4509 - Bridging Mathematics

Lecture 8: Discrete-Time Dynamic Optimization (Euler and Bellman)

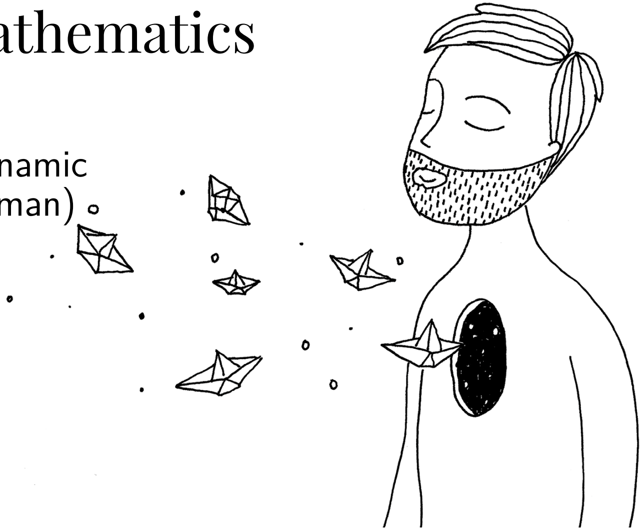

PAULO FAGANDINI

NOVA SCHOOL OF BUSINESS & ECONOMICS

Objectives

- Formulate infinite-horizon sequential optimization and the convergence of discounted sums.
- Derive the **Euler equation** and the **transversality condition** through the variational approach.
- Formulate the **Bellman equation** and prove existence and uniqueness of the value function through the contraction mapping theorem and Blackwell's conditions.

The spine: one infinite-horizon problem, two ways to crack it. The **variational** route perturbs the optimal path and gets the Euler equation, a trade-off between today and tomorrow. The **recursive** route writes the Bellman equation and exists because the Bellman operator is a contraction, Monday's Banach theorem cashed. They give the same answer.

The sequence problem

State space X, discount $\beta \in (0,1)$, feasible transitions $\Gamma : X \rightrightarrows X$:

$$V^*(x_0) = \sup_{\{x_{t+1}\}} \sum_{t=0}^{\infty} \beta^t F(x_t, x_{t+1}) \quad \text{s.t.} \quad x_{t+1} \in \Gamma(x_t),\ x_0 \text{ given.}$$

If $|F| \leq M$ then $\left| \sum \beta^t F \right| \leq \frac{M}{1-\beta} < \infty$: the same $\frac{1}{1-\beta}$ as the Banach bound. Discounting is what makes an infinite sum finite; $\beta < 1$ is doing exactly what it did in the contraction proof.

The variational derivation of the Euler equation

Take an optimal path $\{x_t^*\}$ with interior choices. Perturb only x_τ at one date $\tau \geq 1$; the only terms involving it are $\beta^{\tau-1} F(x_{\tau-1}^*, x_\tau) + \beta^\tau F(x_\tau, x_{\tau+1}^*)$. Optimality sets the derivative to zero and, dividing by $\beta^{\tau-1}$,

$$\boxed{F_2(x_{t-1}^*, x_t^*) + \beta\, F_1(x_t^*, x_{t+1}^*) = 0 \qquad \forall t \geq 1.}$$

Nudge one unit of the state at date t: $F_2 < 0$ is the marginal **cost** of investing today, $\beta F_1 > 0$ the discounted marginal **return** tomorrow. At the optimum they cancel, no free lunch from re-timing. This is the consumption-investment trade-off.

Transversality

The Euler equation is a **second-order** difference equation, so it needs two boundary conditions: the initial x_0, and a condition at infinity,

$$\lim_{T\to\infty} \beta^T F_1(x^*_T, x^*_{T+1})\, x^*_T = 0 \qquad \Big(\text{equivalently } \lim_{T\to\infty} \beta^T V'(x^*_T) x^*_T = 0\Big).$$

Do not leave valuable capital on the table forever (waste), and do not run a Ponzi scheme (borrow against a future you never repay). The present value of the terminal state must vanish.

The Bellman equation

$$V(x) = \max_{y \in \Gamma(x)} \{F(x,y) + \beta V(y)\}, \qquad g(x) = \arg\max_{y \in \Gamma(x)} \{F(x,y) + \beta V(y)\}.$$

Principle of optimality. Under boundedness and continuity, the functional-equation solution V equals the sequence-problem value V^*, and $x_{t+1} = g(x_t)$ achieves the sup.

Break an infinite problem into "today's payoff plus the discounted value of the whole future from tomorrow's state." The tail of an optimal plan is itself optimal from tomorrow's state, and that self-similarity *is* the Bellman equation. Recursion replaces an infinite choice with a single one.

The Bellman operator is a contraction

On $B(X)$, the bounded functions with sup-norm $\|V\| = \sup_x |V(x)|$ (a **complete** space, Lectures 1 and 2), define $(TV)(x) = \max_{y \in \Gamma(x)}\{F(x, y) + \beta V(y)\}$.

Theorem (Blackwell's sufficient conditions)

T is a contraction with modulus β if

1. *$V \leq W \Rightarrow TV \leq TW$ (monotonicity), and*
2. *$T(V + c) \leq TV + \beta c$ for constants $c \geq 0$ (discounting).*

Both are immediate for the Bellman operator: monotonicity because $F + \beta V \leq F + \beta W$, discounting because $T(V + c) = TV + \beta c$.

The keystone callback: Banach cashed

T is a contraction on a complete space, so by **Banach** (Lecture 2) there is a **unique** V^* with $TV^* = V^*$, and from any starting guess V_0, $T^k V_0 \to V^*$.

That iteration is **value-function iteration**, and Banach's error bound

$$\|T^k V_0 - V^*\| \leq \frac{\beta^k}{1-\beta}\|TV_0 - V_0\|$$

says it converges geometrically at rate β.

This is the theorem that lets computational macro just iterate the Bellman equation and trust the answer: existence, uniqueness and a convergent algorithm, all three from Monday.

Worked example: cake-eating — setup

You own a cake of size x_0. Eating c_t today leaves $x_{t+1} = x_t - c_t$ for tomorrow:

$$\max_{\{c_t\}_{t\geq 0}} \sum_{t=0}^{\infty} \beta^t \ln c_t \quad \text{s.t.} \quad x_{t+1} = x_t - c_t, \quad 0 \leq c_t \leq x_t, \quad x_0 \text{ given.}$$

In stacked form the state is the cake x_t and the choice is next period's stock x_{t+1}, with consumption $c_t = x_t - x_{t+1}$ and one-period return

$$F(x_t, x_{t+1}) = \ln(x_t - x_{t+1}).$$

The $\ln$ (Inada: $\ln'(0^+) = \infty$) forces an interior $c_t > 0$, so first-order conditions are necessary and sufficient. We now solve it **two ways**.

Route 1 — Euler: perturb a single date

Write the objective in stacked form, $\sum_{t\geq 0} \beta^t \ln(x_t - x_{t+1})$. Take the optimal path and nudge one stock x_τ ($\tau \geq 1$), holding all other dates fixed. Only **two** terms of the sum contain x_τ:

$$\underbrace{\beta^{\tau-1} \ln(x_{\tau-1} - x_\tau)}_{\text{return at date } \tau-1} + \underbrace{\beta^{\tau} \ln(x_\tau - x_{\tau+1})}_{\text{return at date } \tau}.$$

The stock x_τ plays two roles: it is what is *left over* after date $\tau - 1$, and what is *available* at date τ.

Route 1 — Euler: the first-order condition

Set the derivative with respect to x_τ to zero:

$$\beta^{\tau-1} \cdot \frac{-1}{x_{\tau-1} - x_\tau} \; + \; \beta^\tau \cdot \frac{1}{x_\tau - x_{\tau+1}} \; = \; 0.$$

Divide by $\beta^{\tau-1}$ and use $c_t = x_t - x_{t+1}$:

$$-\frac{1}{c_{\tau-1}} + \beta \, \frac{1}{c_\tau} = 0.$$

Marginal utility lost by saving one unit today = discounted marginal utility gained tomorrow. No re-timing improves the plan.

Route 1 — Euler: the consumption path

Rearrange the first-order condition:

$$\frac{1}{c_{\tau-1}} = \beta \frac{1}{c_\tau} \implies c_\tau = \beta\, c_{\tau-1} \implies \boxed{c_{t+1} = \beta\, c_t.}$$

This first-order linear recursion integrates to

$$c_t = \beta^t c_0.$$

Log utility has unit elasticity, so consumption falls at exactly the rate of impatience β.
One unknown, c_0, is still free.

Route 1 — Euler: unroll the stock

Feed $c_t = \beta^t c_0$ into the law of motion $x_{t+1} = x_t - c_t$ and sum from 0 to $t-1$:

$$x_t = x_0 - \sum_{s=0}^{t-1} c_s = x_0 - c_0 \sum_{s=0}^{t-1} \beta^s = x_0 - c_0 \frac{1-\beta^t}{1-\beta}.$$

The entire trajectory is now pinned down by the single number c_0; we still need one more condition to fix it.

Route 1 — Euler: transversality closes it

Transversality (leave no cake uneaten in the limit) requires $x_t \to 0$. Letting $t \to \infty$ above,

$$x_0 - \frac{c_0}{1-\beta} = 0 \implies c_0 = (1-\beta)\, x_0.$$

Back-substituting, $x_t = \beta^t x_0$ and $c_t = (1-\beta)\beta^t x_0$, i.e. the **policy**

$$\boxed{c_t = (1-\beta)\, x_t, \qquad x_{t+1} = \beta\, x_t.}$$

Check the TVC: $\beta^T F_1(x_T, x_{T+1})\, x_T = \beta^T \frac{x_T}{c_T} = \frac{\beta^T}{1-\beta} \longrightarrow 0.$ ✓

Route 2 — Bellman: the recursive problem

Write the same problem recursively. With next period's stock $y = x_{t+1}$ as the choice,

$$V(x) = \max_{0 \le y \le x} \left\{ \ln(x - y) + \beta V(y) \right\}.$$

$V(x)$ is the best lifetime utility attainable from a cake of size x: today's payoff $\ln(x - y)$ plus the discounted value of continuing optimally from y. We solve this functional equation by **guess and verify**.

Route 2 — Bellman: why guess log-linear?

The guess is not luck; log utility forces it. Scaling the cake by $\theta > 0$ scales every feasible path: from θx the feasible consumptions are exactly θ times those from x. Since $\ln(\theta c) = \ln c + \ln \theta$,

$$\sum_{t\geq 0} \beta^t \ln(\theta c_t) = \sum_{t\geq 0} \beta^t \ln c_t + \Big(\sum_{t\geq 0} \beta^t\Big) \ln \theta.$$

Maximizing both sides over feasible paths, the optimum scales the same way, so

$$V(\theta x) = V(x) + \frac{\ln \theta}{1-\beta}.$$

Set $x = 1$ and rename $\theta \to x$: $\quad V(x) = \underbrace{V(1)}_{=A} + \underbrace{\frac{1}{1-\beta}}_{=B} \ln x.$

Route 2 — Bellman: the guess

The scaling argument already hands us the log-linear form (and even predicts $B = \frac{1}{1-\beta}$). Take it as a guess with undetermined coefficients and verify:

$$V(x) = A + B \ln x, \qquad A, B \text{ unknown.}$$

Substitute the guess on the right-hand side; the object to maximize in y is

$$W(y) = \ln(x - y) + \beta\big(A + B \ln y\big).$$

Plan: maximize W over y, then *require* the result to equal $A + B \ln x$. Consistency fixes A, B (and should return $B = \frac{1}{1-\beta}$).

Route 2 — Bellman: the inner maximization

First-order condition $W'(y) = 0$:

$$-\frac{1}{x-y} + \frac{\beta B}{y} = 0.$$

Cross-multiply: $\beta B(x-y) = y$, so $\beta Bx = y(1+\beta B)$, giving the maximizer

$$y = \frac{\beta B}{1+\beta B}\,x, \qquad c = x - y = \frac{1}{1+\beta B}\,x.$$

The policy is linear in the state; the slope still depends on the unknown B.

Route 2 — Bellman: substitute back

Put the maximizer into W:

$$V(x) = \ln\left(\frac{x}{1+\beta B}\right) + \beta A + \beta B \ln\left(\frac{\beta B x}{1+\beta B}\right).$$

Expand each log with $\ln \frac{u x}{v} = \ln x + \ln u - \ln v$:

$$\begin{aligned} V(x) = {} & \big[\ln x - \ln(1+\beta B)\big] + \beta A \\ & + \beta B\big[\ln x + \ln(\beta B) - \ln(1+\beta B)\big]. \end{aligned}$$

Route 2 — Bellman: match the $\ln x$ coefficient

Group the terms in $\ln x$:

$$V(x) = (1 + \beta B) \ln x + (\text{constant in } x).$$

Consistency with the guess $V(x) = A + B \ln x$ forces the coefficients of $\ln x$ to agree:

$$B = 1 + \beta B \implies B(1 - \beta) = 1 \implies \boxed{B = \frac{1}{1 - \beta}.}$$

The constant terms give a separate equation for A (two slides on).

Route 2 — Bellman: the policy

With $B = \dfrac{1}{1-\beta}$ we get $1 + \beta B = \dfrac{1}{1-\beta}$, hence

$$\frac{\beta B}{1+\beta B} = \beta, \qquad \frac{1}{1+\beta B} = 1-\beta.$$

So the optimal policy is

$$\boxed{x_{t+1} = g(x) = \beta\, x, \qquad c = (1-\beta)\, x,}$$

identical to the Euler solution: same trade-off, reached from the recursive side.

Route 2 — Bellman: the value constant

Matching the constant terms, and using $\ln(1+\beta B) = -\ln(1-\beta)$ and $\ln(\beta B) = \ln\beta - \ln(1-\beta)$, they collapse to

$$A(1-\beta) = \ln(1-\beta) + \frac{\beta}{1-\beta}\ln\beta.$$

Therefore

$$A = \frac{1}{1-\beta}\left[\ln(1-\beta) + \frac{\beta}{1-\beta}\ln\beta\right], \qquad V(x) = \frac{1}{1-\beta}\ln x + A.$$

Not needed for the policy, but it completes the value function.

Route 3 — value-function iteration: iterate the operator

If you cannot guess, **iterate** $V_{k+1} = TV_k$ from any V_0. Start at $V_0 \equiv 0$:

$$V_1(x) = \max_{0 \le y \le x} \{ \ln(x - y) + \beta \cdot 0 \} = \ln x \quad \text{(best to save nothing, } y = 0\text{).}$$

One more step, now using $V_1(y) = \ln y$:

$$V_2(x) = \max_{y} \{ \ln(x - y) + \beta \ln y \} = (1 + \beta) \ln x + \text{const.}$$

The value keeps its log-linear shape; only the coefficient of $\ln x$ changes.

Route 3 — value-function iteration: the inner step

Each iterate keeps the log-linear shape, $V_k(x) = B_k \ln x + \text{const}$, so the maximization inside T is the *same* one as in guess-and-verify ("the inner maximization" slide), now with B_k in place of B. The FOC

$$-\frac{1}{x-y} + \frac{\beta B_k}{y} = 0 \implies y = g_k(x) = \frac{\beta B_k}{1+\beta B_k} x, \qquad c = x - y = \frac{1}{1+\beta B_k} x.$$

Substituting this maximizer back, $V_{k+1}(x) = (1 + \beta B_k) \ln x + \text{const}$, so the coefficient obeys

$$B_{k+1} = 1 + \beta B_k, \qquad B_1 = 1.$$

Route 3 — value-function iteration: convergence

The recursion $B_{k+1} = 1 + \beta B_k$ with $B_1 = 1$ unrolls to

$$B_k = 1 + \beta + \cdots + \beta^{k-1} = \frac{1-\beta^k}{1-\beta} \xrightarrow[k\to\infty]{} \frac{1}{1-\beta},$$

and the policy $g_k(x) = \dfrac{\beta B_k}{1+\beta B_k}\, x \longrightarrow \beta x$, the guess-and-verify answer.

Banach guarantees this: T is a contraction, so from *any* start the iterates converge to the unique V^* and its policy, geometrically at rate β.

Euler vs Bellman: one problem, one policy

Euler (variational)	**Bellman (recursive)**
Perturb the optimal path at one date.	Write $V = TV$, a fixed point of the operator.
Get a 2nd-order difference equation $c_{t+1} = \beta c_t$, closed by x_0 and the TVC.	Solve by guess-and-verify, or iterate $V_{k+1} = TV_k$ (VFI).
Good for analytics and intuition (marginal cost today = discounted return tomorrow).	Good for existence/uniqueness (Banach) and for computation.

Both deliver $\boxed{c_t = (1-\beta)x_t}$: consume the **annuity value** of the cake each period, never the principal at once. In Lecture 9 the continuous-time household does exactly this with lifetime wealth.

References

- N. L. Stokey, R. E. Lucas and E. C. Prescott, *Recursive Methods in Economic Dynamics*, Ch. 2, 3 and 4.
- D. Acemoglu, *Introduction to Modern Economic Growth*, Ch. 5 (Sections 5.1 to 5.4).
- R. K. Sundaram, *A First Course in Optimization Theory*, Ch. 11 and 12.