Graduate Macro Sequence: Dynamic Programming

Hui-Jun Chen

National Tsing Hua University
Department of Economics

February 23, 2026

Dynamic Programming

Lecture 3: What we need for computation

- We want to solve a Bellman equation: $V = TV$.
- Key idea: treat this as a **fixed point problem** for an operator T.
- Two questions:
  1. When does a fixed point exist and is it unique?
  2. If it exists, how do we **find it on a computer**?
- The workhorse theorem: Contraction Mapping Theorem.

Contraction Mapping Theorem

Contraction Mapping Theorem

- **Definition (Contraction Mapping).** Let (S, d) be a metric space and $T : S \to S$ be a mapping of S into itself. T is a contraction mapping with modulus β if for some $\beta \in (0, 1)$,

$$d(Tv_1, Tv_2) \leq \beta\, d(v_1, v_2) \qquad \text{for all } v_1, v_2 \in S.$$

Contraction Mapping Theorem

- **Contraction Mapping Theorem.** Let (S, d) be a **complete** metric space and suppose that $T : S \to S$ is a contraction mapping. Then, T has a **unique** fixed point $v^* \in S$ such that

$$Tv^* = v^* = \lim_{N \to \infty} T^N v_0 \qquad \text{for all } v_0 \in S.$$

- The beauty of CMT is that it is a constructive theorem: it not only tells us existence/uniqueness of v^*, but also shows us how to find it ($v_{n+1} = Tv_n$).

(Optional) Why $T^N v_0$ converges: the key inequality

Let $v_{n+1} = Tv_n$. Because T is a contraction,

$$d(v_{n+1}, v_n) = d(Tv_n, Tv_{n-1}) \leq \beta\, d(v_n, v_{n-1}) \Rightarrow d(v_{n+1}, v_n) \leq \beta^n d(v_1, v_0).$$

Then for any $m \geq 1$,

$$d(v_{n+m}, v_n) \leq \sum_{j=0}^{m-1} d(v_{n+j+1}, v_{n+j}) \leq \sum_{j=0}^{m-1} \beta^{n+j} d(v_1, v_0) \leq \frac{\beta^n}{1-\beta}\, d(v_1, v_0).$$

So $\{v_n\}$ is Cauchy; completeness gives convergence to some $v^* \in S$. (Then show $Tv^* = v^*$, and uniqueness follows from contraction.)

How to use CMT in Dynamic Programming

Step 1: choose the function space

In DP, $v(\cdot)$ is a **function**. A standard choice:

$$S = \mathcal{B}(X) \equiv \{v : X \to \mathbb{R} \ : \ \sup_{x \in X} |v(x)| < \infty\}, \quad d(v, w) = \|v - w\|_\infty .$$

- $\left(\mathcal{B}(X), \|\cdot\|_\infty\right)$ is a **complete** metric space.
- So CMT applies if we can show T is a **contraction** on $\mathcal{B}(X)$.

Step 2: Bellman operator

Consider a discounted DP:

$$(TV)(x) \equiv \max_{a \in \Gamma(x)} \left\{ u(x,a) + \beta\, \mathbb{E}\big[V(x') \mid x, a\big] \right\}, \qquad \beta \in (0,1).$$

- If we can show T is a contraction on $\mathcal{B}(X)$, then CMT gives $\exists!\ V^*$ and $V^* = \lim_{N\to\infty} T^N V_0$.
- Question: **How do we show T is a contraction?**

Blackwell’s sufficient conditions

Intuition before Blackwell: discounting is geometric shrinkage

- Recall the geometric series:

$$1 + \beta + \beta^2 + \cdots = \begin{cases} \dfrac{1}{1-\beta}, & 0 < \beta < 1, \\ \infty, & \beta \geq 1. \end{cases}$$

- Economic meaning: discounting makes the infinite future *finite* because weights shrink geometrically.
- DP intuition: suppose your continuation value is wrong by at most m everywhere:

$$\|V - W\|_\infty \leq m.$$

Then the one-period-ahead value difference coming purely from the future is at most

$$\beta\, \mathbb{E}[V(x') - W(x')] \leq \beta m.$$

- Repeating the logic across iterations gives a geometric tightening:

$$\|T^n V - T^n W\|_\infty \;\leq\; \beta^n \, \|V - W\|_\infty \, .$$

Takeaway: $\beta < 1$ is the "geometric force" behind contraction and convergence.

Blackwell's sufficient conditions (sufficient, not necessary)

- CMT requires T to be a contraction.
- Blackwell provides sufficient conditions to verify contraction easily.
- They are **not necessary**: an operator can be a contraction even if these fail.

Blackwell's sufficient conditions

Let $T : \mathcal{B}(X) \to \mathcal{B}(X)$. Suppose:

1. Monotonicity: $V \leq W \Rightarrow TV \leq TW$.
2. Discounting: there exists $\beta \in (0, 1)$ s.t. for any constant $a \geq 0$,

$$T(V + a) \leq TV + \beta a.$$

Then T is a contraction with modulus β under $||\cdot||_\infty$.

Why Blackwell is useful here: discounting comes from β

For the Bellman operator

$$(TV)(x) = \max_{a \in \Gamma(x)} \{u(x,a) + \beta \mathbb{E}[V(x') \mid x, a]\},$$

- **Monotonicity:** if $V \leq W$, then $\mathbb{E}[V] \leq \mathbb{E}[W]$, so the RHS is smaller for every action, hence $TV \leq TW$.
- **Discounting:** add a constant $c \geq 0$ to V:

$$T(V+c)(x) = \max_a \{u(x,a) + \beta \mathbb{E}[V(x') + c]\} = \max_a \{u(x,a) + \beta \mathbb{E}[V(x')]\} + \beta c = TV(x) + \beta c.$$

So Blackwell implies: $\|TV - TW\|_\infty \leq \beta \|V - W\|_\infty$.

Value Function Iteration and acceleration

Standard Value Function Iteration (VFI)

- Once we know T is a contraction, CMT implies:

$$V^* = \lim_{N \to \infty} T^N V_0 \quad \text{for any initial guess } V_0.$$

- **Algorithm (Standard VFI):**
  1. Choose $V_0 \in \mathcal{B}(X)$.
  2. Iterate $V_{n+1} = TV_n$.
  3. Stop when $||V_{n+1} - V_n||_\infty$ is small.
- When $\beta \approx 1$, VFI can be very slow.

A practical error bound from contraction

For a contraction with modulus β,

$$\|V^* - V_n\|_\infty \;\leq\; \frac{1}{1-\beta}\, \|V_{n+1} - V_n\|_\infty \,.$$

- This gives a stopping rule: to be within ε, it is enough to have $\|V_{n+1} - V_n\|_\infty < (1-\beta)\varepsilon$.

Howard's Policy Iteration

- Howard's policy iteration follows from two key observations about VFI:
  - The **maximization step** is typically much more costly (in computation time) than the **evaluation step**.
  - But the evaluation step uses the updated decision rule $\tilde{s}_n(k_i, z_j)$ for only one period (since decisions after tomorrow are embedded in V_n on the RHS).

- **Policy Iteration:** repeat the evaluation step multiple times between each maximization step.

- **Definition:** for a given value function J and a decision rule w, define **Howard's mapping** $\tilde{T}_w$ as the operator that "plugs in w" into the RHS of the Bellman equation:

$$\tilde{T}_w J(k_i, z_j) \equiv u\Big(\underbrace{e^{z_j} k_i^{\alpha} + (1-\delta)k_i - w(k_i, z_j)}_{\equiv c(k_i, z_j; w)}\Big) + \beta\, \mathbb{E}\Big[J\big(w(k_i, z_j), z'\big) \mid z_j\Big] .$$

Howard's Policy Iteration: Cont'd

- We will be interested in applying the Howard mapping repeatedly, so for $m = 1, \dots, M$, let

$$J_{m+1} \equiv \tilde{T}_w J_m(k_i, z_j)$$

denote the updated value function.

- Howard (1962)'s key insight is that $\tilde{T}_w$ **is also a contraction mapping** with modulus β.

$$\Rightarrow \text{ applying } \tilde{T}_w \text{ repeatedly converges to a fixed point at rate } \beta.$$

- Of course, this fixed point is **not** the solution of the original Bellman equation (since the policy function w is held fixed while only the value is updated).

- But $\tilde{T}_w$ is an operator that is much cheaper to apply, so it is natural to apply it more than once.

Algorithmus: VFI with Policy Iteration

1. Set $n = 0$. Choose an initial guess $V_0 \in \mathcal{S}$.
2. **Maximization Step:** obtain $\tilde{s}_n$ (the updated decision rule) from V_n.
3. **Howard Step:** set $J_0 \equiv V_n$, and iterate on

$$J_{m+1} \equiv \tilde{T}_{\tilde{s}_n} J_m(k_i, z_j), \qquad m = 0, 1, 2, \ldots$$

   then set $V_{n+1} = J^* \equiv \lim_{m\to\infty} J_m$.

4. Stop if convergence criterion is satisfied:

$$\|V_{n+1} - V_n\|_\infty < \text{toler.}$$

   Otherwise, increase n and return to step 2.

- **Note:** it is often possible to obtain the fixed point J^* in a finite number of steps.

VFI with Modified Policy Iteration (MPI) Algorithm

- Modify Step 3 of Howard's algorithm above:
  - **Modified Howard Step:** set $J_0 \equiv V_n$, and iterate on

$$J_{m+1} \equiv \tilde{T}_{\tilde{s}_n} J_m(k_i, z_j), \qquad m = 0, 1, 2, \ldots, M < \infty.$$

    Choose a **moderate** value for M (by experimentation; use smaller M for more challenging problems). Then set $V_{n+1} = J_M$.

- The choice of M is a **key tuning parameter** in practice.

MacQueen–Porteus (MQP) bounds: why we care

- Iterative algorithms need a stopping rule.
- MQP bounds provide upper and lower bounds on V^* (especially clean in discrete state problems).

Let $V_{n+1} = TV_n$ and define increments

$$\Delta_n(x) \equiv V_n(x) - V_{n-1}(x), \quad \underline{\Delta}_n \equiv \min_x \Delta_n(x), \quad \overline{\Delta}_n \equiv \max_x \Delta_n(x).$$

MacQueen–Porteus (MQP) bounds: the bound

Define correction terms

$$\underline{c}_n \equiv \frac{\beta}{1-\beta}\underline{\Delta}_n, \qquad \bar{c}_n \equiv \frac{\beta}{1-\beta}\overline{\Delta}_n.$$

Then (under standard conditions; most transparent in finite/discrete state problems),

$$V_n(x) + \underline{c}_n \;\le\; V^*(x) \;\le\; V_n(x) + \bar{c}_n \qquad \text{for all } x.$$

- Practical stopping: stop when $\bar{c}_n - \underline{c}_n$ is small.

Wrap-up

Wrap-up

- CMT: if T is a contraction on a complete metric space, then a unique fixed point exists
- In DP, the job is to show the Bellman operator T is a contraction.
- Blackwell conditions are a sufficient (not necessary) way to construst contraction: discounting with β.
- Algorithms:
  - VFI (guaranteed but slow when $\beta \approx 1$)
  - Howard / Modified Policy Iteration (speed)
  - MQP bounds (error bounds / stopping rules)