Discrete-Time Dynamic Programming

Alexis Akira Toda*

October 31, 2017

Abstract

This note explains the classic Samuelson (1969) optimal consumption-portfolio problem. Other useful references might be Hakansson (1970) and Toda (2014).

1 Optimal portfolio problem

1.1 Model

Time is denoted by $t = 0, 1, \dots, T$ (maybe $T = \infty$). There are J assets indexed by $j \in J = \{1, \dots, J\}$. (In Samuelson (1969), $J = 2$, a risky and a riskless asset.) Starting from some initial wealth $w_0 > 0$, the investor wants to maximize the lifetime expected utility

$$\mathrm{E}_0 \sum_{t=0}^{T} \beta^t \frac{c_t^{1-\gamma}}{1-\gamma},$$

where $\beta > 0$ is the discount factor, $\gamma > 0$ is the relative risk aversion coefficient, and c_t is consumption. We say that such an investor has an *additive CRRA preference* (CRRA stands for constant relative risk aversion).

Let P_t^j be the per share price of asset j, and D_t^j be the dividend. The gross return of asset j between time t and $t+1$ is

$$R_{t+1}^j = \frac{P_{t+1}^j + D_{t+1}^j}{P_t^j},$$

that is, we adopt the convention that dividends are paid at the beginning of each period and prices are quoted at the end of each period after dividends are paid out. Let $\mathbf{R}_{t+1} = (R_{t+1}^1, \dots, R_{t+1}^J)$ be the vector of asset returns. For simplicity, assume that $\{\mathbf{R}_t\}_{t=0}^T$ is i.i.d. over time. However, the joint distribution of asset returns is arbitrary in the cross-section. Let θ_t^j be the fraction of wealth invested in asset j ($\theta_t^j > 0$ (< 0) corresponds to a long (short) position in asset j), and $\theta_t = (\theta_t^1, \dots, \theta_t^J)$ be the portfolio. By accounting, we have $\sum_{j=1}^J \theta_t^j = 1$.

The timing is as follows. At the beginning of period t, the asset returns $\mathbf{R}_t$ realizes and determines that period's initial wealth, w_t. Given this wealth, the investor chooses consumption c_t and portfolio θ_t. Let

$$R_{t+1}(\theta) = \mathbf{R}_{t+1}'\theta = \sum_{j=1}^J R_{t+1}^j \theta^j$$

*atoda@ucsd.edu

be the gross return on the portfolio between time t and $t+1$. Then the budget constraint is

$$w_{t+1} = R_{t+1}(\theta_t)(w_t - c_t) \geq 0.$$

1.2 Solution: finite horizon

Since we assumed that the returns are i.i.d., the only state variable is wealth.[1] With a slight abuse of notation, let $V_T(w)$ be the value function when T periods are left in the future (the investor has $T+1$ periods to live, including the current period). Then for $T \geq 1$ the Bellman equation is

$$V_T(w) = \max_{c,\theta} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \beta \operatorname{E}[V_{T-1}(w')] \,\middle|\, w' = R(\theta)(w-c) \right\}.$$

If $T = 0$, since the investor has no choice but to consume his wealth, we have

$$V_0(w) = \frac{w^{1-\gamma}}{1-\gamma}.$$

By homotheticity, we can show:

Lemma 1. *For each T, there exists a_T such that $V_T(w) = a_T \frac{w^{1-\gamma}}{1-\gamma}$.*

Proof. Let $c_0, \dots, c_T$ be the optimal consumption starting from wealth w, with value function $V_T(w)$. By the linearity of the budget constraint, if the initial wealth is λw (where $\lambda > 0$), then the consumption $\lambda c_0, \dots, \lambda c_T$ is feasible. By the homotheticity of the utility function, the associated lifetime utility will be $\lambda^{1-\gamma} V_T(w)$. Since the optimal value starting from wealth λw is $V_T(\lambda w)$, it follows that

$$\lambda^{1-\gamma} V_T(w) \leq V_T(\lambda w). \tag{1}$$

To show the reverse inequality, let $w' = \lambda w$ and $\lambda' = 1/\lambda$ in (1). Then we get

$$(1/\lambda')^{1-\gamma} V_T(\lambda' w') \leq V_T(w') \iff (\lambda')^{1-\gamma} V_T(w') \geq V_T(\lambda' w'). \tag{2}$$

Dropping the primes ($'$) in (2) and using (1), we get

$$\lambda^{1-\gamma} V_T(w) \geq V_T(\lambda w). \tag{3}$$

In particular, letting $\lambda = 1/w$, we get

$$V_T(w) = V_T(1) w^{1-\gamma} \equiv a_T \frac{w^{1-\gamma}}{1-\gamma}. \quad \square$$

Using the Lemma and substituting the budget constraint into the Bellman equation, we obtain

$$\begin{aligned} a_T \frac{w^{1-\gamma}}{1-\gamma} &= \max_{c,\theta} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \beta \operatorname{E}\left[a_{T-1} \frac{(R(\theta)(w-c))^{1-\gamma}}{1-\gamma} \right] \right\} \\ &= \max_{c} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \beta a_{T-1} (w-c)^{1-\gamma} \max_{\theta} \frac{1}{1-\gamma} \operatorname{E}[R(\theta)^{1-\gamma}] \right\}. \end{aligned} \tag{4}$$

From (4) we obtain the second result:

[1] See `https://sites.google.com/site/aatoda111/file-cabinet/172B_L08.pdf` for a short note on dynamic programming.

Proposition 2. *The optimal portfolio is* $\theta^* \in \arg\max_\theta \frac{1}{1-\gamma} \mathrm{E}[R(\theta)^{1-\gamma}]$.

For later computations, it is useful to define

$$\rho = \mathrm{E}[R(\theta^*)^{1-\gamma}]^{\frac{1}{1-\gamma}} = \max_\theta \mathrm{E}[R(\theta)^{1-\gamma}]^{\frac{1}{1-\gamma}}.$$

(The second equality uses the fact that $x \mapsto \frac{x^{1-\gamma}}{1-\gamma}$ is monotone.) Substituting the definition of ρ into (4), we get

$$a_T \frac{w^{1-\gamma}}{1-\gamma} = \max_c \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \beta a_{T-1}(w-c)^{1-\gamma} \frac{\rho^{1-\gamma}}{1-\gamma} \right\}. \tag{5}$$

Now the right-hand side is just a maximization in one variable, c. Since the objective function is concave in c, the first-order condition is necessary and sufficient. Therefore

$$\begin{aligned} & c^{-\gamma} - \beta a_{T-1}\rho^{1-\gamma}(w-c)^{-\gamma} = 0 \qquad (6) \\ \iff & c = (\beta a_{T-1}\rho^{1-\gamma})^{-\frac{1}{\gamma}}(w-c) \\ \iff & c = \frac{w}{1+(\beta a_{T-1}\rho^{1-\gamma})^{\frac{1}{\gamma}}}. \qquad (7) \end{aligned}$$

Substituting (7) into (5) and canceling $1-\gamma$, we get

$$\begin{aligned} a_T w^{1-\gamma} &= c^{1-\gamma} + \beta a_{T-1}\rho^{1-\gamma}(w-c)^{1-\gamma} \\ &= c^{1-\gamma} + c^{-\gamma}(w-c) && (\because (6)) \\ &= wc^{-\gamma} = \left(1+(\beta a_{T-1}\rho^{1-\gamma})^{\frac{1}{\gamma}}\right)^\gamma w^{1-\gamma} && (\because (7)) \\ \iff a_T^{1/\gamma} &= 1 + (\beta\rho^{1-\gamma})^{1/\gamma} a_{T-1}^{1/\gamma}. \end{aligned}$$

Letting $b_T = a_T^{1/\gamma}$, we obtain a first-order linear difference equation

$$b_T = 1 + (\beta\rho^{1-\gamma})^{1/\gamma} b_{T-1}.$$

The initial condition is $b_0 = a_0^{1/\gamma} = 1$. Therefore the solution is

$$b_T = \sum_{k=0}^T (\beta\rho^{1-\gamma})^{k/\gamma} = \frac{1-(\beta\rho^{1-\gamma})^{\frac{T+1}{\gamma}}}{1-(\beta\rho^{1-\gamma})^{\frac{1}{\gamma}}}.$$

Using (7), the optimal consumption rule is

$$c = \frac{w}{b_T} = \frac{1-(\beta\rho^{1-\gamma})^{\frac{1}{\gamma}}}{1-(\beta\rho^{1-\gamma})^{\frac{T+1}{\gamma}}} w.$$

Note that

$$1 = b_1 < b_2 < \cdots < b_T < \cdots,$$

so the longer the time horizon, the smaller fraction of wealth ($1/b_T$) you should consume. However, the portfolio is the same over time (at least with i.i.d. assumptions).

1.3 Solution: infinite horizon

The solution for the case with infinite horizon is basically the same. You might guess that all you need to do is to let $T \to \infty$ in the finite horizon, so (assuming $\beta\rho^{1-\gamma} < 1$) the coefficient of the value function is $b = 1/(1-(\beta\rho^{1-\gamma})^{1/\gamma})$ and the consumption rate is $c/w = 1-(\beta\rho^{1-\gamma})^{1/\gamma}$. This guess is correct, but there are technical subtleties.

To address the technical issues, let us consider the following more general problem:

$$\begin{aligned}&\max_{\{c_t\}_{t=0}^{\infty}} \mathrm{E}_0 \sum_{t=0}^{\infty} f_t(c_t, x_t)\\ &\text{subject to } \ c_t \in \Gamma_t(x_t), \ \ x_{t+1} = g_{t+1}(c_t, x_t).\end{aligned}$$

Here x_t is the state variable, c_t is the control variable, $f_t(c_t, x_t)$ is the flow utility, Γ_t is the constraint set, and g_{t+1} is the law of motion for the state variable. A similar (general) problem is discussed in Stokey and Lucas (1989), but since they put strong assumptions on f_t (such as $f_t(c,x) = \beta^t u(c,x)$ with u bounded), their results are practically inapplicable.[2] Clearly the optimal consumption-portfolio problem is a special case by reinterpreting the variables and functions.

I attack this problem as follows. Let

$$V_t^T(x) = \sup_{\{c_{t+s}\}_{s=0}^{T-1}} \mathrm{E}_t \sum_{s=0}^{T-1} f_{t+s}(c_{t+s}, x_{t+s})$$

be the T period value function starting at t and state variable $x = x_t$. Let $V_t^\infty(x) = \limsup_{T\to\infty} V_t^T(x)$ be the infinite horizon value function and

$$V_t^*(x) = \sup_{\{c_{t+s}\}_{s=0}^{\infty}} \mathrm{E}_t \sum_{s=0}^{\infty} f_{t+s}(c_{t+s}, x_{t+s})$$

be the true value function.

Lemma 3. *$V_t^*(x) \le V_t^\infty(x)$ always.*

Proof. Take any feasible consumption plan $\{c_{t+s}\}_{s=0}^{\infty}$ starting from x. Then by the definition of the value function, for any T we have

$$\mathrm{E}_t \sum_{s=0}^{T-1} f_{t+s}(c_{t+s}, x_{t+s}) \le V_t^T(x).$$

By the definition of the infinite horizon utility and infinite horizon value function, letting $T \to \infty$, we get

$$\begin{aligned}\mathrm{E}_t \sum_{s=0}^{\infty} f_{t+s}(c_{t+s}, x_{t+s}) &= \lim_{T\to\infty} \mathrm{E}_t \sum_{s=0}^{T-1} f_{t+s}(c_{t+s}, x_{t+s})\\ &\le \limsup_{T\to\infty} V_t^T(x) = V_t^\infty(x).\end{aligned}$$

Taking the supremum of the left-hand side over $\{c_{t+s}\}_{s=0}^{\infty}$, we get $V_t^*(x) \le V_t^\infty(x)$. ∎

[2] Dynamic programming in infinite horizon is still an active area of research. See Kamihigashi (2014) for recent developments.

I say that *the plan* $\{c_{t+s}\}_{s=0}^\infty$ *is recursively optimal* if it solves

$$V_{t+s}^\infty(x_{t+s}) = \max_{c\in\Gamma(x_{t+s})} \left\{f_{t+s}(c, x_{t+s}) + \mathrm{E}_{t+s} V_{t+s+1}^\infty(g_{t+s+1}(c, x_{t+s}))\right\}$$

for $s = 0, 1, \dots$. To use the results of the finite horizon dynamic programming, we want to show $V_t^*(x) = V_t^\infty(x)$. The following proposition provides a necessary and sufficient condition.

Proposition 4. $V_t^*(x) = V_t^\infty(x)$ *if and only if the transversality condition*

$$\limsup_{T\to\infty} \mathrm{E}_t[V_T^\infty(x_T)] \le 0$$

holds, where x_T *is the state variable obtained from a recursively optimal policy.*

Proof. Take a recursively optimal policy $\{c_{t+s}\}_{s=0}^\infty$. By definition, we have

$$V_t^\infty(x) = \mathrm{E}_t \sum_{s=0}^{T-1} f_{t+s}(c_{t+s}, x_{t+s}) + \mathrm{E}_t[V_T^\infty(x_T)]. \tag{8}$$

Letting $T \to \infty$ we obtain

$$\begin{aligned}
V_t^*(x) &\ge \liminf_{T\to\infty} \mathrm{E}_t \sum_{s=0}^{T-1} f_{t+s}(c_{t+s}, x_{t+s}) \\
&= \liminf_{T\to\infty} [V_t^\infty(x) - \mathrm{E}_t[V_T^\infty(x_T)]] \\
&= V_t^\infty(x) - \limsup_{T\to\infty} \mathrm{E}_t[V_T^\infty(x_T)] \\
&\ge V_t^*(x) - \limsup_{T\to\infty} \mathrm{E}_t[V_T^\infty(x_T)], && (\because \text{ Lemma})
\end{aligned}$$

so $\limsup_{T\to\infty} \mathrm{E}_t[V_T^\infty(x_T)] \ge 0$ always. If $\limsup_{T\to\infty} \mathrm{E}_t[V_T^\infty(x_T)] \le 0$, then actually $\limsup_{T\to\infty} \mathrm{E}_t[V_T^\infty(x_T)] = 0$, so all the above inequalities become equalities. Therefore $V_t^*(x) = V_t^\infty(x)$. Conversely, if $V_t^*(x) = V_t^\infty(x)$, then the recursively optimal policy is also optimal. Therefore letting $T \to \infty$ in (8), we obtain

$$\begin{aligned}
V_t^*(x) &= \lim_{T\to\infty} \mathrm{E}_t \sum_{s=0}^{T-1} f_{t+s}(c_{t+s}, x_{t+s}) + \lim_{T\to\infty} \mathrm{E}_t[V_T^*(x_T)] \\
&= V_t^*(x) + \lim_{T\to\infty} \mathrm{E}_t[V_T^*(x_T)],
\end{aligned}$$

so

$$\limsup_{T\to\infty} \mathrm{E}_t[V_T^\infty(x_T)] = \lim_{T\to\infty} \mathrm{E}_t[V_T^*(x_T)] = 0 \le 0. \quad \square$$

Now we apply this proposition to solve the infinite horizon optimal consumption-portfolio problem. Assuming $\beta\rho^{1-\gamma} < 1$, the infinite horizon value function is

$$V^\infty(w) = a\frac{w^{1-\gamma}}{1-\gamma},$$

where

$$a^{1/\gamma} = b = \frac{1}{1 - (\beta\rho^{1-\gamma})^{1/\gamma}} > 0.$$

The transversality condition $\limsup_{T\to\infty} \mathrm{E}_0[\beta^T V_T^\infty(w_T)] \le 0$ (here there is β^T because it is a discounted problem) is trivial if $\gamma > 1$ because then $V_T^\infty(w_T) \le 0$. If $0 < \gamma < 1$, then by the budget constraint we have

$$w_{t+1} = R_{t+1}(\theta^*)(w_t - c_t) \le R_{t+1}(\theta^*)w_t,$$

so $w_T \le w_0 \prod_{t=1}^T R_t(\theta^*)$. Taking the $(1-\gamma)$-th power and expectations, we get

$$\mathrm{E}_0[w_T^{1-\gamma}] \le w_0^{1-\gamma}\, \mathrm{E}[R(\theta^*)^{1-\gamma}]^T = w_0^{1-\gamma} \rho^{(1-\gamma)T}.$$

Hence

$$\mathrm{E}_0[\beta^T V_T^\infty(w_T)] \le \frac{a w_0^{1-\gamma}}{1-\gamma}(\beta\rho^{1-\gamma})^T \to 0$$

as $T \to \infty$, because $\beta\rho^{1-\gamma} < 1$.

2 Income fluctuation problem

With multiplicative risk (*e.g.*, random asset returns), it is convenient to work with CRRA utilities for tractability. With additive risk (*e.g.*, random labor income), CARA utilities are more convenient.

2.1 Model

Consider an agent with additive CARA utility

$$\mathrm{E}_0 \sum_{t=0}^\infty \beta^t u(c_t), \tag{9}$$

where $u(c) = -\mathrm{e}^{-\gamma c}/\gamma$ with absolute risk aversion $\gamma > 0$.[3] The agent can borrow or save at a gross risk-free rate $R > 1$. The agent is subject to income risk. The income process is given by

$$y_{t+1} = \rho y_t + \varepsilon_{t+1}, \tag{10}$$

where $0 \le \rho < 1$ and the error term ε_{t+1} is i.i.d. over time.[4] Letting w_t be the financial wealth at the beginning of time t (excluding current income), the budget constraint is

$$w_{t+1} = R(w_t - c_t + y_t).$$

The Bellman equation is

$$V(w,y) = \max_c \left\{u(c) + \beta\, \mathrm{E}[V(R(w-c+y), y')] \,|\, y' = \rho y + \varepsilon\right\}. \tag{11}$$

Since the CARA utility is defined on the entire real line, we assume that consumption can be negative.

[3] I focus on CARA preferences because it is tractable with additive shocks (Calvet, 2001; Wang, 2003, 2007; Angeletos and Calvet, 2005, 2006).

[4] Without loss of generality, we may assume that the AR(1) process (10) does not contain a constant term. This is because I have put no structure on the distribution of ε, so if there is a constant term we can always shift the distribution of ε so that the constant term is 0.

2.2 Solution

The following proposition gives a closed-form solution of the income fluctuation problem.

Proposition 5 (Wang, 2003)**.** *The value function and optimal consumption rule are given by*

$$V(w,y) = -\frac{1}{\gamma a}\mathrm{e}^{-\gamma(aw+b+dy)}, \tag{12a}$$

$$c(w,y) = aw + b + dy, \tag{12b}$$

where

$$\begin{aligned}
a &= 1 - 1/R,\\
b &= \frac{1}{\gamma(1-R)}\log \beta R\,\mathrm{E}[\mathrm{e}^{-\gamma\frac{R-1}{R-\rho}\varepsilon}],\\
d &= \frac{R-1)}{R-\rho}.
\end{aligned}$$

Proof. Again we prove by guess-and-verify. Substituting (12a) into the Bellman equation, we obtain

$$-\frac{1}{\gamma a}\mathrm{e}^{-\gamma(aw+b+dy)} = \max_c \left\{-\frac{1}{\gamma}\mathrm{e}^{-\gamma c} - \frac{\beta}{\gamma a}\,\mathrm{E}\left[\mathrm{e}^{-\gamma(aR(w-c+y)+b+dy')}\right]\right\}. \tag{13}$$

The first-order condition with respect to c is

$$\mathrm{e}^{-\gamma c} - \beta R\,\mathrm{E}\left[\mathrm{e}^{-\gamma(aR(w-c+y)+b+dy')}\right] = 0. \tag{14}$$

Substituting (14) into (13), we obtain

$$-\frac{1}{\gamma a}\mathrm{e}^{-\gamma(aw+b+dy)} = -\frac{1}{\gamma a}\left(a + \frac{1}{R}\right)\mathrm{e}^{-\gamma c}. \tag{15}$$

Comparing the coefficients, (15) trivially holds if $a = 1-1/R$ and $c = aw+b+dy$. In this case, $aR(w-c+y) = aw + (1-R)b + (1-R)(d-1)y$, so (14) becomes

$$\begin{aligned}
\mathrm{e}^{-\gamma(aw+b+dy)} &= \beta R\,\mathrm{E}\left[\mathrm{e}^{-\gamma(aw+(1-R)b+(1-R)(d-1)y+b+dy')}\right]\\
\iff \mathrm{e}^{-\gamma dy} &= \beta R\,\mathrm{E}\left[\mathrm{e}^{-\gamma((1-R)b+(1-R)(d-1)y+d(\rho y+\varepsilon)}\right].
\end{aligned} \tag{16}$$

Since (8) is an identity, comparing the coefficients of y, we obtain

$$d = (1-R)(d-1) + \rho d \iff d = \frac{R-1}{R-\rho}.$$

Substituting into (16), we obtain

$$\begin{aligned}
1 &= \beta R\,\mathrm{E}\left[\mathrm{e}^{-\gamma((1-R)b+\frac{R-1}{R-\rho}\varepsilon)}\right]\\
\iff b &= \frac{1}{\gamma(1-R)}\log \beta R\,\mathrm{E}[\mathrm{e}^{-\gamma\frac{R-1}{R-\rho}\varepsilon}].
\end{aligned}$$

□

Remark. Note that $\mathrm{E}[\mathrm{e}^{-\gamma \frac{R-1}{R-\rho}\varepsilon}]$ is the moment generating function $M_\varepsilon(s) = \mathrm{E}[\mathrm{e}^{s\varepsilon}]$ of ε evaluated at $s = -\gamma \frac{R-1}{R-\rho}$.

Remark. We can embed this income fluctuation problem into a general equilibrium model, which is a version of the Huggett (1993) model. Toda (2017) considers such a model with a VAR(1) income dynamics and shows that multiple equilibria are possible (although the equilibrium is unique in the AR(1) case). With multiple equilibria, comparative statics may go in different directions depending on the choice of the equilibrium. Toda (2017) provides an example in which *increasing* income risk is welfare improving!

Remark. For a proof of the transversality condition, see the appendix of Toda (2017).

References

George-Marios Angeletos and Laurent-Emmanuel Calvet. Incomplete-market dynamics in a neoclassical production economy. *Journal of Mathematical Economics*, 41(4-5):407–438, August 2005. doi:10.1016/j.jmateco.2004.09.005.

George-Marios Angeletos and Laurent-Emmanuel Calvet. Idiosyncratic production risk, growth and the business cycle. *Journal of Monetary Economics*, 53: 1095–1115, 2006. doi:10.1016/j.jmoneco.2005.05.016.

Laurent-Emmanuel Calvet. Incomplete markets and volatility. *Journal of Economic Theory*, 98(2):295–338, June 2001. doi:10.1006/jeth.2000.2720.

Nils H. Hakansson. Optimal investment and consumption strategies under risk for a class of utility functions. *Econometrica*, 38(5):587–607, September 1970. doi:10.2307/1912196.

Mark Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5-6):953–969, September-November 1993. doi:10.1016/0165-1889(93)90024-M.

Takashi Kamihigashi. Elementary results on solutions to the Bellman equation of dynamic programming: Existence, uniqueness, and convergence. *Economic Theory*, 56(2):251–273, 2014. doi:10.1007/s00199-013-0789-4.

Paul A. Samuelson. Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics*, 51(3):239–246, August 1969. doi:10.2307/1926559.

Nancy L. Stokey and Robert E. Lucas, Jr. *Recursive Methods in Economic Dynamics*. Harvard University Press, 1989.

Alexis Akira Toda. Incomplete market dynamics and cross-sectional distributions. *Journal of Economic Theory*, 154:310–348, November 2014. doi:10.1016/j.jet.2014.09.015.

Alexis Akira Toda. Huggett economies with multiple stationary equilibria. *Journal of Economic Dynamics and Control*, 84:77–90, November 2017. doi:10.1016/j.jedc.2017.09.002.

Neng Wang. Caballero meets Bewley: The permanent-income hypothesis in general equilibrium. *American Economic Review*, 93(3):927–936, June 2003. doi:10.1257/000282803322157179.

Neng Wang. An equilibrium model of wealth distribution. *Journal of Monetary Economics*, 54(7):1882–1904, October 2007. doi:10.1016/j.jmoneco.2006.11.005.