Topics for Today

1.) Dynamic Programming (without Uncertainty)

- The 'Principle of Optimality' / Dynamic Consistency
- The Hamilton-Jacobi-Bellman (HJB) Equation
- Examples

2.) Dynamic Programming (with Uncertainty)

- Ito's Lemma
- Examples

Dynamic Programming

- The basic problem in finance has two components:

  a.) How much should I consume today vs. How much should I save?

  b.) How should I invest my savings?

- The Key aspect of these problems is that they involve intertemporal trade-offs. I can live like a King today, and suffer the consequences in the future. Or, I can live like a monk today and retire early.

- These are inherently dynamic problems, so we must learn how to solve dynamic optimization problems.

- We can interpret the solution to these problems from either a normative perspective (i.e., providing advice to investors), or from a positive perspective (i.e., as providing predictions about observed data).

- As always, the solutions will depend on the interaction between preferences and market opportunities.

- There has been a lot of interesting work recently on the structure of intertemporal preferences.

- Solving dynamic optimization problems is much more difficult than solving static optimization problems.

- Static optimization problems involve finding a single optimal choice, e.g., Should I buy an iPhone or an Android? They can be solved using calculus.

- Dynamic Optimization problems involve find an entire path. You can't choose myopically because today's choice has consequences for the future! These problems must be solved using functional analysis

calculus

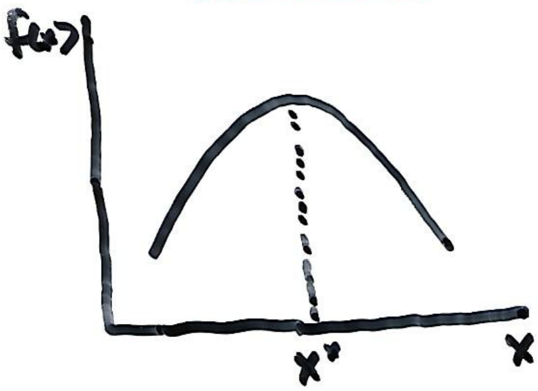

solution is a point

Functional Analysis

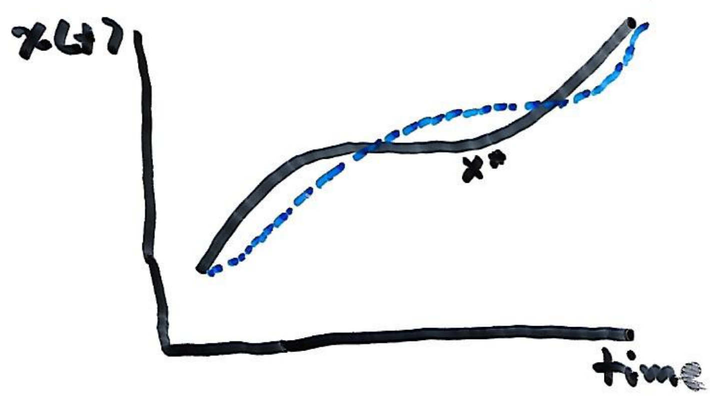

solution is a function

- There are 3 general approaches to these problems:
  1.) Calculus of Variations ⇒ Euler Eq.
  2.) Optimal Control Theory ⇒ Hamiltonian
  3.) Dynamic Programming ⇒ Bellman Eq.

- We are mainly going to use dynamic programming
- DP has distinct advantages in stochastic/uncertain settings, since it involves finding an optimal contingency plan. Future optimal choices are a function of the future state. You solve for the function rather than the choices themselves. [It clearly makes no sense to commit to a certain amount of future consumption until you know what your future income will be!]

- DP exploits the recursive structure of many dynamic optimization problems. When a problem is recursive, you can break it into 2 parts: Today & Tomorrow, exploiting the fact that Tomorrow's problem will be the same as Today's, except with new initial conditions determined by Today's decisions.

- Bellman called this the "Principle of Optimality"

Principle of Optimality

Under certain conditions, optimal paths have the property that whatever the initial conditions and controls were over some initial period, the controls over the remaining periods must be optimal for the "remainder problem" given the state resulting from the earlier controls.

- Or, more succinctly:

  1.) It is optimal to continue optimal paths.

  2.) All parts of an optimal path are optimal.

- By construction, solutions produced by DP are <u>dynamically consistent</u>. It is not necessary to reconsider an optimal policy

- Key condition for DP to produce Optimum:

  1.) Today's actions influence current + <u>future</u> returns, but not past returns

  examples:
  - consumption determines current + future utility, but not past utility
  - investment determines current + future profits, but not past profits

  or equivalently,

  2.) Current returns depend on current + <u>past</u> actions, but not on future actions.

  example: current profits do not depend on future investment

- This condition is often violated in 2-agent/strategic settings, where agents often have incentives to make promises or threats which they may not want to keep ex post.

The HJB Equation

- Although DP is especially useful in stochastic settings, we start with the deterministic case, which is easier.

- As usual, we start in discrete-time, then take continuous-time limits.

Problem

$$\max_u \int_0^T f(x,u,t)\,dt$$

subject to: 1.) $\dot{x} = g(x,u,t)$
2.) $x(0)$, $x(T)$ given

x = state (e.g., wealth)

u = control (e.g., consumption)

$\dot{x} = \frac{dx}{dt}$

g = state transition eq. (e.g., budget constraint)

- Start by defining the value function, $V(t_0, x_0)$, which is the optimized value of the above problem, given the state is x_0 and the time is t_0.

- $$V(t_0, x_0) = \max_u \int_{t_0}^{\infty} f(x,u,t) \quad \text{s.t. } \dot{x} = g(x,u,t), \; x(t_0) = x_0$$

- Note that we can break this integral into 2 pieces:

$$V(t_0, x_0) = \max_u \left\{ \int_{t_0}^{t_0+\Delta t} f\,dt + \int_{t_0+\Delta t}^{T} f\,dt \right\}$$

• If the Principle of Optimality applies,

$$V(t_0, x_0) = \max_{\substack{u \\ t_0 \le t \le t_0+\Delta t}} \left\{ \int_{t_0}^{t_0+\Delta t} f\,dt + \max_{\substack{u \\ t_0+\Delta t \le t \le T}} \left[\int_{t_0+\Delta t}^{T} f\,dt \right] \right\}$$

when choosing u in the future, you can neglect past f's

$\Rightarrow$

$$V(t_0, x_0) = \max_{\substack{u \\ t_0 \le t \le t_0+\Delta t}} \left\{ \int_{t_0}^{t_0+\Delta t} f\,dt + V(t_0+\Delta t, x_0+\Delta x) \right\}$$

For Δt small,

1.) $\int_{t_0}^{t_0+\Delta t} f\,dt \approx f \cdot \Delta t$

2.) $V(t_0+\Delta t, x_0+\Delta x) \approx V(t_0, x_0) + V_t \cdot \Delta t + V_x \cdot \Delta x$ } 1st. order Taylor Series Approx.

Sub-in, divide by Δt, and let $\Delta t \to 0$

1.

$$\boxed{-V_t = \max_u \left\{ f(x,u,t) + V_x(t,x) \cdot g(x,u,t) \right\}}$$ } HJB Eq.

$\Rightarrow u^* = h(x; V)$ } Policy Function (Recursive Representation of optimal Path)

Sub u^* in

2.

$$\boxed{-V_t = f(x, h(x;V), t) + V_x(t,x) \cdot g(x, h(x;V), t)}$$

- Both (1) and (2) are referred to as the Hamiton-Jacobi-Bellman Eq. Note, it is a nonlinear partial differential eq. in V.

- PDEs are the bread + butter of science. Unfortunately, they are notoriously difficult to solve. For linear PDEs, Fourier/Laplace transform methods provide a general strategy. For nonlinear PDEs, the only hope is a separation-of-variables, guess + verify approach.

- In general, there are many solutions to a PDE. Unique solutions are pinned down by boundary conditions.

Example

$$\min_u \int_0^\infty e^{-rt}(ax^2 + bu^2)\,dt$$

$$\text{s.t. } \dot{x} = c\cdot x + u$$

$$-V_t = \min_u \{ e^{-rt}(ax^2 + bu^2) + V_x \cdot (cx + u) \}$$

FOC(u): $2e^{-rt}bu + V_x = 0$

$$\Rightarrow \boxed{u = -\frac{V_x}{2b} e^{rt}}$$

sub the optimal u back into the HJB eq.

$$-V_t = e^{-rt}\left(ax^2 + \frac{1}{4b}V_x^2 e^{2rt}\right) + V_x \cdot \left(cx - \frac{V_x}{2b}e^{rt}\right)$$

$$\boxed{\text{Guess: } V(t,x) = e^{-rt} \cdot A x^2} \Rightarrow V_t = -r e^{-rt} A x^2$$

$$V_x = 2 e^{-rt} A x$$

sub these into HJB

$$r e^{-rt} A x^2 = a e^{-rt} x^2 + \frac{1}{b} e^{-rt} A^2 x^2 + 2c e^{-rt} A x^2 - \frac{2}{b} e^{-rt} A^2 x^2$$

Note: There is a common $e^{-rt}x^2$ term, which can be cancelled out. (This ability to cancel <u>defines</u> a correct guess!)

- We are left with the following quadratic equation for A. (We must select the positive root).

$$\frac{1}{b}A^2 + (r-2c)A - a = 0$$

- We then get the optimal feedback policy

$$\boxed{u = -\frac{V_x}{2b}e^{rt} = -\frac{A}{b}x}$$

A Useful Shortcut

- Many econ/finance problems feature an infinite horizon, (As a useful approximation), and time only enters as an exponential discount factor in the objective function.
- In these problems, time per se doesn't matter. (There is always tomorrow!) As a result, the HJB partial diff. eq. reduces to a much easier ordinary diff. eq.
- Here's how it works,

$$V(x) = \max_u \left\{ f(x,u)\cdot\Delta t + e^{-r\Delta t} V(x+\Delta x) \right\}$$

Note: V & f do not depend on t.

Note: $e^{-r\Delta t} \approx \frac{1}{1+r\cdot\Delta t}$

Multiply both-sides by $(1+r\Delta t)$ + expand $V(x+\Delta x)$

$$V(x)\cdot(1+r\cdot\Delta t) = f(x,u)\cdot\Delta t + f(x,u)\cdot r(\Delta t)^2 + V(x) + V_x\cdot\Delta X$$

Simplify, divide by Δt, let $\Delta t \to 0$ and drop higher-order terms

$$\underbrace{rV(x)}_{\text{riskless return / opportunity cost}} = \max_u \left\{ \underbrace{f(x,u)}_{\text{dividends (flow return)}} + \underbrace{V_x'(x)\cdot g(x,u)}_{\text{capital gain / loss}} \right\}$$

Stationary HJB Eq. (An ODE)

Stochastic Dynamic Programming

- Let's now consider a more realistic situation, where the future evolution of the state is uncertain.

- The previous analysis goes through, with two exceptions:

  1.) Since we don't know the future, we can only optimize expected returns/utility.

  2.) Doing the Taylor series approximation of $V(x)$ when obtaining the HJB eq. is a bit tricky when x is a function of Brownian motion. Since $dx \sim \sqrt{dt}$, we must expand to 2nd-order to get all the dt terms.

Ito's Lemma

Suppose x follows the Brownian motion with drift process,

$$dx = \mu \cdot dt + \sigma \cdot dw$$

Comments

1.) Remember, we cannot write

$$\frac{dx}{dt} = \mu + \sigma \frac{dw}{dt}$$

2.) A more mathematically correct notation would be

$$x_t = x_0 + \mu \int_0^t ds + \sigma \int_0^t dw_s$$

- Suppose we also have some function, $F(t,x)$, which depends on x (and t), and we want to approximate dF = total differential

- Normally, a 1st order approximation would be:

$$dF \approx \frac{\partial F}{\partial x} dx + \frac{\partial F}{\partial t} \cdot dt$$

- However, since $dx \sim \sqrt{dt}$, we must go out to 2nd order

$$dF \approx \frac{\partial F}{\partial x} dx + \frac{\partial F}{\partial t} dt + \frac{1}{2} \frac{\partial^2 F}{\partial x^2} (dx)^2 + \frac{1}{2} \frac{\partial^2 F}{\partial t^2} (dt)^2 + \frac{\partial F}{\partial x \partial t} dx \cdot dt$$

- Note that the final two terms are of order higher than dt, and so can be dropped from our 1st order approx.

- However, note that

$$(dx)^2 = \mu^2 (dt)^2 + 2\mu\sigma \, dt \cdot dw + \sigma^2 (dw)^2$$
$$\approx \sigma^2 dt$$

- Therefore,

$$\boxed{\begin{aligned} dF &= \frac{\partial F}{\partial x} dx + \frac{\partial F}{\partial t} dt + \frac{1}{2} \sigma^2 \frac{\partial^2 F}{\partial x^2} dt \\ &= \left(\frac{\partial F}{\partial x} \cdot \mu + \frac{1}{2} \sigma^2 \frac{\partial^2 F}{\partial x^2} + \frac{\partial F}{\partial t} \right) dt + \sigma \frac{\partial F}{\partial x} dw \end{aligned}}$$

> Ito's Lemma

Examples

① Suppose $dx = \mu\, dt + \sigma\, dW$ and define

$$Z = F(x) = e^x$$

Note: $F_x = e^x = z$

$F_{xx} = e^x = z$

$F_t = 0$

Plugging into Ito's Lemma,

$$dz = dF = F_x \cdot dx + \tfrac{1}{2}\sigma^2 F_{xx}\, dt$$

$$= z(\mu\, dt + \sigma\, dW) + \tfrac{1}{2}\sigma^2 \cdot z\, dt$$

$$= (\mu + \tfrac{1}{2}\sigma^2) z \cdot dt + \sigma Z\, dW$$

Geometric Brownian Motion

② Suppose $dx = \mu x \cdot dt + \sigma x\, dW$

Define $Z = F(x) = \log(x)$

Note: $F_x = \frac{1}{x}$

$F_{xx} = -\frac{1}{x^2}$

$F_t = 0$

Plugging into Ito's Lemma,

$$dz = dF = F_x\, dx + \tfrac{1}{2}\sigma^2 x^2 \left(-\tfrac{1}{x^2}\right) dt$$

$$= \tfrac{1}{x}\left[\mu x\, dt + \sigma x\, dW\right] - \tfrac{1}{2}\sigma^2 dt$$

$$= (\mu - \tfrac{1}{2}\sigma^2) dt + \sigma\, dW$$