ISOMORPHIC DYNAMIC PROGRAMS

JOHN STACHURSKI AND JUNNAN ZHANG

Abstract. We study relationships between dynamic programs by applying conjugacy methods from dynamical systems theory. When two dynamic programs are connected by an order isomorphism, we show that optimality properties transmit from one formulation to the other. We apply these results to Epstein–Zin preferences with time preference shocks, obtaining a sharp characterization of when optimality holds. We also show that multiplicative Kreps–Porteus preferences and risk-sensitive preferences are isomorphic, so that well-known results for the latter carry over to the former. Finally, we demonstrate how isomorphic transformations can improve the numerical accuracy of value function approximations, with gains of two orders of magnitude in a multisector real business cycle model.

1. Introduction

Dynamic programming is a major field of optimization with applications across a wide spectrum of scientific domains, including economics and finance (see, e.g., Bellman (1957), Stokey and Lucas (1989), Puterman (2005), Hernández-Lerma and Lasserre (2012), Bäuerle and Rieder (2011), Bertsekas (2012)). In recent years, variations of the standard model have flourished. This is particularly true in economics and finance, as applied researchers extend dynamic models to get closer to the data. Some of these variations add features to the standard framework, including Epstein–Zin preferences, risk-sensitive preferences, adversarial agents, ambiguity, and hyperbolic discounting.[1] Other variations take an existing model and rearrange the structure and timing. These variations include the Q-factor and exponential risk-sensitive Q-factor

Date: November 5, 2026.

The authors thank Yuchao Li, Shu Hu, and Jingni Yang for valuable comments and suggestions. Financial support from Schmidt Futures is gratefully acknowledged. The second author acknowledges support from the National Natural Science Foundation of China (Projects 72303190 and 72533007).

[1]See, for example, Epstein and Zin (1989), Cagetti et al. (2002), Hansen and Sargent (2011), Bäuerle and Jaśkiewicz (2018), Fedus et al. (2019), Marinacci and Montrucchio (2019), Gao et al. (2021), or de Castro et al. (2022).

formulations from Q-learning, and the expected value function models and integrated value function models from structural estimation.[2]

Given this proliferation of variations and modifications, one significant issue is that each new formulation typically requires its own bespoke analysis to establish desirable optimality properties (e.g., existence of optimal policies, Bellman's principle of optimality, or convergence of Howard policy iteration). In this paper, we address this issue from a new direction, by borrowing ideas from the field of dynamical systems. Since each policy operator defines a discrete-time dynamical system, tools from that field become directly applicable. In particular, we draw on notions of conjugacy. In dynamical systems and chaos theory, topological conjugacy is used to identify families of dynamical systems that have "equivalent" or related dynamics (see, e.g., Smale (1967) or Arnold (2012)). These ideas have proved enormously useful for categorizing different types of systems and allowing researchers to make deductions about a particular system from their knowledge of other related systems.[3]

We transfer these ideas to dynamic programming by studying order-theoretic conjugacy relationships between the policy operators that define a dynamic program. We call two dynamic programs linked in this way *isomorphic*. Under an order-reversing link, we call them *anti-isomorphic*. Our main result shows that verifying this single relationship is enough to transfer an entire package of optimality properties from one program to the other: existence and uniqueness of the value function, Bellman's principle of optimality, existence of optimal policies, and convergence of Howard policy iteration. Solving one program therefore delivers a complete solution to the other, including recovery of the value function and optimal policies.

The benefits of this approach are both theoretical and numerical. On the theoretical side, dynamic programs that are not easily analyzed can often be connected, via an isomorphism, to a program whose optimality properties are already established or more readily verified; the isomorphism then allows those properties to carry over to the original problem. We illustrate this in two settings that are not covered by standard contraction-based methods. First, for a dynamic program with Epstein–Zin utility and time preference shocks, we obtain an exact characterization of when

[2]See, for example, Kochenderfer et al. (2022), Rust (1994), Fei et al. (2021), or Kristensen et al. (2021).

[3]For example, Kennedy and Stockman (2008), Gardini et al. (2009), and Raines and Stockman (2012) characterize chaotic dynamical systems in economic models utilizing the shift homeomorphism. In a similar vein, Flynn and Sastry (2022) establish topological conjugacy between equilibrium economic dynamics and dynamical systems with known chaotic behaviors. Also see Battaglini (2021) and Deng et al. (2022).

optimality holds. Second, we show that multiplicative Kreps–Porteus preferences are isomorphic to risk-sensitive preferences, so that the well-developed contraction theory for the latter yields optimality results for the former. On the numerical side, within an isomorphism class one is free to work with whichever representative is best suited to computation. Exploiting this flexibility in the multisector real business cycle model of Long and Plosser (1983), we obtain solutions that are, on average, two orders of magnitude more accurate than solving the original problem, by choosing a representative within the isomorphism class whose value function has lower curvature.

This paper builds on the abstract dynamic programming framework of Sargent and Stachurski (2025). In that paper, individual dynamic programs are represented as families of policy operators acting on a partially ordered set. When dynamic programs are viewed in this way, studying conjugacy relationships between related families of policy operators becomes a natural idea, since each policy operator identifies a discrete time dynamical system.[4] Whereas Sargent and Stachurski (2025) study a single dynamic program in isolation, we introduce order isomorphisms and anti-isomorphisms between distinct dynamic programs and show that they preserve optimality properties. To apply these results to continuous-state, continuous-action models, we also relax their assumption that every candidate value function admits a maximizing policy in the Bellman equation.

This paper contributes to a large and growing literature that studies the optimality properties of dynamic programs with nonstandard features, such as recursive utility (Marinacci and Montrucchio, 2010; Bloise et al., 2024; Rincón-Zapatero, 2024), risk sensitive preferences (Bäuerle and Jaśkiewicz, 2018), hyperbolic discounting (Bäuerle et al., 2021; Balbus et al., 2022), and Q-factors (Ma et al., 2022). For example, our work relates to the literature on existence and uniqueness in dynamic programs with recursive preferences, such as Bloise and Vailakis (2018), Ren and Stachurski (2021), Stachurski and Zhang (2021), and Bloise et al. (2024). Our approach complements this literature in two ways. First, rather than focusing on specific models, we propose a general framework that, under appropriate transformations, can simplify the analysis of such dynamic programs. Second, we provide exact necessary and sufficient conditions for optimality in a continuous-state dynamic program with Epstein–Zin preferences and state-dependent discounting, whereas existing dynamic programming results for recursive preferences establish only sufficient conditions in similar settings.

[4]The main deviation from traditional dynamical systems methodology is that we use order-theoretic conjugacy relationships, rather than topological ones, since the key objects in optimization are order-theoretic and preserved by this relation.

Our order-theoretic approach also connects to the Bellman existence and uniqueness results of Kamihigashi (2014) and Bloise et al. (2024). Kamihigashi (2014) establishes existence and uniqueness of Bellman solutions and convergence of value iteration under purely order-theoretic conditions, while Bloise et al. (2024) use Tarski's fixed point theorem to establish Bellman solutions for stochastic recursive programs. In our framework, such fixed point results can be applied to a conjugate program where their hypotheses are easier to verify, as we illustrate in our continuous-state Epstein–Zin application using the existence result of Bloise et al. (2024).

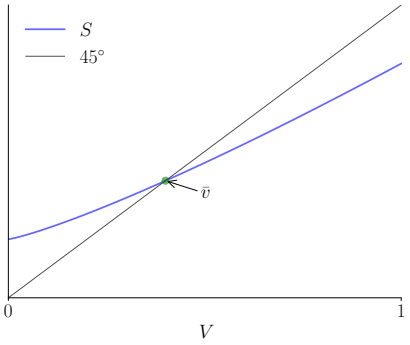
While not the primary focus of this paper, our work also connects to the literature on computational methods for solving dynamic economic models. These include traditional methods such as projection and perturbation (Foerster et al., 2016; Bayer and Luetticke, 2020), endogenous grid method (Carroll, 2006; Barillas and Fernández-Villaverde, 2007), adaptive sparse grids (Brumm and Scheidegger, 2017), and deep learning (Azinovic et al., 2022; Maliar et al., 2021). A notable feature of our framework is that it can complement existing computation methods: any suitable method can be applied to an isomorphic dynamic program to improve numerical performance compared to solving the original problem. Relatedly, Ma and Stachurski (2021) study transformations of Bellman equations that preserve optimality and can improve computational efficiency. Their approach refactors the Bellman equation, while ours relates entire dynamic programs through order conjugacy relationships.

The structure of the paper is as follows. Section 2 introduces order conjugacy and discusses its properties. Sections 3–4 present definitions and optimality results for "abstract" dynamic programs, which are well-suited to our dynamical systems perspective. Section 5 introduces isomorphisms between abstract dynamic programs and shows how isomorphisms preserve optimality properties. Section 6 provides applications and Section 7 concludes.

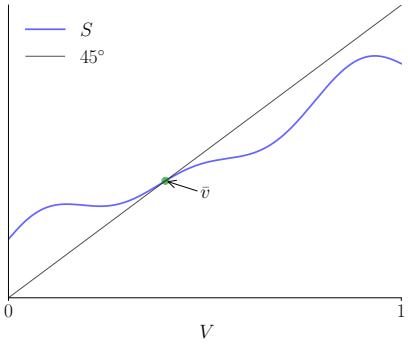
2. Preliminaries

This section introduces the key order-theoretic concepts needed for our analysis of abstract dynamic programs. Additional definitions of posets, dynamical systems, and related background are given in Appendix A. Throughout, if A and B are sets, then the symbol B^A indicates the set of all maps from A to B.

A *partially ordered set* (poset) is a pair $V = (V, \preceq)$ where $\preceq$ is reflexive, antisymmetric, and transitive. Given posets $(V, \preceq)$ and $(W, \trianglelefteq)$, a function $F \colon V \to W$ is called *order-preserving* if $v \preceq w$ implies $Fv \trianglelefteq Fw$, and *order-reversing* if $v \preceq w$ implies $Fw \trianglelefteq Fv$.



(A) Order-preserving



(B) Non-monotone

FIGURE 1. Order stable maps S on $[0, 1]$

A bijection $F\colon V \to \hat{V}$ is an *order isomorphism* if both F and F^{-1} are order-preserving, and an *order anti-isomorphism* if both F and F^{-1} are order-reversing.

Let (V, S) be a dynamical system (a set V with a self-map S) where V is a poset and S has a unique fixed point $\bar{v}$. Following Sargent and Stachurski (2025), (V, S) is called *upward stable* if $v \preceq S\,v$ implies $v \preceq \bar{v}$, *downward stable* if $S\,v \preceq v$ implies $\bar{v} \preceq v$, and *order stable* if both hold.

Figure 1 gives two examples of order stable maps S on $V = [0, 1]$. All points mapped up by S lie below its unique fixed point, while all points mapped down by S lie above its fixed point. The map in panel (b) is non-monotone, illustrating that order stability does not require monotonicity.

Intuitively, order stability says that the fixed point $\bar{v}$ of S separates V into two regions: elements that S maps upward all lie below $\bar{v}$, and elements that S maps downward all lie above $\bar{v}$. This is an order-theoretic analogue of global stability. (A dynamical system (V, S) on a metric space is called *globally stable* if S has a unique fixed point $\bar{v}$ and $S^k v \to \bar{v}$ for all $v \in V$.) Whereas global stability requires a topology and ensures convergence of trajectories, order stability requires only a partial order. Moreover, order stability is implied by global stability when S is order-preserving and the space V has both a topology *and* a partial order:

Example 2.1. Let $(V, \preceq, \rho)$ be a partially ordered space (i.e., a metric space (V, ρ) where $\preceq$ is closed under limits). If (V, S) is globally stable and S is order-preserving, then (V, S) is order stable. To see this, let $\bar{v}$ be the unique fixed point of S in V. Upward stability holds because if $v \in V$ and $v \preceq S\,v$, then, iterating on this inequality and using the fact that S is order-preserving, we have $v \preceq S^k\,v$ for all $k \in \mathbb{N}$. Applying

global stability and taking the limit gives $v \preceq \bar{v}$. Hence upward stability holds. The proof of downward stability is similar.

Two dynamical systems (V, S) and $(\hat{V}, \hat{S})$ on posets are called *order conjugate* under F when they are conjugate (i.e., $F \circ S = \hat{S} \circ F$ for some bijection F) and F is an order isomorphism. Order conjugacy is an equivalence relation that preserves order stability: if (V, S) and $(\hat{V}, \hat{S})$ are order conjugate, then one is order stable if and only if the other is (Lemma A.4 in Appendix A.4).

3. Abstract Dynamic Programs

We follow the abstract dynamic programming framework of Sargent and Stachurski (2025), which was partly inspired by Bertsekas (2022). For us, the benefit of the framework in Sargent and Stachurski (2025) is that dynamic programs are represented as families of policy operators. This abstract representation allows us to apply the notion of order conjugacy to study relationships between dynamic programs.

3.1. Definition and Examples.

Sargent and Stachurski (2025) define an *abstract dynamic program* (ADP) to be a pair $\mathscr{A} = (V, \{T_\sigma\}_{\sigma \in \Sigma})$, where

(i) $V = (V, \preceq)$ is a partially ordered set and
(ii) $\{T_\sigma\}_{\sigma \in \Sigma}$ is a family of order-preserving self-maps on V, indexed by $\sigma \in \Sigma$.

Elements of the index set Σ are referred to as *policies* and elements of $\{T_\sigma\}_{\sigma \in \Sigma}$ are called *policy operators*. When Σ is understood, we often write $\{T_\sigma\}_{\sigma \in \Sigma}$ as $\{T_\sigma\}$. Thus, each $\sigma \in \Sigma$ defines a dynamical system (V, T_σ), whereas $(V, \{T_\sigma\})$ denotes the entire ADP. In all applications, the significance of each policy operator T_σ is that its fixed point, denoted below by v_σ, represents the lifetime value (or cost) of following policy σ.

Example 3.1 (MDPs)**.** Consider a *Markov decision process* (MDP; see, e.g., Puterman (2005)) where the aim is to maximize $\mathbb{E} \sum_{t \geqslant 0} \beta^t r(X_t, A_t)$ when X_t takes values in finite set X (the state space), A_t takes values in finite set A (the action space), Γ is a nonempty correspondence from X to A, the set $\mathsf{G} := \{(x, a) \in \mathsf{X} \times \mathsf{A} : a \in \Gamma(x)\}$ denotes feasible state-action pairs, r is a reward function defined on G, $\beta \in (0, 1)$ is a discount factor, and $P \colon \mathsf{G} \times \mathsf{X} \to [0, 1]$ provides transition probabilities. The Bellman equation for this problem is

$$v(x) = \max_{a \in \Gamma(x)} \left\{ r(x, a) + \beta \sum_{x'} v(x') P(x, a, x') \right\} \qquad (x \in \mathsf{X}). \tag{1}$$

The set of feasible policies is $\Sigma := \{\sigma \in \mathsf{A}^{\mathsf{X}} : \sigma(x) \in \Gamma(x) \text{ for all } x \in \mathsf{X}\}$. We combine $\mathbb{R}^{\mathsf{X}}$ (the set of all real-valued functions on X) with the pointwise partial order $\leqslant$ and, for $\sigma \in \Sigma$ and $v \in \mathbb{R}^{\mathsf{X}}$, define the MDP policy operator

$$(T_\sigma v)(x) = r(x, \sigma(x)) + \beta \sum_{x'} v(x') P(x, \sigma(x), x') \qquad (x \in \mathsf{X}). \tag{2}$$

The pair $(\mathbb{R}^{\mathsf{X}}, \{T_\sigma\})$ is an ADP.

Example 3.2 (Risk-sensitive MDPs)**.** Risk-sensitive MDPs (Howard and Matheson, 1972) modify the MDP model in Example 3.1 so that the policy operators take the form

$$(T_\sigma^\theta v)(x) = r(x, \sigma(x)) + \frac{\beta}{\theta} \ln \left[\sum_{x'} \exp(\theta v(x')) P(x, \sigma(x), x') \right]$$

where θ is some fixed value in $\Theta := \mathbb{R} \setminus \{0\}$. The pair $(\mathbb{R}^{\mathsf{X}}, \{T_\sigma^\theta\})$ is an ADP.

Further examples of ADPs, including Q-factor and risk-sensitive Q-factor formulations from the reinforcement learning literature, are given in Appendix B.

3.2. **Lifetime Values.** The objective of dynamic programming is to maximize or minimize lifetime value. Following Sargent and Stachurski (2025), we identify the lifetime value of any given policy σ as the unique fixed point of T_σ, whenever it exists. When it does exist, we denote this fixed point by v_σ and call it the *σ-value function*.

Example 3.3. Consider the MDP setting of Example 3.1. Let r_σ and P_σ be defined by

$$r_\sigma(x) := r(x, \sigma(x)) \quad \text{and} \quad (P_\sigma v)(x) := \sum_{x'} v(x') P(x, \sigma(x), x'). \tag{3}$$

The standard MDP definition of the lifetime value of policy σ given $X_0 = x$ is $v_\sigma(x) = \mathbb{E} \sum_{t \geqslant 0} \beta^t r(X_t, \sigma(X_t))$, where $(X_t)_{t \geqslant 0}$ is a Markov chain generated by P_σ with initial condition $X_0 = x \in \mathsf{X}$. Pointwise on X, we can express v_σ as

$$v_\sigma = \sum_{t \geqslant 0} (\beta P_\sigma)^t r_\sigma = (I - \beta P_\sigma)^{-1} r_\sigma \tag{4}$$

(see, e.g., Puterman (2005), Theorem 6.1.1, or Kochenderfer et al. (2022), Section 7.2). Equivalently, v_σ is the unique solution to the equation $v = r_\sigma + \beta P_\sigma v$. Inspecting the definition of T_σ in (2), we see this is also equivalent to the statement that v_σ is the unique fixed point of T_σ.

Similarly, for the risk-sensitive MDP of Example 3.2, the operator T_σ^θ is a contraction on $\mathbb{R}^{\mathsf{X}}$ and its unique fixed point represents the lifetime value of following σ.

3.3. **Greedy Policies.** The core idea behind the theory of dynamic programming is that an optimal policy can be obtained by choosing actions that solve a two-period problem involving a Bellman equation (Bellman, 1957). The optimal actions in the two-period problem produce what are typically called "greedy policies." For example, in the context of Example 3.1, a policy σ satisfying

$$\sigma(x) \in \arg\max \left\{ r(x,a) + \beta \sum_{x'} v(x') P(x,a,x') \right\} \qquad \text{for all } x \in \mathsf{X}. \tag{5}$$

is called a "greedy policy with respect to v." See, for example, Kochenderfer et al. (2022), Section 7.3.

Following Sargent and Stachurski (2025), we generalize this idea to the abstract setting. In particular, given an ADP $(V, \{T_\sigma\})$ and an element $v \in V$, a policy σ in Σ is called

- *v-min-greedy* if $T_\sigma v \preceq T_\tau v$ for all $\tau \in \Sigma$, and
- *v-max-greedy* if $T_\tau v \preceq T_\sigma v$ for all $\tau \in \Sigma$.

For example, in the context of Example 3.1, with $\leqslant$ as the pointwise partial order, a policy σ obeying (5) satisfies $T_\tau v \leqslant T_\sigma v$ for all $\tau \in \Sigma$, and hence is v-max-greedy. A v-min-greedy policy can be constructed by replacing arg max in (5) with arg min.

3.4. **Bellman Operators.** For a generic ADP $(V, \{T_\sigma\})$, we respectively define the *Bellman min-operator* and the *Bellman max-operator* via

$$T_\triangledown v := \bigwedge_\sigma T_\sigma v \quad \text{and} \quad T_\triangle v := \bigvee_\sigma T_\sigma v \tag{6}$$

whenever the infimum (resp., supremum) exists. We say that $v \in V$ satisfies the *Bellman max-equation* (resp., the *Bellman min-equation*) if it is a fixed point of $T_\triangle$ (resp., $T_\triangledown$). Notice that $\sigma \in \Sigma$ is

(i) v-max-greedy if and only if $T_\sigma v = T_\triangle v$, and
(ii) v-min-greedy if and only if $T_\sigma v = T_\triangledown v$.

To illustrate, consider the MDP setting of Example 3.1. Traditionally, the Bellman operator for this model is given by

$$(Tv)(x) = \max_{a \in \Gamma(x)} \left\{ r(x,a) + \beta \sum_{x'} v(x') P(x,a,x') \right\} \qquad (x \in \mathsf{X}). \tag{7}$$

(See, e.g., Puterman (2005).) For the ADP $(\mathbb{R}^{\mathsf{X}}, \{T_\sigma\})$ generated by this MDP, this Bellman operator exactly coincides with the Bellman max-operator $T_\triangle$ in (6). Replacing max with min in (7) produces the Bellman min-operator from (6).

3.5. **Properties of ADPs.** To obtain optimality results, we need to place structure on ADPs. Here and below, we call an ADP $\mathscr{A} = (V, \{T_\sigma\})$

- *well-posed* if T_σ has a unique fixed point v_σ in V for each $\sigma \in \Sigma$,
- *order stable* if (V, T_σ) is order stable for each $\sigma \in \Sigma$.

In addition, we let

$$V_\triangle^G := \{v \in V : v \text{ has at least one } v\text{-max-greedy policy}\}$$

and define $V_\triangledown^G$ analogously, with min-greedy in place of max-greedy. Note that $T_\triangle\, v$ is well-defined whenever $v \in V_\triangle^G$, since the supremum in (6) is then attained. We call $\mathscr{A}$ *max-regular* if $V_\triangle^G = V$ and *min-regular* if $V_\triangledown^G = V$.

Our optimality results require the following combination of conditions. We call $\mathscr{A}$

- *max-stable* if $\mathscr{A}$ is order stable and $T_\triangle$ has a fixed point in $V_\triangle^G$, and
- *min-stable* if $\mathscr{A}$ is order stable and $T_\triangledown$ has a fixed point in $V_\triangledown^G$.

Well-posedness is a minimum regularity condition for ADPs. Without it, we cannot be sure that policies have well-defined lifetime values. Well-defined lifetime values are essential because maximizing (or minimizing) lifetime value over the set of all policies is the objective of dynamic programming.

Order stability is a natural regularity assumption for ADPs. To understand this, suppose $\mathscr{A}$ is well-posed and consider upward stability for an arbitrary policy operator T_σ with fixed point v_σ. If $v \preceq T_\sigma v$, then following policy σ for one period offers an improvement in value. Since the problem is stationary, this suggests that following the policy forever will also be an improvement. Thus, we expect $v \preceq v_\sigma$, in which case upward stability holds. Intuition for downward stability is similar.

Example 3.4. The ADP $(\mathbb{R}^{\mathsf{X}}, \{T_\sigma\})$ from the MDP setting of Example 3.1 is well-posed, order stable and max-regular. Well-posedness was shown in Section 3.2, where we saw that each policy operator T_σ has unique fixed point $v_\sigma := (I - \beta P_\sigma)^{-1} r_\sigma$. Regarding order stability, fix $\sigma \in \Sigma$. If $T_\sigma v \geqslant v$, then $r_\sigma + \beta P_\sigma v \geqslant v$ and hence $(I - \beta P_\sigma)v \leqslant r_\sigma$. Since $(I - \beta P_\sigma)^{-1}$ is positive, we get $v \leqslant v_\sigma$. Hence (V, T_σ) is upward stable. A similar proof shows that (V, T_σ) is downward stable, and therefore order stable. Finally, since X and each $\Gamma(x)$ are finite, the maximum in (5) is attained at

each $x \in \mathsf{X}$ for every $v \in \mathbb{R}^{\mathsf{X}}$. Hence every $v \in \mathbb{R}^{\mathsf{X}}$ has a v-max-greedy policy, and the ADP is max-regular.

Example 3.5. The ADP $(\mathbb{R}^{\mathsf{X}}, \{T_\sigma^\theta\})$ generated by the risk-sensitive MDP from Example 3.2 is well-posed, order stable and max-regular. Evidently T_σ^θ is order-preserving on $V = \mathbb{R}^{\mathsf{X}}$. Moreover, (V, T_σ^θ) is globally stable (see, e.g., Bäuerle and Jaśkiewicz (2018)). In particular, T_σ^θ has a unique fixed point in V, so the ADP is well-posed. Global stability also implies that (V, T_σ^θ) is order stable, by Example 2.1. Given $v \in V$, we can construct a v-max-greedy policy σ by setting

$$\sigma(x) \in \arg\max \left\{ r(x,a) + \frac{\beta}{\theta} \ln \left[\sum_{x'} \exp(\theta v(x')) P(x,a,x') \right] \right\}$$

for all $x \in \mathsf{X}$. Hence $(V, \{T_\sigma^\theta\})$ is max-regular.

In Section 4.3 we show that the ADPs in Examples 3.4 and 3.5 are also max-stable.

4. Optimality

In this section we state conditions for optimality in our abstract setting. These conditions are closely related to those in Sargent and Stachurski (2025). Later, in Section 5, we will study how optimality properties are preserved under transformations.

4.1. Max-Optimality.

Let $\mathscr{A}$ be a well-posed ADP. We let $V_\Sigma := \{v_\sigma\}_{\sigma \in \Sigma}$ denote the set of σ-value functions. If V_Σ has a greatest element, then we denote it by v_Δ and call it the *max-value function*. A policy $\sigma \in \Sigma$ is called *max-optimal* for $\mathscr{A}$ if v_σ is a greatest element of V_Σ; that is, if v_Δ exists and $v_\sigma = v_\Delta$.

In the following result, we take $\mathscr{A}$ to be an ADP with Bellman operator T_Δ.

Theorem 4.1 (Max-optimality)**.** *If $\mathscr{A}$ is max-stable, then*

(i) *the max-value function v_Δ exists and lies in V_Δ^G,*
(ii) *v_Δ is the unique solution to the Bellman max-equation in V_Δ^G,*
(iii) *a policy is max-optimal if and only if it is v_Δ-max-greedy, and*
(iv) *at least one max-optimal policy exists.*

The proof of Theorem 4.1 is given in the appendix (page 28).

Theorem 4.1 mildly extends Theorem 5.1 of Sargent and Stachurski (2025) by replacing max-regularity with the weaker requirement that T_Δ have a fixed point in

V_Δ^G. This weaker condition allows the isomorphism results of Section 5 to be applied to continuous-state models, where max-regularity typically fails. If $\mathscr{A}$ is max-regular, then $V_\Delta^G = V$, and part (ii) states that v_Δ is the unique solution to the Bellman max-equation in V.

4.2. **Min-Optimality.** Analogous results hold for minimization. Let $\mathscr{A}$ be a well-posed ADP. We call a policy $\sigma \in \Sigma$ *min-optimal* for $\mathscr{A}$ if v_σ is a least element of V_Σ. When V_Σ has a least element we denote it by v_∇ and call it the *min-value function*.

In the following result, we take $\mathscr{A}$ to be an ADP with Bellman operator T_∇.

Theorem 4.2 (Min-optimality)**.** *If $\mathscr{A}$ is min-stable, then*

(i) *the min-value function v_∇ exists and lies in V_∇^G,*
(ii) *v_∇ is the unique solution to the Bellman min-equation in V_∇^G,*
(iii) *a policy is min-optimal if and only if it is v_∇-min-greedy, and*
(iv) *at least one min-optimal policy exists.*

Theorem 4.2 follows from Theorem 4.1 by order duality. The proof is given in Appendix C.2.

4.3. **Sufficient Conditions.** Theorems 4.1 and 4.2 are high-level results that assume the Bellman operator has a fixed point in V_Δ^G or V_∇^G. Since order stability is often straightforward to check, as Examples 3.4 and 3.5 show, verifying max- or min-stability in applications mainly requires establishing this fixed point condition. Below we discuss two familiar settings in which max- or min-stability can be verified using general results, followed by cases that require more specific metric or order-theoretic arguments.

4.3.1. *Contractions.* A common strategy is to find a subset $V_0 \subseteq V_\Delta^G$ on which T_Δ has a fixed point. By Theorem 4.1(ii), v_Δ then lies in V_0. This is useful because policy evaluation and Bellman optimization often call for different regularity, as illustrated below.

The next result implements this strategy when the policy operators are contractions. It requires the metric ρ to be sup-nonexpansive, a condition defined in Appendix C.3 and satisfied, for example, by the supremum norm on bounded functions with the pointwise order.

Proposition 4.3. *Let $(V, \preceq, \rho)$ be a partially ordered space with ρ complete and sup-nonexpansive, and let V_0 be a closed subset of V such that $V_0 \subseteq V_\Delta^G$ and $T_\Delta V_0 \subseteq V_0$. If each T_σ is a contraction of modulus $\beta \in [0, 1)$ on V, then $\mathscr{A}$ is max-stable and $v_\Delta \in V_0$.*

Proposition 4.3 is a variation on a result in Peng et al. (2026) and is proved in Appendix C.3. Applying it to the dual ADP $\mathscr{A}^\partial$ (Appendix C.2) yields a min version whenever ρ satisfies the analogous condition for infima.

To illustrate, consider a continuous-state version of the MDP in Example 3.1 with compact metric state and action spaces and bounded rewards. Let V be the bounded Borel measurable functions under the supremum norm and V_0 be the subset of continuous functions. Each policy operator is a contraction on V but need not preserve continuity. Under standard regularity conditions on the model primitives and admissible policies, however, T_Δ maps V_0 into itself and every $v \in V_0$ has a measurable greedy policy (see Peng et al. (2026)). Proposition 4.3 then implies that the ADP is max-stable and $v_\Delta \in V_0$, even though all measurable policies, including those with discontinuous value functions, are compared.

4.3.2. *Finite Policy Sets.* Here is another useful condition, which covers the case where state and action spaces are finite.

Proposition 4.4. *Let $\mathscr{A}$ be order stable and suppose the set of policies is finite. In this setting,*

(i) *if $\mathscr{A}$ is max-regular, then $\mathscr{A}$ is max-stable, and*
(ii) *if $\mathscr{A}$ is min-regular, then $\mathscr{A}$ is min-stable.*

Proposition 4.4 follows from Theorem 5.4 of Sargent and Stachurski (2025); for completeness, a proof is given in Appendix C.3.

For example, the ADPs in Examples 3.4 and 3.5 are order stable and max-regular, and their policy sets are finite, since X and the choice sets are all finite. Hence, by Proposition 4.4, both ADPs are max-stable.

4.3.3. *Direct Verification.* Some important models fit neither of the settings above. In such cases, max- or min-stability can be verified directly using fixed point theorems suited to the operators at hand, often after passing to a conjugate representation (Section 5). Although an order isomorphism preserves purely order-theoretic properties such as chain completeness and order continuity, a nonlinear transformation can induce properties that combine the partial order with vector-space or metric structure,

such as concavity, convexity, or contractiveness. For example, in the continuous-state Epstein–Zin model of Section 6.1, the operators need not be contractions, but a conjugate representation makes them concave or convex, allowing us to apply the fixed point theorems of Stachurski et al. (2025). Other fixed point results that can be used to establish Bellman fixed points include the order-theoretic results of Kamihigashi (2014) for additive programs and those of Bloise et al. (2024) for recursive programs.

5. Isomorphic ADPs

In this section we introduce isomorphic relationships between ADPs and explore their implications for optimality. True to their name, isomorphic relationships are symmetric, transitive and reflexive. We show that isomorphic ADPs have identical optimality properties.

5.1. Definition and Properties.

Let $\mathscr{A} = (V, \{T_\sigma\})$ and $\hat{\mathscr{A}} = (\hat{V}, \{\hat{T}_\sigma\})$ be two ADPs. We call $\mathscr{A}$ and $\hat{\mathscr{A}}$ *isomorphic* under F if these two ADPs have the same policy set Σ and F is an order isomorphism from V to $\hat{V}$ such that

$$F \circ T_\sigma = \hat{T}_\sigma \circ F \quad \text{on } V \text{ for all } \sigma \in \Sigma. \tag{8}$$

In other words, (V, T_σ) and $(\hat{V}, \hat{T}_\sigma)$ are order conjugate under F for all $\sigma \in \Sigma$.[5]

An illustration of isomorphic ADPs built from Q-factor formulations is given in Appendix B. We also provide a further illustration involving risk-sensitive and Kreps–Porteus preferences in Section 6.2.

Remark 5.1. For ADPs $\mathscr{A}, \hat{\mathscr{A}}$, let $\mathscr{A} \sim \hat{\mathscr{A}}$ indicate that $\mathscr{A}$ and $\hat{\mathscr{A}}$ are isomorphic. It is elementary to show that the relation $\sim$ is reflexive, symmetric and transitive. Hence $\sim$ is an equivalence relation on the set of all ADPs.

5.2. Isomorphisms and Optimality.

We seek a connection between value functions and optimality properties of isomorphic ADPs. The theory below provides this relationship. For all of this section, we take $\mathscr{A} = (V, \{T_\sigma\})$ and $\hat{\mathscr{A}} = (\hat{V}, \{\hat{T}_\sigma\})$ to be two ADPs with the same policy set. When they exist, we let

- v_σ (resp., $\hat{v}_\sigma$) be the unique fixed point of T_σ (resp., $\hat{T}_\sigma$)
- T_Δ (resp., $\hat{T}_\Delta$) be the Bellman max-operator of $\mathscr{A}$ (resp., $\hat{\mathscr{A}}$)

[5] While the definition requires that the two ADPs have the same policy set Σ, it suffices that the policy sets can be put in one-to-one correspondence with each other. Indeed, given a bijection $h\colon \Sigma \to \hat{\Sigma}$, we can index the second family of policy operators by Σ, writing $\tilde{T}_\sigma := \hat{T}_{h(\sigma)}$.

- T_∇ (resp., $\hat{T}_\nabla$) be the Bellman min-operator of $\mathscr{A}$ (resp., $\hat{\mathscr{A}}$)
- v_Δ (resp., $\hat{v}_\Delta$) be the max-value function of $\mathscr{A}$ (resp., $\hat{\mathscr{A}}$)
- v_∇ (resp., $\hat{v}_\nabla$) be the min-value function of $\mathscr{A}$ (resp., $\hat{\mathscr{A}}$)

Isomorphic ADPs share the same regularity properties:

Theorem 5.1. *If $\mathscr{A}$ and $\hat{\mathscr{A}}$ are isomorphic under F, then*

(i) *$\mathscr{A}$ is well-posed if and only if $\hat{\mathscr{A}}$ is well-posed.*

(ii) *$\mathscr{A}$ is order stable if and only if $\hat{\mathscr{A}}$ is order stable.*

(iii) *$\mathscr{A}$ is max-stable if and only if $\hat{\mathscr{A}}$ is max-stable. In this case,*

$$F \circ T_\Delta = \hat{T}_\Delta \circ F \quad \text{on } V_\Delta^G \quad \text{and} \quad \hat{v}_\Delta = F v_\Delta. \tag{9}$$

Moreover, $\mathscr{A}$ and $\hat{\mathscr{A}}$ have the same max-optimal policies.

(iv) *$\mathscr{A}$ is min-stable if and only if $\hat{\mathscr{A}}$ is min-stable. In this case,*

$$F \circ T_\nabla = \hat{T}_\nabla \circ F \quad \text{on } V_\nabla^G \quad \text{and} \quad \hat{v}_\nabla = F v_\nabla. \tag{10}$$

Moreover, $\mathscr{A}$ and $\hat{\mathscr{A}}$ have the same min-optimal policies.

The proof (page 31) uses condition (8) to show that F maps V_Δ^G onto $\hat{V}_\Delta^G$ and commutes with the Bellman operators on V_Δ^G (and similarly for the min case), and then combines Theorem 4.1 with properties of order conjugate systems.

Remark 5.2 (Monotone comparative statics)**.** Theorem 5.1 also allows monotone comparative statics established in one representation to carry over to another.[6] Let θ range over a poset Θ, and suppose that $\mathscr{A}_\theta$ and $\hat{\mathscr{A}}_\theta$ are max-stable and isomorphic under F_θ for each $\theta \in \Theta$, with a common partially ordered policy set Σ. By Theorem 5.1(iii), $\mathscr{A}_\theta$ and $\hat{\mathscr{A}}_\theta$ have the same max-optimal policies at every θ, so any monotonicity of the optimal-policy correspondence in θ established for one family holds immediately for the other, even when F_θ depends on θ. On the other hand, comparing max-value functions across parameters requires more structure: if $F_\theta = F$ is independent of θ, then $\hat{v}_\Delta = F v_\Delta$ implies that the max-value function is increasing in θ in one representation if and only if it is in the other, whereas a parameter-dependent F_θ requires additional conditions on how F_θ varies with θ.

In applications, monotonicity of optimal policies is often established by verifying lattice-programming conditions on the Bellman objective. Ordinal conditions such as quasisupermodularity and single crossing (Milgrom and Shannon, 1994) are preserved whenever the objective at each state undergoes a strictly increasing transformation,

[6] We thank an anonymous referee for pointing out this connection.

so they hold in one representation if and only if they hold in the other. Cardinal conditions such as supermodularity and increasing differences, by contrast, are not invariant under nonlinear transformations. As with concavity or contractiveness in Section 6, this allows a suitable transformation F_θ to induce supermodularity or increasing differences in $\hat{\mathscr{A}}_\theta$ when they fail in $\mathscr{A}_\theta$, although identifying general conditions that guarantee these properties in the transformed problem is nontrivial.

5.3. **Anti-Isomorphic ADPs.** Let $\mathscr{A} = (V, \{T_\sigma\})$ and $\hat{\mathscr{A}} = (\hat{V}, \{\hat{T}_\sigma\})$ be two ADPs. We call $\mathscr{A}$ and $\hat{\mathscr{A}}$ *anti-isomorphic* under F if they have the same policy set Σ and, in addition, there exists an order anti-isomorphism F from V to $\hat{V}$ such that (8) holds (see Example B.3).

If $\mathscr{A} \overset{a}{\sim} \hat{\mathscr{A}}$ indicates that $\mathscr{A}$ and $\hat{\mathscr{A}}$ are anti-isomorphic, then $\overset{a}{\sim}$ is symmetric but, in general, neither reflexive nor transitive.

Here is an optimality result for anti-isomorphic ADPs that parallels Theorem 5.1.

Theorem 5.2. *If $\mathscr{A}$ and $\hat{\mathscr{A}}$ are anti-isomorphic under F, then*

(i) *$\mathscr{A}$ is well-posed if and only if $\hat{\mathscr{A}}$ is well-posed.*
(ii) *$\mathscr{A}$ is order stable if and only if $\hat{\mathscr{A}}$ is order stable.*
(iii) *$\mathscr{A}$ is max-stable if and only if $\hat{\mathscr{A}}$ is min-stable. In this case,*

$$F \circ T_\triangle = \hat{T}_\triangledown \circ F \quad \text{on } V_\triangle^G \quad \text{and} \quad \hat{v}_\triangledown = F v_\triangle. \tag{11}$$

Moreover, $\sigma \in \Sigma$ is max-optimal for $\mathscr{A}$ if and only if σ is min-optimal for $\hat{\mathscr{A}}$.

The proof of Theorem 5.2 is given in the appendix (page 32).

5.4. **Algorithms.** Next we consider three standard dynamic programming algorithms: value function iteration (VFI), optimistic policy iteration (OPI) and Howard policy iteration (HPI). We show that isomorphic ADPs generate conjugate sequences under each of these algorithms, so that convergence in one representation is equivalent to convergence in the other.

Let $\mathscr{A} = (V, \{T_\sigma\})$ be a max-stable ADP and fix $v \in V_\triangle^G$ and $m \in \mathbb{N}$.

- A *VFI sequence* from v is a sequence $(v_n)_{n\geqslant 0}$ in $V_\triangle^G$ with $v_0 = v$ and $v_{n+1} = T_\triangle v_n$ for all n.
- An *OPI sequence* with step size m from v is a pair of sequences $(v_n)_{n\geqslant 0}$ in $V_\triangle^G$ and $(\sigma_n)_{n\geqslant 0}$ in Σ such that $v_0 = v$, each σ_n is v_n-max-greedy and $v_{n+1} = T_{\sigma_n}^m v_n$ for all n.

- An *HPI sequence* from v is a pair of sequences $(v_n)_{n\geqslant 0}$ in V_Δ^G and $(\sigma_n)_{n\geqslant 0}$ in Σ such that $v_0 = v$, each σ_n is v_n-max-greedy and $v_{n+1} = v_{\sigma_n}$ for all n.[7]

Following Peng et al. (2026), we use an order-theoretic notion of convergence, writing $v_n \uparrow v_\Delta$ if (v_n) is increasing and $\bigvee_n v_n = v_\Delta$ (see Appendix A.1).[8] An HPI sequence *converges in finitely many steps* if $v_n = v_\Delta$ for some n. When $\mathscr{A}$ is order stable and max-regular and Σ is finite, every HPI sequence converges in finitely many steps (see Proposition 4.4, Theorem 5.4 of Sargent and Stachurski (2025), and Appendix C.3).

Proposition 5.3. *Let $\mathscr{A}$ and $\hat{\mathscr{A}}$ be isomorphic under F, and let $\mathscr{A}$ be max-stable. For all $v \in V_\Delta^G$ and $m \in \mathbb{N}$,*

(i) *(v_n) is a VFI sequence for $\mathscr{A}$ from v if and only if (Fv_n) is a VFI sequence for $\hat{\mathscr{A}}$ from Fv,*
(ii) *$((v_n), (\sigma_n))$ is an OPI sequence with step size m for $\mathscr{A}$ from v if and only if $((Fv_n), (\sigma_n))$ is an OPI sequence with step size m for $\hat{\mathscr{A}}$ from Fv, and*
(iii) *$((v_n), (\sigma_n))$ is an HPI sequence for $\mathscr{A}$ from v if and only if $((Fv_n), (\sigma_n))$ is an HPI sequence for $\hat{\mathscr{A}}$ from Fv.*

In each case, $v_n \uparrow v_\Delta$ if and only if $Fv_n \uparrow \hat{v}_\Delta$, and an HPI sequence for $\mathscr{A}$ converges in finitely many steps if and only if the corresponding HPI sequence for $\hat{\mathscr{A}}$ does.

Thus, isomorphic ADPs run the same algorithms in lockstep. In particular, OPI and HPI select the same sequence of policies in both representations. No continuity assumption on F is required. An analogous result for anti-isomorphic ADPs is given with the proofs in Appendix C.4.

6. Applications

In this section we show how isomorphic relationships can simplify or illuminate dynamic programming problems.

[7] If $\mathscr{A}$ is max-regular, then VFI sequences always exist, while OPI and HPI sequences exist for every choice of greedy policies. Without regularity, the iterates can leave V_Δ^G, in which case no such sequence exists from v.

[8] When $v \preceq T_\Delta v$, the sequence (v_n) is increasing under each algorithm whenever it exists.

6.1. **Modified Epstein–Zin Equations.** Consider an Epstein–Zin version of the MDP in Example 3.1 (see, e.g., Epstein and Zin (1989) or Weil (1990)) with X and A compact metric spaces. The Bellman max-equation takes the form

$$v(x) = \max_{a \in \Gamma(x)} \left\{ r(x,a)^\alpha + \beta(x) \left(\int_{\mathsf{X}} v(x')^\gamma P(x,a,\mathrm{d}x') \right)^{\alpha/\gamma} \right\}^{1/\alpha},$$

where γ and α are nonzero parameters. Following much of the recent literature, we have allowed the rate of time preference β to depend on the state (see, e.g., Albuquerque et al. (2016), Schorfheide et al. (2018), de Groot et al. (2018), Gomez-Cram and Yaron (2020)). Let Γ be continuous with nonempty compact values, and let $r\colon \mathsf{G} \to (0,\infty)$ and $\beta\colon \mathsf{X} \to (0,\infty)$ be continuous. Let Σ consist of all Borel measurable selections from Γ. With r_σ as in (3) and P_σ defined by

$$(P_\sigma v)(x) := \int_{\mathsf{X}} v(x') P(x, \sigma(x), \mathrm{d}x'),$$

the policy operator T_σ can be written as

$$T_\sigma v = \left\{ r_\sigma^\alpha + \beta \, (P_\sigma v^\gamma)^{\alpha/\gamma} \right\}^{1/\alpha}, \tag{12}$$

where powers are taken pointwise.

Let $B_b(\mathsf{X})$ and $C(\mathsf{X})$ denote the bounded Borel measurable and continuous real-valued functions on X, respectively. We assume that the transition kernel P is weakly continuous: for every $f \in C(\mathsf{X})$, the map

$$(x,a) \mapsto \int_{\mathsf{X}} f(x') P(x,a,\mathrm{d}x')$$

is continuous on G. Set

$$V := \{ v \in B_b(\mathsf{X}) : \inf_{x \in \mathsf{X}} v(x) > 0 \}, \tag{13}$$

paired with the pointwise order. Under these assumptions, T_σ is an order-preserving self-map on $(V, \leqslant)$, and $\mathscr{A} := (V, \{T_\sigma\})$ is an ADP.

Now set

$$\theta := \frac{\gamma}{\alpha} \quad \text{and} \quad \hat{T}_\sigma v := \left\{ r_\sigma^\alpha + \beta \, (P_\sigma v)^{1/\theta} \right\}^\theta. \tag{14}$$

The pair $\hat{\mathscr{A}} = (V, \{\hat{T}_\sigma\})$ is also an ADP.

Lemma 6.1. *The following relationships hold:*

(i) *If $\gamma > 0$, then $\mathscr{A}$ and $\hat{\mathscr{A}}$ are isomorphic.*
(ii) *If $\gamma < 0$, then $\mathscr{A}$ and $\hat{\mathscr{A}}$ are anti-isomorphic.*

Proof. Let $F\colon V \to V$ be the bijective map $Fv = v^\gamma$ (since powers preserve Borel measurability, boundedness, and a positive infimum). Fixing σ and applying (12) yields

$$F T_\sigma v = (T_\sigma v)^\gamma = \left\{ r_\sigma^\alpha + \beta \, (P_\sigma v^\gamma)^{\alpha/\gamma} \right\}^{\gamma/\alpha} = \left\{ r_\sigma^\alpha + \beta \, (P_\sigma v^\gamma)^{1/\theta} \right\}^\theta .$$

Inspection of (14) shows that $\hat{T}_\sigma F v = \hat{T}_\sigma v^\gamma$ is identical to the last expression in the display above. Hence $F \circ T_\sigma = \hat{T}_\sigma \circ F$ on V. If $\gamma > 0$, then F is order-preserving, so $\mathscr{A}$ and $\hat{\mathscr{A}}$ are isomorphic. If $\gamma < 0$, then F is order-reversing, so $\mathscr{A}$ and $\hat{\mathscr{A}}$ are anti-isomorphic. □

Lemma 6.1 gives us a way to solve for max-optimal policies of $\mathscr{A}$ by studying $\hat{\mathscr{A}}$ (and applying either Theorem 5.1 or Theorem 5.2). This is convenient because $\hat{\mathscr{A}}$ is easier to analyze. The next section illustrates.

6.1.1. *Characterizing Optimality.* Suppose that β depends on x through a purely exogenous state component (as in, say, de Groot et al. (2018) and Schorfheide et al. (2018)). Specifically, $\mathsf{X} = \mathsf{Y} \times \mathsf{Z}$ and $x = (y, z)$, where

$$P(x, a, \mathsf{Y} \times E) = Q(z, E) \quad \text{and} \quad \beta(x) = \beta(z) \tag{15}$$

for every $(x, a) \in \mathsf{G}$ and Borel $E \subseteq \mathsf{Z}$. Here Q is a Markov kernel on Z, which is weakly continuous by the assumptions on P and Γ. For each $z \in \mathsf{Z}$, let $(Z_t(z))_{t \geqslant 0}$ be a Markov chain on Z generated by Q and starting at z. Define

$$\Lambda := \lim_{n \to \infty} \left[\sup_{z \in \mathsf{Z}} \left\{ \mathbb{E} \prod_{t=0}^{n-1} \beta(Z_t(z))^\theta \right\}^{1/\theta} \right]^{1/n} . \tag{16}$$

Set $V_0 := V \cap C(\mathsf{X})$. We can now state the following exact result.

Theorem 6.2. *Under the preceding assumptions, the following are equivalent:*

(i) $\Lambda < 1$*;*
(ii) $\mathscr{A}$ *is well-posed;*
(iii) $\mathscr{A}$ *is max-stable.*

Hence, under $\Lambda < 1$, $v_\triangle \in V_0$ *and all of the optimality results in Theorem 4.1 apply. Conversely, if* $\Lambda \geq 1$, *no* T_σ *has a fixed point in* V.

Theorem 6.2 shows that $\Lambda < 1$ is both necessary and sufficient for well-posedness and optimality across all nonzero α and γ. With time-preference shocks, $\Lambda < 1$ restricts (α, γ) jointly with the shock process through $\theta = \gamma/\alpha$: by the power-mean

inequality, Λ is increasing in θ, as a larger θ places more weight on states with high $\beta(z)$. Here $1-\gamma$ is relative risk aversion (RRA) and $(1-\alpha)^{-1}$ is the elasticity of intertemporal substitution (EIS). In calibrations with RRA and the EIS both above one (Albuquerque et al., 2016; Schorfheide et al., 2018), we have $\alpha > 0 > \gamma$ and $\theta < 0$, so a higher EIS or lower RRA increases θ and tightens $\Lambda < 1$.

Existing dynamic programming results for recursive preferences establish sufficient conditions for optimality under constant discounting (Bloise and Vailakis, 2018; Marinacci and Montrucchio, 2019; Ren and Stachurski, 2021) or state-dependent discounting (Stachurski and Zhang, 2021; Bloise et al., 2024).[9] By contrast, Theorem 6.2 also establishes necessity, which is new. Its proof isolates the linear continuation operator $\beta^\theta P_\sigma$ inside the conjugate ADP $\hat{\mathscr{A}}$, which leads to the necessity of $\Lambda < 1$. For sufficiency, we establish order stability and existence of a Bellman fixed point using the concave and convex fixed point theorems of Stachurski et al. (2025); we also provide an alternative proof of the Bellman fixed point step using the order-theoretic results of Bloise et al. (2024) (Appendix C.5).

6.1.2. *Finite Irreducible Shocks.* When Z is finite and Q is irreducible, the condition has a simpler spectral form. Define the matrix K by $K(z,z') := \beta(z)^\theta Q(z,z')$, and let $\rho(K)$ denote its spectral radius.

Corollary 6.3. *Under the assumptions of Theorem 6.2, if Z is finite and Q is irreducible, then $\Lambda = \rho(K)^{1/\theta}$. Hence $\mathscr{A}$ is max-stable if and only if*

$$\rho(K)^{1/\theta} < 1. \tag{17}$$

Without irreducibility, the identity $\Lambda = \rho(K)^{1/\theta}$ can fail when $\theta < 0$. The proof of the corollary is also in Appendix C.5.

6.2. **Kreps–Porteus vs Risk Sensitive Preferences.** As a further application, we show that the risk-sensitive preferences and a multiplicative variation of Kreps–Porteus preferences, are connected by an isomorphism. This allows optimality results from the well-studied risk-sensitive framework to be transferred to the multiplicative Kreps–Porteus model.

Consider first the finite MDP framework of Example 3.1, using r_σ and P_σ as in (3). Recall the risk-sensitive MDP in Example 3.2, where the policy operators take the

[9] For existence and uniqueness of recursive utility under an exogenous consumption process, Borovička and Stachurski (2020) and Stachurski et al. (2025) obtain necessary and sufficient spectral conditions under irreducibility and compactness assumptions.

form

$$(T_\sigma^\theta v)(x) = r_\sigma(x) + \frac{\beta}{\theta} \ln \left[\sum_{x'} \exp(\theta v(x')) P(x, \sigma(x), x') \right] \tag{18}$$

for $v \in \mathbb{R}^{\mathsf{X}}$ and $\theta \in \mathbb{R} \setminus \{0\}$. We show in Example 3.5 and Section 4.3 that the ADP $(\mathbb{R}^{\mathsf{X}}, \{T_\sigma^\theta\})$ is max-stable and hence all of the optimality results in Theorem 4.1 apply.

An alternative formulation is to replace the entropic certainty equivalent with Kreps–Porteus expectations, leading to the policy operator

$$(K_\sigma v)(x) = r_\sigma(x) + \beta \left\{ (P_\sigma v^\nu)(x) \right\}^{1/\nu} \tag{19}$$

for $v \in (0, \infty)^{\mathsf{X}}$ and $\nu \in \mathbb{R} \setminus \{0\}$, where r is strictly positive. The Kreps–Porteus ADP $((0, \infty)^{\mathsf{X}}, \{K_\sigma\})$ is harder to analyze directly because K_σ is not generally a contraction. There is, however, a multiplicative variation on the Kreps–Porteus ADP that is simple to analyze. The model is obtained by setting

$$(M_\sigma v)(x) = r_\sigma(x) \cdot \left\{ (P_\sigma v^\nu)(x) \right\}^{\beta/\nu} \tag{20}$$

where $v \in (0, \infty)^{\mathsf{X}}$, r is strictly positive, and $\beta \in [0, 1)$. We call $\mathscr{A} := ((0, \infty)^{\mathsf{X}}, \{M_\sigma\})$ the multiplicative Kreps–Porteus (MKP) ADP.

It turns out that the MKP ADP and the risk-sensitive ADP are isomorphic. To see this, we take logs of the Bellman max-equation associated with the MKP ADP, obtaining

$$\ln v(x) = \max_{a \in \Gamma(x)} \left\{ \ln r(x, a) + \frac{\beta}{\nu} \ln \left[\sum_{x'} v(x')^\nu P(x, a, x') \right] \right\}.$$

Setting $\hat{v} = \ln v$ and $\hat{r} = \ln r$ yields

$$\hat{v}(x) = \max_{a \in \Gamma(x)} \left\{ \hat{r}(x, a) + \frac{\beta}{\nu} \ln \left[\sum_{x'} \exp(\nu \hat{v}(x')) P(x, a, x') \right] \right\}.$$

This is the Bellman max-equation for the risk-sensitive MDP in (18), with $\hat{r}$ in place of r and ν in place of θ.

We can formalize this observation. The MKP ADP is $\mathscr{A} = ((0, \infty)^{\mathsf{X}}, \{M_\sigma\})$. Let $F \colon (0, \infty)^{\mathsf{X}} \to \mathbb{R}^{\mathsf{X}}$ be the pointwise logarithm, $Fv = \ln v$, which is an order isomorphism with inverse $F^{-1} = \exp$. Consider the risk-sensitive ADP $\hat{\mathscr{A}} := (\mathbb{R}^{\mathsf{X}}, \{T_\sigma^\nu\})$ with reward $\hat{r} = \ln r$ and risk-sensitivity parameter $\theta = \nu$.

It remains to verify condition (8), i.e., that $M_\sigma = F^{-1} \circ T^\nu_\sigma \circ F$ on $(0,\infty)^{\mathsf{X}}$. Setting $\hat{v} = Fv = \ln v$ for $v \in (0,\infty)^{\mathsf{X}}$, we have

$$\begin{aligned}(F^{-1} \circ T^\nu_\sigma \circ F)(v)(x) &= \exp\left[\ln r_\sigma(x) + \frac{\beta}{\nu} \ln \left[\sum_{x'} \exp(\nu\, \hat{v}(x'))P(x,\sigma(x),x')\right]\right] \\ &= r_\sigma(x) \cdot \left\{\sum_{x'} \exp(\nu\, \ln v(x'))P(x,\sigma(x),x')\right\}^{\beta/\nu} \\ &= r_\sigma(x) \cdot \{(P_\sigma\, v^\nu)(x)\}^{\beta/\nu} = (M_\sigma\, v)(x).\end{aligned}$$

Since F is an order isomorphism and $\hat{\mathscr{A}}$ is max-stable, Theorem 5.1(iii) implies that the MKP ADP $\mathscr{A}$ is also max-stable for every $\nu \neq 0$, the value functions are related by $\hat{v}_\Delta = Fv_\Delta$, and the two ADPs have the same max-optimal policies.

This result provides a practical benefit. Since the risk-sensitive MDP is an additive model with well-developed theory (Bäuerle and Jaśkiewicz, 2018), the isomorphism allows us to immediately establish optimality results for the multiplicative Kreps–Porteus model without the need for separate analysis. The same isomorphism extends naturally to continuous-state settings, where sums are replaced by expectations and $(0,\infty)^{\mathsf{X}}$ by the exponential image of the value space used in Bäuerle and Jaśkiewicz (2018).

6.3. **Curvature and Approximation.** Implementing dynamic programs with continuous state spaces on computers requires approximation of value functions, policy functions, or both. The need for high quality approximations of these functions is amplified by the fact that, in dynamic programming, most computational algorithms use some form of iteration (such as value or Howard policy iteration), and approximation errors tend to compound across iterations, potentially leading to poor final approximations or failures of convergence (see, e.g., Farahmand et al. (2010)).

In general, functions are harder to approximate when they involve high degrees of curvature. Intuitively, isomorphic relationships can improve accuracy of approximations and stabilize iteration by transforming dynamic programs in order to reduce curvature or other complexities in target functions. We explore this idea in the current section, using the multisector real business cycle model of Long and Plosser (1983) as a test case.

In their model, the Bellman max-equation has the form

$$v(y) = \max_{c,X} \left\{u(c) + \beta \int v(f(\lambda, X))\varphi(\mathrm{d}\lambda)\right\}, \tag{21}$$

subject to

$$0 \leqslant c_i, X_{ij}, y_i \quad \text{and} \quad c_j + \sum_{i=1}^{n} X_{ij} = y_j \quad \text{for all } i, j \text{ in } \{1, \ldots, n\}. \tag{22}$$

The function f is a Borel measurable production function taking values in $\mathbb{R}^n_+$, $c = (c_j)$ is an n-vector of consumption quantities across n goods, $y = (y_j)$ is an n-vector of final outputs, $X = (x_{ij})$ is an $n \times n$ matrix of commodity inputs, and φ is a distribution over $\mathbb{R}^n_+$. (The labor-leisure decision is missing from (21) because we adopt a special case of the parameterization in Long and Plosser (1983) that allows us to assume inelastic labor supply.)

Since any strictly increasing function $F\colon \mathbb{R} \to \mathbb{R}$ commutes with maximization, we can equivalently solve the transformed Bellman equation

$$w(y) = \max_{c,X} F\left[u(c) + \beta \int (F^{-1} \circ w)(f(\lambda, X))\varphi(\mathrm{d}\lambda)\right] \tag{23}$$

for $w = Fv$ and recover the value function of the original problem via $v = F^{-1}w$, with the same maximizing choices. (We understand F as both a strictly increasing function on $\mathbb{R}$ and as the induced pointwise order isomorphism onto its image, sending v into $y \mapsto F(v(y))$.) This is advantageous because, for suitably chosen F, solving (23) significantly improves numerical accuracy.

To illustrate this point, we follow Long and Plosser (1983) in assuming that utility obeys $u(c) = \sum_{i=1}^n \theta_i \ln c_i$ and production has the multisector Cobb-Douglas form $f_i(\lambda, X) = \lambda_i \prod_{j=1}^n x_{ij}^{a_{ij}}$. Under these assumptions, an analytical solution to (21) exists, allowing us to compare approximate numerical solutions with exact solutions. The analytical solution has the form $v_0(y) = \sum_{i=1}^n \gamma_i \ln y_i + K$, where γ_i and K depend on the parameters a_{ij} and θ_i. This function has relatively high curvature near zero, with both slope and curvature approaching infinity as output converges to the origin. These features make it challenging to approximate with a high degree of accuracy. In contrast, if we consider an isomorphic transformation (23) with $F(x) = e^{mx+b}$, then the transformed value function is given by

$$(Fv_0)(y) = e^{mK+b} \prod_{i=1}^{n} y_i^{m\gamma_i}. \tag{24}$$

This function has lower curvature near zero, which facilitates approximation and leads to improved accuracy.

To demonstrate, we solve (21) and (23) for the two-sector case under several parameterizations using value function iteration where the value functions are approximated

TABLE 1. Accuracy Gain from Solving the Transformed Problem

Parameterization	$A = ([0.2, 0.7], [0.6, 0.1])$	$A = ([0.1, 0.7], [0.3, 0.1])$
	$\theta_2 = 1.0$	$\theta_2 = 1.0$
$\theta_1 = 0.5$	104.6	68.8
$\theta_1 = 0.6$	153.6	65.3
$\theta_1 = 0.7$	143.6	63.1
$\theta_1 = 0.8$	109.7	62.6
$\theta_1 = 0.9$	173.3	178.9
$\theta_1 = 1.0$	178.7	122.3

by linear interpolation on a common grid. To ensure a fair comparison, we perform a fixed number of iterations ($N = 500$) for both the original and transformed problems. The maximization operations in (21) and (23) are carried out using the Adam optimizer (Kingma and Ba, 2014) for efficient computation. In each maximization operation, we run the optimizer for 2000 iterations following a cosine decay learning rate schedule with a 500-step warmup. The initial, peak, and final learning rates of the schedule are 1×10^{-6}, 1×10^{-4}, and 1×10^{-7}, respectively, chosen through trial and error. We then compute the sup-norm errors for the two problems, $\|v - v_0\|$ and $\|F^{-1}w - v_0\|$. Table 1 shows the improvement in accuracy from solving the transformed problem, measured by $\|v - v_0\| / \|F^{-1} \circ w - v_0\|$.[10]

Although the transformed value function w is not exactly linear for the two-sector case (see (24)), we can make it approximately linear in at least one dimension by choosing F appropriately (e.g., setting $m \approx 1/\gamma_1$). Table 1 shows that this improved approximation of the value function leads to solutions that are, on average, 100 times more accurate across a wide range of parameterizations. The numerical gains come from choosing a representation better suited to linear interpolation, while conjugacy ensures that solving the transformed problem recovers the solution to the original problem.

[10]For each parameterization, we choose the parameters of the exponential transformation $F(x) = e^{mx+b}$ so that Fv_0 is approximately linear in at least one dimension. In particular, in the second column, $m = [0.35, 0.3, 0.3, 0.25, 0.25, 0.25]$ and $b = [75, 60, 65, 60, 60, 75]$; in the third column, $m = [0.9, 0.8, 0.7, 0.7, 0.6, 0.5]$ and $b = [90, 85, 80, 85, 80, 75]$. In all cases, the state space is $[1 \times 10^{-7}, 20]^2$ with 500 grid points along each dimension. The other parameter values are listed in the table.

7. Conclusion

We studied isomorphic and anti-isomorphic relationships between dynamic programs and showed how optimality properties transmit across these relationships. We applied these ideas to Epstein–Zin preferences with time preference shocks, to the connection between multiplicative Kreps–Porteus and risk-sensitive preferences, and to improving numerical accuracy in multisector real business cycle models.

Our research focused on discrete time dynamic programs, which we paired with discrete time concepts of topological and order conjugacy. Continuous time dynamic programs are also important and the concept of topological conjugacy has a well-defined analogy in continuous time dynamics. This suggests that many of our ideas can be carried over to continuous time systems. We leave this task for future research.

Another direction is to use parameterized conjugacies to study monotone comparative statics, particularly in recursive utility models (see Remark 5.2 and, e.g., Flynn et al., 2023). This raises questions about which transformations yield the required cardinal conditions and how parameter-dependent transformations affect comparisons of value functions.

Appendix A. Order-Theoretic Preliminaries

This appendix collects the order-theoretic definitions, dynamical systems background, and auxiliary results that underpin the main text.

A.1. **Posets.** Let $V = (V, \preceq)$ be a partially ordered set, also called a poset. When $A \subset V$, the symbol $\bigvee A$ refers to the supremum of A in V and $\bigwedge A$ is the infimum (see, e.g., Davey and Priestley (2002)). V is called *bounded* if V has a least and greatest element. A sequence (v_n) in $V = (V, \preceq)$ is called *increasing* if $v_n \preceq v_{n+1}$ for all $n \in \mathbb{N}$. If (v_n) is increasing and $\bigvee_n v_n = v$ for some $v \in V$, then we write $v_n \uparrow v$. The set V is called *countably chain complete* if V is bounded and every increasing sequence in V has a supremum in V. A self-map S on V is called *order continuous* on V if $Sv_n \uparrow Sv$ whenever $v_n \uparrow v$.

A.2. **Dynamical Systems.** A *dynamical system* is a pair (V, S) where V is a set and S is a self-map on V. Given $v \in V$, the sequence $(S^n v)_{n \geqslant 1}$ is called the *trajectory* of v under S. When V has a topology, we can study the convergence of trajectories. In particular, in this setting, a system (V, S) is called *globally stable* when S has a unique fixed point $\bar{v}$ in V and $\lim_{k \to \infty} S^k v = \bar{v}$ for all $v \in V$.

Two dynamical systems (V, S) and $(\hat{V}, \hat{S})$ are called *conjugate* under F if there exists a bijection F from V to $\hat{V}$ such that $F \circ S = \hat{S} \circ F$ on V. In this setting, if v is a fixed point of S, then $\hat{S}Fv = FSv = Fv$, so Fv is a fixed point of $\hat{S}$. With slightly more effort, one can show the following.

Lemma A.1. *If (V, S) and $(\hat{V}, \hat{S})$ are conjugate under F and S has a unique fixed point $v \in V$, then Fv is the unique fixed point of $\hat{S}$ in $\hat{V}$.*

For more details see, for example, Sternberg (2014) or Layek (2015).

Dynamical systems (V, S) and $(\hat{V}, \hat{S})$ are called *topologically conjugate* under F if both V and $\hat{V}$ have topologies, the pairs (V, S) and $(\hat{V}, \hat{S})$ are conjugate under F and, in addition, F is a homeomorphism (i.e., continuous bijection with continuous inverse) from V to $\hat{V}$. In this setting we can extend Lemma A.1 as follows.

Lemma A.2. *If (V, S) and $(\hat{V}, \hat{S})$ are topologically conjugate under F, then (V, S) is globally stable if and only if $(\hat{V}, \hat{S})$ is globally stable.*

Lemma A.2 is standard and routinely used in dynamical systems theory. (There are similar results concerning local stability for topologically conjugate systems, as well as additional results concerning other kinds of asymptotic behavior. See, for example, Sternberg (2014) or Layek (2015).)

A.3. **Order Stability.** Let (V, S) be a dynamical system where V is partially ordered by $\preceq$ and S has a unique fixed point $\bar{v}$ in V. In this environment, Sargent and Stachurski (2025) call (V, S)

(i) *upward stable* if $v \in V$ and $v \preceq Sv$ implies $v \preceq \bar{v}$,
(ii) *downward stable* if $v \in V$ and $Sv \preceq v$ implies $\bar{v} \preceq v$, and
(iii) *order stable* if upward and downward stability both hold.

In this paper, we discuss both maximization and minimization. To link these ideas in a poset environment, we use the notion of order duals. In particular, given poset $V = (V, \preceq)$, let $(V, \preceq^\partial)$ be the order dual, so that, for $u, v \in V$, we have $u \preceq^\partial v$ if and only if $v \preceq u$. For convenience, we sometimes denote $(V, \preceq^\partial)$ by V^∂.

Lemma A.3. *(V, S) is order stable if and only if (V^∂, S) is order stable.*

Proof of Lemma A.3. Let (V, S) be order stable. By definition, S has a unique fixed point $\bar{v} \in V$. We claim that (V^∂, S) is upward and downward stable. Regarding upward stability, suppose $v \in V$ and $v \preceq^\partial Sv$. Then $Sv \preceq v$ and hence $\bar{v} \preceq v$, by downward

stability of (V, S). But then $v \preceq^{\partial} \bar{v}$, so (V^{∂}, S) is upward stable. The proof of downward stability is similar. Hence order stability of (V, S) is sufficient for order stability of (V^{∂}, S). Necessity follows from sufficiency, since the dual of (V^{∂}, S) is (V, S). □

A.4. **Order Conjugacy.** Our next step is to replace the concept of topological conjugacy from Section A.2 with a parallel notion for dynamical systems on posets. To this end, we recall that a bijective map $F\colon V \to \hat{V}$ is called an *order isomorphism* if both F and its inverse F^{-1} are order preserving. (We call $F\colon V \to \hat{V}$ an *order anti-isomorphism* if both F and F^{-1} are order reversing.) Let (V, S) and $(\hat{V}, \hat{S})$ be two dynamical systems where V and $\hat{V}$ are partially ordered. We call (V, S) and $(\hat{V}, \hat{S})$ *order conjugate* under F when (V, S) and $(\hat{V}, \hat{S})$ are conjugate under F and, in addition, F is an order isomorphism. It is easy to verify that order conjugacy is an equivalence relation on the set of dynamical systems over partially ordered sets.

Lemma A.4. *If (V, S) and $(\hat{V}, \hat{S})$ are order conjugate under F, then (V, S) is order stable if and only if $(\hat{V}, \hat{S})$ is order stable.*

Proof of Lemma A.4. Let (V, S) and $(\hat{V}, \hat{S})$ be order conjugate under F, and suppose that (V, S) is order stable. By Lemma A.1, the map $\hat{S}$ has a unique fixed point in $\hat{V}$. Let $\hat{w}$ be an element of $\hat{V}$ satisfying $\hat{S}\hat{w} \preceq \hat{w}$. Let v and $\hat{v} := Fv$ be the fixed points of S and $\hat{S}$, respectively. Then $F^{-1}\hat{S}\hat{w} \preceq F^{-1}\hat{w}$ and hence $SF^{-1}\hat{w} \preceq F^{-1}\hat{w}$. But then $v \preceq F^{-1}\hat{w}$, by downward stability of (V, S). Applying F gives $\hat{v} \preceq \hat{w}$. Hence $(\hat{V}, \hat{S})$ is downward stable. Similarly, if $\hat{w}$ is an element of $\hat{V}$ satisfying $\hat{w} \preceq \hat{S}\hat{w}$, then $F^{-1}\hat{w} \preceq F^{-1}\hat{S}\hat{w} = SF^{-1}\hat{w}$. By upward stability of (V, S), we have $F^{-1}\hat{w} \preceq v$. Applying F gives $\hat{w} \preceq \hat{v}$, so $(\hat{V}, \hat{S})$ is upward stable. Together, these results show that $(\hat{V}, \hat{S})$ is order stable. The converse implication follows from symmetry. □

Appendix B. Additional Examples

The following examples extend the ADP framework of Section 3 to Q-factor formulations from the reinforcement learning literature. All primitives are as in Example 3.1.

Example B.1 (Q-factors)**.** The Q-learning literature studies the Q-factor Bellman equation (see, e.g., Kochenderfer et al. (2022), Section 17.2), which is given by

$$f(x, a) = r(x, a) + \beta \sum_{x'} \max_{a' \in \Gamma(x')} f(x', a') P(x, a, x') \qquad ((x, a) \in \mathsf{G}). \tag{25}$$

Here $f \in \mathbb{R}^{\mathsf{G}}$ is called the *Q-factor*. The policy operators over Q-factors take the form

$$(Q_\sigma f)(x,a) = r(x,a) + \beta \sum_{x'} f(x', \sigma(x'))P(x,a,x'), \tag{26}$$

where $f \in \mathbb{R}^{\mathsf{G}}$ and $\sigma \in \Sigma$. If we pair $\mathbb{R}^{\mathsf{G}}$ (the set of all real-valued functions on G) with the pointwise partial order $\leqslant$ and Q_σ as in (26), then $(\mathbb{R}^{\mathsf{G}}, \{Q_\sigma\})$ is an ADP.

Example B.2 (Risk-sensitive Q-factors)**.** It has become popular in reinforcement learning and related fields to extend the Q-factor approach from Example B.1 to risk-sensitive decision processes (see, e.g., Fei et al. (2021)). The corresponding Q-factor Bellman equation is given by

$$f(x,a) = r(x,a) + \frac{\beta}{\theta} \ln \left\{ \sum_{x'} \exp \left[\theta \max_{a' \in \Gamma(x')} f(x', a') \right] P(x,a,x') \right\} \qquad ((x,a) \in \mathsf{G}). \tag{27}$$

The policy operators over risk-sensitive Q-factors take the form

$$(Q^\theta_\sigma f)(x,a) = r(x,a) + \frac{\beta}{\theta} \ln \left[\sum_{x'} \exp \left[\theta f(x', \sigma(x')) \right] P(x,a,x') \right] \tag{28}$$

where $f \in \mathbb{R}^{\mathsf{G}}$ and $\sigma \in \Sigma$. The pair $(\mathbb{R}^{\mathsf{G}}, \{Q^\theta_\sigma\})$ is an ADP.

Example B.3 (Exponential risk-sensitive Q-factors)**.** Fei et al. (2021) work with an "exponential" risk-sensitive Q-factor Bellman equation where the corresponding policy operators have the form

$$(M_\sigma h)(x,a) = \exp \left\{ \theta r(x,a) + \beta \ln \left[\sum_{x'} h(x', \sigma(x')) P(x,a,x') \right] \right\}$$

Here $h \in (0,\infty)^{\mathsf{G}}$ and M_σ maps $(0,\infty)^{\mathsf{G}}$ into itself. All primitives are as in the risk-sensitive Q-factor ADP $\mathscr{A} := (\mathbb{R}^{\mathsf{G}}, \{Q_\sigma\})$ in Example B.2. Since each M_σ is order preserving (under the usual pointwise order), the pair $\hat{\mathscr{A}} := ((0,\infty)^{\mathsf{G}}, \{M_\sigma\})$ is also an ADP. If we take F to be the bijection from $\mathbb{R}^{\mathsf{G}}$ to $(0,\infty)^{\mathsf{G}}$ defined by $(Fh)(x,a) = \exp(\theta h(x,a))$, then, for Q_σ defined in (28), $h \in \mathbb{R}^{\mathsf{G}}$ and $(x,a) \in \mathsf{G}$,

$$(FQ_\sigma h)(x,a) = \exp \left\{ \theta r(x,a) + \beta \ln \left[\sum_{x'} \exp \left[\theta h(x', \sigma(x')) \right] P(x,a,x') \right] \right\}.$$

This is equal to $(M_\sigma F h)(x,a)$, which shows that $F \circ Q_\sigma = M_\sigma \circ F$ on $\mathbb{R}^{\mathsf{G}}$. As a consequence, $\mathscr{A}$ and $\hat{\mathscr{A}}$ are isomorphic when $\theta > 0$ and anti-isomorphic when $\theta < 0$.

Appendix C. Remaining Proofs

C.1. **Proof of Optimality Results.**

Proof of Theorem 4.1. Since $\mathscr{A}$ is max-stable, T_Δ has a fixed point $\bar{v} \in V_\Delta^G$. Since $\bar{v} \in V_\Delta^G$, we can find a $\bar{v}$-max-greedy policy σ, so that $T_\sigma \bar{v} = T_\Delta \bar{v} = \bar{v}$. But T_σ has a unique fixed point in V, equal to v_σ, so $\bar{v} = v_\sigma$. Moreover, if τ is any policy, then $T_\tau \bar{v} \preceq T_\Delta \bar{v} = \bar{v}$ and hence, by downward stability, $v_\tau \preceq \bar{v}$. These facts imply that $\bar{v}$ is the greatest element of V_Σ. Hence v_Δ exists and $v_\Delta = \bar{v} \in V_\Delta^G$, which proves (i). The same argument applies to any fixed point of T_Δ in V_Δ^G, so every such fixed point is equal to v_Δ. This proves (ii).

For (iii), parts (i)–(ii) give $v_\Delta \in V_\Delta^G$ and $T_\Delta v_\Delta = v_\Delta$. Now recall that σ is max-optimal if and only if $v_\sigma = v_\Delta$. Since v_σ is the unique fixed point of T_σ, this is equivalent to $T_\sigma v_\Delta = v_\Delta$. Since $T_\Delta v_\Delta = v_\Delta$, the last statement is equivalent to $T_\sigma v_\Delta = T_\Delta v_\Delta$, which is, in turn, equivalent to the statement that σ is v_Δ-max-greedy.

Part (iv) follows from part (iii) and existence of a v_Δ-max-greedy policy, which holds because $v_\Delta \in V_\Delta^G$. □

C.2. **Min-Optimality.**

In the abstract setting of Section 4, minimization results are readily recovered from maximization results by order duality. Min-optimal policies and the min-value function v_∇ are defined in Section 4.2.

Below, if $\mathscr{A} := (V, \{T_\sigma\})$ is an ADP then its *dual* $\mathscr{A}^\partial$ is the ADP $(V^\partial, \{T_\sigma\})$ where the partial order $\preceq$ on V is replaced with its dual $\preceq^\partial$. In this setting, we let T_Δ^∂ be the Bellman max-operator for $\mathscr{A}^\partial$, v_Δ^∂ be the max-value function for $\mathscr{A}^\partial$, $V_\Delta^{G,\partial}$ be the set of $v \in V$ that have a v-max-greedy policy for $\mathscr{A}^\partial$, and so on. Since $\mathscr{A}$ and $\mathscr{A}^\partial$ have the same policy operators, $\mathscr{A}$ is well-posed if and only if $\mathscr{A}^\partial$ is well-posed, and the two ADPs have the same set of σ-value functions V_Σ.

Lemma C.1. *Let $\mathscr{A}$ be an ADP. Then*

(i) *$\mathscr{A}$ is order stable if and only if $\mathscr{A}^\partial$ is order stable,*
(ii) *for $v \in V$, a policy σ is v-min-greedy for $\mathscr{A}$ if and only if σ is v-max-greedy for $\mathscr{A}^\partial$,*
(iii) *$V_\nabla^G = V_\Delta^{G,\partial}$ and $T_\nabla = T_\Delta^\partial$ on V_∇^G,*
(iv) *if $\mathscr{A}$ is well-posed, then a policy σ is min-optimal for $\mathscr{A}$ if and only if σ is max-optimal for $\mathscr{A}^\partial$, and v_∇ exists if and only if v_Δ^∂ exists, in which case $v_\nabla = v_\Delta^\partial$, and*

(v) *$\mathscr{A}$ is min-stable if and only if $\mathscr{A}^\partial$ is max-stable.*

Proof. Part (i) follows from Lemma A.3, applied to each T_σ. For (ii), fix $v \in V$ and note that σ is v-min-greedy for $\mathscr{A}$ if and only if $T_\sigma v \preceq T_\tau v$ for all $\tau \in \Sigma$, which is equivalent to $T_\tau v \preceq^\partial T_\sigma v$ for all $\tau \in \Sigma$; that is, to σ being v-max-greedy for $\mathscr{A}^\partial$. Part (iii) follows from (ii) and the fact that, for any such σ, we have $T_\nabla v = T_\sigma v = T_\Delta^\partial v$. Part (iv) holds because a least element of V_Σ under $\preceq$ is a greatest element of V_Σ under $\preceq^\partial$, and vice versa. Part (v) follows from (i) and (iii), since (iii) implies that the fixed points of T_∇ in V_∇^G are exactly the fixed points of T_Δ^∂ in $V_\Delta^{G,\partial}$. Each of these statements remains true when the roles of $\mathscr{A}$ and $\mathscr{A}^\partial$ are reversed, since $(\mathscr{A}^\partial)^\partial = \mathscr{A}$. □

Proof of Theorem 4.2. Let $\mathscr{A}$ be min-stable. By Lemma C.1, the dual $\mathscr{A}^\partial$ is max-stable, $V_\Delta^{G,\partial} = V_\nabla^G$, and $T_\Delta^\partial = T_\nabla$ on V_∇^G. In particular, the fixed points of T_Δ^∂ in $V_\Delta^{G,\partial}$ coincide with those of T_∇ in V_∇^G. Applying Theorem 4.1 to $\mathscr{A}^\partial$, we see that v_Δ^∂ exists, lies in V_∇^G, and is the unique fixed point of T_∇ in V_∇^G. By Lemma C.1, v_∇ exists and equals v_Δ^∂, which proves (i)–(ii). Theorem 4.1 also tells us that a policy is max-optimal for $\mathscr{A}^\partial$ if and only if it is v_Δ^∂-max-greedy for $\mathscr{A}^\partial$, and that at least one such policy exists. Since $v_\nabla = v_\Delta^\partial$, parts (ii) and (iv) of Lemma C.1 translate these statements into (iii)–(iv). □

C.3. **Sufficient Conditions.** Recall from Example 2.1 that a partially ordered space is a triple $(V, \preceq, \rho)$, where (V, ρ) is a metric space and $\preceq$ is closed under limits. We call ρ *sup-nonexpansive* if, for any index set A and any families $(v_\alpha)_{\alpha\in A}$ and $(w_\alpha)_{\alpha\in A}$ in V,

$$\rho\left(\bigvee_\alpha v_\alpha, \bigvee_\alpha w_\alpha\right) \leqslant \sup_\alpha \rho(v_\alpha, w_\alpha)$$

whenever both suprema on the left-hand side exist and are attained (i.e., $\bigvee_\alpha v_\alpha = v_{\alpha'}$ and $\bigvee_\alpha w_\alpha = w_{\alpha''}$ for some $\alpha', \alpha'' \in A$). This is weaker than the notion of sup-nonexpansiveness in Peng et al. (2026), which requires the inequality whenever both suprema exist. For example, if V is a set of bounded real-valued functions with the pointwise order and the supremum norm, then ρ is sup-nonexpansive, since attained suprema are pointwise maxima.

Proof of Proposition 4.3. By Banach's fixed point theorem, each (V, T_σ) is globally stable. Since each T_σ is order-preserving, $\mathscr{A}$ is order stable, by Example 2.1. Now fix $v, w \in V_0$. Since $V_0 \subseteq V_\Delta^G$, the suprema defining $T_\Delta v$ and $T_\Delta w$ exist and are

attained (by $T_\sigma v$ and $T_\tau w$, where σ is v-max-greedy and τ is w-max-greedy), so sup-nonexpansiveness gives

$$\rho(T_\Delta v, T_\Delta w) \leqslant \sup_\sigma \rho(T_\sigma v, T_\sigma w) \leqslant \beta\, \rho(v, w).$$

Hence T_Δ is a contraction on V_0. Since V_0 is a closed subset of the complete space V, the operator T_Δ has a fixed point in $V_0 \subseteq V_\Delta^G$. Thus $\mathscr{A}$ is max-stable, and $v_\Delta \in V_0$ by Theorem 4.1(ii). □

To prove Proposition 4.4, we use policy iteration. Let $\mathscr{A} = (V, \{T_\sigma\})$ be an order stable ADP that is max-regular. We define a map H_Δ from V to V_Σ via $H_\Delta v = v_\sigma$ where σ is v-max-greedy. Iterating with H_Δ is an abstraction of Howard policy iteration.[11] We call H_Δ the *Howard max-operator* generated by the ADP. As above, we denote the greatest element of V_Σ by v_Δ whenever it exists.

Lemma C.2. *If $\mathscr{A}$ is order stable and max-regular, then the following statements are true:*

(i) H_Δ *obeys* $v_\sigma \preceq H_\Delta v_\sigma$ *for all* $\sigma \in \Sigma$.
(ii) *If* $\sigma \in \Sigma$ *and* $T_\Delta v_\sigma = v_\sigma$, *then* v_Δ *exists in* V *and* $v_\sigma = v_\Delta$.
(iii) *If* $v \in V$ *and* $H_\Delta v = v$, *then* $v = v_\Delta$ *and* $T_\Delta v_\Delta = v_\Delta$.
(iv) *If* $v \in V$ *and* Σ *is finite, then* v_Δ *exists,* $H_\Delta v_\Delta = v_\Delta$ *and* $(H_\Delta^k v)_{k \geqslant 0}$ *converges to* v_Δ *in finitely many steps.*

Proof. As for (i), fix $\sigma \in \Sigma$ and let τ be such that $H_\Delta v_\sigma = v_\tau$. Since τ is v_σ-greedy, we have $v_\sigma = T_\sigma v_\sigma \preceq T_\Delta v_\sigma = T_\tau v_\sigma$. Upward stability of T_τ gives $v_\sigma \preceq v_\tau = H_\Delta v_\sigma$.

As for (ii), suppose $\sigma \in \Sigma$ and $T_\Delta v_\sigma = v_\sigma$. Fix $\tau \in \Sigma$ and note that $v_\sigma = T_\Delta v_\sigma \succeq T_\tau v_\sigma$. Downward stability of T_τ implies $v_\sigma \succeq v_\tau$. Since $\tau \in \Sigma$ was arbitrary, $v_\sigma = v_\Delta$.

As for (iii), fix $v \in V$ with $H_\Delta v = v$ and let σ be such that $H_\Delta v = v_\sigma$. Then $v_\sigma = v$, and, since σ is v-max-greedy, $T_\sigma v = T_\Delta v$. But then $T_\sigma v_\sigma = T_\Delta v_\sigma$, and, since $v_\sigma = T_\sigma v_\sigma$, we have $v_\sigma = T_\Delta v_\sigma$. Part (ii) now implies $v = v_\sigma = v_\Delta$. This proves the first claim. Regarding the second, substituting $v_\sigma = v_\Delta$ into $v_\sigma = T_\Delta v_\sigma$ yields $v_\Delta = T_\Delta v_\Delta$.

For (iv), it suffices to show that $H_\Delta v_\Delta = v_\Delta$ and there exists a $K \in \mathbb{N}$ such that $H_\Delta^K v = v_\Delta$. To this end, let $v_k = H_\Delta^k v$ and note that $v_k \in V_\Sigma$ for all $k \geqslant 1$. Part (i)

[11] For H_Δ to be well-defined, we must always select the same v-greedy policy when the operator is applied to v. We can use the axiom of choice to assign to each v a designated v-greedy policy, although, in applications, a simple rule usually suffices. For example, if Σ is finite, we can enumerate the policy set Σ and choose the first v-greedy policy.

implies that $v_{k+1} \succeq v_k$ for all $k \in \mathbb{N}$. Since the sequence (v_k) is contained in the finite set V_Σ, it must be that $v_{K+1} = v_K$ for some $K \in \mathbb{N}$ (since otherwise V_Σ contains an infinite sequence of distinct points). But then $H_\Delta v_K = v_{K+1} = v_K$, so v_K is a fixed point of H_Δ. Part (iii) now implies that $v_K = v_\Delta$. □

Proof of Proposition 4.4. If $\mathscr{A}$ is order stable and max-regular and Σ is finite, then, by (iii)–(iv) of Lemma C.2, the point v_Δ is a fixed point of T_Δ, so $\mathscr{A}$ is max-stable. This proves (i) of Proposition 4.4. Part (ii) follows from (i) applied to the dual ADP $\mathscr{A}^\partial$ (see Lemma C.1). □

Notice also that, under the hypotheses of Proposition 4.4(i), if $((v_k), (\sigma_k))$ is any HPI sequence from $v \in V$ (Section 5.4), then $v_k = v_{\sigma_{k-1}}$ for $k \geqslant 1$, and the argument in part (i) of Lemma C.2 shows that $v_k \preceq v_{k+1}$ in the finite set V_Σ for all $k \geqslant 1$. Hence $v_K = v_{K+1} = v_{\sigma_K}$ for some $K \geqslant 1$. Since σ_K is v_K-max-greedy, $T_\Delta v_K = T_{\sigma_K} v_K = v_K$ and part (ii) of Lemma C.2 gives $v_K = v_\Delta$, so every HPI sequence converges in finitely many steps.

C.4. **Proofs of Isomorphism Results.**

Proof of Theorem 5.1. Claims (i)–(ii) follow directly from Lemmas A.1 and A.4, respectively. Regarding (iii), we first show that $F(V_\Delta^G) = \hat{V}_\Delta^G$ and $F \circ T_\Delta = \hat{T}_\Delta \circ F$ on V_Δ^G. To this end, fix $v \in V_\Delta^G$ and let σ be v-max-greedy. Since F is order-preserving, the conjugacy relation (8) gives

$$\hat{T}_\tau Fv = F T_\tau v \preceq F T_\sigma v = \hat{T}_\sigma Fv \quad \text{for all } \tau \in \Sigma.$$

Hence σ is Fv-max-greedy for $\hat{\mathscr{A}}$, so $Fv \in \hat{V}_\Delta^G$, and $\hat{T}_\Delta Fv = \hat{T}_\sigma Fv = F T_\sigma v = F T_\Delta v$. This proves $F(V_\Delta^G) \subseteq \hat{V}_\Delta^G$, as well as the commutation relation. Since $\hat{\mathscr{A}}$ and $\mathscr{A}$ are isomorphic under F^{-1}, the same argument gives $F^{-1}(\hat{V}_\Delta^G) \subseteq V_\Delta^G$. Hence $F(V_\Delta^G) = \hat{V}_\Delta^G$.

Now suppose that $\mathscr{A}$ is max-stable. By (ii), $\hat{\mathscr{A}}$ is order stable. Moreover, if $\bar{v} \in V_\Delta^G$ is a fixed point of T_Δ, then $F\bar{v} \in \hat{V}_\Delta^G$ and $\hat{T}_\Delta F\bar{v} = F T_\Delta \bar{v} = F\bar{v}$. Hence $\hat{\mathscr{A}}$ is max-stable. The reverse implication follows by symmetry, since $\hat{\mathscr{A}}$ and $\mathscr{A}$ are isomorphic under F^{-1}.

Finally, by (i) and Lemma A.1, we have $\hat{v}_\sigma = Fv_\sigma$ for all $\sigma \in \Sigma$. Since F is an order isomorphism, v_σ is a greatest element of V_Σ if and only if Fv_σ is a greatest element of $\hat{V}_\Sigma = F(V_\Sigma)$. Hence $\hat{v}_\Delta = Fv_\Delta$, and $\mathscr{A}$ and $\hat{\mathscr{A}}$ have the same max-optimal policies. We have now confirmed all the claims in (iii).

The proof of (iv) is identical after replacing max with min. (Alternatively, (iv) can be derived from (iii) and duality.) □

Proof of Theorem 5.2. Let $\mathscr{A}$ and $\hat{\mathscr{A}}$ be anti-isomorphic, so that $\mathscr{A}$ is isomorphic to $\hat{\mathscr{A}}^\partial$. If $\mathscr{A}$ is well-posed, then, by Theorem 5.1, $\hat{\mathscr{A}}^\partial$ is well-posed, so $\hat{T}_\sigma$ has a unique fixed point in $\hat{V}$ for all $\sigma \in \Sigma$. This implies that $\hat{\mathscr{A}}$ is likewise well-posed, completing the proof of (i). Similarly, if $\mathscr{A}$ is order stable, then, by Theorem 5.1, $\hat{\mathscr{A}}^\partial$ is order stable, in which case $\hat{\mathscr{A}}$ is order stable, by Lemma A.3. This proves (ii).

Regarding (iii), since $\mathscr{A}$ and $\hat{\mathscr{A}}^\partial$ are isomorphic under F, Theorem 5.1 implies that $\mathscr{A}$ is max-stable if and only if $\hat{\mathscr{A}}^\partial$ is max-stable. Since the dual of $\hat{\mathscr{A}}^\partial$ is $\hat{\mathscr{A}}$, Lemma C.1 implies that $\hat{\mathscr{A}}^\partial$ is max-stable if and only if $\hat{\mathscr{A}}$ is min-stable. Combining these equivalences proves the first claim in (iii).

Now suppose $\mathscr{A}$ is max-stable. By Theorem 5.1 and its proof, F maps $V_\triangle^G$ onto $\hat{V}_\triangle^{G,\partial}$, the set of $\hat{\upsilon} \in \hat{V}$ that have a $\hat{\upsilon}$-max-greedy policy for $\hat{\mathscr{A}}^\partial$, and we have $F \circ T_\triangle = \hat{T}_\triangle^\partial \circ F$ on $V_\triangle^G$ and $\hat{\upsilon}_\triangle^\partial = F\upsilon_\triangle$. By Lemma C.1, we have $\hat{V}_\triangle^{G,\partial} = \hat{V}_\triangledown^G$, $\hat{T}_\triangle^\partial = \hat{T}_\triangledown$ on $\hat{V}_\triangledown^G$, and $\hat{\upsilon}_\triangle^\partial = \hat{\upsilon}_\triangledown$. Hence $F \circ T_\triangle = \hat{T}_\triangledown \circ F$ on $V_\triangle^G$ and $\hat{\upsilon}_\triangledown = F\upsilon_\triangle$, which proves (11).

Finally, Theorem 5.1 tells us that $\mathscr{A}$ and $\hat{\mathscr{A}}^\partial$ have the same max-optimal policies. Applying Lemma C.1, we see that the max-optimal policies of $\mathscr{A}$ are the same as the min-optimal policies of $\hat{\mathscr{A}}$. □

Proof of Proposition 5.3. Let $\mathscr{A}$ and $\hat{\mathscr{A}}$ be isomorphic under F and let $\mathscr{A}$ be max-stable. By Theorem 5.1 and its proof, $\hat{\mathscr{A}}$ is max-stable, $\hat{\upsilon}_\triangle = F\upsilon_\triangle$, $F(V_\triangle^G) = \hat{V}_\triangle^G$, $F \circ T_\triangle = \hat{T}_\triangle \circ F$ on $V_\triangle^G$, and, for $\upsilon \in V_\triangle^G$, a policy σ is υ-max-greedy for $\mathscr{A}$ if and only if it is $F\upsilon$-max-greedy for $\hat{\mathscr{A}}$. Moreover, $F T_\sigma^m = \hat{T}_\sigma^m F$ for all $\sigma \in \Sigma$ and $m \in \mathbb{N}$ by (8), and $\hat{\upsilon}_\sigma = F\upsilon_\sigma$ for all $\sigma \in \Sigma$ by Lemma A.1. Claims (i)–(iii) follow immediately from these identities and the fact that F is a bijection from $V_\triangle^G$ onto $\hat{V}_\triangle^G$.

Since F is an order isomorphism, (υ_n) is increasing if and only if $(F\upsilon_n)$ is increasing, and $\bigvee_n \upsilon_n = \upsilon_\triangle$ if and only if $\bigvee_n F\upsilon_n = F\upsilon_\triangle = \hat{\upsilon}_\triangle$. Hence $\upsilon_n \uparrow \upsilon_\triangle$ if and only if $F\upsilon_n \uparrow \hat{\upsilon}_\triangle$. Similarly, $\upsilon_n = \upsilon_\triangle$ if and only if $F\upsilon_n = \hat{\upsilon}_\triangle$. □

For min-stable ADPs, min-VFI, min-OPI and min-HPI sequences from $\upsilon \in V_\triangledown^G$ are defined analogously, with $T_\triangledown$, $V_\triangledown^G$ and min-greedy policies in place of $T_\triangle$, $V_\triangle^G$ and max-greedy policies. Given a sequence (υ_n) in V and $\upsilon \in V$, we write $\upsilon_n \downarrow \upsilon$ if (υ_n) is decreasing and $\bigwedge_n \upsilon_n = \upsilon$.

Proposition C.3. *Let $\mathscr{A}$ and $\hat{\mathscr{A}}$ be anti-isomorphic under F, and let $\mathscr{A}$ be max-stable. For all $\upsilon \in V_\triangle^G$ and $m \in \mathbb{N}$, statements (i)–(iii) of Proposition 5.3 hold with VFI, OPI and HPI sequences for $\hat{\mathscr{A}}$ replaced by min-VFI, min-OPI and min-HPI sequences. Moreover, $\upsilon_n \uparrow \upsilon_\triangle$ if and only if $F\upsilon_n \downarrow \hat{\upsilon}_\triangledown$, and an HPI sequence for $\mathscr{A}$ converges in finitely many steps if and only if the corresponding min-HPI sequence for $\hat{\mathscr{A}}$ does.*

Proof of Proposition C.3. Since $\mathscr{A}$ and $\hat{\mathscr{A}}^{\partial}$ are isomorphic under F, the result follows by applying Proposition 5.3 to $\mathscr{A}$ and $\hat{\mathscr{A}}^{\partial}$ and using Lemma C.1, together with the fact that a sequence is increasing in $\hat{V}^{\partial}$ with supremum $\hat{w}$ if and only if it is decreasing in $\hat{V}$ with infimum $\hat{w}$. □

C.5. **Proofs of Epstein–Zin Optimality Results.**

Proof of Theorem 6.2. By Lemma 6.1, $Fv = v^{\gamma}$ bijects V onto itself and satisfies $FT_\sigma = \hat{T}_\sigma F$. It also bijects V_0 onto itself (since F and F^{-1} preserve continuity). We first prove that $\mathscr{A}$ is well-posed if and only if $\Lambda < 1$.

Necessity. Let $\mathbb{1}$ be the constant-one function and define the positive linear operator K on $C(\mathsf{Z})$ by

$$(Kf)(z) = \beta(z)^{\theta} \int_{\mathsf{Z}} f(z')Q(z, \mathrm{d}z').$$

Then $K^n\mathbb{1}(z) = \mathbb{E}\prod_{t=0}^{n-1}\beta(Z_t(z))^{\theta}$. Set $c_n = \sup_z (K^n\mathbb{1}(z))^{1/\theta}$. We have

$$K^n\mathbb{1} \le c_n^{\theta}\mathbb{1} \quad (\theta > 0), \qquad K^n\mathbb{1} \ge c_n^{\theta}\mathbb{1} \quad (\theta < 0). \tag{29}$$

Applying K^m and then taking the power $1/\theta$ and the supremum gives $c_{m+n} \le c_m c_n$ in either case. By Fekete's lemma, $\Lambda = \lim_n c_n^{1/n} = \inf_{n\ge 1} c_n^{1/n}$ exists and lies between $\min\beta > 0$ and $\max\beta < \infty$.

Set $A_\sigma = \beta^{\theta} P_\sigma$ on $B_b(\mathsf{X})$. If $\hat{T}_\sigma w = w \in V$, rearranging the fixed-point equation gives

$$\frac{A_\sigma w}{w} = \left(1 - \frac{r_\sigma^{\alpha}}{w^{1/\theta}}\right)^{\theta}.$$

Since $w^{1/\theta} = r_\sigma^{\alpha} + (A_\sigma w)^{1/\theta}$ with both terms positive, the ratio $r_\sigma^{\alpha}/w^{1/\theta}$ is strictly below one; it is also bounded away from zero since r_σ^{α} is bounded away from zero and $w^{1/\theta}$ is bounded above. Choose $\varepsilon \in (0,1)$ with $r_\sigma^{\alpha}/w^{1/\theta} \ge \varepsilon$ and set $\kappa = (1-\varepsilon)^{\theta}$. Then $\kappa^{1/\theta} = 1 - \varepsilon < 1$ and

$$A_\sigma w \le \kappa w \quad (\theta > 0), \qquad A_\sigma w \ge \kappa w \quad (\theta < 0).$$

By (15), $A_\sigma^n\mathbb{1}(y,z) = K^n\mathbb{1}(z)$. Writing $m = \inf w > 0$ and $M = \sup w < \infty$, positivity and linearity allow us to iterate these inequalities:

$$mK^n\mathbb{1} \le A_\sigma^n w \le M\kappa^n\mathbb{1} \quad (\theta > 0),$$
$$MK^n\mathbb{1} \ge A_\sigma^n w \ge m\kappa^n\mathbb{1} \quad (\theta < 0).$$

Taking the $1/\theta$ power in these bounds and then taking the supremum over z gives $c_n \le (M/m)^{1/|\theta|}\kappa^{n/\theta}$. Taking nth roots and letting $n \to \infty$ yields $\Lambda \le \kappa^{1/\theta} < 1$. If $\Lambda \ge 1$, no $\hat{T}_\sigma$ has a fixed point, and $FT_\sigma = \hat{T}_\sigma F$ gives the same conclusion for every T_σ.

Sufficiency. Assume $\Lambda < 1$. The order-preserving map $\hat{T}_\sigma \colon V \to V$ is concave for $\theta < 0$ or $\theta \geq 1$ and convex for $0 < \theta \leq 1$ (see, e.g., Proposition 3.3 of Stachurski et al., 2025). We shall apply Theorem 2.2 and Theorem 2.4 of Stachurski et al. (2025), which give global stability of $\hat{T}_\sigma$ in V, the interior of the positive cone of $B_b(\mathsf{X})$.

Choose N with $c_N < 1$ and define

$$h = \sum_{j=0}^{N-1} K^j \mathbb{1}.$$

Since K preserves $C(\mathsf{Z})$ and is positive, $h \in C(\mathsf{Z})$ and $h \geq \mathbb{1}$. Regard functions on Z as functions on X independent of y. Set $q = (Kh/h)^{1/\theta}$. For $t > 0$,

$$\frac{\hat{T}_\sigma(th)}{th} = \left\{ \frac{r_\sigma^\alpha}{(th)^{1/\theta}} + \left(\frac{Kh}{h} \right)^{1/\theta} \right\}^\theta = \left\{ \frac{r_\sigma^\alpha}{(th)^{1/\theta}} + q \right\}^\theta. \tag{30}$$

Since $Kh \geq K\mathbb{1} = \beta^\theta > 0$, we have $q > 0$. Note that $Kh = h + K^N\mathbb{1} - \mathbb{1}$. By (29), the choice $c_N < 1$ gives $K^N\mathbb{1} < \mathbb{1}$ for $\theta > 0$ and $K^N\mathbb{1} > \mathbb{1}$ for $\theta < 0$. Hence $Kh < h$ in the first case and $Kh > h$ in the second, so $q < 1$ in both. Continuity and compactness bound q away from zero and one.

For $\theta > 0$, (30) tends to infinity as $t \to 0$ and to $q^\theta < 1$ as $t \to \infty$. For $\theta < 0$, the term instead tends to $q^\theta > 1$ and zero. In either case, the right-hand side of (30) exceeds 1 for t sufficiently small and is less than 1 for t sufficiently large. The bounds on r, h and q make these limits uniform across states and policies. Fix $w \in V$. Choose $0 < \ell < \inf(w/h)$ sufficiently small and $U > \sup(w/h)$ sufficiently large. Then $\ell h \leq w \leq Uh$ and, for some $\varepsilon > 0$,

$$\hat{T}_\sigma(\ell h) \geq \ell h + \varepsilon \mathbb{1}, \qquad \hat{T}_\sigma(Uh) \leq Uh - \varepsilon \mathbb{1} \quad \text{for all } \sigma \in \Sigma. \tag{31}$$

In the concave case, $\ell h \in V$ satisfies $\ell h \leq w$ and $\hat{T}_\sigma(\ell h) - \ell h \in V$ by (31), verifying condition (i) of Theorem 2.2 in Stachurski et al. (2025). Likewise, $Uh \in V$ satisfies $w \leq Uh$ and $\hat{T}_\sigma(Uh) \leq Uh$, verifying condition (ii). Since $w \in V$ was arbitrary, Theorem 2.2 applies.

In the convex case, $1/\theta > 0$, so the operator formula (14) extends $\hat{T}_\sigma$ to the non-negative cone of $B_b(\mathsf{X})$ as an order-preserving convex self-map. The extension has $\hat{T}_\sigma 0 = r_\sigma^y \in V$, verifying condition (i) of Theorem 2.4 in Stachurski et al. (2025). For each $w \in V$, the chosen $Uh \in V$ satisfies $w \leq Uh$ and $Uh - \hat{T}_\sigma(Uh) \in V$ by (31), verifying condition (ii).

Thus every policy operator has a unique globally attracting fixed point in V, and Example 2.1 gives order stability. Parts (i)–(ii) of Theorem 5.1 or 5.2 transfer well-posedness and order stability to $\mathscr{A}$. This proves the equivalence of (i) and (ii), as well as the final claim of the theorem.

Optimality. Assume $\Lambda < 1$. We verify max-stability of $\hat{\mathscr{A}}$ when $\gamma > 0$ and min-stability when $\gamma < 0$. For $w \in V_0$, weak continuity of P implies that $(x,a) \mapsto \{r(x,a)^\alpha + \beta(x)(\int_X w(x')P(x,a,\mathrm{d}x'))^{1/\theta}\}^\theta$ is continuous on G. By Berge's maximum theorem (Aliprantis and Border, 2006, Theorem 17.31), the Bellman operators $\hat{T}_\triangle$ and $\hat{T}_\triangledown$ of $\hat{\mathscr{A}}$ map V_0 into itself. Moreover, since G is closed in $\mathsf{X} \times \mathsf{A}$, this function extends continuously to $\mathsf{X} \times \mathsf{A}$ by the Tietze extension theorem (Aliprantis and Border, 2006, Theorem 2.47), and the measurable maximum theorem (Aliprantis and Border, 2006, Theorem 18.19) implies that $V_0 \subseteq \hat{V}_\triangle^G \cap \hat{V}_\triangledown^G$. Both $\hat{T}_\triangle$ and $\hat{T}_\triangledown$ inherit the common bounds in (31). We need a fixed point of $\hat{T}_\triangle$ for $\gamma > 0$ and of $\hat{T}_\triangledown$ for $\gamma < 0$. We shall apply the same fixed-point theorems on $C(\mathsf{X})$, whose positive-cone interior is V_0.

Suppose first that $\gamma \leq \alpha$. If $\gamma > 0$, then $0 < \theta \leq 1$, so $\hat{T}_\triangle$ is convex and extends to the nonnegative cone of $C(\mathsf{X})$, with $\hat{T}_\triangle 0 = \max_\sigma r_\sigma^\gamma$ uniformly positive. Theorem 2.4 of Stachurski et al. (2025) and the common upper bounds give a unique fixed point in V_0. If $\gamma < 0$, then $\theta < 0$ or $\theta \geq 1$, so $\hat{T}_\triangledown$ is concave. Theorem 2.2 of Stachurski et al. (2025) and the common lower and upper bounds give the same conclusion.

When $\alpha < \gamma$, we instead face a maximum of concave operators or a minimum of convex operators. Neither operation preserves the respective curvature in general. Define $Hw = w^{\alpha/\gamma}$ on V and set

$$\tilde{T}_\sigma w := H\hat{T}_\sigma H^{-1} w = r_\sigma^\alpha + \beta(P_\sigma w^\theta)^{1/\theta}. \tag{32}$$

On V_0, the corresponding Bellman maximum $\tilde{T}_\triangle$ is convex when $\alpha > 0$ ($\theta > 1$), and the minimum $\tilde{T}_\triangledown$ is concave when $\alpha < 0$ ($\theta < 1$), by the power-mean inequalities. Since H is an order isomorphism when $\theta > 0$ and an order anti-isomorphism when $\theta < 0$, it conjugates the required Bellman operator to $\tilde{T}_\triangle$ ($\alpha > 0$) or $\tilde{T}_\triangledown$ ($\alpha < 0$). Preservation of V_0, existence of greedy policies on V_0, and the common strict bounds transfer under H, with the lower and upper bounds exchanged if $\theta < 0$. The same applications of Theorems 2.4 and 2.2 in Stachurski et al. (2025) give a unique fixed point in V_0, which Lemma A.1 transfers back under H^{-1}.

Hence $\hat{\mathscr{A}}$ is max-stable when $\gamma > 0$ and min-stable when $\gamma < 0$. Theorem 5.1(iii) or 5.2(iii) gives max-stability of $\mathscr{A}$ and $v_\triangle \in V_0$, since F^{-1} preserves V_0. Theorem 4.1 gives the optimality conclusions. Since (iii) implies (ii), this completes the proof. □

Alternative proof of optimality in Theorem 6.2. Retain the preceding proof except for the Bellman fixed-point step when $\alpha < \gamma$. The common bounds (31) and order stability show that each $\hat{T}_\sigma$ preserves $[\ell h, Uh]$ and has its unique fixed point there.

The action objective in (14) is Borel measurable for Borel continuation values and continuous for continuous ones. On this uniformly positive bounded interval, dominated convergence gives convergence of the objective along monotone sequences. Together with continuity and compactness of Γ, these properties verify the hypotheses of Bloise et al. (2024, Propositions 1–3). For $\gamma > 0$, their result gives a fixed point of $\hat{T}_\Delta$ in V_0. For $\gamma < 0$, apply the same result to $w \mapsto -\hat{T}_\sigma(-w)$ on $[-Uh, -\ell h]$; negating the resulting continuous Bellman fixed point gives a fixed point of $\hat{T}_\nabla$ in V_0.

Preservation of V_0, existence of greedy policies on V_0, and the final application of Theorem 5.1(iii) or 5.2(iii) and Theorem 4.1 are unchanged. □

Proof of Corollary 6.3. Since Q is irreducible and $\beta > 0$, the matrix K is irreducible. Perron–Frobenius theory gives a vector h with $h(z) > 0$ for every z and $Kh = \rho(K)h$. Choose constants $m, M > 0$ such that $mh \leqslant \mathbb{1} \leqslant Mh$. Positivity of K yields

$$m\rho(K)^n h \leqslant K^n \mathbb{1} \leqslant M\rho(K)^n h.$$

Taking nth roots shows that $(K^n \mathbb{1}(z))^{1/n}$ converges to $\rho(K)$ uniformly over z. Hence $\Lambda = \rho(K)^{1/\theta}$ for either sign of θ. The optimality characterization follows from Theorem 6.2. □

References

Albuquerque, R., Eichenbaum, M., Luo, V. X., and Rebelo, S. (2016). Valuation risk and asset pricing. *The Journal of Finance*, 71(6):2861–2904.

Aliprantis, C. D. and Border, K. C. (2006). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag, New York, 3 edition.

Arnold, V. I. (2012). *Geometrical methods in the theory of ordinary differential equations*, volume 250. Springer Science & Business Media.

Azinovic, M., Gaegauf, L., and Scheidegger, S. (2022). Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525.

Balbus, Ł., Reffett, K., and Woźny, Ł. (2022). Time-consistent equilibria in dynamic models with recursive payoffs and behavioral discounting. *Journal of Economic Theory*, page 105493.

Barillas, F. and Fernández-Villaverde, J. (2007). A generalization of the endogenous grid method. *Journal of Economic Dynamics and Control*, 31(8):2698–2712.

Battaglini, M. (2021). Chaos and unpredictability in dynamic social problems. Technical report, National Bureau of Economic Research.

Bäuerle, N. and Jaśkiewicz, A. (2018). Stochastic optimal growth model with risk sensitive preferences. *Journal of Economic Theory*, 173:181–200.

Bäuerle, N., Jaśkiewicz, A., and Nowak, A. S. (2021). Stochastic dynamic programming with non-linear discounting. *Applied Mathematics & Optimization*, 84(3):2819–2848.

Bäuerle, N. and Rieder, U. (2011). *Markov decision processes with applications to finance*. Springer Science & Business Media.

Bayer, C. and Luetticke, R. (2020). Solving discrete time heterogeneous agent models with aggregate risk and many idiosyncratic states by perturbation. *Quantitative Economics*, 11(4):1253–1288.

Bellman, R. (1957). *Dynamic Programming*. Princeton University Press, Princeton, NJ.

Bertsekas, D. (2012). *Dynamic programming and optimal control*, volume 1. Athena Scientific.

Bertsekas, D. P. (2022). *Abstract dynamic programming*. Athena Scientific, 3 edition.

Bloise, G., Le Van, C., and Vailakis, Y. (2024). Do not blame Bellman: It is Koopmans' fault. *Econometrica*, 92(1):111–140.

Bloise, G. and Vailakis, Y. (2018). Convex dynamic programming with (bounded) recursive utility. *Journal of Economic Theory*, 173:118–141.

Borovička, J. and Stachurski, J. (2020). Necessary and sufficient conditions for existence and uniqueness of recursive utilities. *The Journal of Finance*, 75(3):1457–1493.

Brumm, J. and Scheidegger, S. (2017). Using adaptive sparse grids to solve high-dimensional dynamic models. *Econometrica*, 85(5):1575–1612.

Cagetti, M., Hansen, L. P., Sargent, T. J., and Williams, N. (2002). Robustness and pricing with uncertain growth. *The Review of Financial Studies*, 15(2):363–404.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters*, 91(3):312–320.

Davey, B. A. and Priestley, H. A. (2002). *Introduction to lattices and order*. Cambridge University Press.

de Castro, L., Galvao, A. F., and Nunes, D. (2022). Dynamic economics with quantile preferences. *Available at SSRN 4108230*.

de Groot, O., Richter, A. W., and Throckmorton, N. A. (2018). Uncertainty shocks in a model of effective demand: Comment. *Econometrica*, 86(4):1513–1526.

Deng, L., Khan, M. A., and Mitra, T. (2022). Continuous unimodal maps in economic dynamics: On easily verifiable conditions for topological chaos. *Journal of Economic Theory*, 201:105446.

Epstein, L. G. and Zin, S. E. (1989). Risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969.

Farahmand, A.-M., Munos, R., and Szepesvári, C. (2010). Error propagation for approximate policy and value iteration. In Lafferty, J. D., Williams, C. K. I., Shawe-Taylor, J., Zemel, R. S., and Culotta, A., editors, *Advances in Neural Information Processing Systems 23*, pages 568–576.

Fedus, W., Gelada, C., Bengio, Y., Bellemare, M. G., and Larochelle, H. (2019). Hyperbolic discounting and learning over multiple horizons.

Fei, Y., Yang, Z., Chen, Y., and Wang, Z. (2021). Exponential Bellman equation and improved regret bounds for risk-sensitive reinforcement learning. *Advances in Neural Information Processing Systems*, 34:20436–20446.

Flynn, J. P. and Sastry, K. (2022). *The macroeconomics of narratives*. SSRN.

Flynn, J. P., Schmidt, L. D. W., and Toda, A. A. (2023). Robust comparative statics for the elasticity of intertemporal substitution. *Theoretical Economics*, 18(1):231–265.

Foerster, A., Rubio-Ramírez, J. F., Waggoner, D. F., and Zha, T. (2016). Perturbation methods for markov-switching dynamic stochastic general equilibrium models. *Quantitative economics*, 7(2):637–669.

Gao, Y., Lui, K. Y. C., and Hernandez-Leal, P. (2021). Robust risk-sensitive reinforcement learning agents for trading markets. Technical report, arXiv preprint arXiv:2107.08083.

Gardini, L., Hommes, C., Tramontana, F., and De Vilder, R. (2009). Forward and backward dynamics in implicitly defined overlapping generations models. *Journal of Economic Behavior & Organization*, 71(2):110–129.

Gomez-Cram, R. and Yaron, A. (2020). How Important Are Inflation Expectations for the Nominal Yield Curve? *The Review of Financial Studies*, 34(2):985–1045.

Hansen, L. P. and Sargent, T. J. (2011). *Robustness*. Princeton university press.

Hernández-Lerma, O. and Lasserre, J. B. (2012). *Discrete-time Markov control processes: basic optimality criteria*, volume 30. Springer Science & Business Media.

Howard, R. A. and Matheson, J. E. (1972). Risk-sensitive Markov decision processes. *Management Science*, 18(7):356–369.

Kamihigashi, T. (2014). Elementary results on solutions to the Bellman equation of dynamic programming: existence, uniqueness, and convergence. *Economic Theory*,

56:251–273.

Kennedy, J. A. and Stockman, D. R. (2008). Chaotic equilibria in models with backward dynamics. *Journal of Economic Dynamics and Control*, 32(3):939–955.

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kochenderfer, M. J., Wheeler, T. A., and Wray, K. H. (2022). *Algorithms for decision making*. The MIT Press.

Kristensen, D., Mogensen, P. K., Moon, J. M., and Schjerning, B. (2021). Solving dynamic discrete choice models using smoothing and sieve methods. *Journal of Econometrics*, 223(2):328–360.

Layek, G. (2015). *An Introduction to Dynamical Systems and Chaos*. Springer.

Long, J. B. and Plosser, C. I. (1983). Real business cycles. *Journal of Political Economy*, 91(1):39–69.

Ma, Q. and Stachurski, J. (2021). Dynamic programming deconstructed: Transformations of the Bellman equation and computational efficiency. *Operations Research*, 69(5):1591–1607.

Ma, Q., Stachurski, J., and Toda, A. A. (2022). Unbounded dynamic programming via the q-transform. *Journal of Mathematical Economics*, 100:102652.

Maliar, L., Maliar, S., and Winant, P. (2021). Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122:76–101.

Marinacci, M. and Montrucchio, L. (2010). Unique solutions for stochastic recursive utilities. *Journal of Economic Theory*, 145(5):1776–1804.

Marinacci, M. and Montrucchio, L. (2019). Unique tarski fixed points. *Mathematics of Operations Research*, 44(4):1174–1191.

Milgrom, P. and Shannon, C. (1994). Monotone comparative statics. *Econometrica*, 62(1):157–180.

Peng, C., Stachurski, J., and Yang, J. (2026). Abstract dynamic programming on partially ordered spaces. arXiv preprint arXiv:2606.12777.

Puterman, M. L. (2005). *Markov decision processes: discrete stochastic dynamic programming*. Wiley Interscience.

Raines, B. and Stockman, D. (2012). Fixed points imply chaos for a class of differential inclusions that arise in economic models. *Transactions of the American Mathematical Society*, 364(5):2479–2492.

Ren, G. and Stachurski, J. (2021). Dynamic programming with value convexity. *Automatica*, 130:109641.

Rincón-Zapatero, J. P. (2024). Existence and uniqueness of solutions to the bellman equation in stochastic dynamic programming. *Theoretical Economics*, 19(3):1223–1260.

Rust, J. (1994). Structural estimation of Markov decision processes. *Handbook of Econometrics*, 4:3081–3143.

Sargent, T. J. and Stachurski, J. (2025). Dynamic programs on partially ordered sets. *SIAM Journal on Control and Optimization*, 63(2):778–795.

Schorfheide, F., Song, D., and Yaron, A. (2018). Identifying long-run risks: A Bayesian mixed-frequency approach. *Econometrica*, 86(2):617–654.

Smale, S. (1967). Differentiable dynamical systems. *Bulletin of the American Mathematical Society*, 73(6):747–817.

Stachurski, J., Wilms, O., and Zhang, J. (2025). Unique solutions to power-transformed affine systems. *Journal of Mathematical Analysis and Applications*, 550(1):129515.

Stachurski, J. and Zhang, J. (2021). Dynamic programming with state-dependent discounting. *Journal of Economic Theory*, 192:105190.

Sternberg, S. (2014). *Dynamical Systems*. Dover Publications.

Stokey, N. L. and Lucas, R. E. (1989). *Recursive methods in dynamic economics*. Harvard University Press.

Weil, P. (1990). Nonexpected utility in macroeconomics. *The Quarterly Journal of Economics*, 105(1):29–42.

National Graduate Institute for Policy Studies

Email address: john.stachurski@gmail.com

Center for Macroeconomic Research and Wang Yanan Institute for Studies in Economics, Xiamen University, China

Email address: zhangjunnan1224@gmail.com