

Lecture 9: Dynamic Programming and the Bellman Equation

Maria Titova, Vanderbilt University

Fall 2026

In many economic problems, the decision maker faces a sequence of choices, each affecting 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.

1 Sequential problems and their recursive form


1.1 The growth problem

Time is discrete, $t = 0, 1, 2, \dots$. An agent (decision-maker) starts with a stock of capital $k_0 > 0$. Capital k_t produces output $f(k_t)$ at date t , and output is either consumed, c_t , or carried forward as next period's capital, k_{t+1} . The agent values consumption streams by a discounted sum of one-period utilities and solves

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

The number $\beta \in (0, 1)$ is the *discount factor*. The utility function $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ is continuous, strictly increasing, and strictly concave, and $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuous, increasing, and concave with $f(0) = 0$; both are differentiable on $(0, \infty)$. When we differentiate a value function we assume in addition that $u'(c) \rightarrow \infty$ as $c \rightarrow 0$, so that an agent never chooses zero consumption.

Problem (1) is written as a *sequential problem*: the unknown is the whole sequence of consumptions and capital stocks, $c_0, k_1, c_1, k_2, \dots$, chosen at once. Since $c_t = f(k_t) - k_{t+1}$, the capital sequence determines the consumption sequence. We can therefore use $k_{t+1} \in [0, f(k_t)]$ as the choice at date t . Every date has the same structure, drawn in Figure 1.1:



The agent inherits the *state* k_t at date t and chooses k_{t+1} , earning the current payoff $u(f(k_t) - k_{t+1})$. The *transition* makes today's choice tomorrow's state.




Figure 1.1: One date of the growth problem. The state k_t fixes the feasible choices $[0, f(k_t)]$; the choice k_{t+1} produces the current payoff and becomes the state at date $t + 1$.

The state is not simply “a variable indexed by t .” It is the information inherited from the past that matters for the choices still to be made. In the growth problem, once k_t is known, the feasible choices at date t and everything after depend on the past only through k_t : two histories that arrive at the same k_t face the same future. Which sequence of earlier consumptions produced k_t is irrelevant, so it is not part of the state. If instead the payoff at date t were $u(c_t, c_{t-1})$, so that yesterday’s consumption changed the value of today’s, then c_{t-1} would have to be carried in the state along with k_t . A state must therefore contain all information from the history that affects future feasible choices, payoffs, or transitions.

1.2 The principle of optimality and the Bellman equation

Suppose the agent begins a date with capital k and chooses k' . From tomorrow on the agent faces the same continuation problem as (1), with k' in place of k_0 : the same technology, preferences, and infinite horizon. The highest utility attainable from tomorrow on therefore depends only on k' .

Definition 1.1 (Value function). The *value function* $V : \mathbb{R}_+ \rightarrow \mathbb{R}$ assigns to each initial stock k the largest lifetime utility attainable from it,

$$V(k) = \max \sum_{t=0}^{\infty} \beta^t u(c_t) \quad \text{subject to the constraints of (1) with } k_0 = k.$$

Whenever we write a maximum, we assume that the objective is finite and the maximum is attained. Bounded one-period utility is sufficient for convergence: if $|u(c)| \leq M$ for every $c \geq 0$, every feasible stream has discounted utility of absolute value at most $M/(1 - \beta)$ by the geometric-sum calculation of Lecture 8’s Example 3.9. Attainment requires additional conditions supplied by the macroeconomics sequence.

Now split lifetime utility from k into today’s payoff and everything after. If today’s choice is k' , then today yields $u(f(k) - k')$, and the most that tomorrow and later can yield is $V(k')$, discounted once because it starts one period later. An optimal choice of k' makes the total as large as possible:

$$V(k) = \max_{0 \leq k' \leq f(k)} \{u(f(k) - k') + \beta V(k')\} \quad \text{for every } k \geq 0. \quad (2)$$

This is the *Bellman equation* of the growth problem. Its terms correspond to the pieces of Section 1.1:

- 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, which is why the same letter appears inside $V(\cdot)$;
- $V(k')$ is the *continuation value*, the value of behaving optimally from tomorrow on;
- β discounts the continuation value back to today.

The reasoning that produced (2) is 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. In particular, 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.

The unknown in (2) is a *function*. The equation must hold at every $k \geq 0$, and the same V appears on both sides, once evaluated at k and once at k' . Thus solving the Bellman equation means finding a function V that satisfies it at every state. For a fixed k and a fixed candidate V , the maximization inside the braces is a one-variable problem of the kind solved in Lecture 6. The recursive formulation repeats this one-period choice at every state.

The recursive formulation gives more than an equation for V . Under our assumptions, the value function of Definition 1.1 is the unique bounded solution of the Bellman equation when u is bounded. Moreover, a feasible capital sequence is optimal for (1) if and only if, at every date, its k_{t+1} attains the maximum in (2) at the state k_t . We use these facts without proof; Section 3.2 returns to the uniqueness claim.

1.3 Value function and policy function

Solving (2) produces two objects, and it is important to keep them apart. The first is the value function V . The second is the choice that attains the maximum at each state.

Definition 1.2 (Policy function). The *policy function* $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ assigns to each state the choice that attains the maximum in the Bellman equation,

$$g(k) \in \arg \max_{0 \leq k' \leq f(k)} \{u(f(k) - k') + \beta V(k')\}.$$

Under our assumptions the maximizer is unique at every k , so g is a function; we take this without proof. Consumption then follows from the constraint, $c = f(k) - g(k)$. The two functions answer different questions:

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

The policy function describes the agents's choice at each state and generates the whole optimal path from k_0 alone:

$$k_1 = g(k_0), \quad k_2 = g(k_1), \quad k_3 = g(k_2), \dots$$

The Bellman equation determines the value function. Its continuation term $\beta V(k')$ values tomorrow's capital in today's utility. The sequential problem (1) and the recursive problem (2) generate the same optimal paths; the recursive formulation describes them with a rule g rather than a list of numbers.

2 First-order and envelope conditions

The maximization inside the Bellman equation is a one-variable problem. Lecture 6 supplies its first-order condition, and Lecture 8's envelope theorem gives the derivative of the optimized value. In this section we assume, for $k > 0$, that V is differentiable and concave and that the maximizer $g(k)$ is interior, $0 < g(k) < f(k)$, and differentiable in k . Standard dynamic-programming results establish these properties under suitable hypotheses; we impose them rather than prove them.

2.1 The first-order condition

Fix k and let $c = f(k) - k'$. The objective in (2) is $u(f(k) - k') + \beta V(k')$, and Lecture 6's interior first-order condition, differentiating in k' , gives

$$-u'(c) + \beta V'(k') = 0, \quad \text{that is,} \quad u'(c) = \beta V'(k'). \quad (3)$$

Since the objective is concave in k' , the first-order condition is sufficient by Lecture 6's Theorem 3.2: it locates the maximizer $k' = g(k)$. In words, one more unit of capital carried into tomorrow costs $u'(c)$ in utility today and is worth $\beta V'(k')$, the discounted marginal value of capital tomorrow. At the optimum the two are equal.

To eliminate the unknown derivative $V'(k')$ from (3), we use the envelope condition.

2.2 The envelope condition and the Euler equation

The Bellman equation defines $V(k)$ as an optimized value with parameter k . By Lecture 8's Theorem 2.1, part 1, we hold the choice at its optimum and differentiate the objective only with respect to k . The parameter appears only in $u(f(k) - k')$, so

$$V'(k) = u'(f(k) - g(k)) f'(k) = u'(c) f'(k). \quad (4)$$

This is the *envelope condition*. One more unit of capital today raises output by $f'(k)$, and at the margin the extra output is worth $u'(c)$ whether it is consumed or saved, because the first-order condition has equated the two uses; the response of $g(k)$ contributes nothing to first order.

Equation (4) holds at every state, so it holds at tomorrow's state k' with tomorrow's consumption $c' = f(k') - g(k')$: $V'(k') = u'(c') f'(k')$. Substituting into (3),

$$u'(c_t) = \beta u'(c_{t+1}) f'(k_{t+1}), \quad (5)$$

where we have restored time subscripts, $c = c_t$, $k' = k_{t+1}$, and $c' = c_{t+1}$. This is the *Euler equation* of the growth problem. Its left side is the utility cost of saving one more unit today. Its right side is the benefit: the unit becomes $f'(k_{t+1})$ units of output tomorrow, each worth $u'(c_{t+1})$, discounted once. Substitution has eliminated V' and left a relation between consumption at adjacent dates and the model's primitives.

The calculation has three steps that recur throughout macroeconomics:

Bellman equation $\longrightarrow$ first-order condition and envelope condition $\longrightarrow$ Euler equation
--

2.3 The same Euler equation from the sequential problem

We can also derive (5) directly from the sequential formulation. Substitute $c_t = f(k_t) - k_{t+1}$ into (1), so that the objective becomes

$$\sum_{t=0}^{\infty} \beta^t u(f(k_t) - k_{t+1}),$$

a function of the capital sequence alone. The variable k_{t+1} appears in exactly two terms of the sum, dated t and $t + 1$. At an interior optimum, Lecture 6's first-order condition in the single variable k_{t+1} , all other capital stocks held fixed, gives

$$-\beta^t u'(c_t) + \beta^{t+1} u'(c_{t+1}) f'(k_{t+1}) = 0,$$

which is (5) after dividing by β^t . The economics is a one-period perturbation: consume one unit less at t , carry it forward, and consume the proceeds at $t + 1$; at an optimum the change in lifetime utility is zero to first order.

The two derivations produce the same condition because the recursive and sequential formulations describe the same problem. The Bellman equation, its first-order condition, and its envelope condition reproduce the intertemporal optimality condition obtained directly from the sequential problem, using a one-period problem with the same form at every state. The Euler equation is a relation between three consecutive capital stocks (k_t, k_{t+1}, k_{t+2}) ; the initial condition k_0 pins down one end of the path, and the macroeconomics sequence supplies the condition that pins down the other, the *transversality condition*. We do not develop the transversality condition in this course.

2.4 Writing down a Bellman equation

Given a new dynamic problem, the procedure is: identify the state, the choice, the payoff, and the transition; then write today's payoff plus β times the value at tomorrow's state, maximized over today's choice. The consumption–saving problem also shows how two choices can parameterize the same recursion.

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.

- *State.* Wealth a . Once a_t is known, nothing about the past matters for the future.
- *Choice.* Consumption $c \in [0, a]$.
- *Payoff.* $u(c)$.
- *Transition.* $a' = R(a - c)$.

The Bellman equation is

$$V(a) = \max_{0 \leq c \leq a} \{u(c) + \beta V(R(a - c))\}.$$

There is more than one convenient way to parameterize the choice. Since $c = a - a'/R$, the consumer may instead choose tomorrow's wealth $a' \in [0, Ra]$:

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

The two equations have the same solution V ; the economic problem determines the recursion, and the notation is a matter of convenience. In the growth problem we chose k' for the same reason: it keeps the continuation value $V(k')$ simple.

Under the assumptions of Section 2, the first-order condition in c from the first formulation is

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

and the envelope condition, differentiating $u(c) + \beta V(R(a - c))$ in the parameter a at the optimal c , is

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

Applying the envelope condition at tomorrow's state, $V'(a') = u'(c')$, and substituting into the first-order condition gives the Euler equation

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

The consumer equates the marginal utility of a unit consumed today with that of the R units it would become tomorrow, discounted. If $\beta R = 1$, consumption is constant over time; if $\beta R > 1$, marginal utility falls and consumption rises. The first-year consumption theory sequence uses this equation as its starting point.

3 Why the recursion works: finite horizons and the fixed point

Finite horizons make the recursive logic explicit because they provide a last date from which to work backward. For an infinite horizon, the contraction mapping theorem determines the value function without a last date.

3.1 Finite horizons and backward induction

Suppose the world ends after date T : the agent solves

$$\max \sum_{t=0}^T \beta^t u(c_t) \quad \text{subject to} \quad c_t + k_{t+1} = f(k_t), \quad c_t, k_{t+1} \geq 0, \quad k_0 \text{ given}, \quad (6)$$

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:

$$V_T(k) = u(f(k)).$$

At date $T - 1$ the agent chooses k_T knowing its continuation value, $V_T(k_T)$:

$$V_{T-1}(k) = \max_{0 \leq k' \leq f(k)} \{ u(f(k) - k') + \beta V_T(k') \}.$$

Once V_T is known, this is a one-variable problem of Lecture 6, and solving it at every k gives the function V_{T-1} . Then V_{T-2} is obtained from V_{T-1} in the same way, and so on down to V_0 :

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

Proposition 3.1 (Finite-horizon Bellman recursion). *For $t = T - 1, T - 2, \dots, 0$ and every $k \geq 0$,*

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

and a feasible capital sequence solves (6) if and only if at every date $t \leq T - 1$ its k_{t+1} attains the maximum at the state k_t .

Proof (optional). Fix t and k , and write $W(k)$ for the right-hand side of the display. Every feasible plan from date t at state k begins with some $k' \in [0, f(k)]$ and continues with a feasible plan from date $t + 1$ at state k' , whose utility from $t + 1$ on is at most $V_{t+1}(k')$. So every plan earns at most $u(f(k) - k') + \beta V_{t+1}(k') \leq W(k)$, and $V_t(k) \leq W(k)$. Conversely, let k' attain the maximum and follow it with an optimal plan from date $t + 1$ at state k' ; that plan exists by induction downward from date T , where the optimal plan is to consume everything. The combined plan is feasible and earns $W(k)$, so $V_t(k) \geq W(k)$. The two inequalities give the recursion. The same comparison shows that a plan achieves $V_t(k)$ exactly when its first choice attains the maximum and its continuation is optimal. This proves the biconditional. $\square$

Computing $V_T, V_{T-1}, \dots, V_0$ in that order is called *backward induction*: it is Lecture 1's induction, run from the last date toward the first. Each step is a static problem, and each maximum is attained by the Weierstrass theorem of Lecture 2 when the objective is continuous, since $[0, f(k)]$ is compact. The principle of optimality is also established by the argument: an optimal plan's continuation from any date is optimal for the problem starting from the state reached on that date.

Example 3.2 (Two dates of cake eating). Let $u(c) = \sqrt{c}$ and $f(k) = k$. Capital does not reproduce in this special case: k is a stock of cake, and the agent divides it between consumption now and cake carried to the next date. Fix an upper bound $\bar{k} > 0$ and take the state space to be $S = [0, \bar{k}]$, which is invariant because every feasible choice satisfies $0 \leq k' \leq k$. Let $T = 1$, so there are two dates. At date 1 the agent consumes everything:

$$V_1(k) = \sqrt{k}.$$

At date 0 the recursion gives

$$V_0(k) = \max_{0 \leq k' \leq k} \{\sqrt{k - k'} + \beta \sqrt{k'}\}.$$

For $k > 0$ the objective is strictly concave and its one-sided slopes at the two endpoints point toward the interior. Its first-order condition is

$$\frac{1}{2\sqrt{k - k'}} = \frac{\beta}{2\sqrt{k'}} \quad \text{and hence} \quad k' = \frac{\beta^2}{1 + \beta^2} k,$$

the unique maximizer by Lecture 6's Theorem 3.1. The agent carries the fraction $\beta^2/(1 + \beta^2)$ of the cake to date 1 and consumes the rest. Substituting back gives

$$V_0(k) = \sqrt{1 + \beta^2} \sqrt{k}.$$

Thus V_1 and V_0 are both a constant multiple of $\sqrt{k}$.

The pattern continues. If $V_{t+1}(k) = B_{t+1}\sqrt{k}$ with $B_{t+1} > 0$, the same calculation gives

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

Consequently,

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

As the number of remaining dates grows, B_t increases to $1/\sqrt{1 - \beta^2}$, and the fraction carried forward increases to β^2 . Section 4 obtains these limits by solving the infinite-horizon problem directly.

3.2 The infinite horizon as a fixed point

With no last date there is no V_T from which to start. Example 3.2 shows the coefficient on $\sqrt{k}$ and the fraction carried forward approaching limits as dates are added. In a *stationary* problem, one whose payoff, feasible set, and transition are the same at every date, the infinite-horizon value does not depend on the calendar. The recursion of Proposition 3.1 with $V_t = V_{t+1} = V$ is the Bellman equation (2). Whether an infinite-horizon V exists, and whether it is unique, is a fixed-point question. Under bounded rewards, a contraction argument on the space of bounded functions answers it.

To state the contraction result, write a stationary problem in general notation. It has a state s in a state space S , a set $\Gamma(s)$ of feasible actions at s , a reward $r(s, a)$, a transition $s' = F(s, a)$, and a discount factor $\beta \in (0, 1)$; the growth problem is the case $s = k$, $\Gamma(k) = [0, f(k)]$, $a = k'$, $r(k, k') = u(f(k) - k')$, and $F(k, k') = k'$. Table 3.1 lists the correspondence. The Bellman equation is

$$V(s) = \sup_{a \in \Gamma(s)} \{r(s, a) + \beta V(F(s, a))\} \quad \text{for every } s \in S,$$

written with a supremum so that it makes sense before anything is known about attainment.

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: The pieces of the growth problem and their names in a general stationary dynamic problem. In the growth problem the transition is the identity in the choice, which is why k' serves as both the choice and tomorrow's state.

Definition 3.3 (Bellman operator). Let $B(S)$ be the set of bounded functions $v : S \rightarrow \mathbb{R}$ with the sup norm $\|v\|_\infty = \sup_{s \in S} |v(s)|$. The *Bellman operator* $\mathcal{T}$ sends a candidate value function v to the function

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

The number $(\mathcal{T}v)(s)$ is the supremum value of a one-period choice at s when v values tomorrow's state. The Bellman equation says $\mathcal{T}V = V$: the value function is a *fixed point* of $\mathcal{T}$. Backward induction is iteration of $\mathcal{T}$, since $V_t = \mathcal{T}V_{t+1}$ in Proposition 3.1. The operator $\mathcal{T}$ is the function-valued version of the map $T(x) = a + \beta x$ in Lecture 8's Example 3.9: today's reward plus β times the continuation.

Theorem 3.4 (The Bellman operator is a contraction). *Suppose $\Gamma(s)$ is nonempty for every $s \in S$ and there is $M < \infty$ with $|r(s, a)| \leq M$ whenever $a \in \Gamma(s)$. Then $\mathcal{T}$ maps $B(S)$ into $B(S)$, and for all $v, w \in B(S)$,*

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

Proof (optional). Every number whose supremum defines $(\mathcal{T}v)(s)$ lies between $-M - \beta\|v\|_\infty$ and $M + \beta\|v\|_\infty$, so $|(\mathcal{T}v)(s)| \leq M + \beta\|v\|_\infty$ for every s and $\mathcal{T}v \in B(S)$. Now fix s . For every $a \in \Gamma(s)$ the definition of the sup norm gives $v(F(s, a)) \leq w(F(s, a)) + \|v - w\|_\infty$, hence

$$r(s, a) + \beta v(F(s, a)) \leq r(s, a) + \beta w(F(s, a)) + \beta\|v - w\|_\infty.$$

The inequality holds action by action, so it survives the supremum over $a \in \Gamma(s)$: $(\mathcal{T}v)(s) \leq (\mathcal{T}w)(s) + \beta\|v - w\|_\infty$. Exchanging v and w gives the reverse bound, so $|(\mathcal{T}v)(s) - (\mathcal{T}w)(s)| \leq \beta\|v - w\|_\infty$ at every s , and taking the supremum over s finishes the proof. $\square$

The space $B(S)$ is complete in the sup norm, a fact we use without proof. The contraction mapping theorem therefore gives three conclusions. The Bellman equation has *exactly one* bounded solution. Starting from any bounded v_0 , even $v_0 = 0$, the iterates $v_{n+1} = \mathcal{T}v_n$ converge to it in sup norm, with the worst-case error shrinking by the factor β each time. This procedure is called *value function iteration*. For $v_0 = 0$, the iterate v_n is the value of the problem with n dates remaining and no terminal value. Thus the infinite-horizon value is the limit of finite-horizon values, as in Example 3.2. More generally, value iteration can start from any bounded function, and the effect of that starting function vanishes at the geometric rate β .

Discounting supplies the contraction modulus β . As $\beta \rightarrow 1$, the worst-case convergence bound becomes arbitrarily slow; at $\beta = 1$, Lecture 8's map $T(x) = a + x$ has no fixed point when $a \neq 0$: an undiscounted constant stream of nonzero rewards has no finite value. The theorem also requires a bounded reward. In Example 3.2, $S = [0, \bar{k}]$ and $0 \leq \sqrt{k - k'} \leq \sqrt{\bar{k}}$, so the hypothesis holds and the finite-horizon values converge uniformly to the unique bounded fixed point. If instead the state space were all of $\mathbb{R}_+$, the same reward would be unbounded above and this theorem would no longer apply.

4 A Bellman equation solved by hand

Value function iteration is useful for numerical work. When the primitives are special enough, the Bellman equation can be solved analytically by *guess and verify*: conjecture that V has a particular functional form with unknown constants, substitute the conjecture into the Bellman equation, carry out the maximization, and check whether the result has the conjectured form again; if it does, matching constants yields equations for them. The guess proposes a candidate; substitution verifies that the candidate is a fixed point. A separate theorem is still needed to show that this fixed point is the value function and is unique in the relevant class.

Example 4.1 (Square-root cake eating). Continue the cake-eating problem of Example 3.2, with $u(c) = \sqrt{c}$, $f(k) = k$, and $S = [0, \bar{k}]$. Given $k_0 = k \in S$, the agent solves

$$\max_{\{c_t, k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \sqrt{c_t} \quad \text{subject to} \quad c_t + k_{t+1} = k_t, \quad c_t, k_{t+1} \geq 0.$$

Its Bellman equation is

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

The finite-horizon value functions were multiples of $\sqrt{k}$, so we conjecture

$$V(k) = B\sqrt{k}$$

with $B > 0$ and look for a constant B that makes the equation hold at every $k \in [0, \bar{k}]$.

The maximization. With the conjecture substituted, the objective $\sqrt{k - k'} + \beta B\sqrt{k'}$ is strictly concave in k' . For $k > 0$ its unique maximizer is interior and satisfies

$$\frac{1}{2\sqrt{k - k'}} = \frac{\beta B}{2\sqrt{k'}},$$

so

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

For $k = 0$, the only feasible choice is $k' = 0$, which the same formula gives.

The verification. Substituting the maximizer back, the right-hand side of the Bellman equation becomes

$$\frac{\sqrt{k}}{\sqrt{1 + \beta^2 B^2}} + \frac{\beta^2 B^2 \sqrt{k}}{\sqrt{1 + \beta^2 B^2}} = \sqrt{1 + \beta^2 B^2} \sqrt{k}.$$

The conjectured form is preserved. Matching the coefficient on $\sqrt{k}$ gives

$$B = \sqrt{1 + \beta^2 B^2} \iff B = \frac{1}{\sqrt{1 - \beta^2}}.$$

The value and policy. The value function and its implied policy and consumption are therefore

$$V(k) = \frac{\sqrt{k}}{\sqrt{1 - \beta^2}}, \quad g(k) = \beta^2 k, \quad c = (1 - \beta^2)k.$$

The agent carries the fixed fraction β^2 of the remaining cake to the next date. These are the limits found in Example 3.2. The optimal stock path is $k_t = \beta^{2t} k_0$, so the remaining cake converges to zero.

Substituting the policy into the Bellman equation verifies the solution directly:

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

The Euler check. Since $f'(k) = 1$ and $c_{t+1} = \beta^2 c_t$,

$$\beta u'(c_{t+1}) f'(k_{t+1}) = \frac{\beta}{2\sqrt{\beta^2 c_t}} = \frac{1}{2\sqrt{c_t}} = u'(c_t).$$

Finally, V is bounded on $[0, \bar{k}]$ and the reward is bounded by $\sqrt{\bar{k}}$. The contraction theorem therefore shows that this fixed point is the unique bounded solution of the Bellman equation, and value function iteration from the finite-horizon problems converges to it.