Dynamic Programming I

Fixed Points, Contractions, and Optimal Stopping

Sebastian Dyrda, University of Toronto

ECO2101 Lecture 2 | March 8, 2026

Roadmap for Lectures 2–3

Two lectures, two throughline applications:

Lecture 2 (today): A firm with stochastic productivity

- Fixed-point theory, contractions, Markov dynamics, optimal stopping
- Threshold policies, monotonicity, continuation values
- Sargent & Stachurski, *Dynamic Programming*, Chapters 1–4
- $\rightarrow$ Lectures 4–7: Hopenhayn (1992), Khan-Thomas (2008)

Lecture 3 (next week): A household with stochastic income

- MDPs, VFI / HPI / OPI, extreme value shocks, EGM / ECM
- Application: optimal savings—the individual problem inside Aiyagari (1994)
- Sargent & Stachurski, Chapter 5; Carroll (2006); Maliar & Maliar (2013)
- $\rightarrow$ Lectures 8–10: HANK and sequence-space Jacobians

Roadmap

Firm Valuation without Decisions

Fixed-Point Theory

Optimal Stopping: The Exit Option

Firm Valuation with Exit

Monotonicity and Threshold Policies

Continuation Values and Comparative Statics

The Firm Problem — Setup

A firm operates in discrete time with exogenous stochastic productivity.

Primitives:

- A finite set $\mathsf{Z} \subset \mathbb{R}$ of productivity levels, $\mathsf{Z} = \{z_1, \dots, z_n\}$
- Productivity $(Z_t)_{t \geq 0}$ is Q-Markov on Z, with $Q \in \mathcal{M}(\mathbb{R}^{\mathsf{Z}})$
- Profits $\pi_t = \pi(Z_t)$ for a fixed function $\pi \in \mathbb{R}^{\mathsf{Z}}$
- Constant interest rate $r > 0$; discount factor $\beta := 1/(1+r)$

Because Z is finite, π is bounded. With $\beta \in (0,1)$, the discounted sum $\sum_{t \geq 0} \beta^t \pi(Z_t)$ is finite.

Notation: Since Z is finite, the function space $\mathbb{R}^{\mathsf{Z}}$ is identified with $\mathbb{R}^n$ via $\mathbb{R}^{\mathsf{Z}} \ni u \leftrightarrow (u(z_1), \dots, u(z_n)) \in \mathbb{R}^n$. Operations are pointwise, e.g., $(u \vee v)(z) = \max\{u(z), v(z)\}$. We equip $\mathbb{R}^{\mathsf{Z}}$ with the **supremum norm** $\|u\|_\infty := \max_i |u(z_i)|$.

No decisions yet. The firm simply operates and receives profits. **Question:** What is it worth?

Reference: §1.2.4 (Lemma 1.2.4); §3.2.2.2.

The Markov Operator and Conditional Expectations

For $Q \in \mathcal{M}(\mathbb{R}^{\mathsf{Z}})$ and any $h \in \mathbb{R}^{\mathsf{Z}}$, the **Markov operator** Q acts on functions by

$$(Qh)(z) = \sum_{z' \in \mathsf{Z}} h(z')\, Q(z, z') = \mathbb{E}[h(Z_{t+1}) \mid Z_t = z].$$

Powers give multi-step conditional expectations: $(Q^t h)(z) = \mathbb{E}_z[h(Z_t)]$.

Q is a **stochastic matrix**: $Q \geq 0$ and $Q\mathbb{1} = \mathbb{1}$, so each row sums to 1. For nonnegative matrices, the spectral radius is bounded by the minimum/maximum row sums; hence $\rho(Q) = 1$.

The **spectral radius** of any square matrix A is $\rho(A) := \max\{|\lambda| : \lambda$ is an eigenvalue of $A\}$.

A key property: $\rho(\alpha A) = |\alpha|\, \rho(A)$ for any scalar α (Exercise 1.2.10).

Reference: §3.2.1, equations (3.12)–(3.14); §1.2.1.4, equation (1.15); Lemma 2.3.2 / Exercise 2.3.2(ii).

The Neumann Series Lemma

Theorem 1.2.1 (Neumann Series Lemma)

If A is an $n \times n$ matrix with $\rho(A) < 1$, then $I - A$ is nonsingular and

$$(I - A)^{-1} = \sum_{k \geq 0} A^k.$$

The linear system $u = Au + b$ has the unique solution $u^* = (I - A)^{-1} b$.

Gelfand's formula (Lemma 1.2.2): $\left\|A^k\right\|^{1/k} \to \rho(A)$. Hence $\rho(A) < 1 \implies \left\|A^k\right\| \to 0$ and $\sum A^k$ converges.

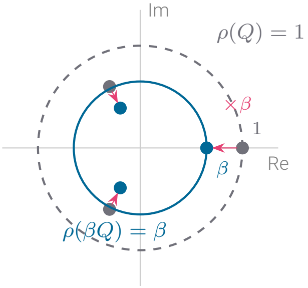

Multiplying by $\beta < 1$ pulls all eigenvalues inside the unit disk.

Reference: Theorem 1.2.1, p. 18; Lemma 1.2.2, p. 19.

Firm Valuation via the Neumann Series

Present value of profits

$$v(z) = \mathbb{E}_z \left[\sum_{t \geq 0} \beta^t \pi(Z_t) \right] = \sum_{t=0}^{\infty} \beta^t (Q^t \pi)(z)$$

Bellman recursion (no decisions)

$$v(z) = \pi(z) + \beta \sum_{z' \in Z} v(z') Q(z, z') = \pi(z) + \beta (Qv)(z)$$

Interpretation
value today = current profit + discounted expected continuation value.
Neumann series gives the closed form

$$v = (I - \beta Q)^{-1} \pi.$$

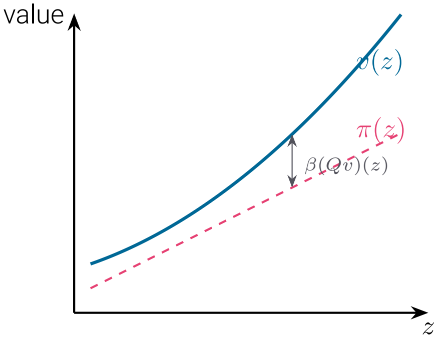

Gap $v(z) - \pi(z)$ is the discounted value of expected future profits.

Monotonicity of Firm Value

Question: If higher z means higher profits *and* predicts higher future z, is the firm worth more?

Definition (§3.2.1.3)

Q is **monotone increasing** if $h \in i\mathbb{R}^{\mathsf{Z}} \implies Qh \in i\mathbb{R}^{\mathsf{Z}}$, where $i\mathbb{R}^{\mathsf{Z}}$ denotes the increasing functions on Z.

Exercise 3.2.6. If $\pi \in i\mathbb{R}^{\mathsf{Z}}$ and Q is monotone increasing, then $v = (I - \beta Q)^{-1}\pi$ is increasing on Z.

Proof. Q monotone increasing implies $Q^t\pi \in i\mathbb{R}^{\mathsf{Z}}$ for all t. Hence $v = \sum_{t \geq 0} \beta^t Q^t \pi$ is increasing. $\square$

Application: If z follows a discretized AR(1) with autocorrelation $\rho \geq 0$ (Tauchen's method, §3.1.3), then Q is monotone increasing (Exercise 3.2.2).

Reference: §3.2.1.3, Exercise 3.2.6, p. 96.

Roadmap

Firm Valuation without Decisions

Fixed-Point Theory

Optimal Stopping: The Exit Option

Firm Valuation with Exit

Monotonicity and Threshold Policies

Continuation Values and Comparative Statics

Contractions and Banach's Fixed-Point Theorem

Definition (§1.2.2.3)

Let $U \subseteq \mathbb{R}^n$ be nonempty and $\|\cdot\|$ a norm. A self-map $T : U \to U$ is a **contraction of modulus** λ if there exists $\lambda \in [0, 1)$ such that

$$\|Tu - Tv\| \leq \lambda \|u - v\| \qquad \text{for all } u, v \in U.$$

Theorem 1.2.3 (Banach's Contraction Mapping Theorem)

If U is closed in $\mathbb{R}^n$ and T is a contraction of modulus λ on U, then T has a **unique fixed point** u^* in U and

$$\left\|T^k u - u^*\right\| \leq \lambda^k \|u - u^*\| \qquad \text{for all } k \in \mathbb{N},\ u \in U.$$

In particular, T is **globally stable**: $T^k u \to u^*$ for any $u \in U$.

Three deliverables: **existence**, **uniqueness**, **constructive algorithm** (successive approximation converges geometrically at rate λ^k).

Banach's Theorem: Geometric Intuition

Proof idea:

1. Fix $u_0 \in U$, set $u_m = T^m u_0$
2. Contraction bound $\implies$ (u_m) is Cauchy
3. Completeness of $\mathbb{R}^n$ $\implies$ limit u^* exists
4. U closed $\implies$ $u^* \in U$
5. Continuity of T $\implies$ $Tu^* = u^*$
6. Uniqueness: if $Tu = u$ and $Tv = v$, then $\|u - v\| = \|Tu - Tv\| \leq \lambda \|u - v\|$ with $\lambda < 1$

For Bellman operators: $\lambda = \beta$. When $\beta \approx 0.99$, convergence is slow ($\beta^{100} \approx 0.37$). Faster algorithms in Lecture 3.

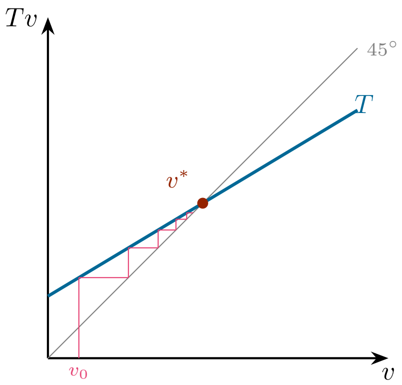

Successive approximation: $v_{k+1} = Tv_k \to v^*$.

Reference: Theorem 1.2.3, p. 23; §1.2.3 (successive approximation), p. 24.

Summary of Part I

Concept	**Firm application**	**Reference**
$\mathbb{R}^{\mathsf{Z}} \leftrightarrow \mathbb{R}^n$, $\|\cdot\|_\infty$	Productivity grid, functions as vectors	§1.2.4
Markov operator $Qh = \sum h(z')Q(z,z')$	Conditional expectations of profits	§3.2.1
Neumann series: $(I - \beta Q)^{-1}\pi$	Firm value without decisions	Lemma 3.2.1
Contraction, Banach FPT	Unique fixed point, λ^k convergence	Thm. 1.2.3

All ingredients are in place. **Next:** give the firm a **choice**—the option to **exit**.
This turns the *linear* fixed-point problem $v = \pi + \beta Q v$ into a *nonlinear* one:

$$v(z) = \max\left\{s,\ \pi(z) + \beta \sum_{z'} v(z')\, Q(z,z')\right\}.$$

The operator $Tv = s \vee (\pi + \beta Q v)$ involves a $\max$, which is nonlinear. Is T still a contraction?

Roadmap

Firm Valuation without Decisions

Fixed-Point Theory

Optimal Stopping: The Exit Option

Firm Valuation with Exit

Monotonicity and Threshold Policies

Continuation Values and Comparative Statics

The Optimal Stopping Problem

Definition (§4.1.1.1)

An **optimal stopping problem** is a tuple $\mathcal{S} = (\beta, P, c, e)$:

- a discount factor $\beta \in (0, 1)$, a Markov operator $P \in \mathcal{M}(\mathbb{R}^{\mathcal{X}})$,
- a **continuation reward** $c \in \mathbb{R}^{\mathcal{X}}$, an **exit reward** $e \in \mathbb{R}^{\mathcal{X}}$.

Timing: Observe state X_t, choose to **continue** (receive $c(X_t)$, proceed) or **stop** (receive $e(X_t)$, terminate).

Policies: $\sigma : \mathcal{X} \to \{0, 1\}$, where $\sigma(x) = 1$ means 'stop." $\Sigma =$ all such maps.

The firm problem is an optimal stopping problem with:

$$\mathcal{X} = \mathsf{Z} \quad P = Q \quad c = \pi \quad e \equiv s \text{ (constant)} \quad \beta = 1/(1 + r)$$

Reference: §4.1.1.1, p. 105–106; §4.1.2.1, p. 111.

Lifetime Values and Policy Operators

Fix $\sigma \in \Sigma$. The σ**-value function** v_σ satisfies

$$v_\sigma(x) = \sigma(x)\, e(x) + (1 - \sigma(x)) \left[c(x) + \beta \sum_{x'} v_\sigma(x')\, P(x, x') \right].$$

Policy operator $T_\sigma : \mathbb{R}^{\mathcal{X}} \to \mathbb{R}^{\mathcal{X}}$: same formula applied to any $v \in \mathbb{R}^{\mathcal{X}}$. Then $v_\sigma = T_\sigma v_\sigma$.

In matrix form: $T_\sigma v = r_\sigma + L_\sigma v$, where

$$r_\sigma(x) := \sigma(x) e(x) + (1 - \sigma(x)) c(x), \qquad L_\sigma(x, x') := \beta(1 - \sigma(x)) P(x, x').$$

Since L_σ is nonnegative and its row sums satisfy $\sum_{x'} L_\sigma(x, x') \le \beta$, we have $\rho(L_\sigma) \le \beta < 1$, so $v_\sigma = (I - L_\sigma)^{-1} r_\sigma$.

Proposition 4.1.1

For any $\sigma \in \Sigma$, T_σ is a **contraction of modulus** β on $\mathbb{R}^{\mathcal{X}}$ under $\|\cdot\|_\infty$.

Reference: §4.1.1.2–4.1.1.3, Proposition 4.1.1, p. 108.

The Upper Envelope Lemma - $\max$ over policies

Lemma 2.2.3 (Upper Envelope Lemma)

Let $\{T_\sigma\}_{\sigma \in \Sigma}$ be a **finite** family of contractions of moduli λ_σ on $V \subseteq \mathbb{R}^{\mathsf{X}}$. Then $T = \bigvee_\sigma T_\sigma$ is a contraction of modulus $\max_\sigma \lambda_\sigma$.

Proof. By Lemma 2.2.2: $|\max f - \max g| \leq \max |f - g|$. Hence

$$\|Tu - Tv\|_\infty \leq \max_\sigma \|T_\sigma u - T_\sigma v\|_\infty \leq \left(\max_\sigma \lambda_\sigma\right) \|u - v\|_\infty \, .$$

.

Pedagogical note: For continuous action sets, the Bellman operator is a $\sup$-envelope and one needs extra conditions to replace $\sup$ by $\max$.

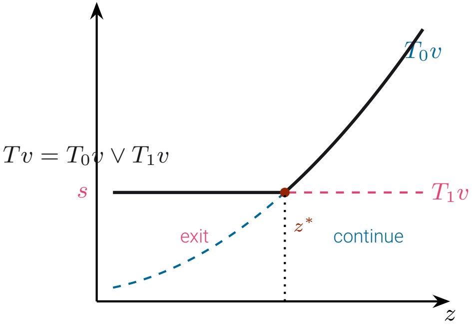

The Bellman operator $Tv = s \vee (\pi + \beta Q v)$ is the upper envelope of the “stop” and “continue” operators.

Reference: Lemma 2.2.3, p. 59; Lemma 2.2.2, p. 58.

The Bellman Operator

The **value function:** $v^*(x) := \max_{\sigma \in \Sigma} v_\sigma(x)$.

The **Bellman operator** $T : \mathbb{R}^{\mathcal{X}} \to \mathbb{R}^{\mathcal{X}}$:

$$(Tv)(x) = \max \left\{ e(x),\ c(x) + \beta \sum_{x'} v(x')\, P(x, x') \right\}.$$

Pointwise: $Tv = e \vee (c + \beta P v)$.

This is the **upper envelope** of two policy operators:

$$(T_1 v)(x) = e(x) \qquad \text{(stop: a constant map, contraction of modulus 0),}$$
$$(T_0 v)(x) = c(x) + \beta (Pv)(x) \qquad \text{(continue: linear contraction of modulus } \beta\text{).}$$

By the upper envelope lemma (Lemma 2.2.3): $T = T_1 \vee T_0$ is a **contraction of modulus** $\max\{0, \beta\} = \beta$.

Reference: §4.1.1.5, equation (4.8), p. 109.

The Fundamental Theorem for Optimal Stopping

Proposition 4.1.2

Let $\mathcal{S}$ be an optimal stopping problem with Bellman operator T and value function v^*. Then:

1. T is a contraction of modulus β on $\mathbb{R}^{\mathcal{X}}$ under $\|\cdot\|_\infty$.
2. The unique fixed point of T in $\mathbb{R}^{\mathcal{X}}$ is v^*.

Proof of (ii): Let $\bar{v}$ be the unique fixed point of T. We show $\bar{v} = v^*$.

Step 1 ($\bar{v} \leq v^*$): Let σ be $\bar{v}$-greedy. Then $T_\sigma \bar{v} = T\bar{v} = \bar{v}$. Since v_σ is the *only* fixed point of T_σ, we get $\bar{v} = v_\sigma \leq v^* = \vee_\sigma v_\sigma$.

Step 2 ($\bar{v} \geq v^*$): For any σ: $T_\sigma v \leq Tv$ for all v. Since T is order-preserving and globally stable, Proposition 2.2.7 gives $v_\sigma \leq \bar{v}$. Taking $\max_\sigma$: $v^* \leq \bar{v}$. □

Reference: Proposition 4.1.2, p. 109–110; Proposition 2.2.7, p. 67.

Bellman's Principle of Optimality

Proposition 4.1.3 (Bellman's Principle)

A policy $\sigma \in \Sigma$ is **optimal** if and only if it is v^*-**greedy**.

A policy σ is v-**greedy** if, for all $x \in \mathcal{X}$,

$$\sigma(x) \in \operatorname*{argmax}_{a \in \{0,1\}} \left\{ a\, e(x) + (1-a) \left[c(x) + \beta \sum_{x'} v(x')\, P(x, x') \right] \right\}.$$

Proof. By Proposition 4.1.2:

$$\sigma \text{ is } v^*\text{-greedy} \iff T_\sigma v^* = T v^* = v^* \iff v_\sigma = v^* \iff \sigma \text{ is optimal.}$$

Reference: Proposition 4.1.3, p. 110; equation (4.9). □

Economic content: To act optimally, one only needs the value function v^*. At each state, compare exit to continuation and pick the higher one.

The Logic of Dynamic Programming

Why Bellman's principle is not circular

Step 1. Define the **Bellman operator**

$$(Tv)(x) = \max \left\{ e(x),\ c(x) + \beta \sum_{x'} v(x')\, P(x, x') \right\}.$$

Step 2. Prove T is a **contraction** on $\mathbb{R}^{\mathcal{X}}$ under $\|\cdot\|_{\infty}$.
Step 3. By Banach's theorem, the fixed point exists and is unique:

$$Tv^* = v^*.$$

Step 4. Then show: **policies greedy w.r.t. v^* are optimal** (*Bellman's principle*).

Key takeaway: the **value function comes first**; the **optimal policy comes after**.

Computation: value iteration + greedy policy recovery

$$v_{k+1} = Tv_k \ \longrightarrow\ v^*, \qquad \sigma^*(x) \in \arg\max_{a \in \{0,1\}} \left\{ a\, e(x) + (1-a) \left[c(x) + \beta \sum_{x'} v^*(x') P(x, x') \right] \right\}.$$

Reference: Banach's theorem (§1.2.2.3) + Proposition 4.1.2–4.1.3.

Roadmap

- Firm Valuation without Decisions
- Fixed-Point Theory
- Optimal Stopping: The Exit Option
- Firm Valuation with Exit
- Monotonicity and Threshold Policies
- Continuation Values and Comparative Statics

The Firm Problem with Exit

Setup (§4.1.2): Same firm as before, now with an **exit option**.

- Each period: **continue** (receive $\pi(Z_t)$) or **exit** (receive scrap value $s > 0$)

The **Bellman equation**:

$$\boxed{v(z) = \max\left\{ s,\ \pi(z) + \beta \sum_{z' \in \mathsf{Z}} v(z')\, Q(z, z') \right\} \qquad (z \in \mathsf{Z}).}$$

Pointwise: $Tv = s \vee (\pi + \beta Q v)$.

By Propositions 4.1.2–4.1.3:

- T is a contraction of modulus β (upper envelope of $T_1 v = s$ and $T_0 v = \pi + \beta Q v$)
- v^* is the unique fixed point of T
- σ is optimal $\iff$ σ is v^*-greedy: $\sigma^*(z) = \mathbb{1}\{s \geq h^*(z)\}$, where $h^*(z) := \pi(z) + \beta(Qv^*)(z)$

Reference: §4.1.2.1, p. 111–112.

Exit Adds Value (Option Value)

Benchmark (no exit). Let

$$w := (I - \beta Q)^{-1}\pi$$

be the value of operating forever (the solution to $w = \pi + \beta Q w$ from Part I).

Key inequality (weak, always true) Exiting is an *option*. Since the policy "never exit" is feasible,

$$v^*(z) \;\geq\; w(z) \qquad \text{for all } z \in \mathsf{Z}.$$

One-line proof. Let $\sigma_0 \equiv 0$ (always continue). Then $v_{\sigma_0} = w$ and

$$v^* = \max_{\sigma \in \Sigma} v_\sigma \;\geq\; v_{\sigma_0} = w.$$

Refinement (optional). If Q is **everywhere positive** and $s > w(z)$ for some z, then

$$v^*(z') > w(z') \quad \text{for all } z' \in \mathsf{Z}$$

(Exercise 4.1.7: the option is strictly valuable everywhere).

Reference: §4.1.2.2; Exercise 4.1.7.

Exit Adds Value

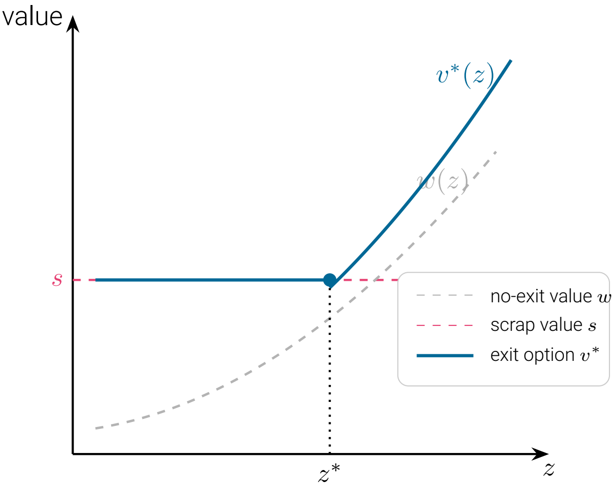

Roadmap

Firm Valuation without Decisions

Fixed-Point Theory

Optimal Stopping: The Exit Option

Firm Valuation with Exit

Monotonicity and Threshold Policies

Continuation Values and Comparative Statics

Monotonicity of the Value Function

Lemma 4.1.4

If $e, c \in i\mathbb{R}^{\mathcal{X}}$ and P is monotone increasing, then v^* and the continuation value h^* are both **increasing**.

Proof. $Tv = e \vee (c + \beta P v)$. Since P monotone increasing, T maps $i\mathbb{R}^{\mathcal{X}}$ into itself (i.e., $i\mathbb{R}^{\mathcal{X}}$ is **invariant** under T).
The set $i\mathbb{R}^{\mathsf{Z}}$ is **closed** in $\mathbb{R}^{\mathsf{Z}}$. Since T is globally stable on the closed set $\mathbb{R}^{\mathsf{Z}}$ and invariant on the closed subset $i\mathbb{R}^{\mathsf{Z}}$, the unique fixed point $v^* \in i\mathbb{R}^{\mathsf{Z}}$ (Exercise 1.2.18). $\square$

For the firm: $e \equiv s$ is constant (weakly increasing) and $c = \pi$. So under $\pi \in i\mathbb{R}^{\mathsf{Z}}$ and Q monotone increasing:

- v^* is increasing in z and $h^*(z) = \pi(z) + \beta(Qv^*)(z)$ is increasing in z

Reference: Lemma 4.1.4, p. 114–115; Example 4.1.3; Exercise 1.2.18, p. 22.

Threshold Policies

Exercise 4.1.9: If e is decreasing and h^* is increasing on $\mathcal{X}$, then σ^* is decreasing.

For the firm: $e \equiv s$ (constant, weakly decreasing) and h^* is increasing. So σ^* is **decreasing**: the firm exits at **low** z and continues at **high** z.

Since Z $\subset \mathbb{R}$ is totally ordered, the policy is of **threshold type**.

Define the reservation productivity on the grid by

$$z^* := \min\{z \in \mathsf{Z} : h^*(z) \geq s\}.$$

(On a finite grid, $h^*(z^*)$ may be strictly larger than s; the figure illustrates the continuous-state geometry.)

Tie-breaking convention: **continue** when $h^*(z) = s$.

$$\sigma^*(z) = \begin{cases} 1 & \text{(exit)} \qquad z < z^* \\ 0 & \text{(continue)} \qquad z \geq z^* \end{cases}$$

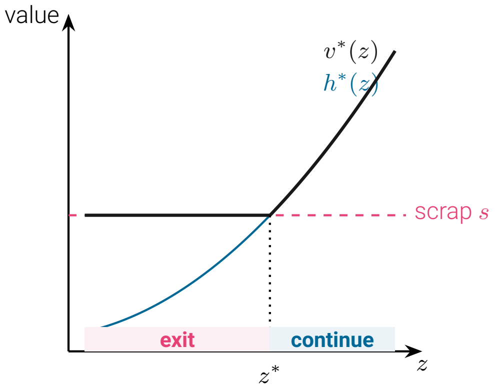

$v^* = s \vee h^*$; exit when $h^*(z) < s$.

Roadmap

Firm Valuation without Decisions

Fixed-Point Theory

Optimal Stopping: The Exit Option

Firm Valuation with Exit

Monotonicity and Threshold Policies

Continuation Values and Comparative Statics

The Continuation Value Operator

Instead of iterating on v via T, we can iterate directly on h^*.

From the Bellman equation: $v^*(x') = e(x') \vee h^*(x')$, so

$$h^*(x) = c(x) + \beta \sum_{x'} \left[e(x') \vee h^*(x')\right] P(x, x').$$

Define the **continuation value operator** $C : \mathbb{R}^{\mathcal{X}} \to \mathbb{R}^{\mathcal{X}}$:

$$(Ch)(x) = c(x) + \beta \sum_{x'} \max\{e(x'),\ h(x')\}\, P(x, x').$$

Proposition 4.1.5

C is a contraction of modulus β on $\mathbb{R}^{\mathcal{X}}$ with unique fixed point h^*.

Advantage: When exit reward contains an IID component, the continuation value has lower dimensionality (§4.1.4.2).

Reference: §4.1.4.1, Proposition 4.1.5, pp. 117–118.

Continuation Values: Dimensionality Reduction for the Firm

Suppose scrap value S_t is **IID** with density φ (e.g., fluctuates with asset prices).

The Bellman operator acts on $\mathsf{Z} \times \mathsf{S}$ space (high-dimensional). But the continuation value h^* depends only on z:

$$(Ch)(z) = \pi(z) + \beta \sum_{z'} \int \max\{s',\, h(z')\}\, \varphi(s')\, ds'\; Q(z, z') \qquad (z \in \mathsf{Z}).$$

$C : \mathbb{R}^{\mathsf{Z}} \to \mathbb{R}^{\mathsf{Z}}$ – dimension $|\mathsf{Z}|$ instead of $|\mathsf{Z}| \times |\mathsf{S}|$.

The optimal policy takes the form $\sigma^*(z, s) = \mathbb{1}\{s \geq h^*(z)\}$: exit when the current scrap value exceeds the continuation value at the current productivity level.

Note: The integration $\int \max\{s', h(z')\}\, \varphi(s')\, ds'$ is a *smoothing* operation. The continuation value h^* is smoother than v^* (which has a kink at the threshold). This aids both analysis and computation.

Reference: §4.1.4.2–4.1.4.3, equations (4.14)–(4.15), Exercise 4.1.12, p. 119.

Comparative Statics

Proposition 2.2.7 (Parametric Monotonicity)

If T dominates S (i.e., $Su \leq Tu$ for all u), and T is order-preserving and globally stable, then the fixed point of T dominates any fixed point of S.

Unconditional dominance cases:

- $s \uparrow$ shifts T up pointwise $\Rightarrow v^* \uparrow$
- $\pi \uparrow$ shifts T up pointwise $\Rightarrow v^* \uparrow$

Caution: claims about β or Q require additional domain restrictions (e.g. $v \geq 0$, or stochastic dominance + monotonicity).

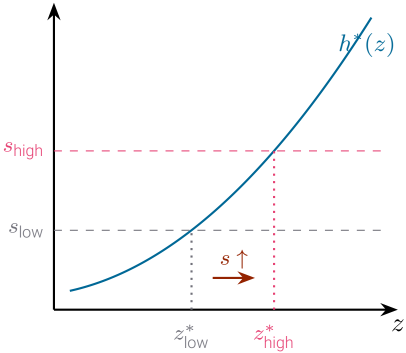

Exit threshold z^* (when h^* shifts pointwise):

- $s \uparrow \implies z^* \uparrow$ (more exit)
- $\pi \uparrow \implies z^* \downarrow$ (less exit)

Reference: Proposition 2.2.7, p. 67; Exercise 4.1.13, p. 119.

Value Function Iteration for Firm Exit: Algorithm

1. Choose an initial guess $v_0 \in \mathbb{R}^{\mathsf{Z}}$

$$v_0 = (I - \beta Q)^{-1} \pi \quad \text{(no-exit value)}$$

2. Iterate

$$v_{k+1}(z) = \max\{s,\ \pi(z) + \beta (Q v_k)(z)\}$$

3. Stop when

$$\|v_{k+1} - v_k\|_\infty \leq \varepsilon$$

4. Compute the greedy policy

$$\sigma_k(z) = \mathbf{1}\{s \geq \pi(z) + \beta (Q v_k)(z)\}$$

5. Exit threshold

$$z^* = \min\{z \in \mathsf{Z} : \sigma_k(z) = 0\}$$

Properties of VFI

- Easy to implement
- Globally convergent
- Robust across models

Reference: §4.1.1.7; Listing 4.1 (firm_exit.jl).

Why VFI Can Be Slow When $\beta \approx 1$

Key fact (Banach). If T is a contraction of modulus β, then

$$\|v_k - v^*\|_\infty \le \beta^k \|v_0 - v^*\|_\infty.$$

Interpretation. VFI is *linear convergence* with rate β. When $\beta \approx 1$, the rate is close to one, so the error shrinks slowly.

Rule of thumb. To shrink the error by a factor κ,

$$k \gtrsim \frac{\ln(\kappa)}{-\ln(\beta)}.$$

(We use this on the next slide.)

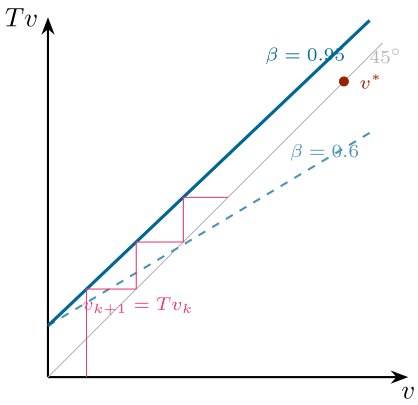

When β is close to 1, T is close to the 45° line, so successive approximation moves slowly.

How Many VFI Iterations? A Back-of-the-Envelope

Convergence bound

$$\|v_k - v^*\|_\infty \leq \beta^k \|v_0 - v^*\|_\infty.$$

Target: shrink error by factor κ

$$\beta^k \leq \frac{1}{\kappa} \quad \Longrightarrow \quad k \geq \frac{\ln(\kappa)}{-\ln(\beta)}.$$

Example. If $\beta = 0.98$ and $\kappa = 10^3$,

$$k \geq \frac{\ln(10^3)}{-\ln(0.98)} \approx \frac{6.91}{0.0202} \approx 342.$$

Practical stopping rule Monitor

$$\|v_{k+1} - v_k\|_\infty.$$

It is cheap to compute and (under contraction) directly tracks progress.
Advanced (Chapter 8). If σ_k is v_k-greedy,

$$\|v^* - v_{\sigma_k}\|_\infty \leq \frac{2\beta}{1-\beta}\|v_k - v_{k-1}\|_\infty.$$

Summary

Concept	Firm application	Reference
$\mathbb{R}^{\mathsf{Z}}$, Markov operator Q, $\|\cdot\|_\infty$	Productivity grid, expectations	§1.2.4, §3.2.1
Neumann series: $(I - \beta Q)^{-1}\pi$	Firm value without exit	Lemma 3.2.1
Banach's FPT, contractions	Unique v^*, geometric convergence	Thm. 1.2.3
Upper envelope lemma	$T = T_1 \vee T_0$ is contraction	Lemma 2.2.3
Optimal stopping (β, P, c, e)	Firm exit: $c = \pi$, $e = s$	§4.1.1
v^* unique FP, Bellman's principle	σ^* optimal $\iff$ v^*-greedy	Props. 4.1.2–3
Monotonicity, threshold z^*	Exit iff $z < z^*$	Lemma 4.1.4
Continuation value operator C	Dimensionality reduction	Prop. 4.1.5
Parametric monotonicity	How z^* moves with β, s, π	Prop. 2.2.7

Forward links:

- **Lectures 4–5 (Hopenhayn):** Embed exit problem in GE with free entry $\to$ stationary firm distribution, exit rate, aggregate productivity
- **Lectures 6–7 (Khan-Thomas):** Add investment choice $\to$ MDP (Lecture 3), lumpy investment with fixed costs
- **Next week (Lecture 3):** Household savings, VFI/HPI/OPI, EGM, extreme value shocks