Dynamic Programming and the Bellman Equation

Lecture 9

Maria Titova

August 18, 2026

Econ 8000 — Math Camp

Motivation

In many dynamic problems, one choice affects the next.

- A household that saves today has more to spend tomorrow.
- A firm that invests today produces more next year.

Today's choice earns a payoff now and changes the situation for every later choice, so the choices at different dates cannot be studied separately.

Dynamic programming is a tool for solving such problems.

Sequential problems and their recursive form

The growth problem

Time is discrete, $t = 0, 1, 2, \ldots$, and the agent starts with capital $k_0 > 0$.

Capital k_t produces output $f(k_t)$, which is consumed or carried forward:

$$\max_{\{c_t,\, k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(c_t) \qquad \text{subject to} \qquad c_t + k_{t+1} = f(k_t), \quad c_t, k_{t+1} \geq 0, \quad k_0 \text{ given.}$$

Here $\beta \in (0,1)$ is the *discount factor*, u is continuous, strictly increasing, and strictly concave, and f is continuous, increasing, and concave with $f(0) = 0$; both are differentiable on $(0, \infty)$.

Choosing the entire sequence $c_0, k_1, c_1, k_2, \ldots$ at once makes this a *sequential problem*.

Every date has the same structure

Since $c_t = f(k_t) - k_{t+1}$, choosing the capital sequence also determines consumption. Thus $k_{t+1} \in [0, f(k_t)]$ is the choice at date t.

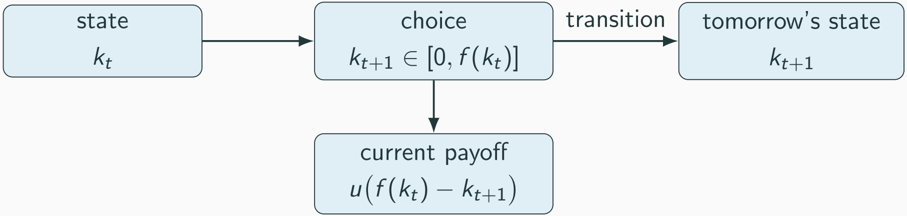

Figure 1.1

What a state is

The state is not simply "a variable indexed by t."

It is the information inherited from the past that matters for future choices.

- Once k_t is known, the past no longer affects feasible choices or payoffs.
- Two histories that arrive at the same k_t face the same future.

If the payoff were $u(c_t, c_{t-1})$, yesterday's consumption would also be a state.

A state contains all history that matters for future choices, payoffs, or transitions.

Value function

Definition 1.1 (Value function)

The *value function* $V(k)$ is the largest lifetime utility attainable from initial stock k:

$$V(k) = \max \sum_{t=0}^{\infty} \beta^t u(c_t) \quad \text{subject to} \quad c_t + k_{t+1} = f(k_t), \quad c_t, k_{t+1} \geq 0, \quad k_0 = k.$$

We assume throughout that the value is finite and a maximizing stream exists.

Bounded utility ensures $|V(k)| \leq M/(1-\beta)$, but not attainment.

The Bellman equation

Starting a date with capital k and choosing k' yields $u(f(k) - k')$ today, and from tomorrow on the agent faces the same problem with k' in place of k_0.

The most that tomorrow and later can yield is therefore $V(k')$, discounted once. An optimal k' makes the total as large as possible:

$$V(k) = \max_{0 \leq k' \leq f(k)} \left\{ u\big(f(k) - k'\big) + \beta V(k') \right\} \qquad \text{for every } k \geq 0.$$

The unknown is a *function*: the equation must hold at every $k \geq 0$, and the same V appears on both sides, once at k and once at k'.

Reading the Bellman equation

$$V(k) = \max_{0 \leq k' \leq f(k)} \left\{ u\big(f(k) - k'\big) + \beta V(k') \right\}$$

- k is the current state;
- k' is the choice, and $[0, f(k)]$ is the set of feasible choices;
- $u(f(k) - k')$ is the current payoff;
- today's choice k' becomes tomorrow's state, so it also appears inside V;
- $V(k')$ is the *continuation value* from behaving optimally tomorrow onward;
- β discounts the continuation value back to today.

The principle of optimality

Bellman's *principle of optimality*: after any feasible first choice, the best continuation is an optimal plan for the problem that starts from the resulting state.

- Every tail of an optimal plan is optimal from the state reached at that date.
- Otherwise, replacing the tail by an optimal continuation would raise lifetime utility without changing any earlier payoff.

Policy function

Definition 1.2 (Policy function)

The *policy function* $g(k)$ is the choice attaining the Bellman maximum at state k:

$$g(k) \in \underset{0 \le k' \le f(k)}{\arg\max} \left\{ u\big(f(k) - k'\big) + \beta V(k') \right\}.$$

We assume the maximizer is unique, so g is single-valued.

Consumption follows from the constraint, $c = f(k) - g(k)$.

The policy function generates the path

The two objects answer different questions:

$$k \xmapsto{V} \text{lifetime utility from } k, \qquad\qquad k \xmapsto{g} \text{tomorrow's capital } k'.$$

From k_0 alone, the policy function produces the whole optimal path:

$$k_1 = g(k_0), \qquad k_2 = g(k_1), \qquad k_3 = g(k_2), \ \ldots$$

The sequential problem and the Bellman equation generate the same optimal paths; the recursive formulation describes them with a rule g rather than a list of numbers.

First-order and envelope conditions

Assumptions for differentiating the Bellman equation

For $k > 0$ we assume:

- V is differentiable and concave;
- the maximizer is interior, $0 < g(k) < f(k)$, and differentiable in k.

We impose these properties rather than prove them.

Step 1: The first-order condition

$$V(k) = \max_{\mathbf{k'} \in [0, f(k)]} \left\{ u\big(f(k) - \mathbf{k'}\big) + \beta V\left(\mathbf{k'}\right) \right\}.$$

Because the Bellman equation maximizes over k', begin with the first-order condition. Hold k fixed and differentiate the objective with respect to k':

$$\frac{\partial}{\partial k'} \left[u\big(f(k) - k'\big) + \beta V(k') \right] = -u'(c) + \beta V'(k') = 0.$$

At the optimum $k' = g(k)$, rearranging gives

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

Step 2: The envelope condition

$$V(\mathbf{k}) = \max_{k' \in [0, f(\mathbf{k})]} \left\{ u\big(f(\mathbf{k}) - k'\big) + \beta V(k') \right\}.$$

The FOC leaves $V'(k')$ unknown, and we need some other equation to eliminate it.

Because $V(k)$ is maximized over k', the envelope theorem holds $k' = g(k)$ fixed:

$$\begin{aligned} V'(k) &= \frac{\partial}{\partial k}\left[u\big(f(k) - k'\big) + \beta V(k')\right]\Bigg|_{k' = g(k)} \\ &= u'\big(f(k) - g(k)\big) f'(k) = \boxed{u'(c) f'(k)}. \end{aligned}$$

Step 3: The Euler equation

At tomorrow's state, the envelope condition is

$$V'(k') = u'(c')f'(k'), \qquad c' = f(k') - g(k').$$

Substitute into the FOC and restore time subscripts to obtain the **Euler equation**:

$$\boxed{u'(c_t) = \beta\, u'(c_{t+1})\, f'(k_{t+1})}.$$

The marginal cost of saving today equals its discounted payoff tomorrow.

From the Bellman equation to the Euler equation

Eliminating V' leaves a relation between adjacent consumptions and the primitives.

$$\boxed{\text{Bellman equation} \;\longrightarrow\; \text{first-order and envelope conditions} \;\longrightarrow\; \text{Euler equation}}$$

These three steps recur throughout macroeconomics.

Writing a Bellman equation

Given a new dynamic problem:

1. identify the state, the choice, the payoff, and the transition;
2. write today's payoff plus β times the value at tomorrow's state;
3. maximize over today's choice.

$$V(\text{state}) = \max_{\text{choice}} \left\{ \text{payoff} + \beta V(\text{tomorrow's state}) \right\}$$

Example 2.1 (Consumption and saving at a fixed interest rate)

A consumer holds wealth $a_t \geq 0$ at the start of date t, consumes $c_t \in [0, a_t]$, and invests the rest at gross return $R > 0$, so that

$$a_{t+1} = R\,(a_t - c_t).$$

Lifetime utility is $\sum_{t=0}^{\infty} \beta^t u(c_t)$, and a_0 is given.

Write this problem's Bellman equation.

Consumption and saving: solution

Example 2.1 (Consumption and saving at a fixed interest rate)

With $c \in [0, a]$, $a' = R(a - c)$ for $R > 0$, and lifetime utility $\sum_t \beta^t u(c_t)$, write the Bellman equation.

Consumption and saving: solution

Example 2.1 (Consumption and saving at a fixed interest rate)

With $c \in [0, a]$, $a' = R(a - c)$ for $R > 0$, and lifetime utility $\sum_t \beta^t u(c_t)$, write the Bellman equation.

Solution.

State	wealth a	*Choice*	$c \in [0, a]$
Payoff	$u(c)$	*Transition*	$a' = R(a - c)$

Therefore

$$V(a) = \max_{0 \le c \le a} \left\{ u(c) + \beta V\bigl(R(a - c)\bigr) \right\}.$$

$\square$

The choice variable is a matter of convenience

Since $c = a - a'/R$, the consumer may instead choose tomorrow's wealth $a' \in [0, Ra]$:

$$V(a) = \max_{0 \leq a' \leq Ra} \left\{ u\left(a - \frac{a'}{R} \right) + \beta V(a') \right\}.$$

Both forms have the same value function; the choice variable is only notation.

The growth problem uses k' because the continuation value is simply $V(k')$.

The Euler equation for consumption and saving

$$V(a) = \max_c \{u(c) + \beta V(R(a-c))\}, \qquad a' = R(a-c).$$

Maximizing over c gives the first-order condition

$$u'(c) = \beta R\, V'(a').$$

Because $V(a)$ is optimized, the envelope condition gives

$$V'(a) = \beta R\, V'(a') = u'(c).$$

At tomorrow's state, $V'(a') = u'(c')$. Substitution yields the **Euler equation**:

$$\boxed{u'(c_t) = \beta R\, u'(c_{t+1})}.$$

Reading the consumption Euler equation

$$u'(c_t) = \beta R \, u'(c_{t+1}).$$

Saving one unit today produces R units tomorrow.

- If $\beta R = 1$, marginal utility and consumption are constant over time.
- If $\beta R > 1$, marginal utility falls and consumption rises.
- If $\beta R < 1$, marginal utility rises and consumption falls.

Why the recursion works

A last date to work back from

Suppose the world ends after date T:

$$\max \sum_{t=0}^{T} \beta^t u(c_t) \qquad \text{subject to} \qquad c_t + k_{t+1} = f(k_t), \quad c_t, k_{t+1} \geq 0, \quad k_0 \text{ given},$$

with no use for capital after date T.

The value function now depends on the date as well as on the state, because the agent at date t has $T - t + 1$ decision dates remaining.

Write $V_t(k)$ for the largest utility from date t on when the date-t state is k.

At the last date the problem is trivial: capital left over is wasted and u is increasing, so the agent consumes all output, $V_T(k) = u(f(k))$.

Finite-horizon Bellman recursion

Proposition 3.1 (Finite-horizon Bellman recursion)

For $t = T-1, T-2, \dots, 0$ and every $k \geq 0$,

$$V_t(k) = \max_{0 \leq k' \leq f(k)} \{u\big(f(k) - k'\big) + \beta V_{t+1}(k')\big\}, \qquad V_T(k) = u\big(f(k)\big),$$

and a feasible capital sequence solves the finite-horizon problem if and only if at every date $t \leq T-1$ its k_{t+1} attains the maximum at the state k_t.

An optimal first choice followed by an optimal continuation proves the result.

Given V_{t+1}, each state k poses a one-variable problem.

Backward induction

Computing the value functions in the order

$$V_T \longrightarrow V_{T-1} \longrightarrow V_{T-2} \longrightarrow \cdots \longrightarrow V_0$$

is called *backward induction*: start from the terminal value and work toward date 0.

- Each step is a static problem.
- Each maximum is attained by the Weierstrass theorem of Lecture 2, since $[0, f(k)]$ is compact and the objective is continuous.

Two dates of cake eating

Example 3.2 (Two dates of cake eating)

$$u(c) = \sqrt{c}, \qquad f(k) = k, \qquad 0 < \beta < 1, \qquad T = 1, \qquad k_0 = k > 0.$$

Cake left after date 1 has no value. The problem is

$$\begin{aligned} \max_{c_0, c_1, k_1, k_2 \geq 0} \quad & \sqrt{c_0} + \beta\sqrt{c_1} \\ \text{subject to} \quad & c_0 + k_1 = k_0, \\ & c_1 + k_2 = k_1. \end{aligned}$$

Find V_1, the policy $k_1 = g_0(k_0)$, and V_0.

Cake eating, step 1: Solve the last date

At date 1, what is the value $V_1(k)$ of a remaining stock k?

Cake eating, step 1: Solve the last date

At date 1, what is the value $V_1(k)$ of a remaining stock k?

There is no date after date 1, so cake carried beyond it is worthless.

Since utility is increasing, consume all remaining cake:

$$c_1 = k, \qquad \boxed{V_1(k) = u(k) = \sqrt{k}}.$$

This terminal value tells us what cake carried out of date 0 will be worth.

Cake eating, step 2: Write the date-0 problem

Given $V_1(k') = \sqrt{k'}$, what problem determines the date-0 value?

Cake eating, step 2: Write the date-0 problem

Given $V_1(k') = \sqrt{k'}$, what problem determines the date-0 value?

At date 0, carrying k' units of cake forward leaves $k - k'$ to consume now.

Current utility plus discounted continuation value is

$$\underbrace{\sqrt{k-k'}}_{\text{utility at date 0}} + \underbrace{\beta V_1(k')}_{\text{value at date 1}} .$$

Substitute $V_1(k') = \sqrt{k'}$:

$$\boxed{V_0(k) = \max_{0 \le k' \le k} \left\{ \sqrt{k-k'} + \beta\sqrt{k'} \right\}}.$$

Backward induction has reduced the dynamic problem to one static maximization.

Cake eating, step 3: Choose how much to carry

Which k' maximizes the date-0 objective?

Cake eating, step 3: Choose how much to carry

Which k' maximizes the date-0 objective?

Strict concavity and the endpoint slopes imply a unique interior optimum.

The first-order condition for the amount k' carried forward is

$$-\frac{1}{2\sqrt{k-k'}} + \frac{\beta}{2\sqrt{k'}} = 0 \quad \iff \quad \sqrt{k'} = \beta\sqrt{k-k'}.$$

Squaring and solving for k' gives the date-0 policy

$$\boxed{g_0(k) = k' = \frac{\beta^2}{1+\beta^2}\,k}.$$

A larger β raises the fraction of cake carried forward.

Cake eating, step 4: Compute the value

Given $k' = \beta^2 k/(1+\beta^2)$, what is $V_0(k)$?

Cake eating, step 4: Compute the value

Given $k' = \beta^2 k/(1+\beta^2)$, what is $V_0(k)$?

The policy implies current consumption

$$k - k' = \frac{k}{1+\beta^2}.$$

Substitute consumption and the policy into the date-0 objective:

$$\begin{aligned} V_0(k) &= \sqrt{\frac{k}{1+\beta^2}} + \beta\sqrt{\frac{\beta^2 k}{1+\beta^2}} \\ &= \boxed{\sqrt{1+\beta^2}\,\sqrt{k}}. \end{aligned}$$

We now know both value functions and the optimal choice at date 0.

More dates repeat the same calculation

If $V_{t+1}(k) = B_{t+1}\sqrt{k}$, what are today's policy and value?

More dates repeat the same calculation

If $V_{t+1}(k) = B_{t+1}\sqrt{k}$, what are today's policy and value?

Today's problem is

$$V_t(k) = \max_{0 \le k' \le k} \left\{ \sqrt{k - k'} + \beta B_{t+1} \sqrt{k'} \right\}.$$

The same first-order condition gives

$$k' = \frac{\beta^2 B_{t+1}^2}{1 + \beta^2 B_{t+1}^2}\, k, \qquad V_t(k) = B_t \sqrt{k}, \qquad B_t^2 = 1 + \beta^2 B_{t+1}^2.$$

Starting from $B_T = 1$ and working backward yields

$$B_t^2 = 1 + \beta^2 + \beta^4 + \cdots + \beta^{2(T-t)}.$$

With more remaining dates, B_t approaches $1/\sqrt{1 - \beta^2}$.

From a finite horizon to an infinite horizon

With no last date there is no V_T from which to start.

In a *stationary* problem, the payoff, feasible set, and transition are the same at every date, so the infinite-horizon value does not depend on the calendar.

Setting $V_t = V_{t+1} = V$ in the finite-horizon recursion returns the Bellman equation

$$V(k) = \max_{0 \leq k' \leq f(k)} \left\{ u\big(f(k) - k'\big) + \beta V(k') \right\}.$$

Existence and uniqueness are therefore fixed-point questions.

The general stationary problem

Object	Growth problem	General stationary problem
State	capital k	$s \in S$
Feasible choices	$k' \in [0, f(k)]$	$a \in \Gamma(s)$
Current payoff	$u(f(k) - k')$	$r(s, a)$
Transition	k'	$s' = F(s, a)$
Bellman equation	$V(k) = \max\{u(f(k) - k') + \beta V(k')\}$	$V(s) = \sup\{r(s, a) + \beta V(F(s, a))\}$
Policy function	$k' = g(k)$	$a = g(s)$

Table 3.1

Bellman operator

Definition 3.3 (Bellman operator)

Let $B(S)$ be the set of bounded functions $v : S \to \mathbb{R}$, equipped with the *sup norm*

$$\|v\|_\infty = \sup_{s \in S} |v(s)|.$$

For $v \in B(S)$, the *Bellman operator* is defined state by state by

$$(\mathcal{T} v)(s) = \sup_{a \in \Gamma(s)} \left\{ r(s, a) + \beta\, v\big(F(s, a)\big) \right\}.$$

The function v gives next-period values; $\mathcal{T} v$ gives optimized current values.

The Bellman operator is a contraction

Theorem 3.4 (The Bellman operator is a contraction)

Suppose $0 < \beta < 1$, $\Gamma(s)$ *is nonempty, and* $|r(s,a)| \le M < \infty$ *for every feasible pair* (s,a).

Then $\mathcal{T}$ *maps* $B(S)$ *into* $B(S)$, *and for all* $v, w \in B(S)$,

$$\|\mathcal{T}v - \mathcal{T}w\|_\infty \le \beta \, \|v - w\|_\infty.$$

- Since $B(S)$ is complete, $\mathcal{T}$ has exactly one fixed point. The equation $\mathcal{T}V = V$ is the Bellman equation, so V is its unique bounded solution.
- From any $v_0 \in B(S)$, value iteration $v_{n+1} = \mathcal{T}v_n$ converges to V, with $\|v_{n+1} - V\|_\infty \le \beta \|v_n - V\|_\infty$.

Value iteration links finite and infinite horizons

Start from zero terminal value and repeatedly apply the Bellman operator:

$$v_0 = 0 \quad \longrightarrow \quad v_1 = \mathcal{T} v_0 \quad \longrightarrow \quad v_2 = \mathcal{T} v_1 \quad \longrightarrow \quad \cdots \quad \longrightarrow \quad V.$$

The function v_n is the value with n dates remaining. Thus value iteration is backward induction with an expanding horizon.

In the cake-eating example, the same limit appears in

$$B_t^2 = 1 + \beta^2 + \cdots + \beta^{2(T-t)} \quad \longrightarrow \quad \frac{1}{1-\beta^2}.$$

A Bellman equation solved by hand

What guess and verify is doing

In the finite-horizon problem, today's value is computed from tomorrow's value:

$$V_t(k) = \max_{0 \le k' \le k} \left\{ \sqrt{k - k'} + \beta V_{t+1}(k') \right\}.$$

In the stationary infinite-horizon problem, the same function appears on both sides:

$$V(k) = \max_{0 \le k' \le k} \left\{ \sqrt{k - k'} + \beta V(k') \right\}.$$

Section 03 gave $V_t(k) = B_t \sqrt{k}$.

Here we try $V(k) = B\sqrt{k}$ and find the B reproduced by maximization.

Square-root cake eating

Example 4.1 (Square-root cake eating)

$$u(c) = \sqrt{c}, \qquad f(k) = k, \qquad 0 < \beta < 1, \qquad S = [0, \bar{k}].$$

Given $k_0 = k \in S$, solve

$$\max_{\{c_t, k_{t+1}\}_{t=0}^{\infty}} \quad \sum_{t=0}^{\infty} \beta^t \sqrt{c_t}$$

$$\text{subject to} \quad c_t + k_{t+1} = k_t, \qquad c_t, k_{t+1} \geq 0.$$

Find the value function and policy.

Square-root cake eating, step 1: Write the Bellman equation

What recursive equation must the value function satisfy?

Square-root cake eating, step 1: Write the Bellman equation

What recursive equation must the value function satisfy?

At state k, choose $k' \in [0, k]$; then current consumption is $k - k'$. Thus

$$\boxed{V(k) = \max_{0 \leq k' \leq k} \left\{ \sqrt{k - k'} + \beta V(k') \right\}}.$$

Square-root cake eating, step 2: Optimize for a given B

Given $V(k') = B\sqrt{k'}$, which k' maximizes the Bellman right-hand side?

Square-root cake eating, step 2: Optimize for a given B

Given $V(k') = B\sqrt{k'}$, which k' maximizes the Bellman right-hand side?

Substitution gives the one-variable problem

$$\max_{0 \le k' \le k} \left\{ \sqrt{k - k'} + \beta B \sqrt{k'} \right\}.$$

Strict concavity and the endpoint slopes make the optimum unique and interior.

The first-order condition gives

$$\frac{1}{2\sqrt{k - k'}} = \frac{\beta B}{2\sqrt{k'}} \quad \Longrightarrow \quad \boxed{k' = \frac{\beta^2 B^2}{1 + \beta^2 B^2}\, k}.$$

Square-root cake eating, step 3: Compute the maximized value

What is the Bellman equation's right-hand side at the maximizing k'?

Square-root cake eating, step 3: Compute the maximized value

What is the Bellman equation's right-hand side at the maximizing k'?

Substitute $k' = \beta^2 B^2 k/(1+\beta^2 B^2)$ into the objective:

$$\begin{aligned}
\sqrt{k-k'} + \beta B \sqrt{k'} &= \sqrt{\frac{k}{1+\beta^2 B^2}} + \beta B \sqrt{\frac{\beta^2 B^2 k}{1+\beta^2 B^2}} \\
&= \frac{1+\beta^2 B^2}{\sqrt{1+\beta^2 B^2}} \sqrt{k} \\
&= \boxed{\sqrt{1+\beta^2 B^2}\,\sqrt{k}}.
\end{aligned}$$

The maximized right-hand side is again a constant times $\sqrt{k}$.

Square-root cake eating, step 4: Impose the fixed point

Which $B > 0$ makes the two sides of the Bellman equation equal?

Square-root cake eating, step 4: Impose the fixed point

Which $B > 0$ makes the two sides of the Bellman equation equal?

The conjecture gives $B\sqrt{k}$ on the left; step 3 gives the maximized right-hand side:

$$\underbrace{B\sqrt{k}}_{V(k)} = \underbrace{\sqrt{1+\beta^2 B^2}\sqrt{k}}_{\text{maximized right-hand side}} .$$

For $k > 0$, cancel $\sqrt{k}$ and solve:

$$B = \sqrt{1+\beta^2 B^2} \quad \Longleftrightarrow \quad \boxed{B = \frac{1}{\sqrt{1-\beta^2}}}.$$

Square-root cake eating, step 5: Recover V and g

How does the solved coefficient B determine the value function and policy?

Square-root cake eating, step 5: Recover V and g

How does the solved coefficient B determine the value function and policy?

The value function follows directly:

$$\boxed{V(k) = B\sqrt{k} = \frac{\sqrt{k}}{\sqrt{1-\beta^2}}}.$$

Substitute $B^2 = 1/(1-\beta^2)$ into the policy from step 2:

$$\begin{aligned} g(k) &= \frac{\beta^2 B^2}{1+\beta^2 B^2}\, k \\ &= \frac{\beta^2/(1-\beta^2)}{1+\beta^2/(1-\beta^2)}\, k = \boxed{\beta^2 k}. \end{aligned}$$

Verification: Substitute into the Bellman equation

Do the proposed V and g satisfy the Bellman equation?

Verification: Substitute into the Bellman equation

Do the proposed V and g satisfy the Bellman equation?

Step 2 showed that $g(k) = \beta^2 k$ is the unique maximizer. Evaluate the Bellman right-hand side at this choice:

$$\begin{aligned}
\sqrt{k - g(k)} + \beta V\big(g(k)\big) &= \sqrt{(1-\beta^2)k} + \frac{\beta\sqrt{\beta^2 k}}{\sqrt{1-\beta^2}} \\
&= \left(\sqrt{1-\beta^2} + \frac{\beta^2}{\sqrt{1-\beta^2}}\right)\sqrt{k} \\
&= \frac{\sqrt{k}}{\sqrt{1-\beta^2}} = V(k).
\end{aligned}$$

Thus the proposed value and policy satisfy the Bellman equation.

Final check: Uniqueness

Why is the verified value function the unique bounded solution?

Final check: Uniqueness

Why is the verified value function the unique bounded solution?

The value function is bounded on the state space:

$$0 \le V(k) = \frac{\sqrt{k}}{\sqrt{1-\beta^2}} \le \frac{\sqrt{\bar{k}}}{\sqrt{1-\beta^2}} \qquad \text{for } k \in [0, \bar{k}].$$

The one-period reward is also bounded:

$$0 \le \sqrt{k - k'} \le \sqrt{\bar{k}}.$$

The contraction theorem therefore gives uniqueness and convergence of value iteration.

Where macroeconomics takes this

This lecture is a first exposure: deterministic, discrete time, one canonical problem.

The first-year sequence develops the omitted theory and extensions:

- Stokey and Lucas with Prescott, *Recursive Methods in Economic Dynamics*;
- Ljungqvist and Sargent, *Recursive Macroeconomic Theory*;
- Acemoglu, *Introduction to Modern Economic Growth*, Chapter 6.