

v1: 18 April 2024

Review Article

Visualizing the Contraction Mapping Theorem

Peer-approved: 18 April 2024

James E. Rauch¹, Alexis Akira Toda²

© The Author(s) 2024. This is an Open Access article under the CC BY 4.0 license.

1. Department of Economics, University of California, San Diego, United States; 2. University of California, San Diego, United States

Qeios, Vol. 6 (2024)
ISSN: 2632-3834

We visualize the process of value function iteration and convergence. We also clarify the conditions under which value function iteration converges to a unique value function, which are often glossed over in practice.

Corresponding authors: James E. Rauch, jrauch@ucsd.edu; Alexis Akira Toda, atoda@ucsd.edu

1. Introduction

In various disciplines we encounter the concept of *value function iteration*. The main goal of this paper and the companion video is to visualize the process of value function iteration and convergence. We also wish to clarify the conditions under which value function iteration converges to a unique value function, which are often glossed over in practice. We use the classic optimal growth model and optimal savings problem for this purpose.

For many of us, visualization is a powerful tool for understanding mathematical concepts (Arcavi, 2003). For example, suppose you had never seen a circle. You could be shown the equation $x^2 + y^2 = 1$ and be told this is the equation for a circle in two-dimensional Euclidean space, with its center at the origin and with a radius of 1. You could also be shown a picture of the circle represented by this equation. Many of us would find that this picture gives us a deeper and more confident understanding of the concept of a circle than the equation by itself.

The audience we have in mind for this paper is mainly first-year graduate students in Economics and their instructors. We hope, however, that our paper will be useful to anyone who wants to take the mystery out of value function iteration and convergence.

The paper is organized as follows. Section 2 discusses basic facts on Banach's contraction mapping theorem. Sections 3 and 4 study the optimal growth model and optimal savings model and visualize the convergence of value function iteration. Section 5 discusses useful tricks in dynamic programming.

2. Preliminaries

2.1. Metric space

We start from a review of basic concepts. Let V be a set. (We use the uncommon notation V because it will later be the set of candidate value functions.) We say that the function $d : V \times V \rightarrow \mathbb{R}$ is a *metric* (or *distance*) if it is nonnegative ($d(v_1, v_2) \geq 0$ for all $v_1, v_2 \in V$ with equality if and only if $v_1 = v_2$), symmetric ($d(v_1, v_2) = d(v_2, v_1)$ for all $v_1, v_2 \in V$), and satisfies the triangle inequality ($d(v_1, v_3) \leq d(v_1, v_2) + d(v_2, v_3)$ for all $v_1, v_2, v_3 \in V$). We call a set V endowed with a metric d a *metric space* and denote by (V, d) . If the metric d is understood, we often just refer to V as the metric space. We say that a sequence $\{v_n\}_{n=1}^{\infty}$ *converges* to v if $d(v_n, v) \rightarrow 0$ as $n \rightarrow \infty$, so the distance between v_n and the limit v tends to zero. We denote convergence by $\lim_{n \rightarrow \infty} v_n = v$ or $v_n \rightarrow v$, etc.

2.2. Complete metric space and Banach space

Sometimes we would like to characterize convergence without specifying the limit. We say that a sequence $\{v_n\}_{n=1}^{\infty}$ is *Cauchy* if the terms approach each other as the indexes tend to infinity, or more formally, for all $\epsilon > 0$ we can take N such that $d(v_m, v_n) < \epsilon$ whenever $m, n > N$. If $\{v_n\}_{n=1}^{\infty}$ is convergent, it is clearly Cauchy. When the converse is also true, i.e., every Cauchy sequence is convergent, we say that the metric space (V, d) is *complete*. Intuitively, a complete metric space is a metric space without “holes”. For instance, both the set of rational numbers \mathbb{Q} and the set of real numbers \mathbb{R} are metric spaces with the metric $d(x, y) = |x - y|$, but \mathbb{R} is complete while \mathbb{Q} is not.

We list a few examples of common complete metric spaces. We omit the proofs as they can easily be found in standard textbooks.

Example 1. The Euclidean space $V_1 = \mathbb{R}^N$ equipped with the usual Euclidean distance is a complete metric space.

Below, let X be a subset of a Euclidean space.

Example 2. Let V_2 be the space of bounded functions $v : X \rightarrow \mathbb{R}$, so $v \in V_2$ if and only if $\sup_{x \in X} |v(x)| < \infty$. For $v_1, v_2 \in V_2$, define the sup metric

$$d(v_1, v_2) = \sup_{x \in X} |v_1(x) - v_2(x)|. \quad (2.1)$$

Then (V_2, d) is a complete metric space.

The space of bounded functions is very large. Sometimes we may want to add more structure such as continuity as in the following example.

Example 3. Let V_3 be the space of bounded continuous functions $v : X \rightarrow \mathbb{R}$ equipped with the sup metric (2.1). Then (V_3, d) is a complete metric space.

Sometimes, imposing boundedness is too strong. If we would like to work with functions that are not bounded but are known to be close to a given function, the following space might be useful.

Example 4. Let $u : X \rightarrow \mathbb{R}$ be given and V_4 be the space of functions whose differences from u are bounded, so

$$V_4 = \left\{ v : X \rightarrow \mathbb{R} : \sup_{x \in X} |v(x) - u(x)| < \infty \right\}. \quad (2.2)$$

If we let d be the sup metric (2.1), then (V_4, d) is a complete metric space.

The Euclidean space as well as V_2, V_3 in Examples 2 and 3 equipped with the norm $\|v\| = \sup_{x \in X} |v(x)|$ are also vector spaces (spaces on which addition and scalar multiplication are defined), which are called *normed spaces*. As V_2, V_3 are complete, they are *complete normed spaces*, a more common name being the *Banach space*. Note that the complete metric space V_4 in Example 4 need not have a vector space structure, so it is generally not a Banach space.³

2.3. Contraction mapping theorem

Let (V, d) be a complete metric space. We say that an operator $T : V \rightarrow V$ is a *contraction* with modulus $\beta \in [0, 1)$ if for all $v_1, v_2 \in V$ we have

$$d(Tv_1, Tv_2) \leq \beta d(v_1, v_2).$$

That is, a contraction is a map such that the distance between two elements shrinks by factor at least $\beta \in [0, 1)$ each time we apply the map. What makes a contraction useful is that it allows us to establish the

existence of a unique fixed point, which is known as the *contraction mapping theorem* or the *Banach fixed point theorem*.

Theorem 1 (Contraction Mapping Theorem). Let (V, d) be a complete metric space and $T : V \rightarrow V$ be a contraction with modulus $\beta \in [0, 1)$. Then the following statements are true.

- i. T has a unique fixed point: there exists a unique $v^* \in V$ such that $Tv^* = v^*$.
- ii. Iterates of T converge to v^* . For any $v \in V$, define the sequence $\{v_n\}_{n=0}^\infty$ by $v_0 = v$ and $v_n = Tv_{n-1} = \dots = T^n v_0$. Then $v_n \rightarrow v^*$, with $d(v_n, v^*) = O(\beta^n)$.

Proof. We omit the proof as it is standard. See Stachurski (2009) for a textbook treatment. \square

Often the contraction mapping theorem is proved under the more restrictive condition that (V, d) is a Banach space. We avoid this restriction so that we can apply the contraction mapping theorem when, for example, V is the space of increasing functions or the space of concave functions, neither of which is Banach.

2.4. Blackwell's sufficient conditions

The contraction mapping theorem allows us to establish the existence and uniqueness of a fixed point of an operator T and a numerical algorithm to approximate the fixed point. To this end, we need to verify that T is indeed a contraction. Blackwell (1965)'s sufficient conditions are very useful in this respect. Let X be a set and V be a space of functions $v : X \rightarrow \mathbb{R}$ equipped with the sup metric (2.1). Let us say that V has the *upward shift property* if for any $v \in V$ and nonnegative constant $\kappa \geq 0$, we have $v + \kappa \in V$, that is, if v is in V , the function obtained by adding a nonnegative constant is also in V . We are deliberately vague in specifying V : depending on the context, V could be a space of bounded functions (V_2 in Example 2), of bounded continuous functions (V_3 in Example 3), or some other space. For our purpose, all that matters is that the distance is the sup metric (2.1).

We say that an operator $T : V \rightarrow V$ is *monotone* if $v_1 \leq v_2$ implies $Tv_1 \leq Tv_2$. More precisely, if $v_1(x) \leq v_2(x)$ for all $x \in X$, then $(Tv_1)(x) \leq (Tv_2)(x)$ for all $x \in X$. We say that T satisfies the *discounting property with modulus* $\beta \in [0, 1)$ if $T(v + \kappa) \leq Tv + \beta\kappa$ for all $v \in V$ and $\kappa \geq 0$.

Proposition 1 (Blackwell's sufficient conditions). If V is a complete metric space of functions $v : X \rightarrow \mathbb{R}$ with upward shift property and $T : V \rightarrow V$ is monotone and satisfies the discounting property with modulus β , then T is a contraction with modulus β .

Proof. Let $v_1, v_2 \in V$. For any $x \in X$, by the definition of the sup metric we have

$$v_1(x) - v_2(x) \leq d(v_1, v_2) =: \kappa.$$

Therefore $v_1 \leq v_2 + \kappa$. By the upward shift property, we have $v_2 + \kappa \in V$. Applying T to both sides and using the monotonicity and the discounting property of T , we obtain

$$Tv_1 \leq T(v_2 + \kappa) \leq Tv_2 + \beta\kappa.$$

Therefore $(Tv_1)(x) - (Tv_2)(x) \leq \beta\kappa$ for all $x \in X$. Changing the role of v_1, v_2 , we obtain $(Tv_2)(x) - (Tv_1)(x) \leq \beta\kappa$, so

$$d(Tv_1, Tv_2) = \sup_{x \in X} |(Tv_1)(x) - (Tv_2)(x)| \leq \beta\kappa = \beta d(v_1, v_2).$$

□

3. Optimal growth model

3.1. Informal description of the problem

Imagine that you are Robinson Crusoe marooned on a desert island. Potatoes grow on the island but each season you need to manage how much to eat and how much to plant for the next season. The problem is how to eat and cultivate potatoes optimally.

More formally, time is discrete and indexed by $t = 0, 1, 2, \dots$. You start with some available resources of potatoes, denoted by $a > 0$. If you consume $0 \leq c \leq a$, then you get utility $u(c)$, where u is a utility function. If you plant $k = a - c$ potatoes, you get a new harvest of $f(k)$ next period, where f is a production function.

Let $a_0 > 0$ be the initial endowment of potatoes and c_0, c_1, c_2, \dots be the consumption over time. At time t , because you cannot consume more than the available resources, denoted by a_t , the consumption c_t must satisfy

$$a_{t+1} = f(a_t - c_t), \quad (3.1a)$$

$$0 \leq c_t \leq a_t. \quad (3.1b)$$

The lifetime utility is then

$$u(c_0) + \beta u(c_1) + \beta^2 u(c_2) + \dots = \sum_{t=0}^{\infty} \beta^t u(c_t), \quad (3.2)$$

where $\beta \in [0, 1)$ is the discount factor. Your goal is to maximize the lifetime utility (3.2) subject to the feasibility constraints (3.1), given the initial endowment a_0 .

This model is often called the *Ramsey model* because Ramsey (1928) introduced a continuous-time version of this model and qualitatively studied its solution using

calculus of variations. Cass (1965) and Koopmans (1965) introduced technological and population growth and so the model is also known as the *optimal growth model*.

3.2. Value function iteration

To solve the optimal growth model, we can apply *value function iteration*, which is based on Bellman's principle of optimality and Banach's contraction mapping theorem.

Given the initial endowment $a_0 = a$, let $V(a)$ be the maximum lifetime utility (the maximum of (3.2) over all possible consumption plans $\{c_t\}_{t=0}^{\infty}$), which is called the *value function* that for now we assume to exist. Imagine what would happen to the lifetime utility if you choose an arbitrary consumption $c_0 = c$ this period but you stick to the optimal plan from the next period on. By choosing $c_0 = c$, you first receive flow utility $u(c)$ and the next period's resource becomes $a' = f(a - c)$ by (3.1a). Since by assumption you stick to the optimal plan from next period on, the sum of the remaining terms in lifetime utility becomes

$$\begin{aligned}\beta u(c_1) + \beta^2 u(c_2) + \beta^3 u(c_3) + \dots &= \beta (u(c_1) + \beta u(c_2) + \beta^2 u(c_3) + \dots) \\ &= \beta V(a') = \beta V(f(a - c)),\end{aligned}$$

because c_1, c_2, \dots are chosen optimally given $a_1 = a' = f(a - c)$. Therefore the lifetime utility under this alternative plan is

$$u(c_0) + \beta u(c_1) + \beta^2 u(c_2) + \dots = u(c) + \beta V(f(a - c)).$$

But $c_0 = c$ is arbitrary, and choosing it optimally leads to the maximum lifetime utility. Thus we obtain

$$V(a) = \max_{0 \leq c \leq a} \{u(c) + \beta V(f(a - c))\}, \quad (3.3)$$

which is called the *Bellman equation*.

For an arbitrary function V defined on the set of nonnegative real numbers $[0, \infty)$, the right-hand side of (3.3),

$$\max_{0 \leq c \leq a} \{u(c) + \beta V(f(a - c))\},$$

defines another function. So the right-hand side of (3.3) can be interpreted as an operation T that acts on the set of functions and outputs a new function TV from an input function V . The formal definition of T , called the *Bellman operator*, is

$$(TV)(a) = \max_{0 \leq c \leq a} \{u(c) + \beta V(f(a - c))\}. \quad (3.4)$$

Using the Bellman operator T , the Bellman equation (3.3) can be compactly written as

$$V = TV. \quad (3.5)$$

Equation (3.5) shows that the value function is a *fixed point* of the Bellman operator T (a function that remains unchanged by applying T). Under certain conditions, the Bellman operator T becomes a contraction, which guarantees the existence and uniqueness of a value function V and the uniform convergence of $V^{(n)} := T^n V^{(0)}$ to V for any initial guess $V^{(0)}$ as the number of iterations n tends to infinity. We summarize the formal result in the following proposition.

Proposition 2. Let $V = bc\mathbb{R}_+$ be the space of bounded continuous functions defined on \mathbb{R}_+ equipped with the supremum norm $\|V\| = \sup_{x \geq 0} |V(x)|$ for $V \in V$. Suppose that

- i. $u \in V$,
- ii. $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuous, and
- iii. $0 \leq \beta < 1$.

Then the following statements are true.

- a. The Bellman operator T is a contraction on V with modulus β .

- b. T has a unique fixed point $V \in V$.

- c. The approximation error $\|V^{(n)} - V\|$ is $O(\beta^n)$.

Proof. It suffices to show that T is a contraction. To this end we verify Blackwell's sufficient conditions (Proposition 1). If $V_1, V_2 \in V$ and $V_1 \leq V_2$, then by the definition of the Bellman operator (3.4) we obtain

$$\begin{aligned}(TV_1)(a) &= \max_{0 \leq c \leq a} \{u(c) + \beta V_1(f(a - c))\} \\ &\leq \max_{0 \leq c \leq a} \{u(c) + \beta V_2(f(a - c))\} = (TV_2)(a),\end{aligned}$$

so $TV_1 \leq TV_2$ and T is monotonic. Clearly V satisfies the upward shift property. If $V \in V$ and $\kappa \geq 0$, we obtain

$$\begin{aligned}(T(V + \kappa))(a) &= \max_{0 \leq c \leq a} \{u(c) + \beta(V(f(a - c)) + \kappa)\} \\ &= \max_{0 \leq c \leq a} \{u(c) + \beta V(f(a - c))\} + \beta \kappa \\ &= (TV)(a) + \beta \kappa,\end{aligned}$$

so T satisfies the discounting property (with equality) with modulus β . \square

Proposition 2 implies that the value function V can be approximated arbitrarily well by starting from any initial guess $V^{(0)}$ and repeatedly applying the Bellman operator T . As an illustration, suppose that the utility function u is increasing and we use the zero function $V^{(0)} \equiv 0$ as the initial value.² Using the definition of the Bellman operator (3.4), after one iteration we obtain

$$\begin{aligned}V^{(1)}(a) &= (TV^{(0)})(a) = \max_{0 \leq c \leq a} \{u(c) + \beta V^{(0)}(f(a - c))\} \\ &= \max_{0 \leq c \leq a} u(c) = u(a),\end{aligned}$$

which is just the utility function. After two iterations, we obtain

$$\begin{aligned}V^{(2)}(a) &= (TV^{(1)})(a) = \max_{0 \leq c \leq a} \{u(c) + \beta V^{(1)}(f(a - c))\} \\ &= \max_{0 \leq c \leq a} \{u(c) + \beta u(f(a - c))\}.\end{aligned}$$

Except for special cases, $V^{(2)}$ (and more generally $V^{(n)}$ for $n \geq 2$) does not admit a closed-form expression and needs to be computed numerically. A standard approach is to define a grid $\{a_g\}_{g=1}^G$ with $a_1 < \dots < a_G$, define $V^{(n-1)}$ on \mathbb{R}_+ by interpolation and extrapolation using the values $\{V^{(n-1)}(a_g)\}_{g=1}^G$, and compute the next values $\{V^{(n)}(a_g)\}_{g=1}^G$ as $V^{(n)}(a_g) = (TV^{(n-1)})(a_g)$ by numerically maximizing the right-hand side of (3.4).

The assumption in Proposition 2 that the utility function u is bounded is often undesirable because it rules out

common utility functions such as $u(c) = \log c$. Although it is not simple to allow functions that are unbounded below (such as $u(c) = \log c$; see Le Van and Morhaim (2002) for a treatment of such cases), unboundedness from above can be easily handled if the production function exhibits a certain type of decreasing returns to scale.

Proposition 3. *Suppose that*

- i. $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ is continuous and bounded below,
- ii. f is increasing and there exists $\bar{k} > 0$ such that $f(k) \leq k$ for all $k \geq \bar{k}$, and
- iii. $0 \leq \beta < 1$.

Take any $\bar{a} \geq \bar{k}$ and let V be the space of bounded continuous functions defined on $[0, \bar{a}]$. Then the conclusions of Proposition 2 hold.

Proof. We only need to verify that $V(f(a - c))$ is well-defined. Since by assumption f is increasing, $f(k) \leq k$ for $k \geq \bar{k}$, $0 \leq c \leq a$, $a \in [0, \bar{a}]$, and $\bar{a} \geq \bar{k}$, we have

$$f(a - c) \leq f(a) \leq f(\bar{a}) \leq \bar{a}.$$

Therefore $f(a - c) \in [0, \bar{a}]$ and $V(f(a - c))$ is well-defined. \square

3.3. Stochastic growth model

The stochastic (optimal) growth model is an extension of the optimal growth model with uncertainty, introduced by Brock and Mirman (1972) and quantitatively studied by Kydland and Prescott (1982). Imagine a situation where Robinson Crusoe's harvest of potatoes depends not only on the amount of potatoes planted but also on other factors outside his control such as rainfall and temperature, or his well-being depends on factors such as temperature and sunshine. For convenience, we suppose that these factors take finitely many values indexed by $z \in Z$. Suppose that this exogenous state evolves according to a Markov chain with transition probability $P(z, z') = \Pr(z_{t+1} = z' \mid z_t = z)$.

In this situation, the utility and value functions u, V may depend on the current exogenous state z , and the production function f may depend on two consecutive states (z, z') . If Robinson Crusoe wishes to maximize the expected utility, then the Bellman equation (3.3) becomes

$$V(a, z) = \tag{3.6}$$

$$\max_{0 \leq c \leq a} \left\{ u(c, z) + \beta \sum_{z' \in Z} P(z, z') V(f(a - c, z, z'), z') \right\},$$

where $u(c, z)$ is the utility function in state z and $f(k, z, z')$ is the production function when transitioning from state z to z' . Propositions 2 and 3 easily generalize to this setting by changing the assumptions on u and f to those on $u(\cdot, z)$ and $f(\cdot, z, z')$, so we omit the precise statement.

3.4. Numerical illustration

As a numerical illustration, we solve the stochastic growth model. Let Z be a finite set. For state $z \in Z$, suppose that the utility function takes the form

$$u(c, z) = \frac{(c + \epsilon)^{1-\gamma(z)}}{1 - \gamma(z)},$$

where $\epsilon > 0$ can be thought of an exogenous source of consumption (e.g., coconuts and fish) and $\gamma(z) > 0$ is the coefficient of relative risk aversion. (The case $\gamma(z) = 1$ corresponds to log utility.) This exogenous consumption prevents the utility function from being unbounded below when $\gamma(z) > 1$. (We can set $\epsilon = 0$ if $\gamma(z) < 1$.) The production function takes the form

$$f(k, z, z') = A(z')k^\alpha + (1 - \delta)k,$$

where $A(z') > 0$ is the productivity in the next state, $\alpha \in (0, 1)$ is the elasticity of output with respect to capital, and $\delta \in (0, 1]$ is the capital depreciation rate.

Because our purpose is to visualize the contraction mapping theorem, we consider a simple specification for the stochastic growth model. We consider a two-state Markov chain with $Z = \{1, 2\}$ with transition probability $P(z, z') = 0.8$ if $z = z'$ and $P(z, z') = 0.2$ if $z \neq z'$. The productivity is $(A(1), A(2)) = (1.1, 0.9)$, so state 1 is the high-productivity state. We set $\alpha = 0.36$ and $\delta = 0.08$, which are standard values. The discount factor is $\beta = 0.95$ and the exogenous consumption is $\epsilon = 0.1$. For the relative risk aversion, we consider two values $\gamma = 1.5$ and $\gamma = 0.5$ because the cases $\gamma \geq 1$ are qualitatively different (we have $u \leq 0$ according as $\gamma \geq 1$).

In this setup, we can easily verify that the assumptions of Proposition 3 are satisfied, so a value function uniquely exists. In particular, solving $f(k, z, z') = k$ for $k > 0$, we obtain $k = (A(z')/\delta)^{\frac{1}{1-\alpha}}$, so we can choose any \bar{a} with $\bar{a} \geq \bar{k} := (A(1)/\delta)^{\frac{1}{1-\alpha}}$. Below, we set $\bar{a} = 2\bar{k}$ and use a 100-point exponential grid on $[0, \bar{a}]$ to numerically solve the stochastic growth model by value function iteration.³

We now illustrate four specifications. The first specification is $\gamma(z) \equiv 1.5$, and we start the value function iteration from the initial guess $V^{(0)} \equiv 0$. Figures 1a and 1b show the value and consumption functions, respectively. Figure 1c shows the evolution of value functions along the iterations (for state 1 only for visibility), where the color changes from light green to blue as we increase the number of iterations n . For this specification, because the utility function is negative, the value function monotonically converges from above.

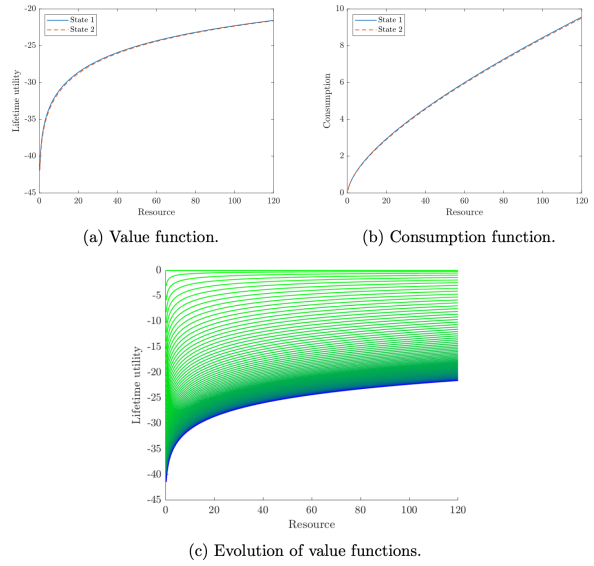


Figure 1. Stochastic growth model with $\gamma = 1.5$ and $V^{(0)} \equiv 0$.

The second specification is the same as the first except that we set the relative risk aversion to $\gamma(z) \equiv 0.5$ (Figure 2). For this specification, because the utility function is positive, the value function monotonically converges from below.

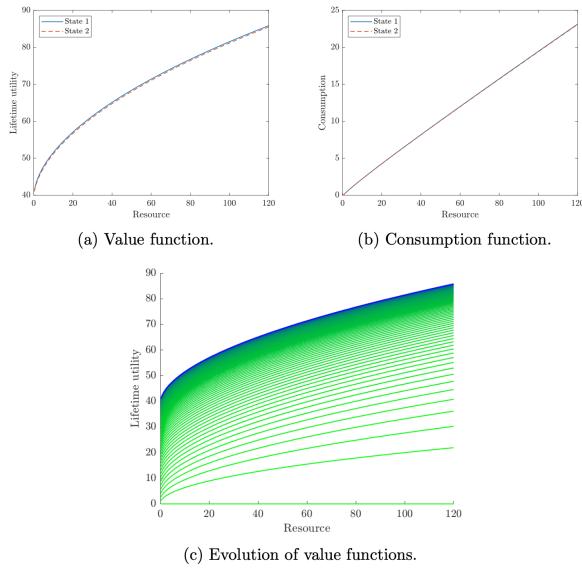


Figure 2. Stochastic growth model with $\gamma = 0.5$ and $V^{(0)} \equiv 0$.

The third specification is the same as the first except that we set the initial guess $V^{(0)}$ to an unnatural function, namely the sine curve flipped upside down (Figure 3). Although the initial guess is artificial (setting $V^{(0)} \equiv 0$ is natural as discussed in Footnote 2), the mathematical theory still applies and the value function converges (but in an erratic manner).




Figure 3. Evolution of value functions with $\gamma = 1.5$ and sine curve $V^{(0)}$.

The fourth specification is an intermediate case of the first and second, where $(\gamma(1), \gamma(2)) = (0.5, 1.5)$. This specification implies that the agent is less risk averse during the high-productivity state (Figure 4). Unlike the first two specifications, when the risk aversion is state-dependent, the value and consumption functions are quite different across states.

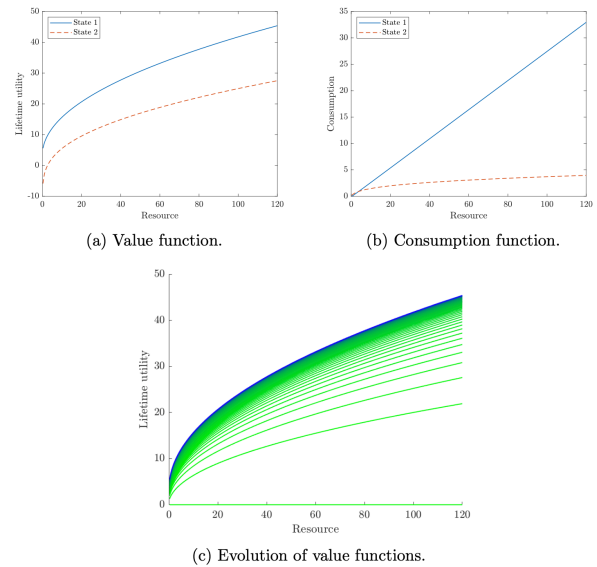


Figure 4. Evolution of value functions. Stochastic growth model with $(\gamma(1), \gamma(2)) = (0.5, 1.5)$ and $V^{(0)} \equiv 0$.

4. Optimal savings problem

4.1. Informal description of the problem

The optimal savings problem is the optimization problem of a single agent that receives income and has

access to the financial market for saving. Standard references for the optimal savings problem are Schechtman and Escudero (1977) and Chamberlain and Wilson (2000), who study the theoretical properties of the optimal consumption assuming a bounded utility function. Mathematically, the optimal savings problem is a special case of the stochastic growth model with production function

$$f(k, z, z') = R(z, z')k + Y(z, z'),$$

where $R(z, z') \geq 0$ is the gross return on savings and $Y(z, z') > 0$ is the non-financial income when transitioning from state z to z' .

4.2. Policy function iteration

Although the optimal savings problem is mathematically a special case of the stochastic growth model, establishing the existence of a solution and studying its properties is not simple when the utility function is unbounded (which is practically almost always the case) and mathematically rigorous results have been obtained only recently by Li and Stachurski (2014) and Ma, Stachurski, and Toda (2020). The reason is that the marginal product of capital $f'(k, z, z') = R(z, z')$ equals the gross return, which could well exceed 1 (imagine a positive interest rate or high stock returns). Then the trick of truncating the state space as in Proposition 3, which relies on marginal product less than 1, is no longer applicable.

To solve the optimal savings problem, we can apply *policy function iteration* instead of value function iteration. We illustrate the idea using the optimal growth model without uncertainty. Consider the Bellman equation (3.3). Assuming that u, f, V are all differentiable and the optimal consumption is interior, the first-order condition for optimality is

$$0 = u'(c) - \beta V'(f(a - c))f'(a - c) = 0.$$

Differentiating both sides of (3.3) with respect to a and applying the envelope theorem, we obtain

$$V'(a) = \beta V'(f(a - c))f'(a - c).$$

Combining (4.1) and (4.2), we obtain

$$u'(c) = V'(a). \quad (4.2)$$

Now let $c = c_t$ and $a = a_t$ be the consumption and resource at time t . Noting that $f(a - c) = f(a_t - c_t) = a_{t+1}$ is the next period's resource (see (3.1a)), combining (4.1) (for $c = c_t$) and (4.2) (for $c = c_{t+1}$), it follows that

$$u'(c_t) = \beta u'(c_{t+1})f'(a_t - c_t), \quad (4.3)$$

which is known as the *Euler equation*. For the stochastic growth model, a similar calculation yields the Euler equation

$$u'(c_t) = E_t[\beta u'(c_{t+1})f'(a_t - c_t, z_t, z_{t+1})], \quad (4.4)$$

where E_t denotes the expectation conditional on time t information.

Coleman (1990) proposed a solution algorithm called *policy function iteration* that exploits the Euler equation (4.4). Suppose that we have a guess of the consumption function $c(a, z)$ and would like to update its value, denoted by ξ . Let $a_t = a$, $z_t = z$, $z_{t+1} = z'$, and $c_t = \xi$. Using the candidate consumption function c and the feasibility constraint (3.1a), we have

$$c_{t+1} = c(a_{t+1}, z_{t+1}) = c(f(a - \xi, z, z'), z').$$

Therefore the Euler equation (4.4) becomes

$$u'(\xi) = E_z[\beta u'(c(f(a - \xi, z, z')), z') f'(a - \xi, z, z')], \quad (4.5)$$

where E_z denotes the expectation conditional on $z_t = z$. Thus given the candidate consumption function $c(a, z)$, we can update it by the value ξ that solves (4.5). Repeating this process until convergence is called policy function iteration.

A key advantage of policy function iteration over value function iteration is that it involves only root-finding, which tends to be numerically more stable than maximization.⁴ A disadvantage is that the Coleman operator (the operation of updating the policy function) is not necessarily a contraction and proving theorems is significantly more challenging than value function iteration; see Mirman et al. (2008) for a rigorous treatment in the context of the stochastic growth model.

However, for the optimal savings problem, the marginal product

$$f'(a - c, z, z') = R(z, z')$$

depends only on the exogenous states and the analysis becomes simpler. Li and Stachurski (2014) apply policy function iteration to the optimal savings problem assuming that the gross return on saving is constant at R . When the utility function satisfies the standard properties such as $u' > 0$ (monotonicity), $u'' < 0$ (concavity), and $u'(0) = \infty$ (Inada condition), they show that the Euler equation (4.5) becomes

$$u'(\xi) = \max \{E_z[\beta R u'(c(R(a - \xi) + Y(z, z'), z'))], u'(a)\}. \quad (4.6)$$

(The reason why we take the maximum with $u'(a)$ is to take into account the possibility that the constraint $\xi \leq a$ binds.) Furthermore, when we define the distance between two candidate consumption functions c_1, c_2 by

$$\rho(c_1, c_2) = \sup_{a, z} |u'(c_1(a, z)) - u'(c_2(a, z))| \quad (4.7)$$

using the marginal utility, they show that the Coleman operator T is a contraction with modulus βR when $\beta R < 1$. Although this approach is specific to the optimal savings problem, the utility function u could be unbounded above and/or below, which is almost always the case in practice.

Stachurski and Toda (2019, 2020) apply policy function iteration to establish a linear lower bound on the consumption function when the utility function exhibits bounded relative risk aversion to show that wealth inherits the tail behavior of income when saving is risk-free as in Aiyagari (1994) models. Ma et al. (2020) generalize the approach of Li and Stachurski (2014) to the case with stochastic returns and discounting. In this case T is not necessarily a contraction but some iterate T^k is under some conditions. Toda (2021) shows that T is a generalization of a contraction called *Perov contraction*, which enables to significantly simplify the proof of Ma et al. (2020). Ma and Toda (2021) apply policy function iteration to prove the asymptotic linearity of consumption functions when the utility function is homothetic, and Ma and Toda (2022) further generalize this result when the marginal utility asymptotically behaves like a power function.

4.3. Numerical illustration

As in the case with the stochastic growth model, we only consider a simple optimal savings problem. The utility function exhibits constant relative risk aversion γ , so the marginal utility is $u'(c) = c^{-\gamma}$, where we set $\gamma = 1.5$. The discount factor is $\beta = 0.95$. We consider a two-state Markov chain with $Z = \{1, 2\}$ with transition probability $P(z, z') = 0.5$ for all (z, z') , so the process is independent and identically distributed over time. We suppose that the agent invests fraction $\theta \in [0, 1]$ of wealth in the stock market with expected return μ and volatility σ , and invests the rest in a risk-free asset with risk-free rate r_f . Therefore we can model the gross return on wealth as

$$R(z, z') = \begin{cases} (1 - \theta)e^{r_f} + \theta e^{\mu - \sigma^2/2 + \sigma} & \text{if } z' = 1 \\ (1 - \theta)e^{r_f} + \theta e^{\mu - \sigma^2/2 - \sigma} & \text{if } z' = 2 \end{cases}$$

We set $r_f = 0.01$, $\mu = 0.05$, $\sigma = 0.2$, and $\theta = 0.5$. Finally, we suppose that income is constant at $Y(z, z') \equiv 1$. Let

B be the 2×2 matrix whose (z, z') -th entry equals $\beta P(z, z')R(z, z')$. Toda (2021, §3.3) shows that if the spectral radius (largest absolute value of all eigenvalues) of B satisfies $\rho(B) < 1$, then the Coleman operator becomes a Perov contraction when we use a (vector-valued) metric similar to (4.7). In our specification we have $\rho(B) = 0.9791 < 1$, so policy function iteration is guaranteed to converge.

Figure 5 shows the consumption function and the evolution of consumption functions along the iterations when we use the initial guess $c^{(0)}(a, z) = a$.⁵ For this specification, the consumption function monotonically converges from above.




Figure 5. Optimal savings problem with $V^{(0)} \equiv 0$.

For policy function iteration, the initial guess $c^{(0)}$ needs to be increasing and satisfy $c^{(0)}(a, z) \leq a$; see Ma et al. (2020). Setting $c^{(0)}(a, z) = a$ is natural⁸ but not necessarily computationally efficient because the asymptotic slope of the true consumption function $c(a, z)$ is far smaller than 1. Ma and Toda (2022) discuss how to choose the initial guess to improve computational efficiency. However, the theory tells us that any initial guess $c^{(0)}$ that is increasing and satisfies $c^{(0)}(a, z) \leq a$ would work. To illustrate this point, we now consider an unnatural initial guess given by $c^{(0)}(a, z) = (\sin(a) + a)/4$. Figure 6 shows that the convergence is non-monotonic.




Figure 6. Evolution of consumption functions with $c^{(0)}(a, z) = (\sin(a) + a)/4$.

5. Some useful tricks

This section discusses various tricks that are useful for studying dynamic programming problems.

5.1. Establishing properties of value function

In many applications, we are not just interested in establishing the existence of a solution to a dynamic programming problem but often would like to establish some properties of the solution. For instance, In Figure 1 we see that the value function is increasing and concave. But how can we establish the monotonicity and concavity of the value function v if we cannot solve for v explicitly? The following proposition is useful in such settings.

Proposition 4. Let (V, d) be a complete metric space and $T : V \rightarrow V$ be a contraction with a unique fixed point $v^* \in V$. If $V_1 \subset V$ is a nonempty closed set and $TV_1 \subset V_1$, then $v^* \in V_1$.

Proof. Since $V_1 \subset V$ is closed, (V_1, d) is a complete metric space. Since $T : V \rightarrow V$ is a contraction and $TV_1 \subset V_1$, T is also a contraction on V_1 . Therefore there exists a unique $v_1^* \in V_1$ such that $Tv_1^* = v_1^*$. Since $V_1 \subset V$, v_1^* is also a fixed point of T in V , and the uniqueness implies $v^* = v_1^* \in V_1$. \square

Although Proposition 4 is almost trivial, it has many applications. Suppose we would like to show that the value function in Figure 1 is increasing. To establish this, we only need to assume that f is increasing.

Proposition 5. Let everything be as in Proposition 2 and suppose f is increasing. Then the value function V is increasing.

Proof. Let $V = bc\mathbb{R}_+$ be the space of bounded continuous functions and $V_1 = \{V \in V : V \text{ is increasing}\}$. Since monotonicity is preserved by taking limits, V_1 is closed. If $V \in V_1$ and $a_1 \leq a_2$, then the definition of the Bellman operator (3.4) implies

$$\begin{aligned} (TV)(a_1) &= \max_{0 \leq c \leq a_1} \{u(c) + \beta V(f(a_1 - c))\} \\ &\leq \max_{0 \leq c \leq a_1} \{u(c) + \beta V(f(a_2 - c))\} \\ &\leq \max_{0 \leq c \leq a_2} \{u(c) + \beta V(f(a_2 - c))\} = (TV)(a_2), \end{aligned}$$

where the first inequality follows from the monotonicity of f and V and the second inequality follows from the fact that taking the maximum on a larger set yields a larger value. Therefore $TV_1 \subset V_1$, and Proposition 4 yields the conclusion. \square

An argument along these lines is used, for example, to show the monotonicity of the consumption and saving functions in Ma et al. (2020, Proposition 2.3). Similarly, suppose that we would like to establish a lower bound

$v \geq \underline{v}$ for the value function. For this purpose we may consider the closed set $V_1 = \{v \in V : v \geq \underline{v}\}$. An application along these lines can be found in Ma and Toda (2021, Theorem 3) for proving the asymptotic linearity of consumption functions and Phelan and Toda (2022, Proposition 3.1) for ranking various value functions.

As another application of Proposition 4, suppose we would like to show that the value function in Figure 1 is concave. To establish this, we only need to assume that u is concave and f is increasing and concave.

Proposition 6. *Let everything be as in Proposition 2 and suppose u is concave and f is increasing and concave. Then the value function V is increasing and concave.*

Proof. Let $V = bc\mathbb{R}_+$ be the space of bounded continuous functions and $V_1 = \{V \in V : V \text{ is increasing and concave}\}$. Since monotonicity and concavity are preserved by taking limits, V_1 is closed. We have already shown that T preserves monotonicity. Therefore it suffices to show that T preserves concavity.

Let $V \in V$, $a_1, a_2 \geq 0$, and $t \in [0, 1]$. Since f is increasing and concave, so is $V \circ f$. To see this, note that the concavity of f implies

$$f((1-t)a_1 + ta_2) \geq (1-t)f(a_1) + tf(a_2),$$

and applying V to both sides and using the monotonicity and concavity of V , we obtain

$$\begin{aligned} V(f((1-t)a_1 + ta_2)) &\geq V((1-t)f(a_1) + tf(a_2)) \\ &\geq (1-t)V(f(a_1)) + tV(f(a_2)). \end{aligned}$$

Fix $c_j \in [0, a_j]$ for $j = 1, 2$ and let $c = (1-t)c_1 + tc_2$ and $a = (1-t)a_1 + ta_2$ for $t \in [0, 1]$. Then the concavity of u and $V \circ f$ implies

$$\begin{aligned} &u(c) + \beta V(f(a-c)) \\ &\geq (1-t)(u(c_1) + \beta V(f(a_1 - c_1))) \\ &\quad + t(u(c_2) + \beta V(f(a_2 - c_2))). \end{aligned}$$

Since $c_j \in [0, a_j]$, we have $c \in [0, a]$. Therefore taking the maximum of the left-hand side over $c \in [0, a]$, we obtain

$$\begin{aligned} (TV)(a) &\geq (1-t)(u(c_1) + \beta V(f(a_1 - c_1))) \\ &\quad + t(u(c_2) + \beta V(f(a_2 - c_2))). \end{aligned}$$

Taking the maximum of the right-hand side over $c_j \in [0, a_j]$, we obtain

$$(TV)((1-t)a_1 + ta_2) \geq (1-t)(TV)(a_1) + t(TV)(a_2),$$

so TV is concave. Therefore T preserves concavity. \square

5.2. Transformation of the Bellman equation

Consider the Bellman equation for a stochastic dynamic programming problem. As a concrete example, consider the Bellman equation for the stochastic growth model (3.6):

$$V(a, z) = \max_{0 \leq c \leq a} \{u(c, z) + \beta \mathbb{E}_z V(f(a - c, z, z'), z')\}, \quad (5.1)$$

where \mathbb{E}_z denotes the expectation conditional on z .
Define the function

$$g(a, c, z) := \beta \mathbb{E}_z V(f(a - c, z, z'), z').$$

Then clearly

$$V(a, z) = \max_{0 \leq c \leq a} \{u(c, z) + g(a, c, z)\}.$$

Changing the notation a, c, z to a', c', z' and setting $a' = f(a - c, z, z')$, it follows from the definition of g that

$$g(a, c, z) = \beta E_z \max_{0 \leq c' \leq f(a-c, z, z')} \{u(c', z') + g(f(a-c, z, z'), c', z')\}. \quad (5.2)$$

Note that the transformed Bellman equation (5.2) now involves only the unknown function g . This kind of transformation may be useful because the expectation has a smoothing effect and g could be better behaved than V . See Ma and Stachurski (2021) and Ma et al. (2022) for more discussion and examples.

Footnotes

¹ A variant of the space V_4 in Example 4 is used to solve the optimal savings problem by policy function iteration as in Section 4; see Li and Stachurski (2014) and Ma et al. (2020).

² Using zero as the initial value is natural because the n -th iterate $V^{(n)} = T^n 0$ is exactly the value function when the agent lives for n periods and the economy ends. Thus, by setting $V^{(0)} = 0$ and iterating the Bellman operator, we would solve the optimal growth model corresponding to various time horizons.

³ See Gouin-Bonenfant and Toda (2023, §4.6) for the specific details on constructing the exponential grid. We use the median grid point $k^*/2$ and spline interpolation for computing value functions off the grid points.

⁴ A variant of policy function iteration that uses a grid on savings $s = a - \xi$ instead of asset a (and hence makes the asset grid endogenous), which is called the *endogenous grid point* method (Carroll, 2006), even avoids root-finding and substantially reduces computing time when the inverse marginal utility function $(u')^{-1}$ is available in closed-form. Examples are the constant relative risk aversion (CRRA) utility $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ or the constant absolute risk aversion (CARA) utility $u(c) = \frac{1}{\gamma} e^{-\gamma c}$, for which $(u')^{-1}(m) = m^{-1/\gamma}$ and $(u')^{-1}(m) = -\frac{1}{\gamma} \log m$, respectively.

⁵ To numerically solve the model, we use a 100-point exponential grid on $[0, 100]$ with a median grid point of 10 and linear interpolation/extrapolation to compute the consumption functions off the grid.

⁶ Because $c = a$ is the optimal consumption when the agent lives for one period, the n -th iterate $c^{(n)} = T^n c^{(0)}$ is exactly the consumption function when the agent lives for $n + 1$ periods by the same reason as in Footnote 2. Thus Figure 5b shows the optimal consumption functions for various time horizons.

References

- S. Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109 (3): 659–684, August 1994. doi: 10.2307/2118417.
- Abraham Arcavi. The role of visual representations in the learning of mathematics. *Educational Studies in Mathematics*, 52 (3): 215–241, 2003. doi: 10.1023/a:1024312321077.
- David Blackwell. Discounted dynamic programming. *Annals of Mathematical Statistics*, 36 (1): 226–235, February 1965. doi: 10.1214/aoms/1177700285.
- William A. Brock and Leonard J. Mirman. Optimal economic growth and uncertainty: The discounted case. *Journal of Economic Theory*, 4 (3): 479–513, June 1972. doi: 10.1016/0022-0531(72)90135-4.
- Christopher D. Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91 (3): 312–320, June 2006. doi: 10.1016/j.econlet.2005.09.013.
- David Cass. Optimum growth in an aggregative model of capital accumulation. *Review of Economic Studies*, 32 (3): 233–240, July 1965. doi: 10.2307/2295827.
- Gary Chamberlain and Charles A. Wilson. Optimal intertemporal consumption under uncertainty. *Review of Economic Dynamics*, 3 (3): 365–395, July 2000. doi: 10.1006/redo.2000.0098.
- Wilbur John Coleman, II. Solving the stochastic growth model by policy-function iteration. *Journal of Business & Economic Statistics*, 8 (1): 27–29, January 1990. doi: 10.1080/07350015.1990.10509769.
- Émilien Gouin-Bonenfant and Alexis Akira Toda. areto extrapolation: An analytical framework for studying tail inequality. *Quantitative Economics*, 14 (1): 201–233, January 2023. doi: 10.3982/QE1817.
- Tjalling C. Koopmans. On the concept of optimal economic growth. In *Semaine d'Étude sur le Rôle de l'Analyse Économétrique dans la Formulation de Plans de Développement*, volume 28, pages 225–300. Pontificiae Academiae Scientiarum Scripta Varia, 1965.
- Finn E. Kydland and Edward C. Prescott. Time to build and aggregate fluctuations. *Econometrica*, 50 (6): 1345–1370, November 1982. doi: 10.2307/1913386.
- Cuong Le Van and Lisa Morhaim. Optimal growth models with bounded or unbounded returns: A unifying approach. *Journal of Economic Theory*, 105 (1): 158–187, July 2002. doi: 10.1006/jeth.2001.2880.
- Huiyu Li and John Stachurski. Solving the income fluctuation problem with unbounded rewards. *Journal*

of *Economic Dynamics and Control*, 45: 353–365, August 2014. doi: 10.1016/j.jedc.2014.06.003.

- Qingyin Ma and John Stachurski. Dynamic programming deconstructed: Transformations of the Bellman equation and computational efficiency. *Operations Research*, 69 (5): 1591–1607, September 2021. doi: 10.1287/opre.2020.2006.
- Qingyin Ma and Alexis Akira Toda. A theory of the saving rate of the rich. *Journal of Economic Theory*, 192: 105193, March 2021. doi: 10.1016/j.jet.2021.105193.
- Qingyin Ma and Alexis Akira Toda. Asymptotic linearity of consumption functions and computational efficiency. *Journal of Mathematical Economics*, 98: 102562, January 2022. doi: 10.1016/j.jmateco.2021.102562.
- Qingyin Ma, John Stachurski, and Alexis Akira Toda. The income fluctuation problem and the evolution of wealth. *Journal of Economic Theory*, 187: 105003, May 2020. doi: 10.1016/j.jet.2020.105003.
- Qingyin Ma, John Stachurski, and Alexis Akira Toda. Unbounded dynamic programming via the Q-transform. *Journal of Mathematical Economics*, 100: 102652, May 2022. doi: 10.1016/j.jmateco.2022.102652.
- Leonard J. Mirman, Olivier F. Morand, and Kevin L. Reffett. A qualitative approach to Markovian equilibrium in infinite horizon economies with capital. *Journal of Economic Theory*, 139 (1): 75–98, March 2008. doi: 10.1016/j.jet.2007.05.009.
- Thomas Phelan and Alexis Akira Toda. Optimal epidemic control in equilibrium with imperfect testing and enforcement. *Journal of Economic Theory*, 206: 105570, December 2022. doi: 10.1016/j.jet.2022.105570.
- Frank P. Ramsey. A mathematical theory of saving. *Economic Journal*, 38 (152): 543–559, December 1928. doi: 10.2307/2224098.
- Jack Schechtman and Vera L. S. Escudero. Some results on “an income fluctuation problem”. *Journal of Economic Theory*, 16 (2): 151–166, December 1977. doi: 10.1016/0022-0531(77)90003-5.
- John Stachurski. *Economic Dynamics: Theory and Computation*. MIT Press, 2 edition, 2009. URL <https://johnstachurski.net/edtc.html>.
- John Stachurski and Alexis Akira Toda. An impossibility theorem for wealth in heterogeneous-agent models with limited heterogeneity. *Journal of Economic Theory*, 182: 1–24, July 2019. doi: 10.1016/j.jet.2019.04.001.
- John Stachurski and Alexis Akira Toda. Corrigendum to “An impossibility theorem for wealth in heterogeneous-agent models with limited heterogeneity” [Journal of Economic Theory 182 (2019) 1–24]. *Journal of Economic Theory*, 188: 105066, July 2020. doi: 10.1016/j.jet.2020.105066.
- Alexis Akira Toda. erov’s contraction principle and dynamic programming with stochastic discounting. *Operations Research Letters*, 49 (5): 815–819, September 2021. doi: 10.1016/j.orl.2021.09.001.

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.