ \frac{\partial V}{\partial t}(s, X_s^{\mathbf{u}}) + \mathcal{A}^{\mathbf{u}} V(s, X_s^{\mathbf{u}}) \right\} ds$$

$$+ \int_t^{t+h} \nabla_x V(s, X_s^{\mathbf{u}}) \sigma^{\mathbf{u}} dW_s,$$

and plug into the formula above.

We obtain

$$E_{t,x}\left[\int_t^{t+h}\left[F\left(s,X_s^{\mathbf{u}},\mathbf{u}_s\right)+\frac{\partial V}{\partial t}(s,X_s^{\mathbf{u}})+\mathcal{A}^{\mathbf{u}}V(s,X_s^{\mathbf{u}})\right]ds\right]\leq 0.$$

Going to the limit:
Divide by h, move h within the expectation and let h tend to zero.
We get

$$F(t,x,u)+\frac{\partial V}{\partial t}(t,x)+\mathcal{A}^{u}V(t,x)\leq 0,$$

Recall

$$F(t,x,u) + \frac{\partial V}{\partial t}(t,x) + \mathcal{A}^u V(t,x) \leq 0,$$

This holds for all $u = \mathbf{u}(t,x)$, with equality if and only if $\mathbf{u} = \hat{\mathbf{u}}$.

We thus obtain the **HJB equation**

$$\frac{\partial V}{\partial t}(t,x) + \sup_{u \in U} \left\{ F(t,x,u) + \mathcal{A}^u V(t,x) \right\} = 0.$$

The HJB equation

Theorem:
Under suitable regularity assumptions the follwing hold:

I: V satisfies the Hamilton–Jacobi–Bellman equation

$$\begin{aligned} \frac{\partial V}{\partial t}(t,x) + \sup_{u \in U} \left\{ F(t,x,u) + \mathcal{A}^u V(t,x) \right\} &= 0, \\ V(T,x) &= \Phi(x), \end{aligned}$$

II: For each $(t,x) \in [0,T] \times R^n$ the supremum in the HJB equation above is attained by $u = \hat{\mathbf{u}}(t,x)$.

Logic and problem

Note: We have shown that **if** V is the optimal value function, and **if** V is regular enough, **then** V satisfies the HJB equation. The HJB eqn is thus derived as a **necessary** condition, and requires strong *ad hoc* regularity assumptions.

Problem: Suppose we have solved the HJB equation. Have we then found the optimal value function and the optimal control law?

Answer: Yes! This follows from the **Verification tehorem**.

The Verification Theorem

Suppose that we have two functions $H(t,x)$ and $g(t,x)$, such that

- H is sufficiently integrable, and solves the HJB equation

$$\left\{\begin{array}{rcl} \dfrac{\partial H}{\partial t}(t,x) + \sup_{u \in U}\left\{F(t,x,u) + \mathcal{A}^u H(t,x)\right\} &=& 0, \\ H(T,x) &=& \Phi(x), \end{array}\right.$$

- For each fixed (t,x), the supremum in the expression

$$\sup_{u \in U}\left\{F(t,x,u) + \mathcal{A}^u H(t,x)\right\}$$

  is attained by the choice $u = g(t,x)$.

Then the following hold.

1. The optimal value function V to the control problem is given by
$$V(t,x) = H(t,x).$$
2. There exists an optimal control law $\hat{\mathbf{u}}$, and in fact
$$\hat{\mathbf{u}}(t,x) = g(t,x)$$

.

Handling the HJB equation

1. Consider the HJB equation for V.
2. Fix $(t,x) \in [0,T] \times R^n$ and solve, the static optimization problem
$$\max_{u \in U} \left[F(t,x,u) + \mathcal{A}^u V(t,x) \right].$$
Here u is the only variable, whereas t and x are fixed parameters. The functions F, μ, σ and V are considered as given.
3. The optimal $\hat{u}$, will depend on t and x, and on the function V and its partial derivatives. We thus write $\hat{u}$ as
$$\hat{\mathbf{u}} = \hat{\mathbf{u}}\,(t,x;V)\,. \tag{1}$$
4. The function $\hat{\mathbf{u}}\,(t,x;V)$ is our candidate for the optimal control law, but since we do not know V this description is incomplete. Therefore we substitute the expression for $\hat{u}$ into the PDE , giving us the PDE
$$\begin{aligned} \frac{\partial V}{\partial t}(t,x) + F^{\hat{\mathbf{u}}}(t,x) + \mathcal{A}^{\hat{\mathbf{u}}}\,(t,x)\,V(t,x) &= 0, \\ V(T,x) &= \Phi(x). \end{aligned}$$
5. Now we solve the PDE above! Then we put the solution V into expression (1). Using the verification theorem we can identify V as the optimal value function, and $\hat{u}$ as the optimal control law.

Making an Ansatz

- The hard work of dynamic programming consists in solving the highly nonlinear HJB equation
- There are no general analytic methods available for this, so the number of known optimal control problems with an analytic solution is very small indeed.
- In an actual case one usually tries to **guess** a solution, i.e. we typically make a parameterized **Ansatz** for V then use the PDE in order to identify the parameters.
- **Hint:** V often inherits some structural properties from the boundary function Φ as well as from the instantaneous utility function F.
- Most of the known solved control problems have, to some extent, been "rigged" in order to be analytically solvable.

The Linear Quadratic Regulator

$$\min_{u \in R^k} \quad E\left[\int_0^T \left\{X_t' Q X_t + u_t' R u_t\right\} dt + X_T' H X_T\right],$$

with dynamics

$$dX_t = \left\{AX_t + Bu_t\right\} dt + C dW_t.$$

We want to control a vehicle in such a way that it stays close to the origin (the terms $x'Qx$ and $x'Hx$) while at the same time keeping the "energy" $u'Ru$ small.

Here $X_t \in R^n$ and $\mathbf{u}_t \in R^k$, and we impose no control constraints on u.

The matrices Q, R, H, A, B and C are assumed to be known. We may WLOG assume that Q, R and H are symmetric, and we assume that R is positive definite (and thus invertible).

Handling the Problem

The HJB equation becomes

$$\left\{\begin{array}{rcl}\dfrac{\partial V}{\partial t}(t,x) & + & \inf_{u \in R^k} \left\{ x'Qx + u'Ru + [\nabla_x V](t,x)\left[Ax + Bu\right] \right\} \\ & + & \frac{1}{2}\sum_{i,j} \frac{\partial^2 V}{\partial x_i \partial x_j}(t,x)\left[CC'\right]_{i,j} = 0, \\ V(T,x) & = & x'Hx. \end{array}\right.$$

For each fixed choice of (t,x) we now have to solve the static unconstrained optimization problem to minimize

$$u'Ru + [\nabla_x V](t,x)\left[Ax + Bu\right].$$

The problem was:

$$\min_{u} \quad u'Ru + [\nabla_x V](t,x)\,[Ax + Bu]\,.$$

Since $R > 0$ we set the gradient to zero and obtain

$$2u'R = -(\nabla_x V)B,$$

which gives us the optimal u as

$$\hat{u} = -\frac{1}{2}R^{-1}B'(\nabla_x V)'.$$

Note: This is our candidate of optimal control law, but it depends on the unkown function V.

We now make an educated guess about the shape of V.

From the boundary function $x'Hx$ and the term $x'Qx$ in the cost function we make the Ansatz

$$V(t,x) = x'P(t)x + q(t),$$

where $P(t)$ is a symmetric matrix function, and $q(t)$ is a scalar function.

With this trial solution we have,

$$\begin{aligned}
\frac{\partial V}{\partial t}(t,x) &= x'\dot{P}x + \dot{q}, \\
\nabla_x V(t,x) &= 2x'P, \\
\nabla_{xx} V(t,x) &= 2P \\
\hat{u} &= -R^{-1}B'Px.
\end{aligned}$$

Inserting these expressions into the HJB equation we get

$$\begin{aligned}
& x'\left\{\dot{P} + Q - PBR^{-1}B'P + A'P + PA\right\}x \\
& + \dot{q} + tr[C'PC] = 0.
\end{aligned}$$

We thus get the following matrix ODE for P

$$\begin{cases} \dot{P} &= PBR^{-1}B'P - A'P - PA - Q, \\ P(T) &= H. \end{cases}$$

and we can integrate directly for q:

$$\begin{cases} \dot{q} &= -tr[C'PC], \\ q(T) &= 0. \end{cases}$$

The matrix equation is a **Riccati equation**. The equation for q can then be integrated directly.

Final Result for LQ:

$$\begin{aligned} V(t,x) &= x'P(t)x + \int_t^T tr[C'P(s)C]ds, \\ \hat{\mathbf{u}}(t,x) &= -R^{-1}B'P(t)x. \end{aligned}$$

2. Portfolio Theory

- Problem formulation.
- An extension of HJB.
- The simplest consutmption-investment problem.
- The Merton fund separation results.

Recap of Basic Facts

We consider a market with n assets.

$$\begin{aligned} S_t^i &= \text{price of asset No } i, \\ h_t^i &= \text{units of asset No } i \text{ in portfolio} \\ w_t^i &= \text{portfolio weight on asset No } i \\ X_t &= \text{portfolio value} \\ c_t &= \text{consumption rate} \end{aligned}$$

We have the relations

$$X_t = \sum_{i=1}^{n} h_t^i S_t^i, \quad w_t^i = \frac{h_t^i S_t^i}{X_t}, \quad \sum_{i=1}^{n} u_t^i = 1.$$

Basic equation:
Dynamics of self financing portfolio in terms of relative weights

$$dX_t = X_t \sum_{i=1}^{n} w_t^i \frac{dS_t^i}{S_t^i} - c_t dt$$

Simplest model

Assume a scalar risky asset and a constant short rate.

$$\begin{aligned} dS_t &= \alpha S_t dt + \sigma S_t dW_t \\ dB_t &= r B_t dt \end{aligned}$$

We want to maximize expected utility over time

$$\max_{w^0, w^1, c} \quad E\left[\int_0^T F(t, c_t)dt + \Phi(X_T)\right]$$

Dynamics

$$dX_t = X_t\left[u_t^0 r + w_t^1 \alpha\right] dt - c_t dt + w_t^1 \sigma X_t dW_t,$$

Constraints

$$\begin{aligned} c_t &\geq 0, \ \forall t \geq 0, \\ w_t^0 + w_t^1 &= 1, \ \forall t \geq 0. \end{aligned}$$

Nonsense!

What are the problems?

- We can obtain umlimited uttility by simply consuming arbitrary large amounts.
- The wealth will go negative, but there is nothing in the problem formulations which prohibits this.
- We would like to impose a constratin of type $X_t \geq 0$ but this is a **state constraint** and DynP does not allow this.

Good News:
DynP can be generalized to handle (some) problems of this kind.

Generalized problem

Let D be a nice open subset of $[0,T] \times R^n$ and consider the following problem.

$$\max_{u \in U} \quad E\left[\int_0^\tau F(s, X_s^{\mathbf{u}}, \mathbf{u}_s) ds + \Phi\left(\tau, X_\tau^{\mathbf{u}}\right)\right].$$

Dynamics:

$$\begin{aligned} dX_t &= \mu\left(t, X_t, u_t\right) dt + \sigma\left(t, X_t, u_t\right) dW_t, \\ X_0 &= x_0, \end{aligned}$$

The **stopping time** τ is defined by

$$\tau = \inf\left\{t \geq 0 \mid (t, X_t) \in \partial D\right\} \wedge T.$$

Generalized HJB

Theorem: Given enough regularity the follwing hold.

1. The optimal value function satisfies

$$\left\{\begin{array}{rcll} \dfrac{\partial V}{\partial t}(t,x) + \sup\limits_{u \in U} \left\{ F(t,x,u) + \mathcal{A}^u V(t,x) \right\} & = & 0, & \forall (t,x) \in D \\ V(t,x) & = & \Phi(t,x), & \forall (t,x) \in \partial D. \end{array}\right.$$

2. We have an obvious verification theorem.

Reformulated problem

$$\max_{c \geq 0,\ w \in R} \quad E\left[\int_0^\tau F(t, c_t)dt + \Phi(X_T)\right]$$

where

$$\tau = \inf\left\{t \geq 0 \ |X_t = 0\right\} \wedge T.$$

with notation:

$$\begin{aligned} w^1 &= w, \\ w^0 &= 1 - w \end{aligned}$$

Thus no constraint on w.

Dynamics

$$dX_t = w_t\left[\alpha - r\right] X_t dt + \left(rX_t - c_t\right) dt + w\sigma X_t dW_t,$$

HJB Equation

$$
\begin{aligned}
\frac{\partial V}{\partial t} + \sup_{c \geq 0, w \in R} \left\{ F(t,c) + wx(\alpha - r)\frac{\partial V}{\partial x} + (rx - c)\frac{\partial V}{\partial x} + \frac{1}{2}x^2 w^2 \sigma^2 \frac{\partial^2 V}{\partial x^2} \right\} &= 0, \\
V(T,x) &= 0, \\
V(t,0) &= 0.
\end{aligned}
$$

We now specialize (why?) to

$$F(t,c) = e^{-\delta t} c^{\gamma},$$

so we have to maximize

$$e^{-\delta t} c^{\gamma} + wx(\alpha - r)\frac{\partial V}{\partial x} + (rx - c)\frac{\partial V}{\partial x} + \frac{1}{2}x^2 w^2 \sigma^2 \frac{\partial^2 V}{\partial x^2},$$

Analysis of the HJB Equation

In the embedde static problem we maximize, over c and w,

$$e^{-\delta t} c^{\gamma} + wx(\alpha - r)\frac{\partial V}{\partial x} + (rx - c)\frac{\partial V}{\partial x} + \frac{1}{2}x^2 w^2 \sigma^2 \frac{\partial^2 V}{\partial x^2},$$

First order conditions:

$$\begin{aligned} \gamma c^{\gamma - 1} &= e^{\delta t} V_x, \\ w &= \frac{-V_x}{x \cdot V_{xx}} \cdot \frac{\alpha - r}{\sigma^2}, \end{aligned}$$

Ansatz:

$$V(t, x) = e^{-\delta t} h(t) x^{\gamma},$$

Because of the boundary conditions, we must demand that

$$h(T) = 0. \tag{2}$$

Given a V of this form we have (using $\cdot$ to denote the time derivative)

$$
\begin{aligned}
\frac{\partial V}{\partial t} &= e^{-\delta t}\dot{h}x^{\gamma} - \delta e^{-\delta t}hx^{\gamma}, \\
\frac{\partial V}{\partial x} &= \gamma e^{-\delta t}hx^{\gamma-1}, \\
\frac{\partial^2 V}{\partial x^2} &= \gamma(\gamma-1)e^{-\delta t}hx^{\gamma-2}.
\end{aligned}
$$

giving us

$$
\begin{aligned}
\hat{\mathbf{w}}(t,x) &= \frac{\alpha - r}{\sigma^2(1-\gamma)}, \\
\hat{\mathbf{c}}(t,x) &= xh(t)^{-1/(1-\gamma)}.
\end{aligned}
$$

Plug all this into HJB!

After rearrangements we obtain

$$x^{\gamma}\left\{\dot{h}(t)+Ah(t)+Bh(t)^{-\gamma/(1-\gamma)}\right\}=0,$$

where the constants A and B are given by

$$\begin{aligned} A &= \frac{\gamma(\alpha-r)^2}{\sigma^2(1-\gamma)}+r\gamma-\frac{1}{2}\frac{\gamma(\alpha-r)^2}{\sigma^2(1-\gamma)}-\delta \\ B &= 1-\gamma. \end{aligned}$$

If this equation is to hold for all x and all t, then we see that h must solve the ODE

$$\begin{aligned} \dot{h}(t)+Ah(t)+Bh(t)^{-\gamma/(1-\gamma)} &= 0, \\ h(T) &= 0. \end{aligned}$$

An equation of this kind is known as a **Bernoulli equation**, and it can be solved explicitly.

We are done.

Merton's Mutal Fund Theorems

1. The case with no risk free asset

We consider n risky assets with dynamics

$$dS_i = S_i \alpha_i dt + S_i \sigma_i dW, \quad i = 1, \dots, n$$

where W is Wiener in R^k. On vector form:

$$dS = D(S)\alpha dt + D(S)\sigma dW.$$

where

$$\alpha = \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix} \quad \sigma = \begin{bmatrix} \sigma_1 \\ \vdots \\ \sigma_n \end{bmatrix}$$

$D(S)$ is the diagonal matrix

$$D(S) = diag[S_1, \dots, S_n].$$

Wealth dynamics

$$dX = Xw'\alpha dt - cdt + Xw'\sigma dW.$$

Formal problem

$$\max_{c,w} \quad E\left[\int_0^{\tau} F(t, c_t)dt\right]$$

given the dynamics

$$dX = Xw'\alpha dt - cdt + Xw'\sigma dW.$$

and constraints

$$e'w = 1, \quad c \geq 0.$$

Assumptions:

- The vector α and the matrix σ are constant and deterministic.
- The volatility matrix σ has full rank so $\sigma\sigma'$ is positive definite and invertible.

Note: S does not turn up in the X-dynamics so V is of the form

$$V(t, x, s) = V(t, x)$$

The HJB equation is

$$
\left\{
\begin{array}{rcl}
\dfrac{\partial V}{\partial t}(t,x) + \sup\limits_{e'w=1,\ c\geq 0} \left\{ F(t,c) + \mathcal{A}^{c,w} V(t,x) \right\} & = & 0, \\
V(T,x) & = & 0, \\
V(t,0) & = & 0.
\end{array}
\right.
$$

where

$$
\mathcal{A}^{c,w} V = x w' \alpha \frac{\partial V}{\partial x} - c \frac{\partial V}{\partial x} + \frac{1}{2} x^2 w' \Sigma w \frac{\partial^2 V}{\partial x^2},
$$

and where the matrix Σ is given by

$$
\Sigma = \sigma \sigma'.
$$

The HJB equation is

$$\begin{cases} V_t(t,x) + \sup_{w'e=1,\ c\geq 0} \left\{ F(t,c) + (xw'\alpha - c)V_x(t,x) + \frac{1}{2}x^2 w'\Sigma w V_{xx}(t,x) \right\} &= 0, \\ V(T,x) &= 0, \\ V(t,0) &= 0. \end{cases}$$

where $\Sigma = \sigma\sigma'$.

If we relax the constraint $w'e = 1$, the Lagrange function for the static optimization problem is given by

$$L = F(t,c) + (xw'\alpha - c)V_x(t,x) + \frac{1}{2}x^2 w'\Sigma w V_{xx}(t,x) + \lambda\left(1 - w'e\right).$$

$$
\begin{aligned}
L &= F(t,c) + (xw'\alpha - c)V_x(t,x) \\
&+ \frac{1}{2}x^2 w'\Sigma w V_{xx}(t,x) + \lambda\,(1 - w'e)\,.
\end{aligned}
$$

The first order condition for c is

$$
\frac{\partial F}{\partial c}(t,c) = V_x(t,x).
$$

The first order condition for w is

$$
x\alpha' V_x + x^2 V_{xx} w'\Sigma = \lambda e',
$$

so we can solve for w in order to obtain

$$
\hat{w} = \Sigma^{-1}\left[\frac{\lambda}{x^2 V_{xx}} e - \frac{x V_x}{x^2 V_{xx}} \alpha\right].
$$

Using the relation $e'w = 1$ this gives λ as

$$
\lambda = \frac{x^2 V_{xx} + x V_x e'\Sigma^{-1}\alpha}{e'\Sigma^{-1}e},
$$

Inserting λ gives us, after some manipulation,

$$\hat{w} = \frac{1}{e'\Sigma^{-1}e}\Sigma^{-1}e + \frac{V_x}{xV_{xx}}\Sigma^{-1}\left[\frac{e'\Sigma^{-1}\alpha}{e'\Sigma^{-1}e}e - \alpha\right].$$

We can write this as

$$\hat{\mathbf{w}}(t) = g + Y(t)h,$$

where the fixed vectors g and h are given by

$$\begin{aligned} g &= \frac{1}{e'\Sigma^{-1}e}\Sigma^{-1}e, \\ h &= \Sigma^{-1}\left[\frac{e'\Sigma^{-1}\alpha}{e'\Sigma^{-1}e}e - \alpha\right], \end{aligned}$$

whereas Y is given by

$$Y(t) = \frac{V_x(t, X(t))}{X(t)V_{xx}(t, X(t))}.$$

We had

$$\hat{\mathbf{w}}(t) = g + Y(t)h,$$

Thus we see that the optimal portfolio is moving stochastically along the one-dimensional "optimal portfolio line"

$$g + sh,$$

in the $(n-1)$-dimensional "portfolio hyperplane" Δ, where

$$\Delta = \{w \in R^n \,| e'w = 1\}.$$

If we fix two points on the optimal portfolio line, say $w^a = g + ah$ and $w^b = g + bh$, then any point w on the line can be written as an affine combination of the basis points w^a and w^b. An easy calculation shows that if $w^s = g + sh$ then we can write

$$w^s = \mu w^a + (1-\mu)w^b,$$

where

$$\mu = \frac{s-b}{a-b}.$$

Mutual Fund Theorem

There exists a family of mutual funds, given by $w^s = g + sh$, such that

1. For each fixed s the portfolio w^s stays fixed over time.

2. For fixed a, b with $a \neq b$ the optimal portfolio $\hat{\mathbf{w}}(t)$ is, obtained by allocating all resources between the fixed funds w^a and w^b, i.e.

$$\hat{\mathbf{w}}(t) = \mu^a(t) w^a + \mu^b(t) w^b,$$

3. The relative proportions (μ^a, μ^b) of wealth allocated to w^a and $w^b$are given by

$$\begin{aligned} \mu^a(t) &= \frac{Y(t) - b}{a - b}, \\ \mu^b(t) &= \frac{a - Y(t)}{a - b}. \end{aligned}$$

The case with a risk free asset

Again we consider the standard model

$$dS = D(S)\alpha dt + D(S)\sigma dW(t),$$

We also assume the risk free asset B with dynamics

$$dB = rBdt.$$

We denote $B = S_0$ and consider portfolio weights $(w_0, w_1, \ldots, w_n)'$ where $\sum_0^n w_i = 1$. We then eliminate w_0 by the relation

$$w_0 = 1 - \sum_1^n w_i,$$

and use the letter w to denote the portfolio weight vector for the risky assets only. Thus we use the notation

$$w = (w_1, \ldots, w_n)',$$

Note: $w \in R^n$ without constraints.

HJB

We obtain

$$dX = X \cdot w'(\alpha - re)dt + (rX - c)dt + X \cdot w'\sigma dW,$$

where $e = (1, 1, \ldots, 1)'$.

The HJB equation now becomes

$$\begin{cases} V_t(t,x) + \sup_{c \geq 0, w \in R^n} \left\{ F(t,c) + \mathcal{A}^{c,w} V(t,x) \right\} &= 0, \\ V(T,x) &= 0, \\ V(t,0) &= 0, \end{cases}$$

where

$$\begin{aligned} \mathcal{A}^c V &= xw'(\alpha - re)V_x(t,x) + (rx - c)V_x(t,x) \\ &+ \frac{1}{2}x^2 w'\Sigma w V_{xx}(t,x). \end{aligned}$$

First order conditions

We maximize

$$F(t,c) + xw'(\alpha - re)V_x + (rx - c)V_x + \frac{1}{2}x^2 w'\Sigma w V_{xx}$$

with $c \geq 0$ and $w \in R^n$.

The first order conditions are

$$\begin{aligned} \frac{\partial F}{\partial c}(t,c) &= V_x(t,x), \\ \hat{w} &= -\frac{V_x}{xV_{xx}}\Sigma^{-1}(\alpha - re), \end{aligned}$$

with geometrically obvious economic interpretation.

Mutual Fund Separation Theorem

1. The optimal portfolio consists of an allocation between two fixed mutual funds w^0 and w^f.

2. The fund w^0 consists only of the risk free asset.

3. The fund w^f consists only of the risky assets, and is given by
$$w^f = \Sigma^{-1}(\alpha - re).$$

4. At each t the optimal relative allocation of wealth between the funds is given by

$$\begin{aligned} \mu^f(t) &= -\frac{V_x(t, X(t))}{X(t)V_{xx}(t, X(t))}, \\ \mu^0(t) &= 1 - \mu^f(t). \end{aligned}$$

3. The Martingale Approach

- Decoupling the wealth profile from the portfolio choice.
- Lagrange relaxation.
- Solving the general wealth problem.
- Example: Log utility.
- Example: The numeraire portfolio.
- Computing the optimal portfolio.
- The Merton fund separation theorems from a martingale perspective..

Problem Formulation

Standard model with internal filtration

$$
\begin{aligned}
dS_t &= D(S_t)\alpha_t dt + D(S_t)\sigma_t dW_t, \\
dB_t &= rB_t dt.
\end{aligned}
$$

Assumptions:

- Drift and diffusion terms are allowed to be arbitrary adapted processes.
- The market is **complete**.
- We have a given initial wealth x_0

Problem:

$$
\max_{h \in \mathcal{H}} \quad E^P\left[\Phi(X_T)\right]
$$

where

$$
\mathcal{H} = \{\text{self financing portfolios}\}
$$

given the initial wealth $X_0 = x_0$.

Some observations

- In a complete market, there is a unique martingale measure Q.

- Every claim Z satisfying the budget constraint

$$e^{-rT}E^{Q}\left[Z\right]=x_0,$$

is attainable by an $h \in \mathcal{H}$ and vice versa.

- We can thus write our problem as

$$\max_{Z} \quad E^{P}\left[\Phi(Z)\right]$$

subject to the constraint

$$e^{-rT}E^{Q}\left[Z\right]=x_0.$$

- We can forget the wealth dynamics!

Basic Ideas

Our problem was

$$\max_{Z} \quad E^P\left[\Phi(Z)\right]$$

subject to
$$e^{-rT}E^Q\left[Z\right] = x_0.$$

Idea I:

We can **decouple** the optimal portfolio problem:

- Finding the optimal wealth profile $\hat{Z}$.
- Given $\hat{Z}$, find the replicating portfolio.

Idea II:

- Rewrite the constraint under the measure P.
- Use Lagrangian techniques to relax the constraint.

Lagrange formulation

Problem:

$$\max_{Z} \quad E^P\left[\Phi(Z)\right]$$

subject to

$$e^{-rT}E^P\left[L_T Z\right] = x_0.$$

Here L is the likelihood process, i.e.

$$L_T = \frac{dQ}{dP}, \quad \text{on } \mathcal{F}_T$$

The Lagrangian of this is

$$\mathcal{L} = E^P\left[\Phi(Z)\right] + \lambda\left\{x_0 - e^{-rT}E^P\left[L_T Z\right]\right\}$$

i.e.

$$\mathcal{L} = E^P\left[\Phi(Z) - \lambda e^{-rT}L_T Z\right] + \lambda x_0$$

The optimal wealth profile

Given enough convexity and regularity we now expect, given the dual variable λ, to find the optimal Z by maximizing

$$\mathcal{L} = E^P \left[\Phi(Z) - \lambda e^{-rT} L_T Z \right] + \lambda x_0$$

over unconstrained Z, i.e. to maximize

$$\int_\Omega \left\{ \Phi(Z(\omega)) - \lambda e^{-rT} L_T(\omega) Z(\omega) \right\} dP(\omega)$$

This is a trivial problem!

We can simply maximize $Z(\omega)$ for each ω separately.

$$\max_z \quad \left\{ \Phi(z) - \lambda e^{-rT} L_T z \right\}$$

The optimal wealth profile

Our problem:

$$\max_{z} \quad \left\{ \Phi(z) - \lambda e^{-rT} L_T z \right\}$$

First order condition

$$\Phi'(z) = \lambda e^{-rT} L_T$$

The optimal Z is thus given by

$$\hat{Z} = G\left(\lambda e^{-rT} L_T\right)$$

where

$$G(y) = \left[\Phi'\right]^{-1}(y).$$

The dual varaiable λ is determined by the constraint

$$e^{-rT} E^P\left[L_T \hat{Z}\right] = x_0.$$

Example – log utility

Assume that

$$\Phi(x) = \ln(x)$$

Then

$$g(y) = \frac{1}{y}$$

Thus

$$\hat{Z} = G\left(\lambda e^{-rT} L_T\right) = \frac{1}{\lambda} e^{rT} L_T^{-1}$$

Finally λ is determined by

$$e^{-rT} E^P\left[L_T \hat{Z}\right] = x_0.$$

i.e.

$$e^{-rT} E^P\left[L_T \frac{1}{\lambda} e^{rT} L_T^{-1}\right] = x_0.$$

so $\lambda = x_0^{-1}$ and

$$\hat{Z} = x_0 e^{rT} L_T^{-1}$$

The optimal wealth process

From general theory we know that the normalized optimal wealth process

$$e^{-rt}\hat{X}_t$$

is a Q-martingale. We thus have

$$e^{-rt}\hat{X}_t = E^Q\left[e^{-rT}\hat{Z} \middle| \mathcal{F}_t \right]$$

so

$$X_t = e^{-r(T-t)}x_0^{-1}e^{rT}E^Q\left[L_T^{-1} \middle| \mathcal{F}_t \right]$$

Since $L = dQ/dP$, we have $L^{-1} = dP/dQ$ so L^{-1} is a Q martingale. We thus obtain

$$\hat{X}_t = x_0^{-1}e^{rt}L_t^{-1}$$

Log utility is myopic

Recall

$$\hat{X}_t = x_0^{-1} e^{rt} L_t^{-1}$$

This shows that the optimal portfolio for log utility does not depend on the choice of time horizon T. This portfolio is also known as the growth optimal portfolio.

Definition:
The **growth optimal portfolio** (GOP) is the portfolio which is optimal for log utility (for arbitrary terminal date T).

The Numeraire Portfolio

Standard approach:

- Choose a fixed numeraire (portfolio) N.
- Find the corresponding martingale measure, i.e. find Q^N s.t.

$$\frac{B}{N}, \quad \text{and} \quad \frac{S}{N}$$

  are Q^N -martingales.

Alternative approach:

- Choose a fixed measure Q.
- Find numeraire N such that $Q = Q^N$.

Special case:

- Set $Q = P$
- Find numeraire N such that $Q^N = P$ i.e. such that

$$\frac{B}{N}, \quad \text{and} \quad \frac{S}{N}$$

  are Q^N -martingales under the **objective** measure P.
- This N is the **numeraire portfolio**.

Log utility and the numeraire portfolio

Theorem:
Assume that X is GOP. Then X is the numeraire portfolio.

Proof:
We have to show that for an arbitrary asset price process S the process

$$Y_t = \frac{S_t}{X_t}$$

is a P martingale. From above we know that

$$X_t = x_0^{-1} e^{rt} L_t^{-1}$$

Thus

$$\frac{S_t}{X_t} = x_0^{-1} e^{-rt} S_t L_t$$

which is a P martingale, since $x_0^{-1} e^{-rt} S_t$ is a Q martingale (use Bayes' formula).

Back to general case: Computing L_T

We recall

$$\hat{Z} = G\left(\lambda e^{-rT} L_T\right).$$

The likelihood process L is computed by using Girsanov. We recall

$$dS_t = D(S_t)\alpha_t dt + D(S_t)\sigma_t dW_t,$$

We know from Girsanov that

$$dL_t = L_t \varphi_t^{\star} dW_t$$

so

$$dW_t = \varphi_t dt + dW_t^Q$$

where W^Q is Q-Wiener.

Thus

$$dS_t = D(S_t)\left\{\alpha_t + \sigma_t \varphi_t\right\} dt + D(S_t)\sigma_t dW_t^Q,$$

Computing L_T, continued

Recall

$$dS_t = D(S_t)\left\{\alpha_t + \sigma_t\varphi_t\right\}dt + D(S_t)\sigma_t dW_t^Q,$$

The kernel φ is determined by the martingale measure condition

$$\alpha_t + \sigma_t\varphi_t = \mathbf{r}$$

where

$$\mathbf{r} = \begin{bmatrix} r \\ \vdots \\ r \end{bmatrix}$$

Market completeness implies that σ_t is invertible so

$$\varphi_t = \sigma_t^{-1}\left\{\mathbf{r} - \alpha_t\right\}$$

and

$$L_T = \exp\left(\int_0^T \varphi_t dW_t - \frac{1}{2}\int_0^T \|\varphi_t\|^2 dt\right)$$

Finding the optimal portfolio

- We can easily compute the optimal wealth profile.
- How do we compute the optimal portfolio?

Recall:

$$dS_t = D(S_t)\alpha_t dt + D(S_t)\sigma_t dW_t,$$

wealth dynamics

$$dX_t = h_t^B dB_t + h_t^S dS_t$$

or

$$dX_t = X_t u_t^B r dt + X_t u_t^S D(S_t)^{-1} dS_t$$

where

$$h^S = (h^1, \dots, h^n), \quad u^S = (u^1, \dots, u^n)$$

We recall:

- $e^{-rt}X_t$ is a Q martingale.
- $X_T = \hat{Z}$

Thus the optimal wealth process is determined by

$$X_t = e^{-r(T-t)}E^Q\left[\hat{Z}\middle|\mathcal{F}_t\right]$$

We can write this as

$$X_t = e^{-r(T-t)}M_t$$

where the Q-martingale M is defined by

$$M_t = E^Q\left[\hat{Z}\middle|\mathcal{F}_t\right]$$

Recall

$$M_t = E^Q\left[\hat{Z}\middle|\mathcal{F}_t\right]$$

The martingale representation theorem gives us

$$dM_t = \xi_t dW_t^Q,$$

which gives us the Q-dynamics of X as

$$dX_t = rX_t dt + e^{-r(T-t)} dM_t$$

so

$$dX_t = rX_t dt + e^{-r(T-t)} \xi_t dW_t^Q.$$

On the other hand we have

$$\begin{aligned} dX_t &= rX_t dt + X_t u_t^S \sigma_t dW_t^Q \\ dX_t &= rX_t dt + h_t^S D(S_t) \sigma_t dW_t^Q \end{aligned}$$

Thus u^S and h_t^S are determined by

$$\begin{aligned} u_t^S &= \frac{e^{-r(T-t)}}{X_t} \xi_t \sigma_t^{-1} \\ h_t^S &= e^{-r(T-t)} \xi_t \sigma_t^{-1} D(S_t)^{-1}. \end{aligned}$$

and

$$u_t^B = 1 - u_t^S \mathbf{e}, \quad h_t^B = X_t - h_t^S S_t$$

How do we find ξ?

Recall

$$
\begin{aligned}
X_t &= e^{-r(T-t)} E^Q \left[\hat{Z} \middle| \mathcal{F}_t \right] \\
dX_t &= rX_t dt + e^{-r(T-t)} \xi_t dW_t^Q .
\end{aligned}
$$

We need to compute ξ.

In a Markovian framework this follows direcetly from the Itô formula.

Recall

$$
\hat{Z} = H(L_T) = G(\lambda L_T)
$$

where

$$
G = [\Phi']^{-1}
$$

and

$$
\begin{aligned}
dL_t &= L_t \varphi_t^\star dW_t, \\
dW &= \varphi dt + dW_t^Q
\end{aligned}
$$

so

$$
dL_t = L_t \|\varphi_t\|^2 dt + L_t \varphi_t^\star dW_t^Q
$$

Finding ξ

If the model is Markovian we have

$$\alpha_t = \alpha(S_t), \quad \sigma_t = \sigma(S_t), \quad \varphi_t = \sigma(S_t)^{-1}\left\{\alpha(S_t) - \mathbf{r}\right\}$$

so

$$\begin{aligned} X_t &= e^{-r(T-t)} E^Q\left[H(L_T) \middle| \mathcal{F}_t \right] \\ dS_t &= D(S_t)\sigma(S_t)dW_t^Q, \\ dL_t &= L_t \|\varphi(S_t)\|^2 dt + L_t \varphi(S_t)^\star dW_t^Q \end{aligned}$$

Thus we have

$$X_t = F(t, S_t, L_t)$$

where F is given as the solution to the Kolmogorov backward equation.

Kolmogorov

Recall

$$
\begin{aligned}
X_t &= F(t, S_t, L_t =) e^{-r(T-t)} E^Q \left[H(L_T) \middle| \mathcal{F}_t \right] \\
dS_t &= D(S_t)\mathbf{r}dt + D(S_t)\sigma(S_t)dW_t^Q, \\
dL_t &= L_t \|\varphi(S_t)\|^2 dt + L_t \varphi(S_t)^\star dW_t^Q
\end{aligned}
$$

$$
\begin{aligned}
F_t + L\|\varphi\|^2 F_L + \frac{1}{2}L^2\|\varphi\|^2 F_{LL} + F_S D(s)\mathbf{r} + \frac{1}{2} tr\left\{C(s)\right\} - rF &= 0, \\
F(T, s, L) &= H(L)
\end{aligned}
$$

where

$$
C(s) = \sigma^\star(s) D(s) F_{ss} D(s) \sigma(s)
$$

Finding ξ, contd.

We had

$$X_t = F(t, S_t, L_t)$$

and Itô gives us

$$dX_t = rX_t dt + \left\{F_S D(S_t)\sigma(S_t) + F_L L_t \varphi^\star(S_t)\right\} dW_t^Q$$

Recall

$$dX_t = rX_t dt + e^{-r(T-t)}\xi_t dW_t^Q$$

Thus

$$e^{-r(T-t)}\xi_t = F_S D(S_t)\sigma(S_t) + F_L L_t \varphi^\star(S_t).$$

and we obtain u^S and h^S from

$$\begin{aligned} u_t^S &= \frac{e^{-r(T-t)}}{X_t}\xi_t \sigma(S_t)^{-1} \\ h_t^S &= e^{-r(T-t)}\xi_t \sigma(S_t)^{-1} D(S_t)^{-1}. \end{aligned}$$

Mutual Funds – Martingale Version

We now assume constant parameters

$$\alpha(s) = \alpha, \quad \sigma(s) = \sigma, \quad \varphi(s) = \varphi$$

We recall

$$\begin{aligned} X_t &= E^Q \left[H(L_T) \middle| \mathcal{F}_t \right] \\ dL_t &= L_t \|\varphi\|^2 dt + L_t \varphi^\star dW_t^Q \end{aligned}$$

Now L is Markov so we have (without any S)

$$X_t = F(t, L_t)$$

Thus

$$e^{-r(T-t)} \xi_t = F_L L_t \varphi^\star, \quad u_t^S = \frac{F_L L_t}{X_t} \varphi^\star \sigma^{-1}$$

and we have fund separation with the fixed risky fund given by

$$w = \varphi^\star \sigma^{-1} = \left\{ \mathbf{r}^\star - \alpha^\star \right\} \left\{ \sigma \sigma^\star \right\}^{-1}.$$

4. Some stuff on Markov processes

Contents

- The infinitesimal generator.
- The Dynkin Theorem.
- The Kolmogorov backward equation.

The infinitesimal generator

Let X be a Markov process on R^n with internal filtration $\mathbf{F} = \mathbf{F}^X$.

Definition: The **domain** $\mathcal{D}$ is the set of bounded continuous mappings $f : R^n \rightarrow R$ such that the limit

$$\lim_{h \to 0} \frac{E_{t,x}\left[f(X_{t+h})\right] - f(x)}{h}$$

exists pointwise for every (t, x).

Definition: The **infinitesimal generator** $\mathcal{G}$ is the mapping $\mathcal{G} : \mathcal{D} \rightarrow C(R_+ \times R^n)$ defined by

$$\left(\mathcal{G}f\right)(t, x) = \lim_{h \to 0} \frac{E_{t,x}\left[f(X_{t+h})\right] - f(x)}{h}$$

Intuition

The infinitesimal generator gives us the "mean derivative" of the process $f(X_t)$. From the definition we have

$$E_{t,x}\left[f(X_{t+h})\right] = f(x) + \mathcal{G}f(t,x)h + o(h)$$

which suggests the informal interpretation

$$E_{t,x}\left[df(X_t)\right] = \mathcal{G}f(t,x)dt$$

Time dependence

For a function $f(t,x)$ which is also time dependent and C^1 in the t-variable we have

$$\lim_{h \to 0} \frac{E_{t,x}\left[f(t+h, X_{t+h})\right] - f(t,x)}{h} = \frac{\partial f}{\partial t}(t,x) + \mathcal{G} f(t,x)$$

or, alternatively,

$$E_{t,x}\left[df(t, X_t)\right] = \left(\frac{\partial f}{\partial t}(t,x) + \mathcal{G} f(t,x)\right) dt$$

The Dynkin Theorem

Theorem: Consider an arbitrary f in the domain of $\mathcal{G}$, and define the process M by

$$M_t = f(X_t) - \int_0^t \mathcal{G} f(X_s) ds$$

alternatively by

$$df(X_t) = \mathcal{G} f(X_t) dt + dM_t.$$

Then the following hold.

- M is a martingale.
- The process $f(X_t)$is a martingale if and only if $\mathcal{G} f = 0$.

Sketch of proof

Since we have

$$E_{t,x}\left[df(X_t)\right] = \mathcal{G}f(X_t)dt$$

it follows from the Markov property that we have

$$E\left[df(X_t) - \mathcal{G}f(X_t)dt \middle| \mathcal{F}_t\right] = 0.$$

Thus

$$df(X_t) - \mathcal{G}f(X_t)dt$$

is a martingale increment, so M is a martingale. The second part of the statement depends on the (deep) result that a martingale with continuous trajectories of bounded variation must be constant.

The Kolmogorov Backward Equation

Consider a mapping $\Phi : R^n \rightarrow R$ and define the function $f(t,x)$ by

$$f(t,x) = E_{t,x}\left[\Phi(X_T)\right]$$

Then f solves the boundary value problem

$$\begin{aligned} \frac{\partial f}{\partial t}(t,x) + \mathcal{G}f(t,x) &= 0, \\ f(T,x) &= \Phi(x) \end{aligned}$$

This is the Kolmogorov Backward Equation.

Proof

From the definition of f and the Markov property we have

$$f(t, X_t) = E\left[\Phi(X_T) \middle| X_t\right] = E\left[\Phi(X_T) \middle| \mathcal{F}_t\right]$$

thus the process $f(t, X_t)$ is a martingale and Kolmogorov now follows from Dynkin.

5. Filtering theory

- An investment problem.
- The non-linear FKK filtering equations.
- The SPDE for the conditional density.
- The Zakai equation for the unnormalized density.
- The Wonham filter.
- The Kalman filter.

An investment problem with stochastic rate of return

Model:

$$dS_t = S_t \alpha(Y_t) dt + S_t \sigma dW_t$$

W is scalar and Y is some factor process. We assume that (S, Y) is Markov and adapted to the filtration $\mathbf{F}$.

Wealth dynamics

$$dX_t = X_t \left[r + u_t \left(\alpha - r \right) \right] dt + u_t X_t \sigma dW_t$$

Objective:

$$\max_u \quad E^P \left[\Phi(X_T) \right]$$

Information structure:

- Complete information: We observe S and Y, so $u \in \mathbf{F}$
- Incomplete information: We only observe S, so $u \in \mathbf{F}^S$. We need **filtering theory**.

Filtering Theory – Setup

Given some filtration $\mathbf{F}$:

$$\begin{aligned} dX_t &= a_t dt + dM_t \\ dZ_t &= b_t dt + dW_t \end{aligned}$$

Here all processes are $\mathbf{F}$ adapted and

$$\begin{aligned} X &= \text{state process,} \\ Z &= \text{observation process,} \\ M &= \text{martingale w.r.t. } \mathbf{F} \\ W &= \text{Wiener w.r.t. } \mathbf{F} \end{aligned}$$

We assume (for the moment) that M and W are **independent**.

Problem:
Compute (recursively) the filter estimate

$$\widehat{X}_t = \Pi_t[X] = E\left[X_t \middle| \mathcal{F}_t^Z \right]$$

Typical example

A very commen example is given by

$$
\begin{aligned}
dX_t &= \mu(t, X_t)dt + \sigma(t, X_t)dV_t, \\
dZ_t &= b(t, X_t)dt + dW_t
\end{aligned}
$$

where W and V are Wiener.

The innovations process

Recall:

$$dZ_t = b_t dt + dW_t$$

Our best guess of b_t is $\hat{b}_t$, so the genuinely new information should be

$$dZ_t - \hat{b}_t dt$$

Definition:
The **innovations process** ν is defined by

$$\nu_t = dZ_t - \hat{b}_t dt$$

Theorem: The process ν is $\mathbf{F}^Z$-Wiener.

Proof: By Levy it is enough to show that

- ν is an $\mathbf{F}^Z$ martingale.
- $\nu_t^2 - t$ is an $\mathbf{F}^Z$ martingale.

I. ν is an $\mathbf{F}^Z$ martingale:

From definition we have

$$d\nu_t = \left(b_t - \widehat{b}_t\right) dt + dW_t \tag{3}$$

so

$$E_s^Z\left[\nu_t - \nu_s\right] = \int_s^t E_s^Z\left[b_u - \widehat{b}_u\right] du + E_s^Z\left[W_t - W_s\right]$$

$$= \int_s^t E_s^Z\left[E_u^Z\left[b_u - \hat{b}_u\right]\right] du + E_s^Z\left[E_s\left[W_t - W_s\right]\right] = 0$$

I. $\nu_t^2 - t$ is an $\mathbf{F}^Z$ martingale:

From Itô we have

$$d\nu_t^2 = 2\nu_t d\nu_t + \left(d\nu_t\right)^2$$

Here $d\nu$ is a martingale increment and from (3) it follows that $\left(d\nu_t\right)^2 = dt$.

Remark 1:
The innovations process gives us a Gram-Schmidt orthogonalization of the increasing family of Hilbert spaces

$$L^2(\mathcal{F}_t^Z); \quad t \geq 0.$$

Remark 2:
The use of Itô above requires general semimartingale integration theory, since we do not know a priori that ν is Wiener.

Filter dynamics

From the X dynamics we guess that

$$d\widehat{X}_t = \hat{a}_t dt + \text{martingale}$$

Definition: $dm_t = d\widehat{X}_t - \hat{a}_t dt$.

Proposition: m is an $\mathcal{F}_t^Z$ martingale.

Proof:

$$\begin{aligned} E_s^Z\left[m_t - m_s\right] &= E_s^Z\left[\hat{X}_t - \hat{X}_s\right] - E_s^Z\left[\int_s^t \widehat{a}_u du\right] \\ &= E_s^Z\left[X_t - X_s\right] - E_s^Z\left[\int_s^t \widehat{a}_u du\right] \\ &= E_s^Z\left[M_t - M_s\right] - E_s^Z\left[\int_s^t \left(a_u - \widehat{a}_u\right) du\right] \\ &= E_s^Z\left[E_s\left[M_t - M_s\right]\right] - E_s^Z\left[\int_s^t E_u^Z\left[a_u - \widehat{a}_u\right] du\right] = 0 \end{aligned}$$

Filter dynamics

We now have the filter dynamics

$$d\widehat{X}_t = \hat{a}_t dt + dm_t$$

where m is an $\mathcal{F}_t^Z$ martingale.

If the **innovations hypothesis**

$$\mathcal{F}_t^Z = \mathcal{F}_t^\nu$$

is true, then the martingale representation theorem would give us an $\mathcal{F}_t^Z$ adapted process h such that

$$dm_t = h_t d\nu_t \qquad (4)$$

The innovations hypothesis is not generally correct but FKK have proved that in fact (4) is always true.

Filter dynamics

We thus have the filter dynamics

$$d\widehat{X}_t = \widehat{a}_t dt + h_t d\nu_t$$

and it remains to determine the gain process h.

Proposition: The process h is given by

$$h_t = \widehat{X_t b_t} - \widehat{X}_t \widehat{b}_t$$

We give a slighty heuristic proof.

Proof sketch

From Itô we have

$$d(X_t Z_t) = X_t b_t dt + X_t dW_t + Z_t a_t dt + Z_t dM_t$$

using

$$d\widehat{X}_t = \widehat{a}_t dt + h_t d\nu_t$$

and

$$dZ_t = \hat{b}_t dt + d\nu_t$$

we have

$$d\left(\widehat{X}_t Z_t\right) = \widehat{X}_t \widehat{b}_t dt + \widehat{X}_t d\nu_t + Z_t \widehat{a}_t dt + Z_t h_t d\nu_t + h_t dt$$

Formally we also should have

$$E\left[d(X_t Z_t) - d\left(\widehat{X}_t Z_t\right) \middle| \mathcal{F}_t^Z \right] = 0$$

which gives us

$$\left(\widehat{X_t b_t} - \widehat{X}_t \widehat{b}_t - h_t\right) dt = 0.$$

The FKK filter equations

For the model

$$\begin{aligned} dX_t &= a_t dt + dM_t \\ dZ_t &= b_t dt + dW_t \end{aligned}$$

where M and W are independent, we have the Fujisaki-Kallianpur-Kunita (FKK) non-linear filter equations

$$\begin{aligned} d\widehat{X}_t &= \widehat{a}_t dt + \left\{ \widehat{X_t b_t} - \widehat{X}_t \widehat{b}_t \right\} d\nu_t \\ d\nu_t &= dZ_t - \widehat{b}_t dt \end{aligned}$$

Remark: It is easy to see that

$$h_t = E\left[\left(X_t - \widehat{X}_t \right) \left(b_t - \widehat{b}_t \right) \Big| \, \mathcal{F}_t^Z \right]$$

The general filter equations

For the model

$$\begin{aligned} dX_t &= a_t dt + dM_t \\ dZ_t &= b_t dt + \sigma_t dW_t \end{aligned}$$

where

- The process σ is $\mathcal{F}_t^Z$ adapted and positive.
- There is no assumption of independence between M and W.

we have the filter

$$\begin{aligned} d\widehat{X}_t &= \widehat{a}_t dt + \left[\widehat{D}_t + \frac{1}{\sigma_t} \left\{ \widehat{X_t b_t} - \widehat{X}_t \widehat{b}_t \right\} \right] d\nu_t \\ d\nu_t &= \frac{1}{\sigma_t} \left\{ dZ_t - \widehat{b}_t dt \right\} \\ dD_t &= \frac{d\langle M, W \rangle_t}{dt} \end{aligned}$$

Comment on $\langle M, W \rangle$

This requires semimartingale theory but there are two simple cases

- If M is Wiener then

$$d\langle M, W \rangle_t = dM_t dW_t$$

with usual multiplication rules.

- If M is a pure jump process then

$$d\langle M, W \rangle_t = 0.$$

Filtering a Markov process

Assume that X is Markov with generator $\mathcal{G}$. We want to compute $\Pi_t [f(X_t)]$, for some nice function f. Dynkin's formula gives us

$$df(X_t) = (\mathcal{G}f)(X_t)dt + dM_t$$

Assume that the observations are

$$dZ_t = b(X_t)dt + dW_t$$

where W is independent of X.

The filter equations are now

$$\begin{aligned} d\Pi_t [f] &= \Pi_t [\mathcal{G}f] dt + \{\Pi_t [fb] - \Pi_t [f] \Pi_t [b]\} d\nu_t \\ d\nu_t &= dZ_t - \Pi_t [b] dt \end{aligned}$$

Remark: To obtain $d\Pi_t [f]$ we need $\Pi_t [\mathcal{G}f]$, $\Pi_t [fb]$, and $\Pi_t [b]$. This leads generically to an infinite dimensional system of filter equations.

On the filter dimension

$$d\Pi_t[f] = \Pi_t[\mathcal{G}f]\,dt + \{\Pi_t[fb] - \Pi_t[f]\,\Pi_t[b]\}\,d\nu_t$$

- To obtain $d\Pi_t[f]$ we need $\Pi_t[\mathcal{G}f]$, $\Pi_t[fb]$, $\Pi_t[b]$.
- We apply the FKK equations to $\mathcal{G}f$, fb, and b.
- This leads to new filter estimates to determine and generically to an **infinite dimensional** system of filter equations.
- The filter equations are really equations for the **entire conditional distribution** of X.
- You can only expect the filter to be finite when the conditional distribution of X is finitely parameterized.
- There are only very few examples of finite dimensional filters.
- The most well known finite filters are the Wonham and the Kalman filters.

The SPDE for the conditional density

Recall the FKK equation

$$d\Pi_t[f] = \Pi_t[\mathcal{G}f]\,dt + \{\Pi_t[fb] - \Pi_t[f]\,\Pi_t[b]\}\,d\nu_t$$

Now **assume** that X has a conditional density process $p_t(x)$, with interpretation

$$p_t(x)dx = E\left[X_t \in dx \middle| \mathcal{F}_t^Z\right]$$

so

$$\Pi_t[f] = E\left[f(X_t)\middle| \mathcal{F}_t^Z\right] = \int_{R^n} f(x)p_t(x)dx$$

Using the pairing $\langle f, g\rangle = \int f(x)g(x)dx$ we can write FKK as

$$d\langle f, p_t\rangle = \langle \mathcal{G}f, p_t\rangle dt + \{\langle fb, p_t\rangle - \langle f, p_t\rangle\langle b, p_t\rangle\}\,d\nu_t$$

Recall

$$d\langle f, p_t \rangle = \langle \mathcal{G} f, p_t \rangle dt + \{ \langle fb, p_t \rangle - \langle f, p_t \rangle \langle b, p_t \rangle \} d\nu_t$$

We can now dualize this to obtain

$$d\langle f, p_t \rangle = \langle f, \mathcal{G}^\star p_t \rangle dt + \{ \langle f, bp_t \rangle - \langle f, p_t \rangle \langle b, p_t \rangle \} d\nu_t$$

Since this holds for all test functions f we have the following result.

Theorem: The density function $p_t(x)$ satisfies the following stochastic partial differential equation (SPDE)

$$dp_t(x) = \mathcal{G}^\star p_t(x) dt + p_t(x) \left\{ b(x) - \int_{R^n} b(y) p_t(y) dy \right\} d\nu_t$$

This SPDE is know as the **Kushner-Stratonovic equation**.

The Zakai equation

We consider the following model under a measure P.

$$\begin{aligned} dX_t &= a(X_t)dt + b(X_t)dV_t, \\ dZ_t &= h(X_t)dt + dW_t \end{aligned}$$

where V and W are independent Wiener processes.

The SPDE for $p_t(x)$ is quite messy. We now present an alternative along the following lines.

- Perform a Girsanov transformation from P to Q so that X and Z are independent under Q.

- Compute filtering estimates under Q. This should be very easy, due to the independence.

- Transform the filter estimetes back from Q to P, using the abstract Bayes Formula.

The Basic Construction

Consider a probability space $(Q, \mathcal{F}, V, Z)$ where V and Z are independent Wiener processes under Q. Define X by

$$dX_t = a(X_t)dt + b(X_t)dV_t$$

and define $\mathbf{F}$ by

$$\mathcal{F}_t = \mathcal{F}_t^Z \vee \mathcal{F}_\infty^V$$

Define the likelihood process L by

$$dL_t = h(X_t)L_t dZ_t, \quad L_0 = 1$$

and define P by $dP = L_t dQ$ on $\mathcal{F}_t$. From Girsanov we deduce that W, defined by

$$dZ_t = h(X_t)dt + dW_t$$

is $(P, \mathbf{F})$-Wiener. In particular it is independent of $\mathcal{F}_0 = \mathcal{F}_\infty^V$, so W and V are P-independent. It is also easy to see (how?) that (X, V) has the same distribution under P as under Q. Under P we now have our standard model.

The unnormalized estimate

Define $\Pi_t[f]$ as usual by

$$\Pi_t[f] = E^P\left[f(X_t) \middle| \mathcal{F}_t^Z \right].$$

We then have, from Bayes,

$$\Pi_t[f] = \frac{E^Q\left[L_t f(X_t) \middle| \mathcal{F}_t^Z \right]}{E^Q\left[L_t \middle| \mathcal{F}_t^Z \right]}$$

Now define $\sigma_t[f]$ by

$$\sigma_t[f] = E^Q\left[L_t f(X_t) \middle| \mathcal{F}_t^Z \right]$$

which gives us the **Kallianpur-Striebel formula**

$$\Pi_t[f] = \frac{\sigma_t[f]}{\sigma_t[1]}$$

We can view $\sigma_t[f]$ as an **unnormalized** filter estimate of $f(X_t)$, and we now define the SDE for $\sigma_t[f]$.

The Zakai Equation

We have

$$\sigma_t[f] = \Pi_t[f] \cdot \sigma_t[1]$$

By FKK we already have an expression for $d\Pi_t[f]$ and one can show that

$$d\sigma_t[1] = \Pi_t[h]\,\sigma_t[1]\,dZ_t$$

From Ito, and after lots of calculations, we have the following result.

Theorem: The unnormalized filter estimate satisfies the **Zakai Equation**

$$d\sigma_t[f] = \sigma_t[\mathcal{G}f]\,dt + \sigma_t[hf]\,dZ_t$$

The SPDE for the unnormalized density

Let us now assume that there exists an unnormalized density process $q_t(x)$ with interpretation

$$\sigma_t\left[f\right] = \int_{R^n} f(x)q_t(x)dx$$

Arguing as before we then obtain the following result.

Theorem: The unnormalized density Q satisfies the SPDE

$$dq_t(x) = \mathcal{G}^\star q_t(x)dt + h(x)q_t(x)dZ_t$$

This is a **much** nicer equation than the corresponding equation for $p_t(x)$, since

- It is linear in q_t whereas the SPDE for p_t is quadratic in p_t.
- The equation for q is driven directly by the observations process Z, rather than by the innovations process ν.

The Wonham filter

Assume that X is a continuous time Markov chain on the state space $\{1, \ldots, n\}$ with (constant) generator matrix H, i.e.

$$P\left(X_{t+h} = j \,|X_t = i\right) = H_{ij}h + o(h),$$

for $i \neq j$ and

$$H_{ii} = -\sum_{j \neq i} H_{ij}$$

Define the indicator processes by

$$\delta_i(t) = I\left\{X_t = i\right\}, \quad i = 1, \ldots, n.$$

Dynkin's Theorem gives us

$$d\delta_t^i = \sum_j H_{ji}\delta_j dt + dM_t^i, \quad i = 1, \ldots, n.$$

The Wonham filter

Recall

$$d\delta_t^i = \sum_j H_{ji}\delta_j dt + dM_t^i, \quad i = 1, \ldots, n.$$

Observations are

$$dZ_t = b(X_t)dt + dW_t.$$

The filter equations are

$$d\Pi_t\left[\delta_i\right] = \sum_j H_{ji}\Pi_t\left[\delta_j\right]dt + \left\{\Pi_t\left[\delta_i b\right] - \Pi_t\left[\delta_i\right]\Pi_t\left[b\right]\right\}d\nu_t$$

$$d\nu_t = dZ_t - \Pi_t\left[b\right]dt$$

We note that

$$b(X_t) = \sum_i b_i \delta_i(t)$$

so

$$\begin{aligned} \Pi_t[\delta_i b] &= b_i \Pi_t[\delta_i], \\ \Pi_t[b] &= \sum_j b_j \Pi_t[\delta_j] \end{aligned}$$

We finally have the Wonham filter

$$\begin{aligned} d\widehat{\delta_i} &= \sum_j H_{ji} \widehat{\delta_j} dt + \left\{ b_i \widehat{\delta_i} - \widehat{\delta_i} \sum_j b_j \widehat{\delta_j} \right\} d\nu_t, \\ d\nu_t &= dZ_t - \sum_j b_j \widehat{\delta_j} dt \end{aligned}$$

The Kalman filter

$$
\begin{aligned}
dX_t &= aX_t dt + c dV_t, \\
dZ_t &= X_t dt + dW_t
\end{aligned}
$$

W and V are independent Wiener

FKK gives us

$$
\begin{aligned}
d\Pi_t[X] &= a\Pi_t[X]\,dt + \left\{\Pi_t\left[X^2\right] - \left(\Pi_t[X]\right)^2\right\} d\nu_t \\
d\nu_t &= dZ_t - \Pi_t[X]\,dt
\end{aligned}
$$

We need $\Pi_t\left[X^2\right]$, so use Itô to get write

$$
dX_t^2 = \left\{2aX_t^2 + c^2\right\} dt + 2cX_t dV_t
$$

From FKK:

$$
\begin{aligned}
d\Pi_t\left[X^2\right] &= \left\{2a\Pi_t\left[X^2\right] + c^2\right\} dt \\
&+ \left\{\Pi_t\left[X^3\right] - \Pi_t\left[X^2\right]\Pi_t[X]\right\} d\nu_t
\end{aligned}
$$

Now we need $\Pi_t\left[X^3\right]$! Etc!

Define the conditional error variance by

$$H_t = \Pi_t \left[(X_t - \Pi_t [X])^2 \right] = \Pi_t \left[X^2 \right] - (\Pi_t [X])^2$$

Itô gives us

$$d (\Pi_t [X])^2 = \left[2a (\Pi_t [X])^2 + H^2 \right] dt + 2\Pi_t [X] H d\nu_t$$

and Itô again

$$\begin{aligned} dH_t &= \left\{ 2aH_t + c^2 - H_t^2 \right\} dt \\ &+ \left\{ \Pi_t \left[X^3 \right] - 3\Pi_t \left[X^2 \right] \Pi_t [X] + 2 (\Pi_t [X])^3 \right\} d\nu_t \end{aligned}$$

In **this particular case** we know (why?) that the distribution of X conditional on Z is Gaussian!

Thus we have

$$\Pi_t \left[X^3 \right] = 3\Pi_t \left[X^2 \right] \Pi_t [X] - 2 (\Pi_t [X])^3$$

so H is deterministic (as expected).

The Kalman filter

Model:

$$
\begin{aligned}
dX_t &= aX_t dt + c dV_t, \\
dZ_t &= X_t dt + dW_t
\end{aligned}
$$

Filter:

$$
\begin{aligned}
d\Pi_t [X] &= a\Pi_t [X] dt + H_t d\nu_t \\
\dot{H}_t &= 2aH_t + c^2 - H_t^2 \\
d\nu_t &= dZ_t - \Pi_t [X] dt
\end{aligned}
$$

$$
H_t = \Pi_t \left[\left(\Pi_t [X_t] - \Pi_t [X] \right)^2 \right]
$$

Remark: Because of the Gaussian structure, the conditional distribution evolves on a two dimensional submanifold. Hence a two dimensional filter.