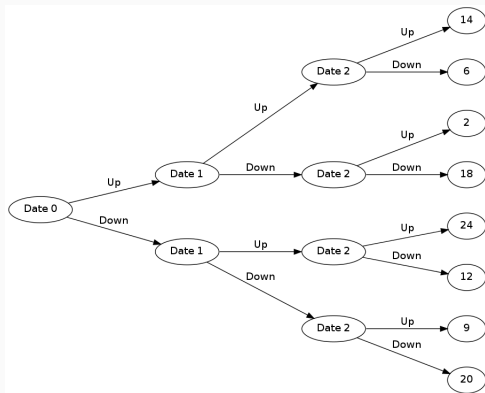


Chapter 9: Dynamic Portfolio Choice

Kerry Back
BUSI 521/ECON 505
Spring 2024
Rice University

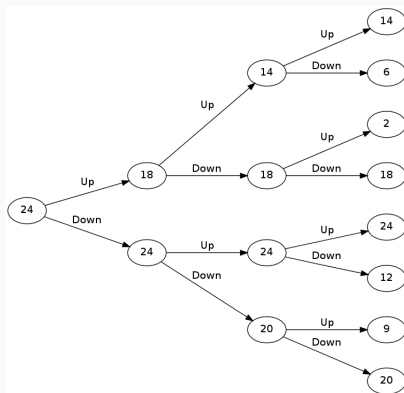
A Dynamic Decision Problem

The problem is to choose Up or Down at each date to maximize the terminal reward. The optimal sequence of decisions is Down-Up-Up.



Value Function

The value at each node is the maximum value that can be reached starting from that node. The value is also the max of the Up and Down values at each node.



Bellman Equation

- Label the nodes at each date t as $x = 1, \dots, 2^t$, starting from the bottom. Label the choice as $\pi = 0$ for Up and $\pi = 1$ for Down. Then

$$x_{t+1} = 2x_t - \pi_t$$

- Let $V_t(x)$ denote the value at node x at date t . The value is the max of the Up and Down values, so

$$V_t(x) = \max_{\pi \in \{0,1\}} V_{t+1}(2x - \pi) \quad (\star)$$

- Equation (\star) is called the Bellman equation.

Intermediate Rewards

- Suppose there is a reward $u_t(x, \pi)$ at each date t , given the node x at date t and the decision π taken at node x at date t . Suppose the objective is to maximize

$$\sum_{t=0}^2 u_t(x_t, \pi_t(x_t)) + u_3(x_3),$$

- Let $V_3(x) = u_3(x)$ and for $t < 3$, define

$$V_t(x) = \max_{\pi \in \{0,1\}} \left\{ u_t(x, \pi) + V_{t+1}(2x - \pi) \right\} \quad (**)$$

- V_t is the value function and $(**)$ is the Bellman equation with intermediate rewards.

Dynamic Programming under Uncertainty

For choice under uncertainty, the Bellman equation is

$$V_t(x) = \max_{\pi} \left\{ u_t(x, \pi) + E[V_{t+1}(X_{t+1}) \mid X_t = x] \right\} \quad (***)$$

Here X_{t+1} denotes the random state (node) at date $t+1$, the distribution of which may depend on the decision π and the state x at date t .

Markovian Model

- Assume X is a stationary Markov process with values in \mathbb{R}^k . This means that
 - the unconditional distribution of any finite set $(X_{t_1}, X_{t_2}, \dots, X_{t_k})$ is the same as the distribution of $(X_{t_1+s}, X_{t_2+s}, \dots, X_{t_k+s})$, for any s
 - the distribution of X_{t+1}, X_{t+2}, \dots conditional on all information at date t is the same as the distribution conditional on X_t only
- Assume the distribution of return vectors R_{t+1}, R_{t+2}, \dots and the distribution of labor income Y_{t+1}, Y_{t+2}, \dots conditional on all information at date t is the same as the distribution conditional on X_t only and are time homogeneous.
- In other words, X_t is a sufficient statistic for forecasting future values of X , R , and Y .

Bellman Equation for Portfolio Choice

- Return vector $R_t \in \mathbb{R}^n$ includes risk-free return if it exists.
- State variables are wealth W and X .
- Choice variables are consumption c and portfolio π . Require $\iota' \pi = 1$.
- Wealth evolves according to the intertemporal budget constraint:

$$W_{t+1} = (W_t - c)\pi' R_{t+1} + Y_{t+1}$$

- Bellman equation: $V_t(x, w) =$

$$\max_{c, \pi} \delta^t u(c) + E \left[V_{t+1} \left(X_{t+1}, (w - c)\pi' R_{t+1} + Y_{t+1} \right) \mid X_t = x \right]$$

Discounting Back Only to Date t

- V_t is the maximum expected sum of utilities at dates $t, t + 1, \dots$ measured in terms of date-0 utility, i.e., discounted to date 0.
- The same expected sum measured in terms of date- t utility is

$$J_t(x, w) = \delta^{-t} V_t(w, x)$$

- Multiplying by δ^{-t} in the Bellman equation gives: $J_t(x, w) =$

$$\max_{c, \pi} u(c) + \delta E \left[J_{t+1} \left(X_{t+1}, Y(w - c)\pi' R_{t+1} + Y_{t+1} \right) \mid X_t = x \right]$$

- If there is an infinite horizon, then there is always an infinite number of periods remaining and J does not depend on t .
- So the Bellman equation is: $J(x, w) =$

$$\max_{c, \pi} u(c) + \delta E \left[J \left(X_{t+1}, (w - c)\pi' R_{t+1} + Y_{t+1} \right) \mid X_t = x \right]$$

- Taking $t = 0$, we can also write it as: $J(x, w) =$

$$\max_{c, \pi} u(c) + \delta E \left[J \left(X_1, (w - c)\pi' R_1 + Y_1 \right) \mid X_0 = x \right]$$

Optimal Portfolio

Optimal Portfolio

- Consider the infinite horizon case for simplicity.
- Let c^* denote optimal consumption given (x, w) . Then, the optimal portfolio π^* solves

$$\max_{\pi} E \left[J \left(X_1, (w - c^*)\pi' R_1 + Y_1 \right) \mid X_0 = x \right]$$

- Compared to a single-period problem, the difference is that the objective function J depends on X_1 in addition to depending on end-of-period wealth.
- The optimal portfolio will usually, to some extent, hedge against adverse changes in X . We discuss “adverse changes” further later.

FOC for Optimal Portfolio

- Recall that we maximize over π subject to $\iota' \pi = 1$. Use a subscript w to denote a partial derivative with respect to wealth.
- FOC is

$$E \left[J_w \left(X_1, (w - c^*) \pi' R_1 + Y_1 \right) R_1 \mid X_0 = x \right] = \lambda \iota$$

where λ is the Lagrange multiplier for the constraint $\iota' \pi = 1$.

- Equivalently, J_w / λ at the optimum is a single-period SDF (for some λ).

Envelope Condition

- Consider the infinite horizon case for simplicity.
- Let $c^*(x, w)$, $\pi^*(x, w)$ achieve the max in the Bellman equation.
- The envelope theorem implies that

$$J_w(x, w) = u'(c^*(x, w))$$

- This is called the envelope condition: the marginal value of wealth equals the marginal utility of consumption at the optimum.

- Recall the Euler equation:

$$\frac{\delta^t u'(C_t)}{u'(C_0)}$$

is an SDF process, where C denotes optimal consumption.

- In the Markovian model, the envelope condition and Euler equation imply

$$\frac{\delta^t J_w(X_t, W_t)}{J_w(X_0, W_0)}$$

is an SDF process, where W denotes optimal wealth.

CRRA Utility

CRRA Utility

- Assume there is no labor income Y . If u is CRRA with risk aversion ρ , then V and J are CRRA in wealth w with the same risk aversion ρ .
- With an infinite horizon,
 - Log utility implies

$$J(x, w) = f(x) + \gamma \log w$$

for a function f and constant γ .

- Power utility implies

$$J(x, w) = f(x) \cdot \frac{1}{1 - \rho} w^{1 - \rho}$$

for a function f .

Example: Maximize Expected Log of Terminal Wealth

- Suppose the objective is to maximize $E[\log W_T]$ with no consumption prior to T . Suppose there is no labor income Y .
- Intertemporal budget constraint implies

$$W_{t+1} = W_t \pi'_t R_{t+1}$$

- Bellman equation is

$$V_t(x, w) = \max_{\pi} E \left[V_{t+1} \left(X_{t+1}, w \pi' R_{t+1} \right) \mid X_t = x \right]$$

- Terminal condition is $V_T(x, w) = \log w$.

Value Function and Optimal Portfolio

- Claim that $V_t(x, w) = f_t(x) + \log w$ for functions f_t .
- Clearly true for $t = T$ (with $f_T = 0$) Suppose it is true at $t + 1, t + 2, \dots, T$ for some date t . Then, the Bellman equation implies

$$\begin{aligned} V_t(x, w) &= \max_{\pi} E [f_{t+1}(X_{t+1}) + \log(w\pi'R_{t+1}) \mid X_t = x] \\ &= \log w + \max_{\pi} E [f_{t+1}(X_{t+1}) + \log(\pi'R_{t+1}) \mid X_t = x] \\ &\stackrel{\text{def}}{=} \log w + f_t(x) \end{aligned}$$

- The maximization in π is equivalent to

$$\max_{\pi} E [\log(\pi'R_{t+1}) \mid X_t = x]$$

- A log investor cares only about the distribution of the return $\pi'R_{t+1}$. He does not hedge against adverse changes in X .

Marginal Value of Wealth

- The “no hedging” result for log utility is due to the fact that J_w does not depend on x :

$$J = f_t(x) + \gamma_t \log w \quad \Rightarrow \quad J_w = \frac{\gamma_t}{w}$$

- Thus, the cross-partial $J_{wx_i} = 0$ for $i = 1, \dots, k$ for log utility. This is a borderline case.
 - For CRRA utility with $\rho < 1$, J_{wx_i} has the same sign as J_{x_i} .
 - For CRRA utility with $\rho > 1$, J_{wx_i} has the opposite sign of J_{x_i} .
- These facts affect hedging demands and risk premia (when investors have CRRA utility).

Interpreting the Cross-Partial

- Example: Suppose $J_{x_i} > 0$. Then an increase in x_i indicates “good times”. The investor gets higher expected future value when x_i is higher.
 - If $\rho < 1$, then $J_{wx_i} > 0$. Thus, J_w is higher when x_i is higher. The marginal value of wealth is higher in good times.
 - If $\rho > 1$, then $J_{wx_i} < 0$. Thus, J_w is higher when x_i is lower. The marginal value of wealth is higher in bad times.
- $\rho > 1$ seems more sensible (on other grounds also).

Proof of the Cross-Partial

- Consider the infinite horizon case for simplicity. Then,

$$J = f(x) \cdot \frac{1}{1 - \rho} w^{1-\rho}$$

$$J_{x_i} = f_{x_i}(x) \cdot \frac{1}{1 - \rho} w^{1-\rho}$$

$$J_{wx_i} = f_{x_i}(x) \cdot w^{-\rho}$$

- The signs of J_{x_i} and J_{wx_i} are the same if and only if $\rho < 1$.

Infinite Horizon

Dynamic Programming with an Infinite Horizon

- Whether the Bellman equation “works” must be verified. Cannot start at the end and work backwards. Cannot use contraction mapping argument (because the utility is unbounded).
- Some facts:
 - The value function always satisfies the Bellman equation.
 - The optimal policy achieves the maximum in the Bellman equation.
- However,
 - There can be other functions that also satisfy the Bellman equation.
 - Achieving the maximum in the Bellman equation (using the true value function) does not always guarantee optimality.

CRRA Utility with an Infinite Horizon and IID Returns

- Assume there are no state variables X . This means that returns are IID.
- Assume no borrowing ($C_t \leq W_t$) to rule out Ponzi schemes.
- We can show (without using the Bellman equation) that $J(w)$ is proportional to

$$\frac{1}{1-\rho} w^{1-\rho} \quad (\star)$$

- Under a parameter restriction (which rules out infinite discounted utility), we can show
 - There is a unique solution of the Bellman equation proportional to (\star) . Hence, that solution is the value function.
 - The policy achieving the max in the Bellman equation is optimal.

Assumptions

- Assume

$$\max_{\pi} E \left[\frac{1}{1-\rho} (\pi' R_t)^{1-\rho} \right] < \infty$$

- Define $B > 0$ by

$$\frac{1}{1-\rho} B^{1-\rho} = \max_{\pi} E \left[\frac{1}{1-\rho} (\pi' R_t)^{1-\rho} \right].$$

- Assume

$$\delta B^{1-\rho} < 1 \quad (**)$$

- Condition $(**)$ holds if $\rho > 1$ and there is a risk-free asset with return $R_f > 1$, because $B \geq R_f$, so $B^{1-\rho} \leq R_f^{1-\rho} < 1$.

Solution

- Optimal consumption is proportional to wealth: $c^*(w) = Aw$ for a constant $A > 0$.
- The optimal portfolio is the same in each period and the same as in a single-period model. It solves

$$\max_{\pi} E \left[\frac{1}{1-\rho} (\pi' R_t)^{1-\rho} \right]$$

- The value function has constant relative risk aversion. It is

$$J(w) = \frac{A^{-\rho}}{1-\rho} w^{1-\rho}$$

- A key step in the proof is to verify the “transversality condition:”

$$\lim_{T \rightarrow \infty} \delta^T E[J(W_T^*)] = 0$$

where W^* is the wealth process obtained by solving the Bellman equation each period.