Journal of Mathematical Economics 111 (2024) 102954

Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

An approximation approach to dynamic programming with unbounded returns☆

G. Bloise[a], C. Le Van[b,c], Y. Vailakis[d,*]

[a] Department of Economics and Finance, University of Rome II, Italy
[b] IPAG Business School, France
[c] Paris School of Economics, TIMAS, CASED, Vietnam
[d] Adam Smith Business School, University of Glasgow, United Kingdom

ARTICLE INFO

Manuscript handled by Editor A Toda

Keywords:
Stochastic recursive utility
Koopmans operator
Dynamic programming
Bellman operator
Multiplicity
Uniqueness

ABSTRACT

We study stochastic dynamic programming with recursive utility in settings where multiplicity of values is only attributed to unbounded returns. That is, we consider Koopmans aggregators that, when artificially restricted to be bounded, satisfy the traditional Blackwell's discounting condition (as it certainly happens with time-additive aggregators). We argue that, when the truncation is removed, the sequence of truncated values converges to the relevant fixed point of the untruncated Bellman operator, whenever it exists, and diverges otherwise. The experiment provides a natural selection criterion, corresponding to an extension of the recursive utility from bounded to unbounded returns.

1. Introduction

A large body of work in macroeconomics and finance models agents' behavior on a recursive utility foundation, instead of using the more restrictive time-additive separable objective. Inspired by the axiomatic contribution of Koopmans (1960), the approach postulates an aggregator function as a primitive representing preferences for current and future utility, as well as the attitude towards uncertainty, and recovers intertemporal utility recursively. Recursive utility broadens the scope of the analysis, encompassing behavioral features such as increasing marginal impatience (Lucas and Stokey, 1984), the distinction of risk attitudes from intertemporal substitution (Epstein and Zin, 1989), preference for early resolution of uncertainty (Kreps and Porteus, 1978), ambiguity aversion (Klibanoff et al., 2009), risk sensitivity and robustness (Hansen and Sargent, 1995, 2001). The wide range of economic applications is manifested in standard textbooks and survey articles (see, among others, Backus et al., 2005; Becker and Boyd, 1997; Hansen and Sargent, 2008; Hansen et al., 2007; Miao, 2014; Skiadas, 2009).

Dynamic programming provides important tools for studying economic models with recursive preferences. At the most fundamental level, the center of attention is on the existence and uniqueness of solutions to Bellman equation. When the value of the recursive program solves this functional equation, the associated policy characterizes the paths that are optimal in the associated sequential program. An important step amounts to show that Bellman's operator has a unique fixed point in a certain class of functions.

The use of general Koopmans aggregators to generate an intertemporal utility complicates the correspondence between sequential and recursive values of a planning program. In traditional dynamic programming, the sequential planning objective is explicitly given as the discounted sum of expected returns, and recursive methods serve to determine this value. With non-linear Koopmans aggregators, instead, the sequential planning objective is only implicitly identified as generated by the aggregator recursively. In principle, without further restrictions, there could be multiple recursively generated utility functions for a given Koopmans aggregator, all of them legitimate planning objectives consistent with the primitives.[1] As recently argued by Bloise et al. (2024), the Bellman operator might be returning this multiplicity and, in fact, this is the *only* source of multiplicity of fixed points of the

☆ We are grateful to Lukasz Balbus, Timothy Christensen, Martin Kaae Jensen, Filipe Martins-da-Rocha, Kevin Reffett, John Stachurski, Lukasz Wozny and Juan Pablo Rincón-Zapatero for helpful comments and observations. We also thank the editor (Alexis Akira Toda) and an anonymous reviewer for their suggestions. Gaetano Bloise acknowledges the financial support of the Italian Ministry of Education (PRIN 2022-PXE3B7). Yiannis Vailakis acknowledges the hospitality of Australian National University and University of Paris-Dauphine.

* Corresponding author.
E-mail addresses: gaetano.bloise@gmail.com (G. Bloise), cuong.le-van@univ-paris1.fr (C.L. Van), yiannis.vailakis@glasgow.ac.uk (Y. Vailakis).

[1] Examples of aggregators at risk of multiplicity appear in Becker and Rincón-Zapatero (2021, 2023), Bloise and Vailakis (2018) and Christensen (2022).

https://doi.org/10.1016/j.jmateco.2024.102954
Received 28 September 2023; Received in revised form 9 January 2024; Accepted 22 January 2024
Available online 2 February 2024

Bellman operator, a reflection of multiple planning objectives implied by the Koopmans operator.

This general framework also accommodates discounted expected utility with unbounded returns. Whenever multiplicity occurs, multiple values of the Bellman operator necessarily lie in an excessively large space in which the time-additive Koopmans operator itself yields multiple utilities. Though generated by the Koopmans operator, these other utilities are obviously spurious, as they do not correspond to the discounted sum of expected returns postulated in the related sequential program. Consequently, the associated fixed points of the Bellman operator can be disregarded because are inconsistent with the primitive objective of the planner. An analogous selection criterion is not immediately available for general recursive utility.

The purpose of this paper is to disentangle unbounded returns as a distinct source of multiplicity by means of a natural counterfactual experiment. More precisely, we restrict attention to Koopmans aggregators that, when artificially restricted to be bounded, satisfy the traditional Blackwell's discounting condition. In this sense we single out all situations in which the multiplicity of values can only be attributed to unbounded returns. Standard techniques, including the Contraction Mapping Theorem, ensure that the fixed point of the Bellman operator is unique for the truncation of the utility aggregator. We prove that, when the artificial truncation is removed, the sequence of truncated values converges to the relevant value of the unbounded recursive program. This experiment provides a natural selection criterion. For time-additive aggregators, in fact, it yields the exact value corresponding to the discounted expected utility.

We study aggregators that are either unbounded from above or unbounded from below, leaving a treatment of fully unbounded aggregators to future research.[2] Perhaps counterintuitively, unbounded-from-below aggregators behave more regularly: we are able to establish the existence of a stationary optimal policy. On the contrary, in general, we do not succeed in proving that the value generated by unbounded-from-above aggregators is upper semicontinuous and that an associated stationary optimal policy exists. However, under further regularity conditions, the discontinuity (if any) occurs for finite-horizon versions of the program and is not intrinsically related to the infinite horizon.

In Bloise et al. (2024) we argue that, whenever all the utilities generated by the Koopmans aggregator are legitimate planning objectives, the *greatest* fixed point of the Bellman operator should be privileged. In this paper, instead, we ostensively select the *least* fixed point for unbounded-from-above aggregators. The inconsistency is only apparent: our selection criterion in this paper implicitly regards all the other utilities generated by the unbounded Koopmans operator as spurious and, hence, as illegitimate planning objectives, because are unapproachable via truncations. For time-additive aggregators, for instance, the only acceptable utility is the discounted sum of expected returns. A similar selection is impracticable for more general recursive utilities studied in Bloise et al. (2024), as the risk of multiplicity persists even after truncating the aggregator (*e.g.*, with Epstein–Zin utility).

The paper is organized as follows. In Section 2, we briefly discuss the related literature. In Section 3, we introduce an abstract recursive program and argue that any value of this program corresponds to a recursively generated utility function. In Section 4, we illustrate, by means of a canonical example, how multiple values to the recursive program obtain under conventional discounting and unbounded returns. In Section 5, we present our major results. In Section 6, we show how our approach applies to programs with quadratic returns and risk-sensitive preferences. Finally, in Section 7, we compare our approach to some recent contributions in the literature.

[2] Notice that the discounted expected utility might be undefined when returns are fully unbounded, because the series might be neither converging nor diverging (or the integral might not exist).

2. Related literature

A well-developed literature has used refinements of the Contraction Mapping Theorem to deal with unbounded Blackwell aggregators. Boyd (1990) introduces the Weighted Contraction approach to economic applications, a method further developed by Alvarez and Stokey (1998), Durán (2000, 2003), Jaśkiewicz and Nowak (2011), Matkowski and Nowak (2011) and more recently by Rincón-Zapatero (2024). An alternative Local Contraction approach is presented by Rincón-Zapatero and Rodríguez-Palmero (2003) and Martins-da-Rocha and Vailakis (2010) to deal with aggregators that allow $-\infty$ as a value. Another method followed by Le Van and Morhaim (2002), Le Van and Vailakis (2005), Kamihigashi (2014) and Wiszniewska-Matyszkiel and Singh (2021) abandons the contraction approach and looks directly for solutions to the Bellman's equation in a suitable space of functions satisfying a sort of transversality condition. Finally, a recent paper by Ma et al. (2022) exploits a transformation of Bellman's operator, along with boundedness of the expected reward, to turn unbounded into bounded programs so that conventional contraction techniques apply.

Dynamic programming with recursive utility was initially approached by Streufert (1990) and Ozaki and Streufert (1996). They introduced the notion of biconvergence, a limiting condition ensuring that returns of any feasible path are sufficiently discounted from above and from below. Further developments along this line appear in Bich et al. (2018) who study deterministic recursive programs under minimal assumptions on primitives. Though foundational, all these biconvergence criteria are not fully operational.

Recent contributions by Balbus (2020), Bloise and Vailakis (2018) and Ren and Stachurski (2021) study dynamic programming with aggregators that fail the conventional sup-norm contractivity property. The approach builds instead around the monotonicity and value concavity properties of the Koopmans aggregator. Marinacci and Montrucchio (2010, 2019), Christensen (2022), Becker and Rincón-Zapatero (2021) and Becker and Rincón-Zapatero (2023) show that these properties are shared by many relevant aggregators.

3. Recursive program

Let X and Z be complete separable metric spaces, and let G be a correspondence from X to Z. We interpret X as the state space, whereas Z is the action space. Feasibility is embedded in the correspondence $G : X \twoheadrightarrow Z$, that is, $G(x) \subset Z$ is the set of admissible actions at state x in X. We use $\Gamma \subset X \times Z$ to denote the graph of the feasibility correspondence. If needed, we also consider a (measurable) Markov transition $\Pi : \Gamma \to \Delta(X)$ governing the evolution of the state over time, that is, $\Pi(x, z)$ is a probability measure on the state space X, endowed with its Borel algebra.

We let $\mathcal{F}$ be the space of all extended real maps $f : X \to \overline{\mathbb{R}}$, endowed with the product topology and the natural ordering, where $\overline{\mathbb{R}}$ is the field of extended reals. We also let $\mathcal{V}$ be the class of measurable maps $v : X \to \overline{\mathbb{R}}$ in $\mathcal{F}$, that is,

$$\mathcal{V} = \{v \in \mathcal{F} : v \text{ is (Borel) measurable}\}.$$

The objective of the planner is given as a bounded-from-below aggregator $W : \Gamma \times \mathcal{V} \to \overline{\mathbb{R}}^{+}$ or a bounded-from-above aggregator $W : \Gamma \times \mathcal{V} \to \overline{\mathbb{R}}^{-}$. The nature of this aggregator will depend on the application of the theory. The most traditional example is given by

$$W(x, z, v) = (1-\delta) u(x, z) + \delta \int v(y) \Pi(x, z)(dy),$$

where δ in $(0, 1) \subset \mathbb{R}^{+}$ is the discount factor and $u : \Gamma \to \overline{\mathbb{R}}$ is the return, or reward, function. Our abstract formulation, inspired by Bertsekas (2018), encompasses many other instances of recursive preferences.

We impose restrictions on fundamentals that are satisfied in typical applications. Assumption 3.1 ensures the applicability of the Maximum

Theorem when values are bounded. In addition, it requires monotonicity. These restrictions will deliver a well-behaved recursive program when values are bounded. Assumption 3.2 reproduces the logic of Levi's Convergence Theorem. This continuity property will be helpful for the limit argument.

Assumption 3.1 (*Basic Properties*). (a) The feasible correspondence $G : X \twoheadrightarrow Z$ is upper hemicontinuous with nonempty compact values. (b) For any v in $\mathcal{V}$, the map $(x,z) \mapsto W(x,z,v)$ is measurable on Γ. (c) For any v in $\mathcal{V}$ that is bounded and upper semicontinuous, the map $(x,z) \mapsto W(x,z,v)$ is upper semicontinuous on Γ. (d) The aggregator is monotone in $\mathcal{V}$, that is,

$$v' \geq v'' \text{ implies } W\left(x,z,v'\right) \geq W\left(x,z,v''\right).$$

Assumption 3.2 (*Monotone Convergence*). For any sequence $\left(v_n\right)_{n\in\mathbb{N}}$ in $\mathcal{V}$ monotonically converging to v in $\mathcal{V}$,

$$\lim_{n\to\infty} W\left(x,z,v_n\right) = W(x,z,v).$$

The *value* of the *recursive program* is given by

$$v(x) = \sup_{z\in G(x)} W(x,z,v).$$

Values of the recursive program obtain as fixed points of the Bellman operator $T : \mathcal{V} \to \mathcal{F}$, given by

$$(Tv)(x) = \sup_{z\in G(x)} W(x,z,v).$$

Notice that, due to well-known issues of measurability, the Bellman operator returns values in the extended space $\mathcal{F}$, though we only consider fixed points which are elements of the allowed space $\mathcal{V}$.

A policy is a measurable map $g : X \to Z$ such that $g(x)$ lies in $G(x)$. Let $\mathcal{G}$ be the space of all such policies. Given a policy g in $\mathcal{G}$, Koopmans operator $T_g : \mathcal{V} \to \mathcal{V}$ is given by

$$\left(T_g v\right)(x) = W(x, g(x), v).$$

A utility function $U : \mathcal{G} \times X \to \overline{\mathbb{R}}$ specifies the utility value $U_g(x)$ in $\overline{\mathbb{R}}$ of a policy g in $\mathcal{G}$ beginning from state x in X. A utility function is recursively generated by the given aggregator whenever

$$U_g(x) = \left(T_g U_g\right)(x).$$

In principle, without further restrictions, there could be multiple recursively generated utility functions for a given aggregator.[3]

We shall devote attention to the least recursively generated utility function, when the aggregator is unbounded-from-above, and the greatest recursively generated utility function, when the aggregator is unbounded-from-below. These selections mimic the logic of the conventional discounted expected utility. In fact, the least and the greatest recursively generated utilities, in both cases, are determined as

$$U_g(x) = \lim_{n\to\infty} \left(T_g^n 0\right)(x).$$

This formula reproduces the idea of a discounted utility,

$$U_g\left(x_0\right) = \lim_{n\to\infty} \mathbb{E}_0 \sum_{t=0}^{n} \delta^t u\left(x_t, g\left(x_t\right)\right),$$

where x_{t+1} in X is conditionally distributed according to the probability $\Pi\left(x_t, g\left(x_t\right)\right)$.

Given a utility function $U : \mathcal{G} \times X \to \overline{\mathbb{R}}$, the *value* of the *sequential program* is given by

$$v(x) = \sup_{g\in\mathcal{G}} U_g(x).$$

Notice that the value of the sequential program is uniquely determined, though it depends on the given utility function. The value of the recursive program, instead, only reflects the aggregator. We shall argue that the value of the sequential program for the chosen utility function can be computed by means of approximations of unbounded aggregators whenever appropriate truncations satisfy the canonical discounting property.

As we extensively argue in Bloise et al. (2024), a multiplicity of fixed points of the Bellman operator is the reflection of a multiplicity of utilities generated by the Koopmans operator on the relevant space. In certain circumstances, some of these utilities are clearly artificial and can be ruled out by an appropriate restriction of the underlying space, as clarified by our example in Section 4. In general, however, many recursively generated utilities are legitimate objectives of the planning program. This feature is annotated in the next proposition.

Proposition 3.1 (*Principle of Optimality*). *Any fixed point $\tilde{v}$ in $\mathcal{V}$ of Bellman operator $T : \mathcal{V} \to \mathcal{F}$, admitting a policy g in $\mathcal{G}$, is the sequential value of some recursively generated utility $\tilde{U} : \mathcal{G} \times X \to \overline{\mathbb{R}}$.*

4. A leading example

Notwithstanding a large literature on value-uniqueness in dynamic programming with economic applications, we find an unexpected paucity of examples of multiplicity. We here present a simple framework with conventional discounting and unbounded returns for which the related Bellman operator admits multiple fixed points. We explain the source of this multiplicity and the emergence of spurious utilities when the underlying space is excessively permissive. We also verify that the fixed point is unique on a suitable space enforcing discounting.

Let $X = Z = \mathbb{R}^+$ and $G(x) = [0,x] \subset \mathbb{R}^+$. Given δ in $(0,1) \subset \mathbb{R}^+$, the utility aggregator is given by

$$W(x,z,v) = u(x,z) + \delta v\left(\frac{x-z}{\delta}\right),$$

where

$$u(x,z) = \sqrt{z+1} - 1.$$

Given these primitives, the induced recursive program is

$$v(x) = \max_{0\leq z\leq x} \sqrt{z+1} - 1 + \delta v\left(\frac{x-z}{\delta}\right).$$

We shall argue that, when unrestricted, the Bellman operator admits multiple fixed points.

The peculiarity of this aggregator is that it induces a natural utility function, namely,

$$\underline{U}_g\left(x_0\right) = \sum_{t=0}^{\infty} \delta^t u\left(x_t, g\left(x_t\right)\right),$$

where $x_{t+1} = \delta^{-1}\left(x_t - g\left(x_t\right)\right)$. The Koopmans aggregator, however, is consistent with other spurious utility functions if its domain is not

[3] The proposed formulation is sufficiently permissive to parsimoniously encompass the more traditional treatments of recursive utility in the literature. Consider, for instance, a deterministic framework with a utility defined on sequences of consumptions. We let $X = \mathbb{T} = \{0,1,2,\ldots,t,\ldots\}$ be the infinite space of periods and $Z = \mathbb{R}^+$ the consumption space. A policy g in $\mathcal{G}$ is thus a consumption plan, with $g(t)$ in $\mathbb{R}^+$ being the consumption level in period t in $\mathbb{T}$. Adapting our notation, a recursive utility satisfies

$$U_g(t) = W\left(t, g(t), U_g\right) = V\left(g(t), U_g(t+1)\right),$$

where $V : \mathbb{R}^+ \times \mathbb{R}^+ \to \mathbb{R}^+$ is the conventional aggregator and $U_g(t)$ in $\mathbb{R}^+$ is the utility value from consumption sequence g in $\mathcal{G}$ beginning from period t in $\mathbb{T}$. Defining $g \circ \sigma(t) = g(t+1)$, this is equivalent to

$$U_g(0) = V\left(g(0), U_{g\circ\sigma}(0)\right),$$

so that $U(0) : \mathcal{G} \to \mathbb{R}^+$ is the intertemporal utility defined on the space of consumption sequences. More generally, by an expansion of the state space from X to $X \times \mathbb{T}$, or even $X^{\mathbb{T}}$, we can accommodate nonstationary policies and history dependence in our analysis.

restricted to an appropriate suitable space. In fact, any other utility will satisfy the recursive condition

$$U_g(x) - \underline{U}_g(x) = \delta\left(U_g\left(\frac{x-g(x)}{\delta}\right) - \underline{U}_g\left(\frac{x-g(x)}{\delta}\right)\right).$$

Thus, the excess utility value will be inflating over time, remaining constant in discounted terms.

It turns out that the value corresponding to the natural utility is

$$\underline{v}(x) = \frac{\sqrt{(1-\delta)x+1}-1}{1-\delta}.$$

To verify this, notice that the recursive program becomes

$$\max_{0\le z\le x} \sqrt{z+1} - 1 + \delta\frac{\sqrt{(1-\delta)\delta^{-1}(x-z)+1}-1}{1-\delta}.$$

A simple first-order condition requires

$$\frac{1}{2\sqrt{z+1}} = \frac{1}{2\sqrt{(1-\delta)\delta^{-1}(x-z)+1}}.$$

This yields $z=(1-\delta)x$ and $x-z=\delta x$, thus confirming our claim.

We now argue that $v(x)=\gamma x$ is also a fixed point of the Bellman operator whenever $1\le 2\gamma$. Indeed, the recursive program reduces to

$$\max_{0\le z\le x} \sqrt{z+1} - 1 + \gamma(x-z).$$

Notice that the first-order condition requires

$$\frac{1}{2\sqrt{z+1}} - \gamma \le 0, \text{ with the equality if } z>0.$$

The maximizer is $z=0$, thus establishing our claim. All these additional fixed points are the artifact of corresponding spurious utilities generated by the Koopmans operator, because $\underline{U}_g(x)=0$ whenever $g(x)=0$ permanently.

The spurious values can be ruled out by means of an appropriate upper bound. Introducing a bounding function $\varphi(x)=\sqrt{x}$, notice that

$$\max_{z\in G(x)} u(x,z) = \sqrt{x+1}-1 \le \sqrt{x} = \varphi(x).$$

Thus, this bounding function is consistent with the growth rate of returns. As a consequence, it also provides an upper bound to the value of the recursive program for the natural utility, as

$$\underline{v}(x) = \frac{\sqrt{(1-\delta)x+1}-1}{1-\delta} \le \frac{\sqrt{x}}{1-\delta} = \left(\frac{1}{1-\delta}\right)\varphi(x).$$

However, all spurious fixed points are precluded by this bound. Indeed, assume the existence of λ in $\mathbb{R}^+$ such that

$$v(x) = \gamma x \le \lambda\sqrt{x} \le \lambda\varphi(x).$$

This yields

$$\lim_{x\to\infty}\frac{v(x)}{\varphi(x)} = \lim_{x\to\infty}\gamma\sqrt{x} \le \lambda,$$

which clearly implies a contradiction.

The relevant properties of the bounding function are that (a) it provides an upper bound on returns and (b) it involves some sort of discounting of values over time, thus ruling out the bubbly excess utility. To see the latter property, notice that

$$\delta\sup_{z\in G(x)}\varphi\left(\frac{x-z}{\delta}\right) \le \delta\sqrt{\frac{x}{\delta}} = \sqrt{\delta}\varphi(x).$$

The left hand-side evaluates, given the bounding function, the maximum discounted value growth permitted by the feasible set, whereas the right hand-side ensures that this value declines over time at geometric rate $\sqrt{\delta}$ in $(0,1)\subset\mathbb{R}^+$. As explained by Bloise et al. (2024, Proposition 5.1), this discounting property for general aggregators is captured by the existence of a monotone sublinear gradient to the Bellman operator whose spectral radius is less than unity.

5. Approximate discounting

5.1. Selection criterion

We develop a selection criterion for situations in which unbounded returns are the *only* source of potential multiplicity. In particular, to single out such circumstances, we consider aggregators satisfying a traditional Blackwell's discounting property for some bounded approximation. Conventional techniques, including the Contraction Mapping Theorem, ensure that the fixed point of the truncated Bellman operator is unique. We argue that, removing the truncation, values converge to the relevant fixed point of the untruncated Bellman operator.

It is worth noticing that restrictions on primitives postulated in Assumption 3.1 are prevalently on bounded values, and serve to establish existence of a fixed point for the bounded approximation of the program, as in conventional dynamic programming. The limit will only exploit a fairly permissive assumption of monotone convergence (Assumption 3.2). In particular, our approach avoids the well-studied issues related to the Feller property for unbounded values. As a matter of fact, we only require a weaker form of Feller property for bounded values.

Our approach separates unbounded from above and unbounded from below aggregators, leaving a full treatment of unbounded returns to future research. The treatment is not symmetric. The theory applies satisfactorily to values that are unbounded from below, permitting to prove the existence of an optimal policy for the untruncated program. Differently, when values are unbounded from above, we only succeed in verifying the existence of an approximate policy. Further restrictions in terms of discounting reveal that a potential failure of existence of the policy occurs in the corresponding finite-horizon programs, rather than to fully expanded infinite-horizon version.

The theory requires the existence of a suitable upper (lower) bound when returns are unbounded from above (below). For the truncated program these bounds are created artificially. For the untruncated program, however, we still need to assume the presence of suitable bounds in order to ensure the convergence of values when the truncation is progressively removed. As a matter of fact, when convergence fails, the original program is misspecified and yields an infinite value. Hence, as a combined effect of all our restrictions, the requirement of bounds is equivalent to the assumption of existence of a finite value for the untruncated program. As an alternative, we could abstain from imposing bounds at all and state all of our results conditional on the existence of a finite value of the underlying program.

5.2. Unbounded-from-above aggregators

We maintain Assumptions 3.1 and 3.2 and restrict attention to positive aggregators, that is, to aggregators of the form $W : \Gamma\times\mathcal{V}\to\overline{\mathbb{R}}^+$. Furthermore, we assume the existence of a *finite* upper bound $\bar{f}$ in $\mathcal{F}$ such that

$$\sup_{z\in G(x)} W\left(x,z,\bar{f}\right) \le \bar{f}(x).$$

The only role of this bound is to ensure convergence of the truncated values. We allow $\bar{f}$ in $\mathcal{F}$ to fail upper semicontinuity and even measurability. The class of all such unbounded-from-above aggregators is denoted by $\mathcal{W}$. Finally, we interpret $\mathcal{V}$ as the space of maps with values in $\overline{\mathbb{R}}^+$.

We say that an aggregator W in $\mathcal{W}$ exhibits the property of *discounting by increasing truncations* if it is the (pointwise) limit of an *increasing* sequence of truncated aggregators $(W_n)_{n\in\mathbb{N}}$ in $\mathcal{W}$ satisfying:

(D-1) each truncated aggregator maps bounded values into bounded values, that is, given a *bounded* v in $\mathcal{V}$,

$$\sup_{(x,z)\in\Gamma}\left|W_n(x,z,v)\right| \text{ is finite;}$$

(D-2) each truncated aggregator satisfies the property of discounting, that is, for some δ_n in $(0,1) \subset \mathbb{R}^+$, given *bounded* v' and v'' in $\mathcal{V}$,

$$W_n\left(x,z,v'\right) - W_n\left(x,z,v''\right) \le \delta_n \sup_{x\in X}\left|v'(x) - v''(x)\right|.$$

As a matter of fact, we postulate that the aggregator can be arbitrarily approximated by truncated aggregators fulfilling the conventional Blackwell's condition. By canonical dynamic programming, this assumption ensures that the truncated Bellman operator admits a unique fixed point v_n in $\mathcal{V}$. Furthermore, as the truncation is relaxed, values increase and, hence, $v_n \le v_{n+1}$. Since $v_n \le \bar{f}$, the sequence converges to v^* in $\mathcal{V}$ because it is the pointwise limit of an increasing sequence of measurable maps. We argue that the limit is the least fixed point of the Bellman operator for the untruncated recursive program.

Proposition 5.1 (*Limit*). *Under the maintained Assumptions 3.1–3.2, if the aggregator satisfies the property of discounting by increasing truncations, the induced limit v^* in $\mathcal{V}$ is the least fixed point of the untruncated Bellman operator $T : \mathcal{V} \to \mathcal{F}$.*

The limit value of truncated recursive programs is indeed the sequential value for the least utility function $U : \mathcal{G}\times X \to \overline{\mathbb{R}}^+$ generated by the aggregator, which is determined as

$$U_g(x) = \lim_{j\to\infty}\left(T_g^j 0\right)(x).$$

For time-additive aggregators, this utility corresponds to the discounted sum of expected returns over the infinite horizon. Thus, our method naturally extends the traditional approach to discounted expected utility with unbounded-from-above returns.

Proposition 5.2 (*Principle of Optimality*). *Under the maintained Assumptions 3.1–3.2, if the aggregator satisfies the property of discounting by increasing truncations, the induced limit v^* in $\mathcal{V}$ satisfies*

$$v^*(x) = \sup_{g\in\mathcal{G}} U_g(x),$$

where $U : \mathcal{G}\times X \to \overline{\mathbb{R}}^+$ is the least utility function recursively generated by the aggregator.

In general, we cannot establish the existence of a policy, unless we further restrict the fundamentals. Here is a practicable set of additional restrictions. Thought it is admittedly far from the weakest set of restrictions, it is an acceptable compromise in practice. In addition to an appropriate form of discounting, we require a well-behaved program over any finite horizon, that is, the value of the program over a finite horizon of length n in $\mathbb{N}$, $(T^n 0)$, is upper semicontinuous. With the purpose of capturing discounting on the relevant space, we introduce the space $\mathcal{L}\left(\bar{f}\right)$ defined as

$$\mathcal{L}\left(\bar{f}\right) = \left\{v\in\mathcal{F} : |v| \le \bar{f}\right\}.$$

This is the space of all values involving a growth rate bounded by $\bar{f}$ in $\mathcal{F}$.

Assumption 5.1 (*Discounting*). There exists a monotone sublinear operator $D : \mathcal{L}\left(\bar{f}\right) \to \mathcal{L}\left(\bar{f}\right)$ such that, for every v' and v'' in $\mathcal{V}\cap\mathcal{L}\left(\bar{f}\right)$,

$$W\left(x,z,v'\right) - W\left(x,z,v''\right) \le D\left(v'-v''\right)(x).$$

Furthermore, each $D^n\left(\bar{f}\right)$ in $\mathcal{L}\left(\bar{f}\right)$ is upper semicontinuous and, for every x in X,

$$\lim_{n\to\infty} D^n\left(\bar{f}\right)(x) = 0.$$

Proposition 5.3 (*Limit Strengthened*). *Under the maintained Assumptions 3.1–3.2–5.1, if the aggregator satisfies the property of discounting by increasing truncations, the induced limit v^* in $\mathcal{V}$ is the only fixed point of the untruncated Bellman operator $T : \mathcal{V}\to\mathcal{F}$ in the relevant space $\mathcal{L}\left(\bar{f}\right)$. Furthermore, a policy g in $\mathcal{G}$ exists, provided that, for every n in $\mathbb{N}$, $(T^n 0)$ in $\mathcal{V}$ is upper semicontinuous.*

5.3. Unbounded-from-below aggregators

We now reverse the logic of the truncation approach in order to encompass unbounded-from-below aggregators. We maintain Assumptions 3.1 and 3.2 and restrict attention to negative aggregators, that is, to aggregators of the form $W : \Gamma\times\mathcal{V}\to\overline{\mathbb{R}}^-$. Furthermore, we assume the existence of a *finite* lower bound $\underline{f}$ in $\mathcal{F}$ such that

$$\sup_{z\in G(x)} W\left(x,z,\underline{f}\right) \ge \underline{f}(x).$$

The only role of this bound is to ensure convergence of the truncated values. We allow $\underline{f}$ in $\mathcal{F}$ to fail upper semicontinuity and even measurability. The class of all such unbounded-from-below aggregators is denoted by $\mathcal{W}$. Finally, we interpret $\mathcal{V}$ as the space of maps with values in $\overline{R}^-$.

We say that an aggregator W in $\mathcal{W}$ exhibits the property of *discounting by decreasing truncations* if it is the (pointwise) limit of a *decreasing* sequence of truncated aggregators $\left(W_n\right)_{n\in\mathbb{N}}$ in $\mathcal{W}$ satisfying:

(D-1) each truncated aggregator maps bounded values into bounded values, that is, given a *bounded* v in $\mathcal{V}$,

$$\sup_{(x,z)\in\Gamma}\left|W_n(x,z,v)\right| \text{ is finite;}$$

(D-2) each truncated aggregator satisfies the property of discounting, that is, for some δ_n in $(0,1)\subset\mathbb{R}^+$, given *bounded* v' and v'' in $\mathcal{V}$,

$$W_n\left(x,z,v'\right) - W_n\left(x,z,v''\right) \le \delta_n \sup_{x\in X}\left|v'(x) - v''(x)\right|.$$

As in our previous analysis, by canonical dynamic programming, this assumption ensures that the truncated Bellman operator admits a unique fixed point v_n in $\mathcal{V}$. Furthermore, as the truncation is relaxed, values decrease and, hence, $v_n \ge v_{n+1}$. Since $v_n \ge \underline{f}$, the sequence converges to v^* in $\mathcal{V}$ because it is the pointwise limit of a decreasing sequence of measurable maps. We argue that the limit is the greatest fixed point of the Bellman operator of the untruncated program.

Unbounded-from-below aggregators behave more regularly than those which are unbounded from above. Indeed, they preserve upper semicontinuity even when values are unbounded. We establish this relevant property in a preliminary lemma.

Lemma 5.1 (*Upper Semicontinuity*). *Under the maintained Assumptions 3.1–3.2, for every upper semicontinuous v in $\mathcal{V}$, the map $(x,z)\mapsto W(x,z,v)$ is upper semicontinuous on Γ.*

Proposition 5.4 (*Limit*). *Under the maintained Assumptions 3.1–3.2, if the aggregator satisfies the property of discounting by decreasing truncations, the induced limit v^* in $\mathcal{V}$ is the greatest fixed point of the untruncated Bellman operator $T : \mathcal{V}\to\mathcal{F}$. Furthermore, an optimal policy g in $\mathcal{G}$ exists.*

As in our previous analysis, the limit value of truncated recursive programs is indeed the sequential value for the greatest utility function $U : \mathcal{G}\times X\to\overline{\mathbb{R}}^-$ generated by the aggregator, which is determined as

$$U_g(x) = \lim_{j\to\infty}\left(T_g^j 0\right)(x).$$

This is established in our last proposition.

Proposition 5.5 (*Principle of Optimality*). *Under the maintained Assumptions 3.1–3.2, if the aggregator satisfies the property of discounting by decreasing truncations, the induced limit v^* in $\mathcal{V}$ satisfies*

$$v^*(x) = \sup_{g\in\mathcal{G}} U_g(x),$$

where $U : \mathcal{G}\times X\to\overline{\mathbb{R}}^-$ is the greatest utility function recursively generated by the aggregator.

5.4. A comparison with Rincón-Zapatero and Rodríguez-Palmero (2003)

Rincón-Zapatero and Rodríguez-Palmero (2003, Section 3.2) also propose a truncation approach to unbounded returns in deterministic recursive programs with traditional discounted utility. Under certain restrictions on the feasible correspondence $G : X \twoheadrightarrow Z$, they construct an increasing sequence of truncated feasible correspondences and prove the existence of a single fixed point of the Bellman operator along the sequence of truncations. They so establish that, under additional assumptions, the limit of such truncated values is the value of the untruncated program.

Differently from Rincón-Zapatero and Rodríguez-Palmero (2003), our truncation involves no manipulation of the feasible correspondence, as it only acts on the planning objective. As a matter of fact, our assumptions (D-1)–(D-2) are totally innocuous in the case of traditional discounted utility, as they obtain by a straightforward truncation of the return (namely, $u_n(x,z) = u(x,z) \wedge n$ or $u_n(x,z) = u(x,z) \vee (-n)$). Our alternative truncation seems more convenient operationally, because any implementation necessarily requires finite bounds for the numerical computation. Furthermore, we dispense with restrictions (DP3)–(DP4) in Rincón-Zapatero and Rodríguez-Palmero (2003, Theorem 5) on the feasible correspondence, and simply rely on more minimal conditions ensuring the applicability of the Maximum Theorem.

For the unbounded-from-above returns studied in Section 5.2, the additional requirements guaranteeing convergence in their and our approaches are basically comparable. In particular, Assumption (i) in Rincón-Zapatero and Rodríguez-Palmero (2003, Theorem 5), as also clarified in their Remark 5(ii), would be satisfied by

$$D(v) = \delta \max_{z \in G(x)} \int |v(y)| \, \Pi(x,z)(dy),$$

provided that the spectral radius of $D : \mathcal{L}(\bar{f}) \to \mathcal{L}(\bar{f})$ is such that $\rho(D) < 1$. We do not impose Assumption (ii) in Rincón-Zapatero and Rodríguez-Palmero (2003, Theorem 5) because our aggregator is bounded-from-below. We remark that, differently from Rincón-Zapatero and Rodríguez-Palmero (2003, Theorem 5), no further assumptions are needed for the unbounded-from-below aggregators of Section 5.3.

To complete the comparison, we notice that, in Proposition 5.3, we also require that $(T^n 0)$ in $\mathcal{V}$ be upper semicontinuous for every n in $\mathbb{N}$. This is unnecessary in Rincón-Zapatero and Rodríguez-Palmero (2003, Theorem 5) because their framework is deterministic. In stochastic applications, instead, a discontinuity might arise because of a failure of the Feller property due to unbounded returns. Our further requirement prevents such an occurrence.

6. Applications

6.1. Quadratic returns

We apply our method to quadratic returns under conventional discounting. Thus, the aggregator is given as

$$W(x,z,v) = (1-\delta)\, u(x,z) + \delta \int v(y)\, \Pi(x,z)(dy),$$

where δ in $(0,1) \subset \mathbb{R}^+$ is the discount factor. The return takes the quadratic form

$$u(x,z) = x\Phi_x x + z\Phi_z z,$$

where Φ_x is a negative definite matrix on $X = \mathbb{R}^n$ and Φ_z is a negative definite matrix on $Z = \mathbb{R}^m$. Furthermore, we postulate that the Markov transition $\Pi : \Gamma \to \Delta(X)$ satisfies the canonical Feller property, that is, for every bounded and continuous map v in $\mathcal{V}$, the map $(\Pi v) : \Gamma \to \mathbb{R}$ is bounded and continuous, where

$$(\Pi v)(x,z) = \int v(y)\, \Pi(x,z)(dy).$$

This quadratic program satisfies all the assumptions for unbounded-from-below returns (Section 5.3). In fact, a natural truncation is given by $u_n(x,z) = u(x,z) \vee (-n)$ with n in $\mathbb{N}$. Therefore, to apply our theory, we only have to verify the existence of a policy g in $\mathcal{G}$ yielding a finite value, that is, such that

$$\lim_{n\to\infty} \left(T_g^n 0\right)(x) \text{ is finite.}$$

Whenever such a policy does not exist, the corresponding sequential program admits no finite value.

6.2. Risk-sensitive preferences

Our approach encompasses risk-sensitive preferences. The aggregator is given by

$$W(x,z,v) = (1-\delta)\, u(x,z) - \frac{\delta}{\theta} \log\left(\int e^{-\theta v(y)} \right) \Pi(x,z)(dy),$$

where δ in $(0,1) \subset \mathbb{R}^+$ is the discount factor and θ in $\mathbb{R}^{++}$ is the risk-sensitive parameter. We assume that the return $u : \Gamma \to \mathbb{R}^-$ is bounded from above (and a similar treatment applies to bounded-from-below returns).

We naturally gain bounded returns by means of the truncation $u_n(x,z) = u(x,z) \vee (-n)$. To verify the discounting property for these bounded truncations, we exploit a canonical argument (as in Christensen, 2022). In particular,

$$\begin{aligned}
&W_n\left(x,z,v'\right) - W_n\left(x,z,v''\right) \\
&= -\frac{\delta}{\theta} \log\left(\frac{\int e^{-\theta v'(y)} \Pi(x,z)(dy)}{\int e^{-\theta v''(y)} \Pi(x,z)(dy)} \right) \\
&= -\frac{\delta}{\theta} \log\left(\frac{\int e^{-\theta v''(y)} e^{-\theta(v'(y)-v''(y))} \Pi(x,z)(dy)}{\int e^{-\theta v''(y)} \Pi(x,z)(dy)} \right) \\
&\le -\frac{\delta}{\theta} \log\left(\frac{\int e^{-\theta v''(y)} e^{-\theta \|v'-v''\|_\infty} \Pi(x,z)(dy)}{\int e^{-\theta v''(y)} \Pi(x,z)(dy)} \right) \\
&= -\frac{\delta}{\theta} \log\left(e^{-\theta \|v'-v''\|_\infty} \right) \\
&= \delta \left\| v' - v'' \right\|_\infty,
\end{aligned}$$

thus confirming our claim. Hence, Propositions 5.4–5.5 apply.

In fact, by the same argument, risk-sensitive preferences cannot sustain distinct fixed points v' and v'' in $\mathcal{V}$ such that $\|v' - v''\|_\infty > 0$ is finite, even though they are unbounded. Indeed, previous inequalities would deliver

$$\left|v'(x) - v''(x)\right| \le \left| \left(Tv'\right)(x) - \left(Tv''\right)(x) \right| \le \delta \left\|v' - v''\right\|_\infty,$$

thus implying a contradiction. This is consistent with the uniqueness established in Ma et al. (2022, Theorem 5.3) by means of their Q-transform. In both cases, however, the existence of other fixed points with unbounded uniform distance is not ruled out.

6.3. Blackwell nonlinear aggregators

Our approach applies to any nonlinear aggregator satisfying, for some δ in $(0,1) \subset \mathbb{R}^+$,

$$W\left(x,z,v'\right) - W\left(x,z,v''\right) \le \delta \sup_{z \in G(x)} \int \left|v'(y) - v''(y)\right| \Pi(x,z)(dy).$$

Examples of such aggregators are studied by Marinacci and Montrucchio (2010). An instance is the KDW aggregator (Koopmans et al., 1964) given by

$$W(x,z,v) = \frac{1}{\theta} \log\left(1 + u(x,z) + \beta \int v(y)\, \Pi(x,z)(dy) \right),$$

where $u : \Gamma \to \mathbb{R}^+$ is the return and $\theta > \beta > 0$ (see Marinacci and Montrucchio, 2010, Example 3).

Assuming that the aggregator is unbounded from above, $W : \Gamma \times \mathcal{V} \to \overline{\mathbb{R}}^+$, a natural truncation is given by $W_n(x,z,v) = W(x,z,v) \wedge n$. It is easy to verify that, for all bounded v' and v'' in $\mathcal{V}$ such that $W_n(x,z,v'') < W_n(x,z,v')$,[4]

$$\begin{aligned} W_n(x,z,v') - W_n(x,z,v'') &\leq W(x,z,v') - W(x,z,v'') \\ &\leq \delta \sup_{z\in G(x)} \int \left|v'(y) - v''(y)\right| \Pi(x,z)(dy) \\ &\leq \delta \sup_{x\in X} \left|v'(x) - v''(x)\right|. \end{aligned}$$

Therefore, the property of discounting by increasing truncations is satisfied.

7. Comments

We compare our findings in this paper with the most recent literature on unbounded returns with time-additive aggregators of the form

$$W(x,z,v) = (1-\delta)u(x,z) + \delta \int v(y)\,\Pi(x,z)(dy).$$

Ma et al. (2022) present an innovative approach to unbounded returns based on the idea of a Q-transform. Rincón-Zapatero (2024) studies unbounded returns under minimal assumptions by means of a contraction-type method. We begin with a comment on this latter contribution.

The approach in Rincón-Zapatero (2024) is a slightly more general version of our discounting assumption (Assumption 5.1). In particular, his restrictions imply the existence of a finite upper bound $\bar{f}$ in $\mathcal{F}$ satisfying

$$|u(x,z)| \leq \bar{f}(x) \text{ and } D\left(\bar{f}\right)(x) \leq \bar{f}(x),$$

where $\lim_{n\to\infty} D^n\left(\bar{f}\right) = 0$.[5] This ensures that the Bellman operator, up to measurability, maps $\mathcal{L}\left(\bar{f}\right)$ into itself. In turn, the fixed point is unique in this space under the maintained assumptions. Furthermore, in addition to the assumption of discounting, and irrespectively of unbounded returns, Rincón-Zapatero (2024) decomposes the state space as $X = S \times Z$ and only considers Markov transitions of the form $\Pi : S \to \Delta(S)$, that is, such that the transition moves from state $x = (s^-, z^-)$ into state $y = (s,z)$, given the action z in $G(s^-, z^-)$, and the probabilities are independent of the current action. This allows him to substantially weaken the restrictions of upper semicontinuity for the application of the Maximum Theorem, including the validity of the Feller property for unbounded values.

The idea in Ma et al. (2022) consists in transforming the aggregator to obtain, under additional restrictions, bounded returns. Their translation can be understood as

$$\tilde{W}(x,z,\tilde{v}) = (1-\delta)u(x,z) - \varphi(x) + \delta \int (\varphi(y) + \tilde{v}(y))\,\Pi(x,z)(dy),$$

where $\varphi(x) = \sup_{z\in G(x)}(1-\delta)u(x,z)$. This operation preserves bounded returns as long as there exists a finite η in $\mathbb{R}^+$ such that

$$\delta \left| \int \varphi(y)\,\Pi(x,z)(dy) \right| \leq \eta.$$

Indeed, noticing that $(1-\delta)u(x,z) \leq \varphi(x)$,

$$\sup_{z\in G(x)} \tilde{W}(x,z,\tilde{v}) \leq \sup_{z\in G(x)} \delta \int (\varphi(y) + \tilde{v}(y))\,\Pi(x,z)(dy) \leq \eta + \delta \|\tilde{v}\|_\infty.$$

and

$$\sup_{z\in G(x)} \tilde{W}(x,z,\tilde{v}) \geq \sup_{z\in G(x)} (1-\delta)u(x,z) - \varphi(x) - \eta - \delta\|\tilde{v}\|_\infty \geq -\eta - \delta\|\tilde{v}\|_\infty.$$

Conventional techniques apply to the transformed aggregator and deliver the existence of a unique fixed point.

CRediT authorship contribution statement

G. Bloise: Conceptualization, Formal analysis, Investigation, Methodology, Writing – original draft, Writing – review & editing. **C. Le Van:** Conceptualization, Formal analysis, Investigation, Methodology, Writing – original draft, Writing – review & editing. **Y. Vailakis:** Conceptualization, Formal analysis, Investigation, Methodology, Writing – original draft, Writing – review & editing.

Declaration of competing interest

Declarations of interest: none

Data availability

No data was used for the research described in the article.

Appendix. Proofs

Proof of Proposition 3.1. Notice that, for every policy g in $\mathcal{G}$, $\left(T_g\tilde{v}\right) \leq \tilde{v}$. Therefore, for every x in X, $\left(\left(T_g^n\tilde{v}\right)(x)\right)_{n\in\mathbb{N}}$ in $\overline{\mathbb{R}}$ is a decreasing sequence. As a consequence, we can consider the utility

$$\tilde{U}_g(x) = \lim_{n\to\infty}\left(T_g^n\tilde{v}\right)(x) \leq \tilde{v}(x).$$

This utility satisfies the Koopmans equation, as

$$\tilde{U}_g(x) = \lim_{n\to\infty}\left(T_g^{n+1}\tilde{v}\right)(x) = \lim_{n\to\infty} W\left(x, g(x), \left(T_g^n\tilde{v}\right)\right) = W\left(x, g(x), \tilde{U}_g\right).$$

Furthermore, as the fixed point $\tilde{v}$ in $\mathcal{V}$ admits a policy g in $\mathcal{G}$, we obviously have

$$\tilde{v}(x) = \sup_{g\in\mathcal{G}} \tilde{U}_g(x).$$

This establishes our claim. □

Proof of Proposition 5.1. For every n in $\mathbb{N}$, we let $T_n : \mathcal{V} \to \mathcal{F}$ be the Bellman operator corresponding to truncated aggregator W_n in $\mathcal{W}$. Notice that, as the sequence of truncated aggregators is increasing, $(Tv) \geq \left(T_n v\right)$ for every v in $\mathcal{V}$. We first argue that the limit is indeed a fixed point of the untruncated Bellman operator.

Obviously, $(Tv^*) \geq \left(Tv_n\right) \geq \left(T_n v_n\right) = v_n$, thus yielding $(Tv^*) \geq v^*$. Suppose that, for some x in X, there exists z in $G(x)$ such that

$$W(x,z,v) > v^*(x).$$

By monotone convergence (Assumption 3.2), for any sufficiently large n in $\mathbb{N}$,

$$W\left(x,z,v_n\right) > v^*(x).$$

By monotone convergence of the truncated aggregator, for any sufficiently large n in $\mathbb{N}$, we also obtain

$$W_n\left(x,z,v_n\right) > v^*(x).$$

As $v^*(x) \geq v_n(x)$, this contradicts that $\left(T_n v_n\right)(x) = v_n(x)$, thus confirming our claim.

We now prove that the limit is the least fixed point of the untruncated Bellman operator. Indeed, letting $\underline{v}$ in $\mathcal{V}$ be this least fixed point, for any j in $\mathbb{N}$,

$$\left(T_n^j 0\right) \leq \left(T_n^j \underline{v}\right) \leq \left(T^j \underline{v}\right) = \underline{v}.$$

As the truncated Bellman operator is a contraction, $v_n = \lim_{j\to\infty}\left(T_n^j 0\right) \leq \underline{v}$. This establishes our claim. □

[4] In fact, notice that

$$W_n\left(x,z,v''\right) < W_n\left(x,z,v'\right) \text{ implies } W_n\left(x,z,v''\right) < n.$$

Therefore, we obtain $W_n\left(x,z,v''\right) = W\left(x,z,v''\right)$ and, consequently,

$$W_n\left(x,z,v'\right) - W_n\left(x,z,v''\right) \leq W\left(x,z,v'\right) - W\left(x,z,v''\right),$$

as claimed.

[5] To properly compare with Rincón-Zapatero (2024, Assumption (B6)), see the Appendix.

Proof of Proposition 5.2. Given n in $\mathbb{N}$, by the Maximum Theorem (along with the Measurable Selection Theorem in Brown and Purves, 1973, Corollary 1), there exists a policy g in $\mathcal{G}$ such that

$$U_{n,g}(x) = v_n(x),$$

where $U_{n,g}$ in $\mathcal{V}$ is the truncated utility, that is, the only bounded fixed point of $T_{n,g} : \mathcal{V} \to \mathcal{V}$. We conclude that

$$U_g(x) = \lim_{j\to\infty} \left(T_g^j 0\right)(x) \geq \lim_{j\to\infty} \left(T_{n,g}^j 0\right)(x) = v_n(x),$$

which suffices to establish that $\sup_{g\in\mathcal{G}} U_g(x) \geq v^*(x)$. On the other side, given any policy g in $\mathcal{G}$,

$$\left(T_g v^*\right)(x) \leq v^*(x),$$

which implies that $U_g(x) \leq v^*(x)$ because

$$U_g(x) = \lim_{j\to\infty} \left(T_g^j 0\right)(x) \leq \lim_{j\to\infty} \left(T_g^j v^*\right)(x) \leq v^*(x).$$

This is sufficient to prove our claim. □

Proof of Proposition 5.3. To the end of establishing uniqueness, suppose there exists another fixed point v^{**} in $\mathcal{V} \cap \mathcal{L}(\bar{f})$. Arguing by induction, Assumption 5.1 implies

$$0 \leq v^{**} - v^* \leq \lim_{n\to\infty} D^n\left(\bar{f}\right) = 0,$$

where we have assumed that $0 \leq v^{**} - v^* \leq \bar{f}$ at no loss of generality. To prove the other claims, notice that, by Assumption 5.1,

$$v^* = (T^n 0) + \left((T^n v^*) - (T^n 0)\right) \leq (T^n 0) + D^n\left(v^*\right) \leq (T^n 0) + D^n\left(\bar{f}\right).$$

Therefore, for any sequence $\left(x_j\right)_{j\in\mathbb{N}}$ in X converging to x in X,

$$\begin{aligned}\limsup_{j\to\infty} v^*\left(x_j\right) &\leq \limsup_{j\to\infty} (T^n 0)\left(x_j\right) + \limsup_{j\to\infty} D^n\left(\bar{f}\right)\left(x_j\right)\\ &\leq (T^n 0)(x) + D^n\left(\bar{f}\right)(x).\end{aligned}$$

Taking the limit with respect to n in $\mathbb{N}$, and exploiting Assumption 5.1,

$$\limsup_{j\to\infty} v^*\left(x_j\right) \leq \lim_{n\to\infty} (T^n 0)(x) + \lim_{n\to\infty} D^n\left(\bar{f}\right)(x) = v^*(x),$$

thus proving that v^* in $\mathcal{V}$ is upper semicontinuous. A policy exists because of Brown and Purves (1973, Corollary 1). □

Proof of Lemma 5.1. Notice that, fixing n in $\mathbb{N}$, the map $(x,z) \mapsto W(x,z,v\vee(-n))$ is upper semicontinuous by Assumption 3.1, because $v\vee(-n)$ in $\mathcal{V}$ is a bounded and upper semicontinuous map. Taking the limit with respect to n in $\mathbb{N}$, noticing that the sequence of maps is decreasing by monotonicity and invoking Assumption 3.2 (Monotone convergence), we prove our claim. □

Proof of Proposition 5.4. Obviously, $(Tv^*) \leq \left(Tv_n\right) \leq v_n$, thus yielding $(Tv^*) \leq v^*$. Fixing x in X, notice that, for every n in $\mathbb{N}$, there exists z_n in $G(x)$ such that

$$W_n\left(x, z_n, v_n\right) = v_n(x) \geq v^*(x).$$

By compactness, we can assume that a subsequence $\left(z_{n(j)}\right)_{j\in\mathbb{N}}$ converges to z in $G(x)$. For large enough j in $\mathbb{N}$, by monotonicity, we have

$$W_m\left(x, z_{n(j)}, v_n\right) \geq v^*(x),$$

because $v_n \geq v_{n(j)}$ and $W_m \geq W_{n(j)}$ for fixed n and m in $\mathbb{N}$ when j in $\mathbb{N}$ is sufficiently large. By upper semicontinuity, this implies

$$W_m\left(x, z, v_n\right) \geq v^*(x).$$

By monotone convergence (Assumption 3.2),

$$W_m\left(x, z, v^*\right) \geq v^*(x).$$

And, finally, by convergence of the sequence of truncated aggregators,

$$W\left(x, z, v^*\right) \geq v^*(x),$$

so proving that $(Tv^*) \geq v^*$. Hence, the limit is a fixed point of the untruncated Bellman operator.

We now argue that $v^* \geq \bar{v}$, where $\bar{v}$ in $\mathcal{V}$ is the greatest fixed point of the untruncated Bellman operator, so in fact establishing coincidence. Indeed, for any j in $\mathbb{N}$,

$$\left(T_n^j 0\right) \geq \left(T_n^j \bar{v}\right) \geq \left(T^j \bar{v}\right) = \bar{v}.$$

As the truncated Bellman operator is a contraction, $v_n = \lim_{j\to\infty}\left(T_n^j 0\right) \geq \bar{v}$. This establishes our claim.

Notice that v^* in $\mathcal{V}$ is an upper semicontinuous map, being the limit of a decreasing sequence of upper semicontinuous maps. To establish the existence of an optimal policy, we prove that the map $(x,z) \mapsto W(x,z,v^*)$ is upper semicontinuous (and then apply Brown and Purves, 1973, Corollary 1). To this end, apply Lemma 5.1. □

Proof of Proposition 5.5. Given optimal policy g in $\mathcal{G}$,

$$U_g(x) = \lim_{j\to\infty}\left(T_g^j 0\right)(x) \geq \lim_{j\to\infty}\left(T_g^j v^*\right)(x) = v^*(x).$$

This shows that

$$\sup_{g\in\mathcal{G}} U_g(x) \geq v^*(x).$$

We now prove the opposite property.

Consider

$$\tilde{v}(x) = \sup_{g\in\mathcal{G}} U_g(x).$$

At no loss of generality, we can assume that $\tilde{v}$ is a upper semicontinuous map in $\mathcal{V}$ (if not, just consider its upper semicontinuous envelope). Observe that, for any policy g in $\mathcal{G}$,

$$(T\tilde{v})(x) \geq \left(T_g\tilde{v}\right)(x) \geq \left(T_g U_g\right)(x) = U_g(x),$$

thus implying

$$(T\tilde{v})(x) \geq \sup_{g\in\mathcal{G}} U_g(x).$$

As $(T\tilde{v})$ is upper semicontinuous by the Maximum Theorem, we conclude that $(T\tilde{v}) \geq \tilde{v}$. Therefore, the interval $[\tilde{v}, 0] \subset \mathcal{V}$ is invariant for Bellman operator $T : \mathcal{V} \to \mathcal{F}$ and, by Bloise et al. (2024, Proposition 4), the greatest fixed point $\bar{v}$ in $\mathcal{V}$ fulfills $\bar{v} \geq \tilde{v} \geq v^*$, contradicting Proposition 5.4. □

A comparison with Rincón-Zapatero (2024). Consider

$$D(f) = \sup_{z\in G(x)} \int f(y)\,\Pi(x,z)(dy).$$

As in Rincón-Zapatero (2024, Assumption (B6)), assuming that f in $\mathcal{F}$ is measurable and that the series $\bar{f} = \sum_{n=0}^{\infty} D^n(f)$ is convergent, this gives the existence of a measurable $\bar{f}$ in $\mathcal{F}$ such that $D\left(\bar{f}\right) \leq \bar{f}$. Indeed, letting $\bar{f}_n = \sum_{j=0}^{n} D^j(f)$, by sublinearity,

$$D\left(\bar{f}_n\right) \leq f + D\left(\bar{f}_n\right) \leq f + D(f) + \cdots + D^{n+1}(f) = \bar{f}_{n+1} \leq \bar{f},$$

which yields $D\left(\bar{f}\right) \leq \bar{f}$. Furthermore, again by sublinearity,

$$\bar{f}_n + D^{n+1}\left(\bar{f}_m\right) \leq \bar{f}_{n+m+1} \leq \bar{f},$$

which, taking the limit with respect to m in $\mathbb{N}$, yields

$$D^{n+1}\left(\bar{f}\right) \leq \bar{f} - \bar{f}_n,$$

thus confirming that $\lim_{n\to\infty} D^n\left(\bar{f}\right) = 0$ pointwisely, as required in our Assumption 5.1. It should be noted that Rincón-Zapatero (2024, Assumption (B6)) postulates a stronger form of convergence of the series $\sum_{n=0}^{\infty} D^n(f)$, namely, for every compact set $K \subset X$,

$$\sum_{n=0}^{\infty} \sup_{x\in K} \left|D^n(f)(x)\right| \text{ is finite.}$$

As a consequence, if each $D^n(f)$ in $\mathcal{F}$ is continuous, so is $\bar{f}$ in $\mathcal{F}$. □

References

Alvarez, F., Stokey, N.L., 1998. Dynamic programming with homogeneous functions. J. Econom. Theory 82 (1), 167–189.

Backus, D.K., Routledge, B.R., Zin, S.E., 2005. Exotic preferences for macroeconomics. In: NBER Macroeconomic Annual 2004, vol. 19. MIT Press, Cambridge, London, pp. 319–414.

Balbus, L., 2020. On recursive utilities with non-affine aggregator and conditional certainty equivalent. Econom. Theory 70 (2), 551–577.

Becker, R.A., Boyd, III, J.H., 1997. Capital Theory, Equilibrium Analysis and Recursive Utility. Basil Blackwell Publisher.

Becker, R.A., Rincón-Zapatero, J.P., 2021. Thompson aggregators, Scott continuous Koopmans operators, and least fixed point theory. Math. Social Sci. 112 (5), 84–97.

Becker, R.A., Rincón-Zapatero, J.P., 2023. Distinct solutions to Koopmans' equation for Thompson aggregators represent distinct preference orders. mimeo.

Bertsekas, D.P., 2018. Abstract Dynamic Programming, second ed. Athena Scientific.

Bich, P., Drugeon, J.P., Morhaim, L., 2018. On temporal aggregators and dynamic programming. Econom. Theory 66, 787–817.

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

Bloise, G., Vailakis, Y., 2018. Convex dynamic programming with (bounded) recursive utility. J. Econom. Theory 173, 118–141.

Boyd, III, J.H., 1990. Recursive utility and the Ramsey problem. J. Econom. Theory 50 (2), 326–345.

Brown, L.D., Purves, R., 1973. Measurable selections of extrema. Ann. Statist. 1 (5), 902–913.

Christensen, T.M., 2022. Existence and uniqueness of recursive utilities without boundedness. J. Econom. Theory 200, 105413.

Durán, J., 2000. On dynamic programming with unbounded returns. Econom. Theory 15 (2), 339–352.

Durán, J., 2003. Discounting long run average growth in stochastic dynamic programs. Econom. Theory 22 (2), 395–413.

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

Hansen, L.P., Heaton, J.C., Lee, N., Roussanov, N., 2007. Intertemporal substitution and risk aversion. In: Handbook of Econometrics, vol. 6A, Chapter 61. Elsevier, Amsterdam, pp. 3967–4056.

Hansen, L.P., Sargent, T.J., 1995. Discounted linear exponential quadratic gaussian control. IEEE Trans. Automat. Control 40 (5), 968–971.

Hansen, L.P., Sargent, T.J., 2001. Robust control and model uncertainty. Amer. Econ. Rev. 91 (2), 60–66.

Hansen, L.P., Sargent, T.J., 2008. Robustness. Princeton University Press, Princeton, N.J..

Jaśkiewisz, A., Nowak, A.S., 2011. Discounted dynamic programming with unbounded returns: Application to economic models. J. Math. Anal. Appl. 378, 450–462.

Kamihigashi, T., 2014. Elementary results on solutions to the Bellman equation of dynamic programming: Existence, uniqueness, and convergence. Econom. Theory 56 (2), 251–273.

Klibanoff, P., Marinacci, M., Mukerji, S., 2009. Recursive smooth ambiguity preferences. J. Econom. Theory 144 (3), 930–976.

Koopmans, T.C., 1960. Stationary ordinal utility and impatience. Econometrica 28 (2), 287–309.

Koopmans, T.C., Diamond, P.A., Williamson, R.E., 1964. Stationary utility and time perspective. Econometrica 32 (1), 82–100.

Kreps, D.M., Porteus, E.L., 1978. Temporal resolution of uncertainty and dynamic choice theory. Econometrica 46 (1), 185–200.

Le Van, C., Morhaim, L., 2002. Optimal growth models with bounded or unbounded returns: A unifying approach. J. Econom. Theory 105 (1), 158–187.

Le Van, C., Vailakis, Y., 2005. Recursive utility and optimal growth with bounded or unbounded returns. J. Econom. Theory 123 (2), 187–209.

Lucas, R.E., Stokey, N.L., 1984. Optimal growth with many consumers. J. Econom. Theory 32 (1), 139–171.

Ma, Q., Stachurski, J., Toda, A.A., 2022. Unbounded dynamic programming via the Q-transform. J. Math. Econom. 100, 102652.

Marinacci, M., Montrucchio, L., 2010. Unique solutions for stochastic recursive utilities. J. Econom. Theory 145 (5), 1776–1804.

Marinacci, M., Montrucchio, L., 2019. Unique tarski fixed points. Math. Oper. Res. 44 (4), 1174–1191.

Martins-da-Rocha, V.F., Vailakis, Y., 2010. Existence and uniqueness of a fixed-point for local contractions. Econometrica 78 (3), 1127–1141.

Matkowski, J., Nowak, A., 2011. On discounted dynamic programming with unbounded returns. Econom. Theory 46 (3), 455–474.

Miao, J., 2014. Economic Dynamics in Discrete Time. The MIT Press, Cambridge, London.

Ozaki, H., Streufert, P.A., 1996. Dynamic programming for non-additive stochastic objectives. J. Math. Econom. 25 (4), 391–442.

Ren, G., 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. Theor. Econ. (forthcoming).

Rincón-Zapatero, J.P., Rodríguez-Palmero, C., 2003. Existence and uniqueness of solutions to the Bellman equation in the unbounded case. Econometrica 71 (5), 1519–1555.

Skiadas, C., 2009. Asset Pricing Theory. Princeton Series in Finance, Princeton University Press.

Streufert, P.A., 1990. Stationary recursive utility and dynamic programming under the assumption of biconvergence. Rev. Econom. Stud. 57 (1), 79–97.

Wiszniewska-Matyszkiel, A., Singh, R., 2021. Necessity of the terminal condition in the infinite horizon dynamic optimization problems with unbounded payoff. Automatica 123, 109332.