Practice Questions - Dynamic Programming

Thomas Walsh / Investment, Finance & Asset Pricing, 2025/6

1 Basic Cake-Eating Setup

(a) Sequential optimization problem:

The consumer's problem is:

$$\max_{\{c_1,c_2,c_3\}} \left\{\ln(c_1) + 0.9\ln(c_2) + 0.9^2\ln(c_3)\right\}$$

subject to:

$$\begin{aligned} W_2 &= W_1 - c_1 \\ W_3 &= W_2 - c_2 \\ 0 &= W_3 - c_3 \\ W_1 &= 100 \text{ (given)} \end{aligned}$$

Or equivalently, subject to the single resource constraint:

$$c_1 + c_2 + c_3 = W_1 = 100$$

The Sequential Lagrangian becomes:

$$\mathcal{L} = \ln(c_1) + \beta\ln(c_2) + \beta^2\ln(c_3) - \lambda_1(W_2 + c_1 - W_1) - \beta\lambda_2(W_3 + c_2 - W_2) - \beta^2\lambda_3(0 + c_3 - W_3)$$

(b) Recursive (Bellman) formulation:

- **State variable:** W_t (wealth at the beginning of period t)
- **Control variable:** c_t (consumption in period t) and W_{t+1} (next period wealth)
- **Value function:** Let $V^T(W)$ denote the value of having wealth W with T periods remaining. Then:

$$V^T(W) = \max_{c,W'} \left\{\ln(c) + \beta V^{T-1}(W')\right\}$$

subject to: $W' = W - c$, for $T = 3, 2, 1$

- With **terminal condition**: $V^0(W) = 0$ for all W (no value with zero periods left).
- The **terminal constraint** is: $W' = 0$ when $T = 1$ (end with no wealth).
- Alternatively, substituting the constraint:

$$V^T(W) = \max_{W'} \left\{ \ln(W - W') + \beta V^{T-1}(W') \right\}$$

- **Backward solution:**
  - Period $T = 1$ (one period left): $V^1(W) = \ln(W)$ (consume everything)
  - Period $T = 2$ (two periods left): $V^2(W) = \max_{W'}\{\ln(W - W') + \beta V^1(W')\}$
  - Period $T = 3$ (three periods left): $V^3(W) = \max_{W'}\{\ln(W - W') + \beta V^2(W')\}$
- With **terminal condition**: $V_4(W_4) = 0$ and $W_4 = 0$.

  The Recursive Lagrangian is:

$$\mathcal{L} = \ln(c) + \beta V^{T-1}(W') - \lambda(W' + c - W)$$

(c) Optimal consumption path using Euler equation:

FOCs:

$$[c]: \quad 1/c = \lambda \tag{1}$$

$$[W']: \quad \beta \frac{\partial V^{T-1}}{\partial W'} - \lambda = 0 \tag{2}$$

$$[EC]: \quad V_W = \lambda \Rightarrow V_{W'} = \lambda' \tag{3}$$

The Euler equation is:

$$\frac{1}{c} = \beta \frac{1}{c'} \implies c' = \beta c$$

For $u(c) = \ln(c)$, we have $u_c(c) = 1/c$.

This gives us:

$$c_2 = \beta c_1 = 0.9 c_1$$

$$c_3 = \beta c_2 = 0.9^2 c_1$$

Using the resource constraint:

$$c_1 + c_2 + c_3 = 100$$

$$c_1 + 0.9c_1 + 0.81c_1 = 100$$

$$c_1(1 + 0.9 + 0.81) = 100$$

$$c_1 = \frac{100}{2.71} = 36.90$$

Therefore:

$$c_1 = 36.90$$

$$c_2 = 0.9 \times 36.90 = 33.21$$

$$c_3 = 0.81 \times 36.90 = 29.89$$

(d) Verification:

$$c_1 + c_2 + c_3 = 36.90 + 33.21 + 29.89 = 100.00$$

The solution satisfies the resource constraint. Note that consumption starts above 100/3 and declines over time because the consumer is impatient ($\beta < 1$).

2 Envelope Condition Application

(a) First-order condition with respect to c:

Set up the Lagrangian:

$$\mathcal{L} = u(c) + \beta V(W') - \lambda(c + W' - W)$$

Taking the derivative with respect to c:

$$\frac{\partial \mathcal{L}}{\partial c} = u_c(c) - \lambda = 0$$

Therefore:

$$u_c(c) = \lambda$$

(b) Envelope condition:

The envelope theorem states that at the optimum, we can ignore the indirect effects of W through the choice variables. Taking the derivative of the value function with respect to W:

$$V_W(W) = \frac{\partial \mathcal{L}}{\partial W} = \lambda$$

Combining with part (a):

$$V_W(W) = u_c(c)$$

Intuition: The marginal value of wealth equals the marginal utility of consumption. An extra unit of wealth is optimally consumed, yielding $u_c(c)$ in utility.

(c) Euler equation:

By stationarity (the value function doesn't depend on time), rolling one period forward:

$$V_{W'}(W') = u_c(c')$$

From the FOC with respect to W':

$$\frac{\partial \mathcal{L}}{\partial W'} = \beta V_{W'}(W') - \lambda = 0$$

Therefore:

$$\lambda = \beta V_{W'}(W') = \beta u'(c')$$

Combining with $\lambda = u_c(c)$ from part (a):

$$u_c(c) = \beta u_c(c')$$

This is the **Euler equation**.

(d) Economic interpretation:

The envelope condition $V_W(W) = u_c(c)$ tells us that the marginal value of wealth is equal to the marginal utility of consumption at the optimum. This makes sense because:

– If you have one more unit of wealth, the optimal thing to do with it is consume it (or save it optimally)

- At the margin, the value of that extra wealth is exactly what you get from consuming it: $dV = V_W \times dW = u_c(c) \times dW$
- This only holds because the **consumer is already optimizing**—hence the "envelope" nature of the result

3 Investment with Productivity Shocks

(a) State and control variables:

State variables:

- θ: current productivity shock (exogenous state)
- K: current capital stock (endogenous state, chosen last period, so is sometimes called predetermined. It was a choice that was already made, so from todays perspective is taken as given.)

Control variable:

- K': next period capital
- I: investment (constraint forces them to obey law of motion, but from the perspective of the problem, I, K' are independent choices until the Lagrangian forces them to comply.

(b) Expectation given $\theta = \theta_L$:

The expectation is taken over the distribution of θ' conditional on current state $\theta = \theta_L$:

$$\mathbb{E}[V(\theta', K')|\theta = \theta_L] = \sum_{\theta' \in \{\theta_L, \theta_H\}} P(\theta'|\theta_L) \cdot V(\theta', K')$$

Using the first row of the transition matrix:

$$\begin{aligned}\mathbb{E}[V(\theta', K')|\theta_L] &= P(\theta_L|\theta_L) \cdot V(\theta_L, K') + P(\theta_H|\theta_L) \cdot V(\theta_H, K') \\ &= 0.7 \cdot V(\theta_L, K') + 0.3 \cdot V(\theta_H, K')\end{aligned}$$

(c) First-order condition:

Taking the derivative with respect to K':

$$\frac{\partial V(\theta, K)}{\partial K'} = -1 - \phi(K' - (1-\delta)K) + \beta\,\mathbb{E}[V_{K'}(\theta', K')] = 0$$

Rearranging:

$$1 + \phi(K' - (1-\delta)K) = \beta\,\mathbb{E}[V_{K'}(\theta', K')]$$

Or in terms of investment $I = K' - (1-\delta)K$:

$$1 + \phi I = \beta\,\mathbb{E}[V_{K'}(\theta', K')]$$

Defining marginal Q as $q = \beta\,\mathbb{E}[V_{K'}(\theta', K')]$:

$$q = 1 + \phi I$$

Interpretation: The marginal cost of investment (purchase cost plus marginal adjustment cost) equals the expected discounted marginal value of installed capital.

(d) Role of adjustment cost parameter ϕ:

The adjustment cost parameter ϕ creates a wedge between the price of capital goods and the marginal value of that capital:

- **When $\phi > 0$:** Investment > 0 and is smooth in economic conditions;
- **Higher ϕ:** Makes investment more costly, leading to:
  * Slower adjustment to shocks
  * Smoother investment over time, small adjustments prefered over large
- The term ϕI represents the marginal cost of adjusting the capital stock quickly

Economically, this captures frictions like installation costs, learning costs, or disruption to production when investment occurs rapidly.

4 Comparing Solution Methods

(a) Backward induction approach:

Step 1: Period $T = 10$ (terminal period)

- In the last period, consume everything: $c_{10} = W_{10}$
- Value function: $V_{10}(W_{10}) = \sqrt{W_{10}}$

Step 2: Period $T - 1 = 9$

- Solve: $V_9(W_9) = \max_{W_{10}}\{\sqrt{W_9 - W_{10}} + 0.95\sqrt{W_{10}}\}$
- FOC gives optimal policy $W_{10}(W_9)$
- Substitute back to get $V_9(W_9)$

Step 3: Continue backward

- For each $t = 8, 7, \ldots, 1$: solve $V_t(W_t) = \max_{W_{t+1}}\{\sqrt{W_t - W_{t+1}} + 0.95V_{t+1}(W_{t+1})\}$
- Use $V_{t+1}(\cdot)$ from previous step

Step 4: Forward simulation

- Given $W_1 = 50$, use policy $W_2 = h_1(W_1)$
- Continue forward: $W_3 = h_2(W_2)$, etc.

(b) Recursive vs Sequential Lagrangian:

Sequential Lagrangian:

- Choose *entire sequence* $\{c_1, c_2, \ldots, c_{10}\}$ simultaneously
- Set up one large Lagrangian with constraints for all periods
- Solve system of T Euler equations plus constraints
- Gives all choices at once

Recursive Approach:

- Decompose into sequence of *two-period* problems
- At each stage, choose current action given continuation value
- Solve backward from terminal condition
- Builds up solution iteratively

Key difference: Sequential approach solves for all periods jointly; recursive approach solves one period at a time using the solution from future periods.

(c) Advantage as $T \to \infty$:

As $T \to \infty$, the recursive approach becomes particularly advantageous because:

1. **Stationarity:** In infinite horizon, the value function and policy functions don't depend on calendar time: $V_t(W) = V_{t+1}(W) = V(W)$ for all t
2. **Single functional equation:** We solve *one* Bellman equation rather than infinitely many Euler equations
3. **Fixed point problem:** Find V^* such that $V^* = \mathcal{T}[V^*]$ where $\mathcal{T}$ is the Bellman operator
4. **Computational feasibility:** Sequential approach with $T = \infty$ is intractable; recursive approach is well-defined

In contrast, the sequential Lagrangian with $T = \infty$ requires solving an infinite system of equations simultaneously, which is not practically feasible.

(d) Principle of Optimality:

The **Principle of Optimality** (Bellman, 1957) states:

> *An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision.*

Justification for recursive formulation:

- Suppose you have an optimal sequence $\{c_1^*, c_2^*, \ldots, c_T^*\}$ starting from W_1
- After choosing c_1^*, you're left with $W_2 = W_1 - c_1^*$
- The principle says: $\{c_2^*, c_3^*, \ldots, c_T^*\}$ must be optimal starting from W_2
- Otherwise, you could improve the original sequence—contradiction!
- This justifies decomposing: $V_1(W_1) = u(c_1^*) + \beta V_2(W_2)$ where V_2 solves the remaining problem optimally

This is why we can write the problem recursively: the value from today onward equals the payoff today plus the optimally-solved continuation value.

5 Stationarity and Time Consistency

(a) Stationarity of value function:

A value function is **stationary** if it does not depend on calendar time. Mathematically:

$$V_t(W) = V_{t+s}(W) = V(W) \quad \forall t, s, W$$

In the infinite horizon problem with time-invariant parameters, the value function is stationary because:

- The problem looks identical from any period t
- There's no terminal date making "time until end" matter
- Parameters (β, utility function) don't change over time

Formally, if the Bellman equation is:

$$V(W) = \max_{W'}\{u(W - W') + \beta V(W')\}$$

then $V(\cdot)$ is the same function for all periods—it depends only on the state W, not on the time period t.

(b) Why the same policy function works:

Starting with $W_0 = 10$:

- Applying $h(\cdot)$: $W_1 = h(10)$
- Applying $h(\cdot)$ again: $W_2 = h(W_1) = h(h(10))$

This is optimal because:

1. The policy function $h(W)$ solves the Bellman equation for *any* starting wealth W
2. By the principle of optimality, after choosing $W_1 = h(W_0)$, the remaining problem is to maximize value from state W_1
3. Since the value function is stationary, the optimal policy from W_1 is given by the *same* function $h(\cdot)$
4. Stationarity means the decision rule doesn't depend on calendar time—only on the current state

Formally: If $h(\cdot)$ is the solution to:

$$h(W) = \arg\max_{W'} \{u(W - W') + \beta V(W')\}$$

then $h(\cdot)$ is the optimal policy for *any* period and *any* state W. This is time consistency: the plan made today remains optimal tomorrow.

6 Marginal Q and Investment

(a) Economic interpretation of marginal Q:

Marginal Q, denoted q, represents the **shadow value of installed capital**. Specifically:

$q =$ Marginal value to the firm of one additional unit of installed capital

Components of interpretation:

- q is the present discounted value of the marginal product of capital over its lifetime
- It includes both direct productivity benefits and the value of non-depreciated capital in future periods
- When $q > 1$: marginal value exceeds purchase cost $\rightarrow$ invest
- When $q < 1$: marginal value below purchase cost $\rightarrow$ don't invest (or disinvest if possible)
- When $q = 1$: marginal benefit equals marginal cost $\rightarrow$ indifferent

From the equation:

$$q = \beta\, \mathbb{E}[\pi_{K'}(\theta', K') + (1 - \delta)q']$$

- $\pi_{K'}(\theta', K')$: additional profit from marginal unit of capital next period
- $(1 - \delta)q'$: resale value of non-depreciated capital (valued at next period's shadow price)
- $\beta\, \mathbb{E}[\cdot]$: discounting and expectation over uncertain productivity

(b) Optimal investment calculation:

Given $q = 1.2$ and $\phi = 2$, use:

$$q = 1 + \phi I$$

Solving for I:

$$1.2 = 1 + 2 \cdot I$$
$$0.2 = 2I$$
$$I = 0.1$$

Interpretation: The firm invests 0.1 units of capital. The marginal unit of capital is worth 20% more than its purchase cost ($q = 1.2$), which justifies incurring the marginal adjustment cost of $\phi I = 2 \times 0.1 = 0.2$.

(c) Effect of high expected productivity:

If $\theta' = \theta_H$ with probability 1, then:

$$q = \beta[\pi_{K'}(\theta_H, K') + (1 - \delta)q']$$

Effects on current investment:

1. **Direct effect:** $\pi_{K'}(\theta_H, K')$ is higher because high productivity increases marginal product of capital
2. **Continuation value effect:** q' will also be higher if high productivity is persistent
3. **Overall:** q increases, which means:
$$I = \frac{q - 1}{\phi}$$
increases, so the firm invests more today

Economic intuition:

- Expecting high productivity tomorrow makes capital more valuable
- The firm wants to have more capital in place to take advantage of favorable conditions
- With adjustment costs, it's optimal to invest *today* in anticipation of high future productivity
- This creates "forward-looking" investment behavior: current investment responds to expected future conditions

(d) Investment when $q < 1$:

From $I = (q - 1)/\phi$:

If $q < 1$, then $I < 0$ (negative investment = disinvestment)

Is this economically sensible?

- **Interpretation:** The firm wants to reduce its capital stock
- $q < 1$ means marginal capital is worth less than its replacement cost
- Capital is too high relative to productivity/demand
- **Mechanism:** The firm lets capital depreciate and doesn't fully replace it

Model limitations:

- In this model, negative investment means: $K' < (1-\delta)K$
- This is possible by simply not replacing depreciated capital
- If we wanted active disinvestment (selling capital): need to allow $K' < (1-\delta)K$ more explicitly, possibly with asymmetric adjustment costs
- In reality: ϕ might differ for investment vs. disinvestment (easier to not buy than to sell)

When does $q < 1$ occur?

- Persistently low productivity shocks
- Expectation of low future profits
- Firm has over-invested in the past
- Industry decline

7 Discretization for Numerical Solution

(a) Grid creation:

MATLAB:

```
k_min = 0.01;  % avoid zero for numerical stability
k_max = 50;
N_k = 100;
k_grid = linspace(k_min, k_max, N_k);
```

Python (NumPy):

```
import numpy as np

k_min = 0.01

k_max = 50

N_k = 100

k_grid = np.linspace(k_min, k_max, N_k)
```

This creates a uniformly spaced grid with 100 points from 0.01 to 50.

Alternative: Use non-uniform grids (e.g., more points near 0) for better accuracy:

```
k_grid = k_min + (k_max - k_min) * (linspace(0,1,N_k).^2);
```

(b) Value function iteration algorithm:

Algorithm Structure:

Step 1: Initialize

- Set parameters: $\alpha = 0.3$, $\delta = 0.1$, $\beta = 0.96$, $k_{max} = 50$
- Create grid: $k_{grid} \in [0.01, 50]$ with $N_k = 100$ points
- Initialize value function: $V_0(k) = 0$ for all $k \in k_{grid}$ (or use better guess)
- Set tolerance: $tol = 10^{-6}$, max iterations: $maxiter = 1000$

Step 2: Iterate until convergence

While $\|V_{n+1} - V_n\| > tol$ and $iter < maxiter$:

For each $k_i \in k_{grid}$:

- Initialize: $V_{best} = -\infty$
- For each $k'_j \in k_{grid}$:
  * Compute flow payoff: $\pi = k_i^\alpha - k'_j + (1-\delta)k_i$
  * If $\pi < 0$: skip (infeasible choice)
  * Compute continuation: $V_{cont} = \beta V_n(k'_j)$
  * Compute total value: $V_{temp} = \pi + V_{cont}$
  * If $V_{temp} > V_{best}$:
    · Update: $V_{best} = V_{temp}$
    · Store policy: $g(k_i) = k'_j$
- Set: $V_{n+1}(k_i) = V_{best}$

Update: $V_n \leftarrow V_{n+1}$

Increment: $iter \leftarrow iter + 1$

Step 3: Return solution

- Value function: $V^*(k)$
- Policy function: $k'(k)$

(c) Checking for convergence:

Several convergence criteria can be used:

1. Supremum norm (most common):

$$\|V_{n+1} - V_n\|_\infty = \max_i |V_{n+1}(k_i) - V_n(k_i)| < tol$$

In code:

```
diff = max(abs(V_new - V_old));
if diff < tol
    break;
end
```

2. Euclidean norm:

$$\|V_{n+1} - V_n\|_2 = \sqrt{\sum_i (V_{n+1}(k_i) - V_n(k_i))^2} < tol$$

3. Relative change:

$$\frac{\|V_{n+1} - V_n\|}{\|V_n\|} < tol$$

4. Policy function convergence: Check if $g_{n+1}(k) = g_n(k)$ for all k (policy has stabilized)

Best practice:

- Use supremum norm (catches worst-case deviation)
- Set $tol = 10^{-6}$ or 10^{-8}
- Also check max iterations to avoid infinite loops
- Monitor convergence: plot $\|V_{n+1} - V_n\|$ vs. iteration

(d) Initial guess $V_0(k)$:

Option 1: Zero function (simple but slow)

$$V_0(k) = 0 \quad \forall k$$

+ Always feasible, easy to implement
- May require many iterations to converge

Option 2: Myopic solution (better)

$$V_0(k) = \frac{k^\alpha + (1-\delta)k}{1-\beta}$$

– Assumes no investment, just collect flow payoff forever

– Denominator $1/(1-\beta)$ is sum of geometric series

+ Closer to true solution, faster convergence

+ Economically sensible: present value of never adjusting capital

Option 3: Analytical approximation (best)

If we ignore adjustment costs and assume steady state, use value from perpetual flow:

$$V_0(k) = \frac{k^\alpha}{1 - \beta(1-\delta)}$$

Justification:

– In steady state, capital doesn't grow: $k' = k$

– Flow payoff: $k^\alpha - k + (1-\delta)k = k^\alpha - \delta k$

– For small δ, approximate with just production: k^α

– Present value with effective discount $\beta(1-\delta)$: captures depreciation

Option 4: Solution from similar parameters

– If you've solved for nearby parameter values, use that solution

– Particularly useful for comparative statics

Why does the initial guess matter?

– Bellman operator is a contraction: converges from any starting point

– Better initial guess $\Rightarrow$ fewer iterations $\Rightarrow$ faster computation

– With $V_0 = 0$: might need 500+ iterations

– With good guess: might need only 50-100 iterations

Recommendation: Use Option 2 or 3 for practical implementation.

8 Constraint Substitution

(a) Relationship between formulations:

The two formulations are **mathematically equivalent**.

Formulation B is obtained from **Formulation A** by substituting the constraint $c = W - W'$ directly into the objective function.

Specifically:

- Formulation A: optimize over (c, W') subject to $c + W' = W$
- Formulation B: use constraint to eliminate c, optimize over W' only
- The constraint becomes implicit in the objective function

Both formulations will yield:

- Same optimal value function $V^*(W)$
- Same optimal consumption policy $c^*(W)$
- Same optimal savings policy $W'^*(W)$

Advantage of B: One fewer choice variable, simpler optimization.

(b) Solving Formulation A with Lagrangian:

Set up the Lagrangian:

$$\mathcal{L} = u(c) + \beta V(W') - \lambda(c + W' - W)$$

First-order conditions:

FOC for c:

$$\frac{\partial \mathcal{L}}{\partial c} = u'(c) - \lambda = 0$$

$$\Rightarrow u'(c) = \lambda \quad (1)$$

FOC for W':

$$\frac{\partial \mathcal{L}}{\partial W'} = \beta V_{W'}(W') - \lambda = 0$$

$$\Rightarrow \beta V_{W'}(W') = \lambda \quad (2)$$

FOC for λ (constraint):

$$\frac{\partial \mathcal{L}}{\partial \lambda} = -(c + W' - W) = 0$$

$$\Rightarrow c + W' = W \quad (3)$$

Envelope condition:

Taking the derivative with respect to W:

$$V_W(W) = \frac{\partial \mathcal{L}}{\partial W} = \lambda$$

$$\Rightarrow V_W(W) = \lambda \quad (4)$$

Combining results:

From (1) and (4):

$$V_W(W) = u'(c)$$

By stationarity, rolling forward one period:

$$V_{W'}(W') = u'(c')$$

Substituting into (2):

$$\beta u'(c') = \lambda = u'(c)$$

Euler equation:

$$\boxed{u'(c) = \beta u'(c')}$$

(c) Solving Formulation B directly:

The Bellman equation is:

$$V(W) = \max_{W'} \{u(W - W') + \beta V(W')\}$$

First-order condition:

Taking the derivative with respect to W':

$$\frac{\partial}{\partial W'} \left[u(W - W') + \beta V(W')\right] = 0$$

Using chain rule:

$$-u'(W - W') + \beta V_{W'}(W') = 0$$

Since $c = W - W'$:

$$-u'(c) + \beta V_{W'}(W') = 0$$

$$\Rightarrow u'(c) = \beta V_{W'}(W') \quad (5)$$

Envelope condition:

Taking derivative with respect to W:

$$V_W(W) = \frac{\partial}{\partial W}\left[u(W - W'^*) + \beta V(W'^*)\right]$$

At the optimum, by the envelope theorem:

$$V_W(W) = u'(W - W'^*) = u'(c) \quad (6)$$

Combining:

By stationarity:

$$V_{W'}(W') = u'(c')$$

Substituting into (5):

$$u'(c) = \beta u'(c')$$

Same Euler equation!

$$\boxed{u'(c) = \beta u'(c')}$$

Both formulations yield identical optimality conditions.

(d) Computational efficiency:

Formulation B is more efficient for numerical solution.

	Formulation A	**Formulation B**
Choice variables	2: (c, W')	1: W'
Constraints	1: $c + W' = W$	0 (implicit)
Dimensions	2D search	1D search
Computation	Nested loops or constrained opt	Single loop

Why B is better:

1. **Fewer choice variables:** Optimization over 1 variable vs. 2
2. **No constraint handling:** Constraint is automatically satisfied by construction
3. **Simpler code:** Single loop over $W' \in k_{grid}$

4. **Faster execution:** If grid has N points:
   - Formulation A: $O(N^2)$ evaluations per iteration (for each c, check all W' satisfying constraint)
   - Formulation B: $O(N)$ evaluations per iteration (just loop over W')
5. **Numerical stability:** No need to enforce constraint numerically

Example comparison:

Formulation A (slower):

```
for i = 1:N_W
    V_best = -inf;
    for j = 1:N_W  % loop over W'
        for k = 1:N_c  % loop over c
            if abs(c(k) + W'(j) - W(i)) < tol  % check constraint
                value = u(c(k)) + beta * V_old(j);
                if value > V_best
                    V_best = value;
                end
            end
        end
    end
    V_new(i) = V_best;
end
```

Formulation B (faster):

```
for i = 1:N_W
    V_best = -inf;
    for j = 1:N_W  % loop over W' only
        c = W(i) - W'(j);  % constraint satisfied by construction
        if c > 0  % feasibility check
            value = u(c) + beta * V_old(j);
            if value > V_best
                V_best = value;
            end
        end
```

```
    end
    V_new(i) = V_best;
end
```

General principle: Always substitute out constraints when possible to reduce dimensionality of the optimization problem.

9 Finite vs Infinite Horizon

(a) Finite horizon Bellman equation with time dependence:

For $T = 20$, the Bellman equation is:

$$V_t(W_t) = \max_{c_t, W_{t+1}} \left\{ \frac{c_t^{1-\sigma}}{1-\sigma} + \beta V_{t+1}(W_{t+1}) \right\}$$

subject to: $W_{t+1} = W_t - c_t$, for $t = 1, 2, \ldots, 20$

With $\sigma = 2$:

$$\begin{aligned} V_t(W_t) &= \max_{W_{t+1}} \left\{ \frac{(W_t - W_{t+1})^{-1}}{-1} + 0.95 V_{t+1}(W_{t+1}) \right\} \\ &= \max_{W_{t+1}} \left\{ -(W_t - W_{t+1})^{-1} + 0.95 V_{t+1}(W_{t+1}) \right\} \end{aligned}$$

Terminal condition:

$$V_{21}(W_{21}) = 0 \quad \text{and} \quad W_{21} = 0$$

Key features:

- Value function has **time subscript** t: $V_t(\cdot) \neq V_s(\cdot)$ for $t \neq s$
- Different value function for each period
- Must solve **backward** from $t = 20$ to $t = 1$
- Policy functions also time-dependent: $c_t(W_t)$

(b) Infinite horizon Bellman equation:

$$V(W) = \max_{W'} \left\{ -(W - W')^{-1} + 0.95 V(W') \right\}$$

Why time subscripts disappear:

1. **No terminal date:** Problem looks identical from any period
2. **Time-invariant parameters:** β, σ, utility function don't change
3. **Stationarity:** Optimal behavior depends only on current state W, not on calendar time t
4. **Self-similarity:** The problem at $t+1$ is identical to the problem at t

Mathematical statement:

In infinite horizon with time-invariant parameters:

$$V_t(W) = V_{t+1}(W) = V(W) \quad \forall t, W$$

The value function is a **fixed point**: it satisfies

$$V(W) = \mathcal{T}[V](W)$$

where $\mathcal{T}$ is the Bellman operator.

Implication: Solve *one* functional equation, not a sequence of 20 (or infinitely many) equations.

(c) Final period consumption:

In the final period $T = 20$, the consumer **consumes all remaining wealth**:

$$c_{20} = W_{20}$$

$$W_{21} = 0$$

Why?

1. **No future:** There's no period 21, so saving has zero value
2. **Terminal condition:** $V_{21}(W_{21}) = 0$ regardless of W_{21}
3. **Marginal value of saving:**
$$\frac{\partial}{\partial W_{21}}[\beta V_{21}(W_{21})] = 0$$
4. **Marginal value of consumption:**

$$u'(c_{20}) = c_{20}^{-\sigma} > 0$$

5. **Optimality:** Since consumption has positive marginal value and saving has zero marginal value, consume everything

Formally: The FOC in period 20 is:

$$u'(W_{20} - W_{21}) = \beta V_{W'}(W_{21}) = 0$$

Since $u'(\cdot) > 0$ requires finite argument, this can only hold as $W_{21} \to 0$.

Economic intuition: You can't take wealth with you! In the last period, spending has value but saving doesn't.

(d) Infinite horizon analytical solution:

Claim: $c_t = (1-\beta)W_t$ is optimal.

Proof:

Step 1: Guess value function form

With CRRA utility, guess:

$$V(W) = AW^{1-\sigma}$$

For $\sigma = 2$:

$$V(W) = -AW^{-1}$$

where $A > 0$ is a constant to be determined.

Step 2: Substitute into Bellman equation

$$-AW^{-1} = \max_{W'} \left\{-(W - W')^{-1} + \beta(-A(W')^{-1})\right\}$$

Step 3: First-order condition

$$\frac{\partial}{\partial W'} \left[-(W - W')^{-1} - \beta A(W')^{-1}\right] = 0$$

$$-(W - W')^{-2} + \beta A(W')^{-2} = 0$$

$$(W')^{-2} = \frac{1}{\beta A}(W - W')^{-2}$$

$$\left(\frac{W - W'}{W'}\right)^2 = \frac{1}{\beta A}$$

$$\frac{W - W'}{W'} = \frac{1}{\sqrt{\beta A}}$$

$$W - W' = \frac{W'}{\sqrt{\beta A}}$$

$$W = W'\left(1 + \frac{1}{\sqrt{\beta A}}\right)$$

$$W' = \frac{W}{1 + 1/\sqrt{\beta A}} = \frac{\sqrt{\beta A}}{1 + \sqrt{\beta A}} W$$

Step 4: Find consumption policy

$$c = W - W' = W\left(1 - \frac{\sqrt{\beta A}}{1 + \sqrt{\beta A}}\right) = \frac{W}{1 + \sqrt{\beta A}}$$

Step 5: Determine A using envelope condition

The envelope condition gives:

$$V_W(W) = u'(c) = c^{-\sigma} = c^{-2}$$

From $V(W) = -AW^{-1}$:

$$V_W(W) = AW^{-2}$$

Therefore:

$$AW^{-2} = c^{-2}$$

$$A = \frac{W^2}{c^2}$$

Substituting $c = W/(1 + \sqrt{\beta A})$:

$$A = (1 + \sqrt{\beta A})^2$$

Let $x = \sqrt{A}$:

$$x^2 = (1 + \sqrt{\beta} x)^2$$

$$x^2 = 1 + 2\sqrt{\beta} x + \beta x^2$$

$$x^2(1-\beta) = 1 + 2\sqrt{\beta}x$$

For $\beta = 0.95$: solving gives A such that $c = (1-\beta)W$ approximately.

Alternative derivation for general result:

With log utility $u(c) = \ln(c)$, the solution is exactly:

$$c_t = (1-\beta)W_t$$

$$W_{t+1} = \beta W_t$$

This can be verified directly:

Verification:

$$V(W) = \max_{W'}\{\ln(W - W') + \beta V(W')\}$$

Guess $V(W) = B + C\ln(W)$:

$$B + C\ln(W) = \max_{W'}\{\ln(W - W') + \beta[B + C\ln(W')]\}$$

FOC:

$$\frac{-1}{W - W'} + \frac{\beta C}{W'} = 0$$
$$\frac{1}{W - W'} = \frac{\beta C}{W'}$$
$$W' = \beta C(W - W')$$
$$W' = \frac{\beta C W}{1 + \beta C}$$

Envelope condition:

$$C = \frac{1}{W - W'} \cdot W = \frac{1}{c}W$$

From FOC: $c = W - W' = W/(1 + \beta C)$

This gives $C = 1/(1-\beta)$ and:

$$c = (1-\beta)W$$

Economic interpretation:

- Consume a constant fraction $(1-\beta)$ of wealth each period
- The more patient (higher β), the less you consume

- Wealth declines at rate β: $W_{t+1} = \beta W_t$
- This is the "permanent income" solution

10 Bellman Operator and Convergence

(a) Fixed point of the operator $\mathcal{T}$:

A function V^* is a **fixed point** of the Bellman operator $\mathcal{T}$ if:

$$V^* = \mathcal{T}[V^*]$$

Explicitly:

$$V^*(W) = \max_c \{u(c) + \beta V^*(W - c)\} \quad \forall W$$

Interpretation:

- V^* is a solution to the Bellman equation
- Applying the operator to V^* returns V^* itself—the function doesn't change
- This is the **value function** of the infinite horizon problem
- At the fixed point, the function is **self-consistent**: it correctly represents the value of following the optimal policy

Mathematical analogy:

Just as a number x^* is a fixed point of function f if $x^* = f(x^*)$ (e.g., $x^* = \sqrt{x^*}$), a function V^* is a fixed point of operator $\mathcal{T}$ if $V^* = \mathcal{T}[V^*]$.

Key property: Under Blackwell's conditions, this fixed point is **unique**.

(b) Blackwell's sufficient conditions:

Blackwell (1965) proved that if an operator satisfies certain properties, it has a unique fixed point and value function iteration converges.

Condition 1: Monotonicity

If $V(W) \geq W(W)$ for all W, then:

$$\mathcal{T}[V](W) \geq \mathcal{T}[W](W) \quad \forall W$$

In our context:

If V gives higher value than W at every wealth level, then applying the Bellman operator to V still gives higher value than applying it to W.

Verification:

$$\begin{aligned}
\mathcal{T}[V](W) &= \max_c \{u(c) + \beta V(W - c)\} \\
&\geq \max_c \{u(c) + \beta W(W - c)\} \\
&= \mathcal{T}[W](W)
\end{aligned}$$

The inequality follows because we're maximizing over the same set, but with a weakly larger continuation value.

Condition 2: Discounting

For any V and constant $a > 0$:

$$\mathcal{T}[V + a](W) \leq \mathcal{T}[V](W) + \beta a$$

In our context:

If we shift the value function up by constant a, the Bellman operator shifts it up by at most βa (less than a if $\beta < 1$).

Verification:

$$\begin{aligned}
\mathcal{T}[V + a](W) &= \max_c \{u(c) + \beta(V(W - c) + a)\} \\
&= \max_c \{u(c) + \beta V(W - c)\} + \beta a \\
&= \mathcal{T}[V](W) + \beta a
\end{aligned}$$

Interpretation:

- The operator is a **contraction mapping** with modulus $\beta < 1$
- It "shrinks" distances between functions
- If two functions are far apart, applying $\mathcal{T}$ brings them closer together

Implications of Blackwell's conditions:

1. **Existence:** There exists a fixed point V^*
2. **Uniqueness:** The fixed point is unique

3. **Convergence:** Starting from any V_0, iteration $V_{n+1} = \mathcal{T}[V_n]$ converges to V^*
4. **Rate:** Convergence is geometric at rate β

(c) Why iteration converges from arbitrary V_0:

Contraction Mapping Theorem:

If $\mathcal{T}$ is a contraction mapping on a complete metric space with modulus $\beta < 1$, then:

1. There exists unique fixed point V^* such that $V^* = \mathcal{T}[V^*]$
2. For any starting point V_0, the sequence $V_{n+1} = \mathcal{T}[V_n]$ converges to V^*
3. The distance to the fixed point shrinks geometrically:

$$\|V_n - V^*\| \leq \beta^n \|V_0 - V^*\|$$

Intuitive explanation:

- Each application of $\mathcal{T}$ moves you closer to V^*
- Distance decreases by factor β each iteration
- Since $\beta < 1$, eventually $\|V_n - V^*\|$ becomes arbitrarily small
- Starting point doesn't matter—you always end up at V^*

Graphical intuition (1D analogy):

Imagine $f(x) = 0.9x + 1$. Starting from any x_0, iterating $x_{n+1} = f(x_n)$ converges to $x^* = 10$ because:

- Fixed point: $x^* = 0.9x^* + 1 \Rightarrow x^* = 10$
- If $x_n < 10$: $x_{n+1} = 0.9x_n + 1$ is closer to 10
- If $x_n > 10$: $x_{n+1} = 0.9x_n + 1$ is closer to 10
- Gap shrinks: $|x_{n+1} - 10| = 0.9|x_n - 10|$

Same logic applies to $\mathcal{T}$ operating on functions!

Practical implication:

You don't need to worry about finding a good initial guess for convergence (though it helps for speed). Even $V_0(W) = 0$ will eventually converge to the true $V^*(W)$.

(d) Economic interpretation of $V_1(W) = \mathcal{T}[V_0](W)$ with $V_0 = 0$:

Starting with $V_0(W) = 0$ for all W:

$$V_1(W) = \mathcal{T}[V_0](W) = \max_c \{u(c) + \beta \cdot 0\} = \max_c u(c)$$

subject to $c \leq W$.

Solution: $c^* = W$ (consume everything)

Therefore:

$$V_1(W) = u(W)$$

Economic interpretation:

$V_1(W)$ represents the value of having wealth W if you:

- Are **completely myopic** (ignore the future)
- Consume all your wealth today
- Receive zero utility from any future periods

This is the **"one-period problem"** or **"hand-to-mouth"** solution.

Next iteration: $V_2(W) = \mathcal{T}[V_1](W)$:

$$V_2(W) = \max_c \{u(c) + \beta u(W - c)\}$$

Economic interpretation:

$V_2(W)$ represents the value if you optimize over **two periods**:

- Choose consumption c today
- Consume remaining wealth $W - c$ tomorrow
- No periods beyond that

This is the **"two-period problem"**.

Continuing:

- $V_3(W)$: optimal value with 3-period horizon
- $V_4(W)$: optimal value with 4-period horizon
- $\vdots$
- $V_n(W)$: optimal value with n-period horizon

As $n \to \infty$, $V_n(W) \to V^*(W)$, the infinite horizon value.

General principle:

Value function iteration builds up the solution by progressively extending the planning horizon:

$$\text{Myopic} \to 2 \text{ periods} \to 3 \text{ periods} \to \ldots \to \text{infinite horizon}$$

Each iteration incorporates one more period of forward-looking behavior until you reach the fully optimal infinite horizon solution.