



Research article

On solutions of the distributional Bellman equation

Julian Gerstenberg*, Ralph Neininger* and Denis Spiegel

Goethe University Frankfurt, Institute of Mathematics, Germany

* **Correspondence:** Email: neinigr@math.uni-frankfurt.de.

Abstract: In distributional reinforcement learning (RL), not only expected returns but the complete return distributions of a policy are taken into account. The return distribution for a fixed policy is given as the solution of an associated distributional Bellman equation. In this note, we consider general distributional Bellman equations and study the existence and uniqueness of their solutions, as well as tail properties of return distributions. We give necessary and sufficient conditions for the existence and uniqueness of return distributions and identify cases of regular variation.

We link distributional Bellman equations to multivariate affine distributional equations. We show that any solution of a distributional Bellman equation can be obtained as the vector of marginal laws of a solution to a multivariate affine distributional equation. This makes the general theory of such equations applicable to the distributional reinforcement learning setting.

Keywords: distributional reinforcement learning; distributional Bellman equation; random difference equation; perpetuity; Markov decision process; regular variation; machine learning

1. Introduction

The objective in RL is to teach an agent that sequentially interacts with an environment to choose “good” actions. For each action, the agent receives an immediate real-valued reward. The rewards are accumulated over time, resulting in the so-called *return*, which describes the overall performance of the agent. Randomness may be involved at all levels of this interaction: in choosing actions, the environment reacting to actions, and/or in the rewards received by the agent. Hence, the return is to be considered random as well. In more classical approaches to RL problems, the randomness is averaged out, and only the expected return is considered when evaluating the performance of an agent. In [1], not only the expectation but the complete distribution of the return was considered, introducing what is now known as distributional RL, see [2].

Mathematically, an RL problem is typically modeled by a Markov decision process (MDP), that is, as a particular type of a discrete-time stochastic control problem. An overview of the classical MDP

theory and its applications is presented in [3]. For more details on distributional RL using notations similar to here, we refer to [2, 4].

For any measurable space \mathcal{X} , we write $\mathcal{P}(\mathcal{X})$ for the set of probability distributions on \mathcal{X} . The distribution (law) of a \mathcal{X} -valued random variable (RV) X is denoted by $\mathcal{L}(X) = \mathbb{P}[X \in \cdot] \in \mathcal{P}(\mathcal{X})$. We also write short $X \sim \mu$, if $\mathcal{L}(X) = \mu \in \mathcal{P}(\mathcal{X})$. In case \mathcal{X} is countable, discrete distributions $\nu \in \mathcal{P}(\mathcal{X})$ can be identified with functions $\nu : \mathcal{X} \rightarrow [0, 1]$ satisfying $\sum_{x \in \mathcal{X}} \nu(x) = 1$. For a random variable X and an event A with $\mathbb{P}[A] > 0$, we write $\mathcal{L}(X|A) \in \mathcal{P}(\mathcal{X})$ for the conditional distribution of X given A .

Let \mathcal{S} and \mathcal{A} be non-empty finite sets. A MDP on states \mathcal{S} and actions \mathcal{A} is a function

$$\rho : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{P}(\mathcal{S} \times \mathbb{R}), \quad (s, a) \mapsto \rho_{(s,a)}.$$

The interpretation is as follows: \mathcal{S} is the set of states an environment can occupy and \mathcal{A} the set of possible actions the agent can perform. If in state $s \in \mathcal{S}$ the agent performs action $a \in \mathcal{A}$, the environment reacts with a (possibly random) successor state S together with a (possibly random) real-valued reward R having joint distribution $(S, R) \sim \rho_{(s,a)} \in \mathcal{P}(\mathcal{S} \times \mathbb{R})$. How an agent chooses its actions can be modeled by a (stationary) *policy*

$$\pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A}), \quad s \mapsto \pi_s.$$

An agent that is in state s and acts according to policy π chooses a (possibly random) action A distributed as $A \sim \pi_s \in \mathcal{P}(\mathcal{A})$.

Suppose an agent acts according to π . Starting at time $t = 0$ from a possibly random state $S^{(0)}$, the following dynamic is defined inductively: at time $t \in \mathbb{N}_0$, the agent finds itself in state $S^{(t)}$, chooses action $A^{(t)} \sim \pi_{S^{(t)}}$ (independent of the past given $S^{(t)}$) and the environment reacts with the next state $S^{(t+1)}$ and immediate reward $R^{(t)}$ jointly distributed as $(S^{(t+1)}, R^{(t)}) \sim \rho_{(S^{(t)}, A^{(t)})}$ (independent of the past given $(S^{(t)}, A^{(t)})$). The resulting stochastic process

$$(S^{(t)}, A^{(t)}, R^{(t)})_{t \in \mathbb{N}_0} \tag{1.1}$$

is called a *Markov reward process* in [2] or a *full MDP* in [4]. To emphasize that the distribution of the stochastic process (1.1) depends on π , we write \mathbb{P}^π instead of \mathbb{P} .

A suitable way to judge the overall performance of the agent is to choose a *discount factor* $\gamma \in (0, 1)$ and to consider the discounted accumulated rewards, the so-called *return*:

$$G = R^{(0)} + \gamma R^{(1)} + \gamma^2 R^{(2)} + \dots = \sum_{t=0}^{\infty} \gamma^t R^{(t)}.$$

The return G is defined as an infinite series of random variables and its existence as a \mathbb{R} -valued random variable is not automatically guaranteed, such as if rewards are (extremely) heavy-tailed. Theorem 1 below provides a complete characterization of when the return G exists almost surely as a \mathbb{R} -valued random variable starting from any state $S^{(0)} = s$. The relation of Theorem 1 to earlier results is discussed in Remarks 2, 3 and 4. In various applications, heavy-tailed reward distributions are not encountered. For instance, if rewards are obtained as a deterministic function of states and actions, the rewards become bounded, and the existence of the return G easily follows from a geometric series argument.

However, unbounded and even heavy-tailed reward distributions are of interest in various applications: fields in which RL approaches have been considered and heavy-tailed distributions play a crucial role include insurance and finance. See [5] and [6] for RL approaches to pricing and trading and [7] for the role of heavy-tailed distributions in that field. Recently, another research area in which heavy-tailed reward distributions have been considered is the study of multi-armed bandit models, which corresponds to an MDP with only a single state, $|\mathcal{S}| = 1$. In [8], the authors report on several studies in this direction, including linear bandits [9, 10], pure exploration [11], Lipschitz-bandits [12], Bayesian optimization [13], and Thompson sampling [14]. Notably, in the latter work, a particular emphasis is placed on α -stable reward distributions, which is an assumption also covered by our analysis, see Theorem 4 and Example 3. It is worth noting that all of these works consider finite horizon non-discounted returns rather than infinite discounted returns. Our analysis offers a theoretical justification for the inclusion of heavy-tailed reward distributions in infinite discounted reward scenarios.

Remark 1. The case $\gamma = 0$ is sometimes considered in RL literature, resulting in $G = R^{(0)}$ and trivializing many of the research questions we investigate here. Although we do not discuss this case further, it should be noted that several parts of our results remain applicable even when $\gamma = 0$.

Suppose the return G exists as a \mathbb{R} -valued RV starting from any state s . Of interest in *distributional policy evaluation* are the distributions of the return starting from given states. The *state(-action) return distributions* are defined by

$$\begin{aligned}\eta_s^\pi &= \mathbb{P}^\pi[G \in \cdot | S^{(0)} = s], \quad s \in \mathcal{S}, \\ \eta_{(s,a)}^\pi &= \mathbb{P}^\pi[G \in \cdot | S^{(0)} = s, A^{(0)} = a], \quad (s, a) \in \mathcal{S} \times \mathcal{A}.\end{aligned}$$

In [2], the collection $\eta^\pi = (\eta_s^\pi)_{s \in \mathcal{S}}$ is called a *return distribution function*. The state(-action) return distributions can be found as solutions to the *distributional Bellman equations*, which we now explain. For $r \in \mathbb{R}$, $\gamma \in (0, 1)$ let $f_{r,\gamma} : \mathcal{P}(\mathbb{R}) \rightarrow \mathcal{P}(\mathbb{R})$ be the map that sends $\nu = \mathcal{L}(X)$ to $f_{r,\gamma}(\nu) = \mathcal{L}(r + \gamma X)$. Using the recursive structure of the return, $G = R^{(0)} + \gamma G'$ with $G' = \sum_{t=0}^{\infty} \gamma^t R^{(t+1)}$, and Markov properties of (1.1), it is seen that state and state-action return distributions are related by

$$\begin{aligned}\eta_s^\pi &= \sum_{a \in \mathcal{A}} \pi_s(a) \eta_{(s,a)}^\pi, \quad s \in \mathcal{S}, \\ \eta_{(s,a)}^\pi &= \int_{\mathcal{S} \times \mathbb{R}} f_{r,\gamma}(\eta_{s'}^\pi) d\rho_{(s,a)}(s', r), \quad (s, a) \in \mathcal{S} \times \mathcal{A}.\end{aligned}$$

Substituting the formulas in one another yields the distributional Bellman equations, which comes in two forms: one for states and one for state-actions:

$$\begin{aligned}\eta_s^\pi &= \sum_{a \in \mathcal{A}} \int_{\mathcal{S} \times \mathbb{R}} \pi_s(a) f_{r,\gamma}(\eta_{s'}^\pi) d\rho_{(s,a)}(s', r), \quad s \in \mathcal{S}, \\ \eta_{(s,a)}^\pi &= \int_{\mathcal{S} \times \mathbb{R}} \sum_{a' \in \mathcal{A}} \pi_{s'}(a') f_{r,\gamma}(\eta_{(s',a')}^\pi) d\rho_{(s,a)}(s', r), \quad (s, a) \in \mathcal{S} \times \mathcal{A}.\end{aligned}\tag{1.2}$$

The distributional Bellman equations are a system of $|\mathcal{S}|$ resp. $|\mathcal{S} \times \mathcal{A}|$ one-dimensional distributional equations in \mathbb{R} . The right hand side of the equations can be used to introduce the *distributional Bellman operator*, which for the state return distributions is defined as:

$$\mathcal{T}^\pi : \mathcal{P}(\mathbb{R})^{\mathcal{S}} \longrightarrow \mathcal{P}(\mathbb{R})^{\mathcal{S}}, \quad \eta = (\eta_s)_{s \in \mathcal{S}} \longmapsto \mathcal{T}^\pi(\eta) = \left(\mathcal{T}_s^\pi(\eta) \right)_{s \in \mathcal{S}}$$

with s -th component function $\mathcal{T}_s^\pi(\eta) = \sum_{a \in \mathcal{A}} \int_{\mathcal{S} \times \mathbb{R}} \pi_s(a) f_{r,\gamma}(\eta_{s'}) d\rho_{(s,a)}(s', r). \quad (1.3)$

No assumption on the MDP ρ nor the policy π is needed to define (1.3), and the relation to the state return distributions is as follows: if the return G , starting from any state s , converges almost surely in \mathbb{R} then the return distribution function $\eta^\pi = (\eta_s^\pi)_{s \in \mathcal{S}}$ is a fixed point of \mathcal{T}^π , that is $\eta^\pi = \mathcal{T}^\pi(\eta^\pi)$. We explore this connection in more detail.

We write $\log^+(x) = \log(x)$ for $x \geq 1$ and $\log^+(x) = 0$ for $x < 1$. The definition of “essential state” is given before Theorem 2. One result of this note is the following:

Theorem 1. *The following are equivalent:*

- (i) \mathcal{T}^π has a fixed point $\eta \in \mathcal{P}(\mathbb{R})^{\mathcal{S}}$,
- (ii) $G = \sum_{t=0}^{\infty} \gamma^t R^{(t)}$ converges $\mathbb{P}^\pi[\cdot | S^{(0)} = s]$ -almost surely in \mathbb{R} for every initial state $s \in \mathcal{S}$,
- (iii) $\int_{\mathcal{S} \times \mathbb{R}} \log^+(|r|) d\rho_{(s,a)}(s', r) < \infty$ for every pair $(s, a) \in \mathcal{S} \times \mathcal{A}$ that is essential with respect to the law of the Markov chain $(S^{(t)}, A^{(t)})_{t \in \mathbb{N}_0}$ under \mathbb{P}^π ,
- (iv) For every $\nu \in \mathcal{P}(\mathbb{R})^{\mathcal{S}}$ as $n \rightarrow \infty$ the sequence $(\mathcal{T}^\pi)^{on}(\nu)$ of n -th iterations of \mathcal{T}^π converges weakly in the product space $\mathcal{P}(\mathbb{R})^{\mathcal{S}}$.

If these hold the fixed point of \mathcal{T}^π is unique, given by the return distribution function η^π and $(\mathcal{T}^\pi)^{on}(\nu) \rightarrow \eta^\pi$ weakly as $n \rightarrow \infty$ for every $\nu \in \mathcal{P}(\mathbb{R})^{\mathcal{S}}$.

A theorem that reads analogously can be formulated for the distributional Bellman operator $\mathcal{T}^\pi : \mathcal{P}(\mathbb{R})^{\mathcal{S} \times \mathcal{A}} \rightarrow \mathcal{P}(\mathbb{R})^{\mathcal{S} \times \mathcal{A}}$ defined in terms of state-actions. We cover both cases simultaneously by introducing simplified notations in Section 1.2 that allow to analyze the distributional Bellman operator, its fixed point equations and its solutions more conveniently. In Section 2, we present our main results in these notations, Theorem 1 will follow from the more general Theorem 2 presented there. For connections to a model and results in [15] see Remark 2.

Besides giving necessary and sufficient conditions for the existence of solutions to the distributional Bellman equations, we also study properties of their solutions, that is, of the return distribution function η^π (and also state-action return distributions). We consider tail probability asymptotics $\eta_s^\pi((-\infty, -x))$ and $\eta_s^\pi((x, \infty))$ as $x \rightarrow \infty$. In Theorem 3, we identify cases of exponential decay and existence of p -th moments and in Theorem 4 we cover regular variation. A possible application of the latter is in distributional dynamic programming (see Chapter 5 in [2]), see Remark 6 of the present note.

Building upon our simplified notation, we explore the connection of the distributional Bellman equations to multivariate distributional equations of the form $\mathbf{X} \stackrel{d}{=} \mathbf{A}\mathbf{X} + \mathbf{B}$ in Section 2.1. This connection seems to have not been noticed in the literature so far, and can be used to apply available results in that field to the distributional RL setting; we do that when proving Theorem 4 by applying results presented in [16].

Before we switch to the simplified notation and present the main results, we shortly present the ordinary Bellman equations for convenience.

1.1. The ordinary Bellman equations

The following is well-known and the standard setting in much of the distributional RL literature: if reward distributions $\rho_{(s,a)}(\mathcal{S} \times \cdot) \in \mathcal{P}(\mathbb{R})$ have finite expectations for all (s, a) , the return converges almost surely and has finite expectations as well; this also follows from Theorems 2 and 3 presented in this note. In many classical approaches to RL, only expected returns are used for policy evaluation. The *state values* and *state-action values* are defined by

$$v_s^\pi = \mathbb{E}^\pi[G|S^{(0)} = s] = \int_{\mathbb{R}} g \, d\eta_s^\pi(g),$$

$$q_{(s,a)}^\pi = \mathbb{E}^\pi[G|S^{(0)} = s, A^{(0)} = a] = \int_{\mathbb{R}} g \, d\eta_{(s,a)}^\pi(g).$$

The collection $v^\pi = (v_s^\pi)_{s \in \mathcal{S}}$ is called a state value function and q^π a state-action value function. They can be found as solutions to the (ordinary) *Bellman equations*, which can be derived from their distributional counterparts (1.2) using linearity of expectation. For action a in state s , let $r_{(s,a)} \in \mathbb{R}$ be the expected reward and $p_{(s,a)} \in \mathcal{P}(\mathcal{S})$ the distribution of the next state, that is

$$r_{(s,a)} = \int r \, d\rho_{(s,a)}(s', r) = \mathbb{E}^\pi[R^{(t)}|S^{(t)} = s, A^{(t)} = a],$$

$$p_{(s,a)} = \rho_{(s,a)}(\cdot \times \mathbb{R}) = \mathbb{P}^\pi[S^{(t+1)} \in \cdot | S^{(t)} = s, A^{(t)} = a].$$

The ordinary Bellman equations are a system of $|\mathcal{S}|$ (resp. $|\mathcal{S} \times \mathcal{A}|$) linear equations in the same number of unknowns, and read as follows:

$$v_s^\pi = \sum_{a \in \mathcal{A}} \pi_s(a) \left[r_{(s,a)} + \gamma \sum_{s' \in \mathcal{S}} p_{(s,a)}(s') v_{s'}^\pi \right], \quad s \in \mathcal{S},$$

$$q_{(s,a)}^\pi = r_{(s,a)} + \gamma \sum_{s' \in \mathcal{S}} p_{(s,a)}(s') \sum_{a' \in \mathcal{A}} \pi_{s'}(a') q_{(s',a')}^\pi, \quad (s, a) \in \mathcal{S} \times \mathcal{A}. \quad (1.4)$$

As in the distributional setting, one can use the right-hand side of these equations to define the (ordinary) Bellman operators ($\mathbb{R}^{\mathcal{S}} \rightarrow \mathbb{R}^{\mathcal{S}}$ for states) such that (1.4) are equivalent to the associated fixed point equation of these operators.

When judging a policy based on expected returns, π is considered to be at least as good as some other policy $\tilde{\pi}$ if $v^\pi \geq v^{\tilde{\pi}}$ pointwise at each s . Classical MDP theory shows that there always exists a policy at least as good as every other policy. Being able to calculate expected returns is not just of interest for evaluating π , but also to improve it: the policy π' defined by $\pi'_s = \delta_{\arg\max_{a \in \mathcal{A}} q_{(s,a)}^\pi}$ is at least as good as π . This leads to an iterative algorithm known as *policy iteration*: starting from an initial policy $\pi^{(0)}$ and letting $\pi^{(k+1)} = (\pi^{(k)})'$ it can be shown that $v^{\pi^{(k)}}$ converges as $k \rightarrow \infty$ to the state value function of an optimal policy.

1.2. Simplified notations

We introduce simplified notations and define the distributional Bellman operator in this setting. Let $d \in \mathbb{N}$ and $(R_1, J_1), \dots, (R_d, J_d)$ be random pairs such that each R_i is real-valued and J_i is discrete taking values in $[d] = \{1, \dots, d\}$. For each $i \in [d]$, it is $\mathcal{L}(R_i, J_i) \in \mathcal{P}(\mathbb{R} \times [d])$ the joint distribution of the

random pair (R_i, J_i) . The random pairs $(R_1, J_1), \dots, (R_d, J_d)$ together with some constant $\gamma \in (0, 1)$ define (a version of) the *distributional Bellman operator* as follows:

$$\mathcal{T} : \mathcal{P}(\mathbb{R})^d \rightarrow \mathcal{P}(\mathbb{R})^d, \quad \eta = \begin{pmatrix} \eta_1 \\ \vdots \\ \eta_d \end{pmatrix} \mapsto \mathcal{T}(\eta) = \begin{pmatrix} \mathcal{T}_1(\eta) \\ \vdots \\ \mathcal{T}_d(\eta) \end{pmatrix},$$

the i -th coordinate function $\mathcal{T}_i : \mathcal{P}(\mathbb{R})^d \rightarrow \mathcal{P}(\mathbb{R})$ being

$$\mathcal{T}_i(\eta) = \mathcal{T}_i(\eta_1, \dots, \eta_d) = \mathcal{L}(R_i + \gamma G_{J_i}), \quad (1.5)$$

where on the right hand side $\mathcal{L}(G_j) = \eta_j$ for each $j \in [d]$ and $G_1, \dots, G_d, (R_i, J_i)$ are independent.

Example 1 (State return). Let ρ be a MDP and π a stationary policy. Let $d = |\mathcal{S}|$ and choose an arbitrary enumeration $i = s$ to identify $[d]$ with \mathcal{S} . Let $i = s, j = s'$ and $B \subseteq \mathbb{R}$ be measurable. Consider the joint distribution

$$\begin{aligned} \mathbb{P}[R_i \in B, J_i = j] &= \mathbb{P}^\pi[R^{(t)} \in B, S^{(t+1)} = s' | S^{(t)} = s] \\ &= \sum_{a \in \mathcal{A}} \pi_s(a) \rho_{(s,a)}(\{s'\} \times B). \end{aligned}$$

In this case \mathcal{T} in (1.5) is the same as \mathcal{T}^π in (1.3).

Example 2 (State-action return). Let ρ be a MDP and π a stationary policy. Let $d = |\mathcal{S} \times \mathcal{A}|$ and choose an arbitrary enumeration $i = (s, a)$ to identify $[d]$ with $\mathcal{S} \times \mathcal{A}$. Let $i = (s, a), j = (s', a')$ and $B \subseteq \mathbb{R}$ measurable. Consider the joint distribution

$$\begin{aligned} \mathbb{P}[R_i \in B, J_i = j] &= \mathbb{P}^\pi[R^{(t)} \in B, S^{(t+1)} = s', A^{(t+1)} = a' | S^{(t)} = s, A^{(t)} = a] \\ &= \rho_{(s,a)}(\{s'\} \times B) \pi_{s'}(a'). \end{aligned}$$

In this case, the distributions $\mathcal{L}(R_i, J_i), i \in [d]$ completely encode both ρ and π and, provided existence, state-action return distributions $(\eta_{(s,a)}^\pi)_{(s,a) \in \mathcal{S} \times \mathcal{A}}$ are fixed points of \mathcal{T} .

The connection to distributional RL shows that it is of interest to understand the fixed points of \mathcal{T} , that is probability distributions $\eta \in \mathcal{P}(\mathbb{R})^d$ satisfying

$$\eta = \mathcal{T}(\eta). \quad (1.6)$$

Let G_1, \dots, G_d be real-valued random variables with $\eta_i = \mathcal{L}(G_i)$. Then, (1.6) is equivalent to the distributional Bellman equations, which is a system of d one-dimensional distributional equations

$$G_i \stackrel{d}{=} R_i + \gamma G_{J_i}, \quad i \in [d], \quad (1.7)$$

where $\stackrel{d}{=}$ (also denoted by $=_d$) denotes equality in distribution and in the i -th equation $G_1, \dots, G_d, (R_i, J_i)$ are independent, $i \in [d]$. For $i, j \in [d]$ let

$$p_{ij} = \mathbb{P}[J_i = j] \quad \text{and} \quad \mu_{ij} = \mathcal{L}(R_i | J_i = j),$$

where in case $p_{ij} = 0$ we set $\mu_{ij} = \delta_0$, the Dirac measure in 0. It holds $\mathbb{P}[R_i \in \cdot, J_i = j] = p_{ij}\mu_{ij}(\cdot)$. Let $F_i(x) = \mathbb{P}[G_i \leq x]$ be the cumulative distribution function (cdf) of real-valued random variable G_i . A third equivalent formulation of the distributional Bellman equations (1.7) in terms of cdf's is

$$F_i(\cdot) = \sum_{j=1}^d p_{ij} \int_{-\infty}^{\infty} F_j\left(\frac{\cdot - r}{\gamma}\right) d\mu_{ij}(r), \quad i \in [d].$$

The distributional Bellman equations in terms of cdfs appeared in the literature before, for instance, in [17] or, in a more specific form, in [18].

We study the following aspects regarding solutions to the distributional Bellman equations, i.e., fixed points of \mathcal{T} :

- Existence and uniqueness for which conditions being necessary and sufficient are given in Theorem 2, see also Remark 3.
- Tail probability asymptotics, that is, if η is a fixed point of \mathcal{T} with i -th component $\eta_i = \mathcal{L}(G_i)$ we study the asymptotic behavior of $\mathbb{P}[G_i > x]$ and $\mathbb{P}[G_i < -x]$ as $x \rightarrow \infty$. Depending on properties of the distributions $\mathcal{L}(R_1, J_1), \dots, \mathcal{L}(R_d, J_d)$, we identify cases of
 - exponential decay and existence of p -th moments, see Theorem 3.
 - regular variation, see Theorem 4.

In Section 2.1, we explain how our findings relate to the area of multivariate distributional fixed point equations of the form

$$\mathbf{X} \stackrel{d}{=} \mathbf{A}\mathbf{X} + \mathbf{B}, \quad (1.8)$$

in which \mathbf{X}, \mathbf{B} are random d -dimensional vectors, \mathbf{A} a random $d \times d$ matrix, and \mathbf{X} and (\mathbf{A}, \mathbf{B}) are independent. Such equations arise in many different applications, such as in the context of perpetuities and random difference equations. In the above notation, the distributional Bellman equations (1.7) are a system of d *one-dimensional distributional equations* which is less restrictive as a *single distributional equation in \mathbb{R}^d* such as (1.8), but more complicated than a single one-dimensional distributional equation. However, using Theorem 2, we show how general results from the multivariate setting become applicable to study distributional Bellman equations, see Corollary 1.

In case $d = 1$, the random pairs $(R_1, J_1), \dots, (R_d, J_d)$ reduce to a single real-valued random variable (R_1, J_1) with $J_1 \equiv 1$ and the distributional Bellman equations (1.7) reduce to a single distributional equation in \mathbb{R} :

$$G_1 \stackrel{d}{=} \gamma G_1 + R_1,$$

with G_1 independent of R_1 . This is a simple special case of a one-dimensional distributional equation $X \stackrel{d}{=} AX + B$, also called perpetuity equation, for which the existence and uniqueness of solutions as well as their properties have been studied in great detail, see, e.g., [19, 20] or [21]. A comprehensive discussion is given in Chapter 2 in [16].

2. Results

Let $\gamma \in (0, 1)$ and $(R_1, J_1), \dots, (R_d, J_d)$ be the $\mathbb{R} \times [d]$ -valued random pairs defining the Bellman operator $\mathcal{T} : \mathcal{P}(\mathbb{R})^d \rightarrow \mathcal{P}(\mathbb{R})^d$ as in (1.5). Some results can be formulated and proven within a special

construction based on random variables

$$(I_t)_{t \in \mathbb{N}_0}, (R_{ijt})_{(i,j) \in [d]^2, t \in \mathbb{N}_0}, \quad (2.1)$$

in which

- $(I_t)_{t \in \mathbb{N}_0}$ is a $[d]$ -valued homogeneous Markov chain having transition probabilities $(p_{ij})_{(i,j) \in [d]^2}$ and a uniform starting distribution, that is, for any path $i_0, i_1, \dots, i_T \in [d]$ it holds that

$$\mathbb{P}[I_0 = i_0, \dots, I_T = i_T] = d^{-1} p_{i_0 i_1} p_{i_1 i_2} \cdots p_{i_{T-1} i_T},$$

- $(R_{ijt})_{(i,j) \in [d]^2, t \in \mathbb{N}_0}$ forms an array of independent real-valued random variables with

$$\mathcal{L}(R_{ijt}) = \mathcal{L}(R_i | J_i = j) = \mu_{ij},$$

in particular, for each pair (i, j) , the sequence $(R_{ijt})_{t \in \mathbb{N}_0}$ is independent identically distributed (iid) with distribution μ_{ij} ,

- $(I_t)_{t \in \mathbb{N}_0}$ and $(R_{ijt})_{(i,j) \in [d]^2, t \in \mathbb{N}_0}$ are independent.

From now on, we omit quantification of indices in array-like quantities once the index-ranges have been introduced.

Using RVs (2.1), an explicit description of the n -th iteration of \mathcal{T} defined inductively as $\mathcal{T}^{on} = \mathcal{T}^{o(n-1)} \circ \mathcal{T}$ can be given: the i -th component of \mathcal{T}^{on} is the map $\mathcal{T}_i^{on} : \mathcal{P}(\mathbb{R})^d \rightarrow \mathcal{P}(\mathbb{R})$ defined by

$$\mathcal{T}_i^{on}(\eta_1, \dots, \eta_d) = \mathcal{L}\left(\sum_{t=0}^{n-1} \gamma^t R_{I_t, I_{t+1}, t} + \gamma^n G_{I_n} \mid I_0 = i\right),$$

with $\mathcal{L}(G_j) = \eta_j$ and $G_1, \dots, G_d, (R_{ijt})_{ijt}, (I_t)_t$ independent.

To introduce some basic definitions from Markov chain theory, let $i, j \in [d]$ be states. We write $i \rightarrow j$ if there exists $t \in \mathbb{N}_0$ with $\mathbb{P}[I_t = j | I_0 = i] > 0$. State i is *essential* if for every state j the implication $i \rightarrow j \Rightarrow j \rightarrow i$ holds. Since the state space $[d]$ is finite, a state i is essential if and only if it is *recurrent*, that is, $\mathbb{P}[I_t = i \text{ for some } t > 0 | I_0 = i] = 1$, which implies $\mathbb{P}[I_t = i \text{ for } \infty\text{-many } t | I_0 = i] = 1$.

Theorem 2. *The following conditions are equivalent:*

- \mathcal{T} has a (unique) fixed point,
- $\mathbb{E}[\log^+ |R_i|] < \infty$ for all essential i ,
- The infinite series $G = \sum_{t=0}^{\infty} \gamma^t R_{I_t, I_{t+1}, t}$ is almost surely (absolutely) convergent,
- $\mathcal{T}_i^{on}(v)$ converges weakly as $n \rightarrow \infty$ for each $i \in [d]$ for some (all) $v \in \mathcal{P}(\mathbb{R})^d$.

If one and hence all of (i)–(iv) hold, then the unique fixed point $\eta \in \mathcal{P}(\mathbb{R})^d$ of \mathcal{T} has i -th component

$$\eta_i = \mathcal{L}(G | I_0 = i) = \mathcal{L}\left(\sum_{t=0}^{\infty} \gamma^t R_{I_t, I_{t+1}, t} \mid I_0 = i\right).$$

Remark 2. In [15], a model of iterations of random univariate affine linear maps within a Markovian environment is introduced and its stability is studied. This model is similar to our version of the distributional Bellman operator in (1.5). In [15], the Markov chain is allowed to have countable state space but required to be ergodic whereas the present Markov chain in the distributional RL setting has finite state space and no further restrictions. The results in Section 3 of [15] partly imply claims of the present Theorem 2 with similar underlying techniques based on [19, 21].

Note that for a real-valued random variable X , the following much stronger properties each imply $\mathbb{E}[\log^+ |X|] < \infty$:

- 1) $|X|$ is bounded: $|X| \leq K$ for some constant $K \geq 0$,
- 2) $|X|$ has a finite exponential moment: $\mathbb{E}[\exp(\beta|X|)] < \infty$ for some $\beta > 0$,
- 3) $|X|$ has a finite p -th moment: $\mathbb{E}[|X|^p] < \infty$ for some $p \in [1, \infty)$.

Of course, we have $1) \Rightarrow 2) \Rightarrow 3)$. Hence, by Theorem 2, a sufficient condition for the existence of a fixed point of \mathcal{T} is that every $R_j, j \in [d]$ satisfies one of the latter three properties. Moreover, supposing \mathcal{T} has a fixed point η with i -component $\eta_i = \mathcal{L}(G_i)$, these properties are transferred from $R_j, j \in [d]$ to G_i , the subsequent theorem presents a concise summary of these well-known facts:

Theorem 3. *Let $i \in [d]$. If for all j with $i \rightarrow j$*

- 1) $|R_j| \leq K$ then $|G_i| \leq \frac{1}{1-\gamma} K$,
- 2) $\mathbb{E}[\exp(\beta|R_j|)] < \infty$ then $\mathbb{E}[\exp(\beta|G_i|)] < \infty$,
- 3) $\mathbb{E}[|R_j|^p] < \infty$ then $\mathbb{E}[|G_i|^p] < \infty$.

The first property transfer is fundamental in MDP theory, while property transfers of the second type are well-established in perpetuity equation theory (cf. Theorem 2.4.1 in [16]). The third is common throughout the Distributional RL literature (cf. Remark 3). For completeness, a short proof of Theorem 3 is presented in Section 3.2.

Remark 3. For $p \geq 1$, let $\mathcal{P}_p(\mathbb{R}) \subset \mathcal{P}(\mathbb{R})$ be the subset of p -integrable distributions on \mathbb{R} , that is, $\mathcal{P}_p(\mathbb{R}) = \{\mu \in \mathcal{P}(\mathbb{R}) \mid \int_{\mathbb{R}} |x|^p d\mu(x) < \infty\}$. Let $\text{pr}_i : \mathbb{R}^d \rightarrow \mathbb{R}$ be the i -th projection. The p -th Wasserstein distance on $\mathcal{P}_p(\mathbb{R})$ is defined as

$$d_p(\mu_1, \mu_2) = \inf \left\{ \int_{\mathbb{R} \times \mathbb{R}} |x - y|^p d\psi(x, y) \mid \psi \in \mathcal{P}(\mathbb{R}^2) \text{ with } \psi \circ \text{pr}_i^{-1} = \mu_i, i = 1, 2 \right\}^{1/p}.$$

The space $(\mathcal{P}_p(\mathbb{R}), d_p)$ is a complete metric space and so is the product $(\mathcal{P}_p(\mathbb{R})^d, d'_p)$ with $d'_p(\eta, \nu) = \max_{i \in [d]} d_p(\eta_i, \nu_i)$. In [1], the following was shown by extending techniques of [22]: if for all $j \in [d]$ it holds that $\mathbb{E}[|R_j|^p] < \infty$, that is $\mathcal{L}(R_j) \in \mathcal{P}_p(\mathbb{R})$, then the restriction $\mathcal{T}_{|\mathcal{P}_p(\mathbb{R})^d}$ of \mathcal{T} to $\mathcal{P}_p(\mathbb{R})^d$ maps to $\mathcal{P}_p(\mathbb{R})^d$ and is a γ -contraction with respect to d'_p . Banach's fixed point theorem yields that the operator $\mathcal{T}_{|\mathcal{P}_p(\mathbb{R})^d} : \mathcal{P}_p(\mathbb{R})^d \rightarrow \mathcal{P}_p(\mathbb{R})^d$ has a unique fixed point $\eta \in \mathcal{P}_p(\mathbb{R})^d$. The latter also follows from our results: if $\mathbb{E}[|R_j|^p] < \infty$ for all j then $\mathbb{E}[\log^+ |R_j|] < \infty$ for all j and, by Theorem 2, there exists a unique fixed point $\eta \in \mathcal{P}(\mathbb{R})^d$. Now $\mathcal{L}(R_j) \in \mathcal{P}_p(\mathbb{R})$ for all j implies $\eta \in \mathcal{P}_p(\mathbb{R})^d \subset \mathcal{P}(\mathbb{R})^d$ by Theorem 3. Note, however, that our Theorem 2 improves upon results of [1] based on the d_p metric twofold: First, Theorem 2 has necessary and sufficient conditions for the existence and uniqueness of fixed points of \mathcal{T} without moment assumption such as $\mathbb{E}[|R_j|^p] < \infty$ for some $p \geq 1$. Second, the uniqueness of the fixed point is shown in the full space $\mathcal{P}(\mathbb{R})^d$, whereas the earlier results show uniqueness only in the subspace $\mathcal{P}_p(\mathbb{R})^d \subset \mathcal{P}(\mathbb{R})^d$.

Remark 4. Theorem 2 shows that, if it exists, the fixed point η of \mathcal{T} can be obtained as the weak limit of the iterates $\mathcal{T}^{on}(\nu)$ as $n \rightarrow \infty$. This was observed in the distributional RL literature before, see Proposition 4.34 in Chapter 4 of [2], and the proof of this part of the theorem shares many similarities with their arguments.

Next, we focus on asymptotic behavior of tail probabilities $\mathbb{P}[X > x]$ and $\mathbb{P}[X < -x]$ as $x \rightarrow \infty$. Each of the three properties stated in Theorem 3 above directly yields asymptotic bounds, in the latter two cases by applying Markov's inequality:

- 1) if $|X| \leq K$ then $\mathbb{P}[X > x] = \mathbb{P}[X < -x] = 0$ for all $x > K$,
- 2) if $\mathbb{E}[\exp(\beta|X|)] < \infty$ for some $\beta > 0$ then

$$\limsup_{x \rightarrow \infty} \frac{\log \mathbb{P}[X > x]}{x} \leq -\beta \quad \text{and} \quad \limsup_{x \rightarrow \infty} \frac{\log \mathbb{P}[X < -x]}{x} \leq -\beta,$$

- 3) if $\mathbb{E}[|X|^p] < \infty$ for some $p \in [1, \infty)$ then $\mathbb{P}[X > x], \mathbb{P}[X < -x]$ are $O(x^{-p})$ as $x \rightarrow \infty$.

With this at hand, in each of the three cases presented in Theorem 3, one can obtain asymptotic bounds for the tail probabilities $\mathbb{P}[G_i > x]$ and $\mathbb{P}[G_i < -x]$ as $x \rightarrow \infty$.

Besides the three classical properties discussed so far, a further interesting and important property of (the distribution of) a random variable X implying $\mathbb{E}[\log^+ |X|] < \infty$ is that of *regular variation*. Regularly varying distributions represent a broad and adaptable class of heavy-tailed distributions, containing several important distribution families as specific cases, see Example 3. As mentioned in the introduction, previous applied works have investigated scenarios that involve heavy-tailed reward distributions. Consequently, an in-depth understanding of how heavy-tailed reward behavior is reflected in the returns is of interest, a question that Theorem 4 answers in the case of regular variation.

Regular variation is directly defined in terms of the asymptotic behavior of the tail probabilities $\mathbb{P}[X > x]$ and $\mathbb{P}[X < -x]$ as $x \rightarrow \infty$. First, a function $f : (0, \infty) \rightarrow [0, \infty)$ with $f(x) > 0$ for all $x > x_0$ is called *regularly varying with index* $\beta \in \mathbb{R}$ if it is measurable and

$$\lim_{x \rightarrow \infty} \frac{f(tx)}{f(x)} = t^\beta \quad \text{for every } t > 0.$$

A function $L : (0, \infty) \rightarrow [0, \infty)$ is called *slowly varying* if it is regularly varying with index $\beta = 0$, so $L(tx)/L(x) \rightarrow 1$ for every $t > 0$. Every regularly varying function f with index β is of the form $f(x) = x^\beta L(x)$ for some slowly varying function L , see [23].

Following [16], a real-valued random variable X is called *regularly varying with index* $\alpha > 0$ if

- a) the function $x \mapsto \mathbb{P}[|X| > x]$ is regularly varying with index $-\alpha$ and
- b) left and right tail probabilities are *asymptotically balanced*: for some $q \in [0, 1]$

$$\lim_{x \rightarrow \infty} \frac{\mathbb{P}[X > x]}{\mathbb{P}[|X| > x]} = q \quad \text{and} \quad \lim_{x \rightarrow \infty} \frac{\mathbb{P}[X < -x]}{\mathbb{P}[|X| > x]} = 1 - q.$$

An equivalent definition reads as follows: a real-valued random variable X is regularly varying with index $\alpha > 0$ if and only if there exists a slowly varying function L and some $q \in [0, 1]$ such that

$$\lim_{x \rightarrow \infty} \frac{\mathbb{P}[X > x]}{x^{-\alpha} L(x)} = q \quad \text{and} \quad \lim_{x \rightarrow \infty} \frac{\mathbb{P}[X < -x]}{x^{-\alpha} L(x)} = 1 - q. \quad (2.2)$$

In this situation, we have $\mathbb{P}[|X| > x] = \mathbb{P}[X > x] + \mathbb{P}[X < -x] \sim x^{-\alpha} L(x)$ as $x \rightarrow \infty$, so $\mathbb{P}[|X| > x]$ is regularly varying of index $-\alpha$, and the tail probabilities are asymptotically balanced.

Similar to the other three properties, regular variation transfers from rewards to returns in the following sense:

Theorem 4. Let $i \in [d]$, $\alpha > 0$ and L a slowly varying function. Suppose for all j with $i \rightarrow j$, there exist $q_j \in [0, 1]$ and $c_j \in [0, \infty)$, with $c_j > 0$ for at least one j , such that

$$\lim_{x \rightarrow \infty} \frac{\mathbb{P}[R_j > x]}{x^{-\alpha} L(x)} = q_j c_j \quad \text{and} \quad \lim_{x \rightarrow \infty} \frac{\mathbb{P}[R_j < -x]}{x^{-\alpha} L(x)} = (1 - q_j) c_j.$$

Then with

$$w_{ij} = \sum_{t=0}^{\infty} (1 - \gamma^\alpha) \gamma^{\alpha t} \mathbb{P}[I_t = j | I_0 = i],$$

that is $w_{ij} = \mathbb{P}[I_N = j | I_0 = i]$ with $N \sim \text{Geo}(1 - \gamma^\alpha)$ independent of $(I_t)_t$, it holds that

$$\lim_{x \rightarrow \infty} \frac{\mathbb{P}[G_i > x]}{x^{-\alpha} L(x)} = \frac{\sum_{j=1}^d w_{ij} q_j c_j}{1 - \gamma^\alpha} \quad \text{and} \quad \lim_{x \rightarrow \infty} \frac{\mathbb{P}[G_i < -x]}{x^{-\alpha} L(x)} = \frac{\sum_{j=1}^d w_{ij} (1 - q_j) c_j}{1 - \gamma^\alpha}. \quad (2.3)$$

In particular, G_i is regularly varying with index α .

Remark 5. For $d = 1$ Theorem 4 is equivalent to Theorem 2.4.3 (2) in [16]. For $d = 1$, the tail asymptotic of $\mathbb{P}[|G_1| > x]$ as $x \rightarrow \infty$ could also be directly obtained from [24], Theorem 2.3.

Example 3. A regularly varying random variable R is called Pareto-like if $\mathbb{P}[R > x] \sim q c x^{-\alpha}$ and $\mathbb{P}[R < -x] \sim (1 - q) c x^{-\alpha}$ for some constants $c > 0, q \in [0, 1], \alpha > 0$, that is, if in (2.2) one can choose L to be constant. Examples of Pareto-like distributions include Pareto, Cauchy, Burr, and α -stable distributions, $\alpha < 2$, see [7]. If there exists a non-empty subset $J \subseteq [d]$ and $\alpha > 0$ such that for each $j \in J$ R_j is Pareto-like with index α and for each $j \notin J$ $\mathbb{P}[|R_j| > x] = o(x^{-\alpha})$, then Theorem 4 applies by choosing $L \equiv 1$ and yields that G_i is also Pareto-like.

Remark 6 (Application in distributional dynamic programming). The goal in distributional dynamic programming is to perform distributional policy evaluation on a computer. A major problem to solve is that probability distributions $\eta_s \in \mathcal{P}(\mathbb{R})$ are infinite dimensional objects and hence some form of approximation has to be applied in practice. Several approaches are discussed in Chapter 5 of [2]. Additional issues arise in cases of Pareto-like tail behaviors with unbounded support as in Example 3: it is then known that (some of) the true state-value distributions are Pareto-like, which may be an information desirable to include in evaluating the policy. Theorem 4 justifies modeling the tails of the state-value distributions a priori as a Pareto-like distribution, the correct choice of asymptotic parameters is given by (2.3). We plan to report on such issues in future work.

2.1. Connection to multivariate fixed point equations

A prevalent method for examining probabilistic behaviors of systems involves coupling techniques, wherein multiple dependent versions of a system of interest are constructed and their collective behavior is analyzed. In this section, we adapt this approach to our specific context, demonstrating how it allows for the application of results from the well-established theory of multivariate distributional fixed point equations in the study of distributional Bellman equations.

A first key insight is that the operator $\mathcal{T} : \mathcal{P}(\mathbb{R})^d \rightarrow \mathcal{P}(\mathbb{R})^d$, defined in (1.5), depends on the random pairs $(R_1, J_1), \dots, (R_d, J_d)$ only through their marginal laws

$$(\mathcal{L}(R_1, J_1), \dots, \mathcal{L}(R_d, J_d)) \in \mathcal{P}(\mathbb{R} \times [d])^d,$$

recall the definition of $\mathcal{T} = (\mathcal{T}_1, \dots, \mathcal{T}_d)$ being

$$\mathcal{T} : \mathcal{P}(\mathbb{R})^d \longrightarrow \mathcal{P}(\mathbb{R})^d, \quad \mathcal{T}(\mathcal{L}(G_1), \dots, \mathcal{L}(G_d)) = (\mathcal{L}(R_1 + \gamma G_{J_1}), \dots, \mathcal{L}(R_d + \gamma G_{J_d})),$$

where in the i -th component the random variables $G_1, \dots, G_d, (R_i, J_i)$ are independent. However, the second key insight is that the independence of G_1, \dots, G_d is not necessary for the definition: any random vector $(\tilde{G}_1, \dots, \tilde{G}_d)$ having the same marginal laws as (G_1, \dots, G_d) and being independent of (R_i, J_i) satisfies $\mathcal{L}(R_i + \gamma \tilde{G}_{J_i}) = \mathcal{L}(R_i + \gamma G_{J_i})$.

By a *coupling* of $(R_1, J_1), \dots, (R_d, J_d)$ we understand any way the random pairs could be defined on a common probability space, thus having a joint distribution

$$\mathcal{L}((R_1, J_1), \dots, (R_d, J_d)) \in \mathcal{P}((\mathbb{R} \times [d])^d). \quad (2.4)$$

A coupling (2.4) induces an operator \mathcal{S} on $\mathcal{P}(\mathbb{R}^d)$ given by

$$\mathcal{S} : \mathcal{P}(\mathbb{R}^d) \longrightarrow \mathcal{P}(\mathbb{R}^d), \quad \mathcal{S}(\mathcal{L}(G_1, \dots, G_d)) = \mathcal{L}(R_1 + \gamma G_{J_1}, \dots, R_d + \gamma G_{J_d}),$$

where on the right hand side (G_1, \dots, G_d) and $((R_1, J_1), \dots, (R_d, J_d))$ are independent. The two key insights mentioned above yield that \mathcal{S} and \mathcal{T} are related by

$$\mathcal{S}(\zeta) \circ \text{pr}_i^{-1} = \mathcal{T}_i(\zeta \circ \text{pr}_1^{-1}, \dots, \zeta \circ \text{pr}_d^{-1}), \quad i \in [d], \zeta \in \mathcal{P}(\mathbb{R}^d), \quad (2.5)$$

with $\text{pr}_i : \mathbb{R}^d \rightarrow \mathbb{R}$ being the i -th coordinate projection and hence $\zeta \circ \text{pr}_i^{-1} \in \mathcal{P}(\mathbb{R})$ the i -th marginal law of $\zeta \in \mathcal{P}(\mathbb{R}^d)$. An immediate consequence of (2.5) of interest to us is

$$\begin{aligned} \mathcal{L}(G_1, \dots, G_d) = \zeta \in \mathcal{P}(\mathbb{R}^d) \text{ is a fixed point of } \mathcal{S} \\ \implies (\mathcal{L}(G_1), \dots, \mathcal{L}(G_d)) = (\zeta \circ \text{pr}_1^{-1}, \dots, \zeta \circ \text{pr}_d^{-1}) \in \mathcal{P}(\mathbb{R})^d \text{ is a fixed point of } \mathcal{T}. \end{aligned} \quad (2.6)$$

Corollary 1 below implies the following converse of this implication: if \mathcal{T} has a (unique) fixed point $(\mathcal{L}(G_1), \dots, \mathcal{L}(G_d))$ then for *any* coupling (2.4) the induced operator \mathcal{S} has a unique fixed point which, due to (2.6), takes the form of a coupling $\mathcal{L}(G_1, \dots, G_d)$ of G_1, \dots, G_d . This justifies the following approach to study (fixed points of) \mathcal{T} : *choose a convenient coupling (2.4) and study its induced operator \mathcal{S} , in particular the marginal laws of a fixed point of \mathcal{S} .* We use this approach proving Theorem 4, where it is most convenient to consider $(R_1, J_1), \dots, (R_d, J_d)$ to be independent.

One advantage of this coupling approach is that it enables access to numerous results from the field of multivariate fixed point equations, as we explain in the following.

Let \mathcal{S} be induced by a coupling (2.4). Define the random d -dimensional vector

$$\mathbf{R} = (R_1, \dots, R_d)^T$$

and the random $d \times d$ matrix

$$\mathbf{J} = (\mathbf{J}_{ij})_{(i,j) \in [d] \times [d]} \text{ with } \mathbf{J}_{ij} = \gamma \cdot 1(J_i = j).$$

The coupling is then completely encoded in the joint distribution $\mathcal{L}(\mathbf{J}, \mathbf{R}) \in \mathcal{P}(\mathbb{R}^{d \times d} \times \mathbb{R}^d)$. Moreover, for any random vector $\mathbf{G} = (G_1, \dots, G_d)^T$ it holds

$$\mathbf{JG} + \mathbf{R} = (R_1 + \gamma G_{J_1}, \dots, R_d + \gamma G_{J_d})^T$$

and hence, if \mathbf{G} is independent of (\mathbf{J}, \mathbf{R}) , it holds that

$$\mathcal{S}(\mathcal{L}(\mathbf{G})) = \mathcal{L}(\mathbf{J}\mathbf{G} + \mathbf{R})$$

and $\mathcal{L}(\mathbf{G}) \in \mathcal{P}(\mathbb{R}^d)$ is a fixed point of \mathcal{S} if and only if

$$\mathbf{G} \stackrel{d}{=} \mathbf{J}\mathbf{G} + \mathbf{R}. \quad (2.7)$$

Note that (2.7) is a single distributional equation in \mathbb{R}^d , whereas the distributional Bellman equations, see (1.7), are a system of d distributional equations in \mathbb{R} . Multivariate fixed point equations like (2.7) have long been studied for general joint laws $\mathcal{L}(\mathbf{J}, \mathbf{R}) \in \mathcal{P}(\mathbb{R}^{d \times d} \times \mathbb{R}^d)$ under various types of assumptions, some of them directly applicable to the situation presented here. A comprehensive overview of the theory of multivariate distributional fixed point equations and its applications, in particular to stochastic difference equations and many important stochastic time-series models, is presented in [16]. Theorem 2 can be used to show the following results, in which $(\mathbf{J}^{(t)}, \mathbf{R}^{(t)})_{t \in \mathbb{N}_0}$ are iid copies of (\mathbf{J}, \mathbf{R}) .

Corollary 1. *For any operator \mathcal{S} induced by a coupling (2.4), the following statements are equivalent*

- (i) \mathcal{S} has a (unique) fixed point,
- (ii) $\mathcal{S}^n(\zeta)$ converges weakly as $n \rightarrow \infty$ for some (all) $\zeta \in \mathcal{P}(\mathbb{R}^d)$,
- (iii) The infinite series $\mathbf{G} = \sum_{t=0}^{\infty} \left[\prod_{s=0}^{t-1} \mathbf{J}^{(s)} \right] \mathbf{R}^{(t)}$ is almost surely (absolutely) convergent,
- (iv) \mathcal{T} has a (unique) fixed point,
- (v) $\mathbb{E}[\log^+ |R_i|] < \infty$ for all essential i .

If one, and hence all, of these hold, then the unique fixed point $\zeta \in \mathcal{P}(\mathbb{R}^d)$ of \mathcal{S} is given by the law of the infinite vector-valued series $\mathcal{L}(\mathbf{G})$ in (iii) and the unique fixed point of \mathcal{T} is the vector of marginal laws of ζ .

Before we present a short proof, we explain how this result fits the known theory of multivariate fixed points equations. Let $|\cdot|$ be the euclidean norm on \mathbb{R}^d . We also write $|\cdot|$ for the induced operator norm on $\mathbb{R}^{d \times d}$. A result by [25] shows that for any joint law $\mathcal{L}(\mathbf{J}, \mathbf{R}) \in \mathcal{P}(\mathbb{R}^{d \times d} \times \mathbb{R}^d)$ such that $\mathbb{E}[\log^+ |\mathbf{J}|], \mathbb{E}[\log^+ |\mathbf{R}|] < \infty$ and such that (the distribution of) \mathbf{J} has a strictly negative top Lyapunov exponent, that is

$$\inf_{t \geq 1} \frac{1}{t} \mathbb{E}[\log \left| \prod_{s=0}^{t-1} \mathbf{J}^{(s)} \right|] < 0,$$

the sum in (iii) converges almost surely and its law is the unique solution to (2.7). We refer to Appendix E in [16] for more information about top Lyapunov exponents. In our case, having $\mathbf{J}_{ij} = \gamma 1(J_i = j)$, and hence $|\prod_{s=0}^{t-1} \mathbf{J}^{(s)}| = \gamma^t$, the top Lyapunov exponent of \mathbf{J} equals $\log(\gamma) < 0$, so it is strictly negative, and $|\mathbf{J}| = \gamma$ implies $\mathbb{E}[\log^+ |\mathbf{J}|] = 0 < \infty$. Thus, the implication (v) \Rightarrow (iii) in Corollary 1 says that the condition $\mathbb{E}[\log^+ |\mathbf{R}|] < \infty$, which is equivalent to $\mathbb{E}[\log^+ |R_i|] < \infty$ for all $i \in [d]$, can be relaxed to $\mathbb{E}[\log^+ |R_i|] < \infty$ for all essential i to obtain almost sure convergence of the infinite series in (iii) in our case.

The implication (ii) \Rightarrow (iii) can directly be obtained from Theorem 2.1 in [26] since $|\prod_{s=0}^{t-1} \mathbf{J}^{(s)}| = \gamma^t \rightarrow 0$ (almost surely) in our case. To deduce almost sure convergence from weak convergence in Theorem 2 (implication (ii) \Rightarrow (iii) there), we present Proposition 1 in Section 3, the proof of which is based on ideas presented in [27] to prove a version of Kolmogorov's three Series theorem also involving

distributional convergence.

Besides the top Lyapunov exponent another important notion in multivariate fixed point equations is that of *irreducibility*: the joint law $\mathcal{L}(\mathbf{J}, \mathbf{R}) \in \mathcal{P}(\mathbb{R}^{d \times d} \times \mathbb{R}^d)$ is called irreducible if the only affine linear subspace $\mathcal{H} \subseteq \mathbb{R}^d$ that fulfills $\mathbb{P}[\mathbf{J}\mathcal{H} + \mathbf{R} \subseteq \mathcal{H}] = 1$ is the complete subspace $\mathcal{H} = \mathbb{R}^d$. In case (\mathbf{J}, \mathbf{R}) is irreducible, the implication (i) \Rightarrow (iii) can be obtained directly from Theorem 2.4 in [28]. In our situation, irreducibility of $\mathcal{L}(\mathbf{J}, \mathbf{R})$ is not given in any situation and Corollary 1 shows that it is also not needed to obtain the implication (i) \Rightarrow (iii).

Proof of Corollary 1. (iv) \iff (v), see Theorem 2.

(iv) \implies (iii). By Theorem 2, the infinite series $\sum_{t=0}^{\infty} \gamma^t R_{I_t, I_{t+1}, t}$ is almost surely (absolute) convergent. For each $i \in [d]$, the i -th component of the vector-valued infinite series $\sum_{t=0}^{\infty} \left[\prod_{s=0}^{t-1} \mathbf{J}^{(s)} \right] \mathbf{R}^{(t)}$ has the same law as $\mathcal{L}\left(\sum_{t=0}^{\infty} \gamma^t R_{I_t, I_{t+1}, t} | I_0 = i\right)$, hence every component converges almost surely (absolutely).

(iii) \implies (ii). Let $\zeta = \mathcal{L}(\mathbf{G}_0)$ with \mathbf{G}_0 independent of $(\mathbf{J}^{(t)}, \mathbf{R}^{(t)})_t$. The n -th iteration of \mathcal{S} can be represented as

$$\mathcal{S}^{on}(\zeta) = \mathcal{L}\left(\sum_{t=0}^{n-1} \left[\prod_{s=0}^{t-1} \mathbf{J}^{(s)} \right] \mathbf{R}^{(t)} + \left[\prod_{s=0}^{n-1} \mathbf{J}^{(s)} \right] \mathbf{G}_0\right).$$

Now $\left[\prod_{s=0}^{n-1} \mathbf{J}^{(s)} \right] \mathbf{G}_0 \rightarrow 0$ almost surely and $\sum_{t=0}^{n-1} \left[\prod_{s=0}^{t-1} \mathbf{J}^{(s)} \right] \mathbf{R}^{(t)}$ converges almost surely by assumption. So $\mathcal{S}^{on}(\zeta)$ converges in distribution, namely to the law of the infinite sum in (iii).

(ii) \implies (i). Assume $\mathcal{S}^{on}(\zeta) \rightarrow \zeta_0$ converges weakly for some ζ . The map $\mathcal{S} : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathcal{P}(\mathbb{R}^d)$ is continuous with respect to the topology of weak convergence, hence $\mathcal{S}(\zeta_0) = \lim_n \mathcal{S}(\mathcal{S}^{on}(\zeta)) = \lim_n \mathcal{S}^{o(n+1)}(\zeta) = \zeta_0$, so ζ_0 is a fixed point of \mathcal{S} .

(i) \implies (iv). See (2.5). □

3. Proofs

Let $m \in \mathbb{N}$ and $\gamma \in (0, 1)$ be fixed. Consider random variables

$$(K_t)_{t \in \mathbb{N}_0}, (X_{lt})_{l \in [m], t \in \mathbb{N}_0} \tag{3.1}$$

such that $(K_t)_t$ is an $[m] = \{1, \dots, m\}$ -valued homogeneous Markov chain with an arbitrary starting distribution, $(X_{lt})_t$ is an array of independent real-valued random variables such that for each $l \in [m]$ the variables $(X_{lt})_t$ are iid and $(X_{lt})_t$ and $(K_t)_t$ are independent. A state l is called *accessible* if $\mathbb{P}[K_t = l] > 0$ for some $t \in \mathbb{N}_0$. Note that we allow any starting distribution in our formulation, hence this is not automatically satisfied. A state l is *accessible essential* if and only if $\mathbb{P}[K_t = l \text{ for infinitely many } t \in \mathbb{N}_0] > 0$ and this probability equals 1 if conditioned on the event of at least one visit in l . In light of Theorem 2, which we prove in the next subsection, any result shown within the setting (3.1) can be applied to (2.1) by considering

$$[m] \simeq [d]^2, l = (i, j), X_{lt} = R_{ijt}, K_t = (I_t, I_{t+1})$$

and in this case working under the condition $\mathbb{P}[\cdot | I_0 = i]$ is equivalent to consider the starting distribution $\mathbb{P}[K_0 = (i', j)] = 1(i' = i)p_{ij}$.

3.1. Existence and uniqueness: Proof of Theorem 2

The proof of Theorem 2 is based on the following proposition, which shares similarities with Kolmogorov's three-series theorem presented Theorem 5.18 in [27].

Proposition 1. For $n \in \mathbb{N}$ let $S_n = \sum_{t=0}^{n-1} \gamma^t X_{K_t,t}$. Then the following are equivalent

- (i) S_n converges in distribution as $n \rightarrow \infty$.
- (ii) S_n converges almost surely as $n \rightarrow \infty$.
- (iii) $\mathbb{E}[\log^+ |X_{l,t}|] < \infty$ for each accessible essential state $l \in [L]$.

Proof. Every finite state space Markov chain with probability one finally takes values in one of its irreducible components, hence to show (iii) \Leftrightarrow (ii), we can reduce to the case that $(K_t)_t$ is irreducible, in particular, every state is accessible essential in this case.

(iii) \Rightarrow (ii): We assume $\mathbb{E}[\log^+ |X_{l,t}|] < \infty$ for all l . Let $0 < c < \log(1/\gamma)$ and $\tilde{c} = \log(1/\gamma) - c > 0$. For each $t \in \mathbb{N}_0$, let

$$A_t = \{\gamma^t |X_{K_t,t}| \geq \exp(-ct)\} = \{\log^+ |X_{K_t,t}| \geq \tilde{c}t\}.$$

We show $\sum_{t=0}^{\infty} \mathbb{P}[A_t] < \infty$ and apply the Borel–Cantelli lemma. Conditioning on $K_t = l$ and bounding $\mathbb{P}[K_t = l] \leq 1$ yields

$$\begin{aligned} \sum_{t=0}^{\infty} \mathbb{P}[A_t] &= \mathbb{P}[A_0] + \sum_{t=1}^{\infty} \mathbb{P}\left[\frac{\log^+ |X_{K_t,t}|}{\tilde{c}} \geq t\right] \leq \mathbb{P}[A_0] + \sum_{l=1}^L \sum_{t=1}^{\infty} \mathbb{P}\left[\frac{\log^+ |X_{l,t}|}{\tilde{c}} \geq t\right] \\ &\leq \mathbb{P}[A_0] + \sum_{l=1}^L \mathbb{E}\left[\frac{\log^+ |X_{l,t}|}{\tilde{c}}\right] < \infty. \end{aligned}$$

The Borel–Cantelli lemma implies that $A_t^c = \{\gamma^t |X_{K_t,t}| < \exp(-ct)\}$ occurs for all but finitely many t almost surely, and hence, almost surely $\sum_{t=0}^{\infty} \gamma^t |X_{K_t,t}| < \infty$, since the tail of the infinite sum is a.s. finally dominated by that of the geometric series $\sum_t \exp(-ct) < \infty$. Hence, S_n is absolutely convergent and thus is convergent almost surely.

(ii) \Rightarrow (iii): We assume there is an accessible essential state l such that $\mathbb{E}[\log^+ |X_{l,t}|] = \infty$ and show that S_n diverges almost surely. We closely follow the idea in the proof of [16], Theorem 2.1.3 (also [19], proof of Lemma 1.7) and apply the root test, which states that for any real-valued sequence $(a_t)_{t \in \mathbb{N}_0}$

$$\limsup_{t \in \mathbb{N}_0} |a_t|^{1/t} > 1 \implies \sum_{t=0}^{\infty} a_t \text{ diverges,}$$

where “diverges” means not having a finite limit. Let $\tilde{C} > 1$, $C = \tilde{C}/\gamma > 1$ and

$$\begin{aligned} B_t &= \left\{ \left| \gamma^t X_{K_t,t} \right|^{1/t} > \tilde{C} \right\} \\ &= \left\{ |X_{K_t,t}|^{1/t} > C \right\} \\ &= \left\{ \log^+ |X_{K_t,t}| > t \log(C) \right\}. \end{aligned}$$

If we can show $\mathbb{P}[\limsup_{t \in \mathbb{N}_0} B_t] = 1$, the root test applies and yields the almost sure divergence of S_n and hence not (ii). We want to apply the second Borel–Cantelli lemma, but since the events B_t , $t \in \mathbb{N}_0$

are not independent, we have to take a little higher effort. For $l \in [m]$, $j \in \mathbb{N}$, let T_{lj} be the time of the j -th visit of $(K_t)_t$ in l with $T_{lj} = \infty$ in case no such visit occurs. With this, it holds that

$$\limsup_{t \in \mathbb{N}_0} B_t = \bigcup_{l=1}^m \limsup_{j \in \mathbb{N}} \{\log^+ |X_{l,T_{lj}}| > T_{lj} \log(C)\}.$$

For calculating the probability of the right-hand side, we can replace $X_{l,T_{lj}}$ with X_{lj} since $(X_t)_t$ and $(T_{lj})_{lj}$ are independent and $(X_t)_t$ is iid for each l . Therefore, by defining the event

$$C_{lj} = \{\log^+ |X_{lj}| > T_{lj} \log(C)\}$$

we have

$$\mathbb{P}[\limsup_{t \in \mathbb{N}_0} B_t] = \mathbb{P}\left[\bigcup_{l'=1}^L \limsup_{j \in \mathbb{N}} C_{l'j}\right] \geq \mathbb{P}[\limsup_{j \in \mathbb{N}} C_{lj}],$$

where on the right hand side the state l is such that $\mathbb{E}[\log^+ |X_t|] = \infty$ which exists by assumption. We show that the right-hand side is one. Using independence of $(X_t)_t$ and $(T_{lj})_j$ Fubini's theorem implies

$$\mathbb{P}[\limsup_{j \in \mathbb{N}} C_{lj}] = \int \mathbb{P}\left[\limsup_{j \in \mathbb{N}} \{\log^+ |X_{lj}| > t_{lj} \log(C)\}\right] d\mathbb{P}^{(T_{lj})_j}((t_{lj})_j), \quad (3.2)$$

where the integral on the right hand side is with respect to the distribution of $(T_{lj})_j$ under \mathbb{P} . Since $(K_t)_t$ is assumed to be an irreducible Markov chain it holds that $\frac{1}{j}T_{lj} \rightarrow c_l > 0$ almost surely as $j \rightarrow \infty$ (where c_l is the inverse of the unique stationary distribution probability for state l). Let $(t_{lj})_{lj}$ be a realization of $(T_{lj})_{lj}$ with $t_{lj} \sim jc_l$ as $j \rightarrow \infty$. If we can show that for any such realization it holds that $\sum_{j=1}^{\infty} \mathbb{P}[\log^+ |X_{lj}| > t_{lj} \log(C)] = \infty$, we can apply the second Borel–Cantelli lemma to conclude $\mathbb{P}[\limsup_{j \in \mathbb{N}} \{\log^+ |X_{lj}| > t_{lj} \log(C)\}] = 1$. Formula (3.2) and the fact that almost every realization $(t_{lj})_j$ of $(T_{lj})_j$ satisfies $t_{lj} \sim jc_l$, then would yield $\mathbb{P}[\limsup_{j \in \mathbb{N}} C_{lj}] = 1$ as desired.

Since $t_{lj} \sim jc_l$ for every $\varepsilon > 0$ there is some j_0 such that $t_{lj} < (1 + \varepsilon)jc_l$ for every $j > j_0$ which allows to conclude

$$\sum_{j=1}^{\infty} \mathbb{P}[\log^+ |X_{lj}| > (1 + \varepsilon)jc_l \log(C)] = \infty \implies \sum_{j=1}^{\infty} \mathbb{P}[\log^+ |X_{lj}| > t_{lj} \log(C)] = \infty$$

With $C' = (1 + \varepsilon)c_l \log(C) > 0$, we have

$$\sum_{j=1}^{\infty} \mathbb{P}[\log^+ |X_{lj}| > (1 + \varepsilon)jc_l \log(C)] \geq \mathbb{E}\left[\left\lfloor \frac{\log^+ |X_{lj}|}{C'} \right\rfloor\right] = \infty,$$

since by assumption $\mathbb{E}[\log^+ |X_{lj}|] = \infty$.

(ii) \implies (i). Obvious

(i) \implies (iii). We show that if there is some accessible essential state l with $\mathbb{E}[\log^+ |X_{l,t}|] = \infty$, then S_n does not converge in distribution. For that, we use two technical Lemmas from [27] which were used to prove a version of Kolmogorov's three series theorem. Let $(X'_t)_t$ be an array of independent RVs, independent of $(X_t)_t$, $(K_t)_t$ and distributed as $(X_t)_t$. Let $\tilde{X}_t = X_t - X'_t$. In [27] \tilde{X}_t is called *symmetrization* of X_t . Let $S'_n = \sum_{t=0}^{n-1} \gamma^t X'_{K_t,t}$ and $\tilde{S}_n = S_n - S'_n = \sum_{t=0}^{n-1} \gamma^t \tilde{X}_{K_t,t}$. Note that \tilde{S}_n is not a

symmetrization of S_n , because S_n, S'_n are not independent as they both depend on $(K_t)_t$. But for any fixed realization $(k_t)_t$ we have $\sum_{t=0}^{n-1} \gamma^t \tilde{X}_{k_t,t}$ a symmetrization of $\sum_{t=0}^{n-1} \gamma^t X_{k_t,t}$, which is enough for our reasoning.

From $\mathbb{E}[\log^+ |X_{t,t}|] = \infty$, one can conclude that $\mathbb{E}[\log^+ |\tilde{X}_{t,t}|] = \infty$, and hence the already established equivalence (ii) \iff (iii) shows that \tilde{S}_n does not converge almost surely, in fact we saw that \tilde{S}_n diverges almost surely. Thus, using independence of $(K_t)_t$ and $(\tilde{X}_{kt})_{kt}$, Fubini's theorem implies

$$1 = \mathbb{P}[\tilde{S}_n \text{ diverges as } n \rightarrow \infty] = \int \mathbb{P}\left[\sum_{t=0}^{n-1} \gamma^t \tilde{X}_{k_t,t} \text{ diverges as } n \rightarrow \infty\right] d\mathbb{P}^{(K_t)_t}((k_t)_t),$$

hence for almost every realization $(k_t)_t$ of $(K_t)_t$ the sequence $\sum_{t=0}^{n-1} \gamma^t \tilde{X}_{k_t,t}$ diverges almost surely. Theorem 5.17 in [27] about sums of independent symmetric random variables applies and shows that $|\sum_{t=0}^{n-1} \gamma^t \tilde{X}_{k_t,t}| \rightarrow \infty$ in probability for almost every $(k_t)_t$. Hence, for almost every realization $(k_t)_t$ and any constant $C > 0$ it holds that

$$\lim_{n \rightarrow \infty} \mathbb{P}\left[\left|\sum_{t=0}^{n-1} \gamma^t \tilde{X}_{k_t,t}\right| > C\right] = 1.$$

Since $\sum_{t=0}^{n-1} \gamma^t \tilde{X}_{k_t,t}$ is a symmetrization of $\sum_{t=0}^{n-1} \gamma^t X_{k_t,t}$, Lemma 5.19 in [27] applies and yields the following bound

$$\mathbb{P}\left[\left|\sum_{t=0}^{n-1} \gamma^t \tilde{X}_{k_t,t}\right| > C\right] \leq 2\mathbb{P}\left[\left|\sum_{t=0}^{n-1} \gamma^t X_{k_t,t}\right| > C\right].$$

Hence, for almost every realization $(k_t)_t$ of $(K_t)_t$ it holds that

$$\liminf_{n \rightarrow \infty} \mathbb{P}\left[\left|\sum_{t=0}^{n-1} \gamma^t X_{k_t,t}\right| > C\right] \geq \frac{1}{2}.$$

Fatou's lemma together with Fubini's theorem yields

$$\liminf_{n \rightarrow \infty} \mathbb{P}\left[|S_n| > C\right] \geq \frac{1}{2}.$$

This holds for every $C > 0$ and hence the sequence $\mathcal{L}(S_n), n \in \mathbb{N}$ is not tight which shows that it does not converge weakly. \square

Proof of Theorem 2. Using the special construction let

$$S_n = \sum_{t=0}^{n-1} \gamma^t R_{I_t, I_{t+1}, t}.$$

For any $\nu \in \mathcal{P}(\mathbb{R})^d$ with i -th component $\nu_i = \mathcal{L}(W_i)$, let $W_1, \dots, W_d, (R_{ijt})_{ijt}, (I_t)_t$ be independent. Then the representation

$$\mathcal{T}_i^{\circ n}(\nu) = \mathcal{L}(S_n + \gamma^n W_{I_n} | I_0 = i)$$

holds. (i) \iff (iv). Let $\eta \in \mathcal{P}(\mathbb{R})^d$ be a fixed point of \mathcal{T} , hence with i -th component $\eta_i = \mathcal{L}(G_i)$ it holds that

$$\eta_i = \mathcal{T}_i^{\circ n}(\eta) = \mathcal{L}(S_n + \gamma^n G_{I_n} | I_0 = i).$$

Since $\gamma^n G_{I_n} \rightarrow 0$ almost surely, this implies that conditioned on $I_0 = i$ the sequence S_n converges to η_i in distribution. In particular, if a fixed point exists, its i -th component has to be the distributional limit of S_n under $I_0 = i$, and thus, fixed points are unique. If a fixed point exists, then $\mathcal{T}^{\circ n}(\nu)$ converges to the unique fixed point η for any ν , since $\gamma^n W_{I_n} \rightarrow 0$ almost surely holds for any random variables W_1, \dots, W_d . Now suppose that $\mathcal{T}^{\circ n}(\nu) \rightarrow \eta$ holds for some ν . The map \mathcal{T} is continuous in the product topology on $\mathcal{P}(\mathbb{R}^d)$ and hence $\mathcal{T}(\eta) = \mathcal{T}(\lim_n \mathcal{T}^{\circ n}(\nu)) = \lim_n \mathcal{T}(\mathcal{T}^{\circ n}(\nu)) = \lim_n \mathcal{T}^{\circ(n+1)}(\nu) = \eta$. So a fixed point exists which finished the equivalence (i) \iff (iv).

(iv) \implies (iii). Supposing (iv) yields that conditioned on $I_0 = i$ the sequence S_n converges in distribution. Applying Proposition 1 by considering $l = (i, j)$ and $X_{lt} = R_{ijt}$ and $K_t = (I_t, I_{t+1})$ with $\mathbb{P}[K_0 = (i', j)] = 1(i' = i)p_{ij}$ yields the almost sure convergence. Since it holds conditioned for every i , it also holds unconditioned, hence (iii).

(iii) \implies (iv). Almost sure convergence holds also conditioned on $I_0 = i$, one can now apply Proposition 1.

(iii) \iff (ii). Again, we apply the equivalence of (ii) and (iii) from Proposition 1. We only note that in the situation $l = (i, j)$ and $K_t = (I_t, I_{t+1})$ a state $l = (i, j)$ is essential with respect to $(K_t)_t$ iff i is essential with respect to $(I_t)_t$ and $p_{ij} > 0$. Also any state l with $p_{ij} > 0$ is accessible when starting in $I_0 = i$. Hence, one obtains that the infinite sum (iii) in Theorem 2 converges almost surely iff $\mathbb{E}[\log^+ |R_{ijt}|] < \infty$ for each essential pair (i, j) , that is for each essential i and j with $p_{ij} > 0$. Since we defined $\mu_{ij} = \delta_0$ for $p_{ij} = 0$ this is equivalent to $\mathbb{E}[\log^+ |R_i|] < \infty$ for every essential i . \square

3.2. Property transfer: Proof of Theorem 3

Proposition 2. *Let $\beta > 0$. If $\mathbb{E}[\exp(\beta|X_{lt}|)] < \infty$ for each accessible state l then the infinite sum $S = \sum_{t=0}^{\infty} \gamma^t X_{K_t,t}$ exists almost surely and it holds that $\mathbb{E}[\exp(\beta|S|)] < \infty$.*

Proof. The infinite sum exists almost surely due to Proposition 1, as the existence of exponential moments for each accessible state clearly yields the existence of the logarithmic moment. Using the triangle inequality, that $x \mapsto \exp(\beta x)$ is continuous non-decreasing and monotone convergence of expectations yields

$$\mathbb{E}[\exp(\beta|S|)] \leq \lim_{n \rightarrow \infty} \mathbb{E}[\exp(\beta \sum_{t=0}^n \gamma^t |X_{K_t,t}|)].$$

Conditioning on (K_0, \dots, K_n) and using independence of $(K_t)_t$ and $(X_{lt})_t$ yields

$$\mathbb{E}[\exp(\beta \sum_{t=0}^n \gamma^t |X_{K_t,t}|)] \leq \mathbb{E}\left[\prod_{t=0}^n \mathbb{E}[\exp(\beta|X_{lt}|)^{\gamma^t}]_{l=K_t}\right]. \quad (3.3)$$

For each $t \in \mathbb{N}_0$, we have $\gamma^t \in (0, 1)$, and hence $x \mapsto x^{\gamma^t}$ is a concave function on $[0, \infty)$. For each accessible state l we have

$$\mathbb{E}[\exp(\beta|X_{lt}|)^{\gamma^t}] \leq \mathbb{E}[\exp(\beta|X_{lt}|)]^{\gamma^t} < \infty.$$

Let $c(\beta) = \max\{\mathbb{E}[\exp(\beta|X_{lt}|)] \mid l \in [m] \text{ accessible}\}$. We have $c(\beta) < \infty$. One can bound the right hand side of (3.3) by

$$\prod_{t=0}^n c(\beta)^{\gamma^t} = c(\beta)^{\sum_{t=0}^n \gamma^t}.$$

Putting all together yields

$$\mathbb{E}[\exp(\beta|S|)] \leq \lim_{n \rightarrow \infty} c(\beta)^{\sum_{t=0}^n \gamma^t} = c(\beta)^{1/(1-\gamma)} < \infty.$$

□

Proof of Theorem 3. Consider $G_i = \sum_t \gamma^t R_{I_t, I_{t+1}, t}$ starting with $I_0 = i$. With probability one (I_t, I_{t+1}) realizes to a pair (j, j') with $i \rightarrow j$ and $p_{jj'} > 0$.

1) We have $\mathbb{P}[|R_j| \leq K] = \sum_{j'} p_{jj'} \mathbb{P}[|R_{jj't}| \leq K]$. Hence, if $\mathbb{P}[|R_j| \leq K] = 1$ for all j with $i \rightarrow j$, then $\mathbb{P}[|R_{jj't}| \leq K] = 1$ for all (j, j') with $i \rightarrow j$ and $p_{jj'} > 0$. This yields $\mathbb{P}[|R_{I_t, I_{t+1}, t}| \leq K | I_0 = i] = 1$ for every t and hence $|G_i| \leq \sum_t \gamma^t |R_{I_t, I_{t+1}, t}| \leq K/(1-\gamma)$ almost surely.

2) Due to Theorem 2, we can consider the situation of Proposition 2 via $l = (i, j)$, $X_{lt} = R_{ijt}$ and $K_t = (I_t, I_{t+1})$ having starting distribution $\mathbb{P}[K_0 = (i, j)] = p_{ij}$. In this situation, a state $l = (i', j)$ is accessible with respect to $(K_t)_t$ if and only if $i \rightarrow i'$ and $p_{i'j} > 0$. So, G_i has the same distribution as the almost sure limit $\sum_{t=0}^{\infty} \gamma^t X_{K_t, t}$ and one obtains $\mathbb{E}[\exp(\beta|G_i|)] < \infty$.

3) We have $\mathbb{E}[|R_j|^p] = \sum_{j'} p_{jj'} \mathbb{E}[|R_{jj't}|^p]$ hence $\mathbb{E}[|R_j|^p] < \infty$ for all j with $i \rightarrow j$ implies $\mathbb{E}[|R_{jj't}|^p] < \infty$ for all (j, j') with $i \rightarrow j$ and $p_{jj'} > 0$. Let $\|R_{jj't}\|_p = \mathbb{E}[|R_{jj't}|^p]^{1/p}$ and $K = \max_{(j, j')} \|R_{jj't}\|_p$ where the maximum runs over all relevant pairs (j, j') , hence $K < \infty$. Let the Markov chain $(I_t)_t$ be started in state $I_0 = i$. Then $G_{i, T} = \sum_{t=0}^T \gamma^t R_{I_t, I_{t+1}, t} \rightarrow G_i$ almost surely as $T \rightarrow \infty$ and the process $(I_t, I_{t+1})_t$ only visits pairs (j, j') with $\|R_{jj't}\|_p \leq K < \infty$. Since $\|\cdot\|_p$ is a norm, for $T' > T$

$$\|G_{i, T'} - G_{i, T}\|_p = \left\| \sum_{t=T+1}^{T'} \gamma^t R_{I_t, I_{t+1}, t} \right\|_p \leq \sum_{t=T+1}^{T'} \|\gamma^t R_{I_t, I_{t+1}, t}\|_p \leq K \cdot \sum_{t=T+1}^{T'} \gamma^t \rightarrow 0 \text{ as } \min\{T, T'\} \rightarrow \infty,$$

hence $(G_{i, T})_{T \in \mathbb{N}}$ is a Cauchy sequence with respect to $\|\cdot\|_p$. The associated L_p -space is complete, hence there is a random variable \tilde{G}_i in L_p , that is $\|\tilde{G}_i\|_p < \infty$, such that $\|G_{i, T} - \tilde{G}_i\|_p \rightarrow 0$. This mode of convergence also implies $G_{i, T} \rightarrow \tilde{G}_i$ in probability. The almost sure convergence $G_{i, T} \rightarrow G_i$ implies that $G_{i, T} \rightarrow G_i$ in probability. Limits with respect to convergence in probability are almost surely unique, hence $\tilde{G}_i = G_i$ with probability one and $\|G_i\|_p = \|\tilde{G}_i\|_p < \infty$. □

3.3. Regularly varying tails: Proof of Theorem 4

Our second main theorem is the result about regular variation of fixed points of \mathcal{T} , see Theorem 4. As noted before, our proof makes use of the well-developed theory for multivariate fixed point equations

$$\mathbf{G} \stackrel{d}{=} \mathbf{J}\mathbf{G} + \mathbf{R}, \quad \mathbf{G} \text{ and } (\mathbf{J}, \mathbf{R}) \text{ independent.} \quad (3.4)$$

In particular, a notion of regular variation for *multivariate* distributions, extending the one-dimensional case, has been explored in the context of equations such as (3.4) under various assumptions on the joint law $\mathcal{L}(\mathbf{J}, \mathbf{R}) \in \mathcal{P}(\mathbb{R}^{d \times d} \times \mathbb{R}^d)$. Below, we use the coupling approach explained in Section 2.1 to apply some of the available results in proving Theorem 4.

First, we introduce the multivariate notion of regular variation in random vectors and state some facts we use in the proof. We refer the reader to Appendix C in [16] for more details on the basic properties of multivariate regular variation. Let $\bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$ be the extended real numbers and

let $\bar{\mathbb{R}}_0^d = \bar{\mathbb{R}}^d \setminus \{0\}$. A \mathbb{R}^d -valued random vector \mathbf{X} is called *regularly varying* if there exists a non-null Radon measure μ on $\bar{\mathbb{R}}_0^d$, that does not charge infinite points, such that

$$\lim_{x \rightarrow \infty} \frac{\mathbb{P}[x^{-1}\mathbf{X} \in C]}{\mathbb{P}[|\mathbf{X}| > x]} = \mu(C) \quad \text{for every measurable } \mu\text{-continuity set } C \subseteq \bar{\mathbb{R}}_0^d, \quad (3.5)$$

note that $|\cdot|$ denotes the euclidean norm on \mathbb{R}^d . The measure μ is called the *limit measure* of regular variation of the random vector \mathbf{X} . The following are well-known facts about regularly varying random vectors, again we refer to Appendix C in [16]

- If μ is the limit measure of regular variation of a random vector \mathbf{X} , then there is a unique $\alpha > 0$, called the *index of regular variation*, such that $\mu(tC) = t^{-\alpha}\mu(C)$ for every μ -continuity set C and $t > 0$.
- If (3.5) is satisfied with index of regular variation $\alpha > 0$, then the set $C = \{\mathbf{x} \in \mathbb{R}^d \mid |\mathbf{x}| > 1\}$ is a μ -continuity set, and hence, because $\{x^{-1}\mathbf{X} \in C\} = \{|\mathbf{X}| > x\}$, it holds $\mu(C) = 1$ and

$$x \mapsto \mathbb{P}[|\mathbf{X}| > x]$$

is a regularly varying function with index $-\alpha$. Further, the set $\{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{w}^T \mathbf{x} > 1\}$ is a μ -continuity set for each $\mathbf{w} \in \mathbb{R}^d$,

- \mathbf{X} is regularly varying with index $\alpha > 0$ if and only if there exists a distribution $\Xi \in \mathcal{P}(\mathbb{S}^{d-1})$ on the unit sphere $\mathbb{S}^{d-1} = \{\mathbf{x} \in \mathbb{R}^d \mid |\mathbf{x}| = 1\}$ such that for every $t > 0$

$$\frac{\mathbb{P}[|\mathbf{X}| > tx, \frac{\mathbf{X}}{|\mathbf{X}|} \in \cdot]}{\mathbb{P}[|\mathbf{X}| > x]} \xrightarrow{w} t^{-\alpha}\Xi(\cdot) \quad \text{as } x \rightarrow \infty,$$

where \xrightarrow{w} is weak convergence of finite measures.

The following theorem can be obtained easily as a special case of Theorem 4.4.24 in [16]; note that the empty product of (random) matrices is the identity matrix.

Theorem 5. *Let $(R_1, J_1), \dots, (R_d, J_d)$ be a coupling, $\mathbf{R} = (R_1, \dots, R_d)^T$ and $\mathbf{J} = (\mathbf{J}_{ij})_{ij}$ with $\mathbf{J}_{ij} = \gamma 1(J_i = j)$, see Section 2.1. Suppose \mathbf{R} is regularly varying with index $\alpha > 0$ and limit measure $\mu_{\mathbf{R}}$. Then the multivariate distributional fixed point equation (3.4) has a unique solution \mathbf{G} which is regularly varying with index α . In particular, with $(\mathbf{J}^{(s)})_{s \in \mathbb{N}_0}$ being iid copies of \mathbf{J} , it holds*

$$\lim_{x \rightarrow \infty} \frac{\mathbb{P}[x^{-1}\mathbf{G} \in C]}{\mathbb{P}[|\mathbf{R}| > x]} = \sum_{t=0}^{\infty} \mathbb{E}[\mu_{\mathbf{R}}(\{\mathbf{x} \in \mathbb{R}^d \mid (\prod_{s=0}^{t-1} \mathbf{J}^{(s)})\mathbf{x} \in C\})] \quad (3.6)$$

for every $\mu_{\mathbf{R}}$ -continuity set C and

$$\mathbb{P}[|\mathbf{G}| > x] \sim c \cdot \mathbb{P}[|\mathbf{R}| > x] \text{ as } x \rightarrow \infty, \quad (3.7)$$

where $c > 0$ equals the right hand side of (3.6) evaluated at $C = \{\mathbf{x} \in \mathbb{R}^d \mid |\mathbf{x}| > 1\}$.

Proof. Since $C = \{\mathbf{x} \in \mathbb{R}^d \mid |\mathbf{x}| > 1\}$ is a μ -continuity set that satisfies $\mathbb{P}[|\mathbf{G}| > x] = \mathbb{P}[x^{-1}|\mathbf{G}| \in C]$ the asymptotic expression (3.7) follows from (3.6). The remaining part of the theorem is a direct consequence of Theorem 4.4.24 from [16] by noticing that the additional necessary assumptions stated in that theorem, $\mathbb{E}[|\mathbf{J}|^\alpha] < 1$ and $\mathbb{E}[|\mathbf{J}|^{\alpha+\delta}] < \infty$ for some $\alpha > 0$ and $\delta > 0$, are satisfied in our situation of interest as $\mathbf{J}_{ij} = \gamma 1(J_i = j)$ and hence $|\mathbf{J}| = \gamma < 1$. \square

In Theorem 5, an arbitrary coupling $(R_1, J_1), \dots, (R_d, J_d)$ is considered and the (multivariate) regular variation of $\mathbf{R} = (R_1, \dots, R_d)^T$, which is an assumption there, transfers over to the (multivariate) regular variation of $\mathbf{G} = (G_1, \dots, G_d)^T$, whose marginal laws are the solutions to the distributional Bellman equations of interest in Theorem 4. Moreover, multivariate regular variation of \mathbf{G} can be used to show univariate regular variation of G_1, \dots, G_d by testing (3.5) on sets of the form $C = \{\mathbf{x} \in \mathbb{R}^d \mid \pm \mathbf{e}_j^T \mathbf{x} > 1\}$ with $\mathbf{e}_j \in \mathbb{R}^d$ being unit vectors. Note that Theorem 4 concerns solutions of the distributional Bellman operator \mathcal{T} . Hence, the result of that theorem only depends on the marginal laws $(\mathcal{L}(R_1, J_1), \dots, \mathcal{L}(R_d, J_d))$ and hence Theorem 5, which assumes presence of a coupling satisfying certain properties, is *not* directly applicable there. However, as explained in Section 2.1, results about coupled situations can become useful by choosing a convenient coupling $\mathcal{L}((R_1, J_1), \dots, (R_d, J_d))$ that makes certain arguments work out. In the following proof, this will be the case when an independence coupling is considered: independence of R_1, \dots, R_d together with the assumptions of Theorem 4 implies that the random vector $\mathbf{R} = (R_1, \dots, R_d)^T$ is regularly varying in the multivariate sense, so Theorem 5 becomes applicable. We now present technical details:

Proof of Theorem 4. Let $(R_1, J_1), \dots, (R_d, J_d)$ be independent. Let $\mathbf{e}_j \in \mathbb{R}^d$ be the j -th unit vector and

$$\mathbf{R} = (R_1, \dots, R_d)^T = \sum_{j=1}^d R_j \cdot \mathbf{e}_j,$$

that is, we express \mathbf{R} as a sum of independent random vectors. In the following, we rely on well-established results from the theory of (multivariate) regular variation which show that regular variation is transferred to sums of *independent* random variables (vectors) each having this property.

First, we determine the asymptotic behavior of $\mathbb{P}[\|\mathbf{R}\| > x]$ as $x \rightarrow \infty$: let $i \in [d]$ be fixed. Only those pairs (R_j, J_j) with $i \rightarrow j$ are relevant for the distribution of G_i , hence we assume w.l.o.g. that $i \rightarrow j$ holds for all j , that is for every $j \in [d]$ there exists constants $q_j \in [0, 1]$ and $c_j \in [0, \infty)$ such that

$$\lim_{x \rightarrow \infty} \frac{\mathbb{P}[R_j > x]}{x^{-\alpha} L(x)} = q_j c_j \quad \text{and} \quad \lim_{x \rightarrow \infty} \frac{\mathbb{P}[R_j < -x]}{x^{-\alpha} L(x)} = (1 - q_j) c_j, \quad (3.8)$$

where $\alpha > 0$ and L is a slowly varying function. This implies for every j

$$c_j = \lim_{x \rightarrow \infty} \frac{\mathbb{P}[|R_j| > x]}{x^{-\alpha} L(x)} = \lim_{x \rightarrow \infty} \frac{\mathbb{P}[R_j^2 > x^2]}{x^{-\alpha} L(x)} = \lim_{y \rightarrow \infty} \frac{\mathbb{P}[R_j^2 > y]}{y^{-\alpha/2} \tilde{L}(y)} \quad \text{with} \quad \tilde{L}(y) = L(\sqrt{y}).$$

The function \tilde{L} is slowly varying which shows that R_j^2 is regularly varying with index $\alpha/2$ in case $c_j > 0$ and $\mathbb{P}[R_j^2 > x] = o(x^{-\alpha/2})$ as $x \rightarrow \infty$ in case $c_j = 0$. In such a situation, and since R_1^2, \dots, R_d^2 are independent by choice of coupling, standard results from the theory of regular variation apply, see Section 1.3.1 in [29], and show

$$\mathbb{P}[R_1^2 + \dots + R_d^2 > y] \sim \mathbb{P}[R_1^2 > y] + \dots + \mathbb{P}[R_d^2 > y] \quad \text{as } y \rightarrow \infty$$

and hence, by re-substituting x^2 for y ,

$$\mathbb{P}[R_1^2 + \dots + R_d^2 > x^2] \sim \mathbb{P}[R_1^2 > x^2] + \dots + \mathbb{P}[R_d^2 > x^2] \sim (c_1 + \dots + c_d) x^{-\alpha} L(x) \quad \text{as } x \rightarrow \infty,$$

that is

$$\lim_{x \rightarrow \infty} \frac{\mathbb{P}[|\mathbf{R}| > x]}{x^{-\alpha} L(x)} = c_1 + \dots + c_d. \quad (3.9)$$

Next, we give an argument to show that \mathbf{R} is regularly varying with index α (in the multivariate sense): let $\text{sgn}(R_j) \in \{-1, 1\}$ be the sign of R_j , so that $|R_j \mathbf{e}_j| = |R_j|$ and $\frac{R_j \mathbf{e}_j}{|R_j \mathbf{e}_j|} = \text{sgn}(R_j) \mathbf{e}_j$ in case $R_j \neq 0$. If $c_j > 0$ then $R_j \mathbf{e}_j$ is regularly varying at index $\alpha > 0$ in the multivariate sense, the spectral measure $\Xi_j \in \mathcal{P}(\mathbb{S}^{d-1})$ is given by

$$\Xi_j = q_j \delta_{\mathbf{e}_j} + (1 - q_j) \delta_{-\mathbf{e}_j} = \lim_{x \rightarrow \infty} \mathbb{P}[\text{sgn}(R_j) \mathbf{e}_j \in \cdot \mid |R_j| > x].$$

Applying standard arguments, see Lemma C.3.1 in [16], it is easily seen that $\mathbf{R} = \sum_{j=1}^d R_j \mathbf{e}_j$ (a sum of independent random vectors) is regularly varying at index α . Let $\mu_{\mathbf{R}}$ be the associated measure. We are now able to apply Theorem 5 to the solution $\mathbf{G} = (G_1, \dots, G_d)^T$ of the distributional equation (3.4). First, we evaluate the limit measure $\mu_{\mathbf{R}}$ at some sets of interest: Combining (3.8) and (3.9) yields for every $j \in [d]$

$$\mu_{\mathbf{R}}(\{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{e}_j^T \mathbf{x} > 1\}) = \lim_{x \rightarrow \infty} \frac{\mathbb{P}[R_j > x]}{\mathbb{P}[|\mathbf{R}| > x]} = \frac{q_j c_j}{c_1 + \dots + c_d}.$$

As in Theorem 5, let $(\mathbf{J}^{(s)})_{s \in \mathbb{N}_0}$ be a sequence of iid copies of \mathbf{J} . For every $t \in \mathbb{N}_0$ the entry of the random matrix $\prod_{s=0}^{t-1} \mathbf{J}^{(s)}$ at position $(i, j) \in [d]^2$ equals $\gamma^t \cdot 1(I_{it} = j)$, where $(I_{it})_{i \in [d], t \in \mathbb{N}_0}$ are $[d]$ -valued random variables such that for every $i \in [d]$ the sequence $(I_{it})_t$ has the same law as $(I_t)_t$ under $\mathbb{P}[\cdot \mid I_0 = i]$.

By Theorem 5 for every $i \in [d]$ it holds that

$$\lim_{x \rightarrow \infty} \frac{\mathbb{P}[G_i > x]}{\mathbb{P}[|\mathbf{R}| > x]} = \sum_{t=0}^{\infty} \mathbb{E} \left[\mu_{\mathbf{R}}(\{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{e}_i^T (\prod_{s=0}^{t-1} \mathbf{J}^{(s)}) \mathbf{x} > 1\}) \right].$$

Using $\mu_{\mathbf{R}}(tC) = t^{-\alpha} \mu_{\mathbf{R}}(C)$ for every $t \in \mathbb{N}_0$ it holds

$$\begin{aligned} \mathbb{E} \left[\mu_{\mathbf{R}}(\{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{e}_i^T (\prod_{s=0}^{t-1} \mathbf{J}^{(s)}) \mathbf{x} > 1\}) \right] &= \mathbb{E} \left[\mu_{\mathbf{R}}(\{\mathbf{x} \in \mathbb{R}^d \mid \gamma^t \mathbf{e}_{I_{it}}^T \mathbf{x} > 1\}) \right] \\ &= \gamma^{\alpha t} \sum_{j \in [d]} \mathbb{P}[I_t = j \mid I_0 = i] \mu_{\mathbf{R}}(\{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{e}_j^T \mathbf{x} > 1\}) \\ &= \gamma^{\alpha t} \sum_{j \in [d]} \mathbb{P}[I_t = j \mid I_0 = i] \frac{q_j c_j}{c_1 + \dots + c_d} \end{aligned}$$

and hence

$$\begin{aligned} \lim_{x \rightarrow \infty} \frac{\mathbb{P}[G_i > x]}{\mathbb{P}[|\mathbf{R}| > x]} &= \sum_{t=0}^{\infty} \gamma^{\alpha t} \sum_{j \in [d]} \mathbb{P}[I_t = j \mid I_0 = i] \frac{q_j c_j}{c_1 + \dots + c_d} \\ &= \frac{1}{1 - \gamma^\alpha} \sum_{j \in [d]} \mathbb{P}[I_N = j \mid I_0 = i] \frac{q_j c_j}{c_1 + \dots + c_d}, \end{aligned}$$

where $N \sim \text{Geo}(\gamma^\alpha)$ is independent of $(I_t)_t$. With $w_{ij} = \mathbb{P}[I_N = j \mid I_0 = i]$ and (3.9) we have

$$\lim_{x \rightarrow \infty} \frac{\mathbb{P}[G_i > x]}{x^{-\alpha} L(x)} = \frac{\sum_{j \in [d]} w_{ij} q_j c_j}{1 - \gamma^\alpha}.$$

The tail asymptotic of $\frac{\mathbb{P}[G_i < -x]}{x^{-\alpha}L(x)}$ can be calculated the same way using

$$\mu_{\mathbf{R}}(\{\mathbf{x} \in \mathbb{R}^d \mid \mathbf{e}_j^T \mathbf{x} < -1\}) = \lim_{x \rightarrow \infty} \frac{\mathbb{P}[R_j < -x]}{\mathbb{P}[|\mathbf{R}| > x]} = \frac{(1 - q_j)c_j}{c_1 + \dots + c_d}.$$

To see that G_i is regularly varying with index α note that $\mathbb{P}[|G_i| > x] = \mathbb{P}[G_i > x] + \mathbb{P}[G_i < -x] \sim (\sum_{j=1}^d w_{ij}c_j)x^{-\alpha}L(x)$ is a regularly varying function of index α and the balance equation hold:

$$\begin{aligned} \lim_{x \rightarrow \infty} \frac{\mathbb{P}[G_i > x]}{\mathbb{P}[|G_i| > x]} &= \frac{\sum_{j=1}^d q_j w_{ij} c_j}{\sum_{j=1}^d w_{ij} c_j} \quad \text{and} \\ \lim_{x \rightarrow \infty} \frac{\mathbb{P}[G_i < -x]}{\mathbb{P}[|G_i| > x]} &= \frac{\sum_{j=1}^d (1 - q_j) w_{ij} c_j}{\sum_{j=1}^d w_{ij} c_j} = 1 - \frac{\sum_{j=1}^d q_j w_{ij} c_j}{\sum_{j=1}^d w_{ij} c_j}. \end{aligned}$$

□

Use of AI tools declaration

The authors have not used Artificial Intelligence (AI) tools in the creation of this article.

Acknowledgments

We thank three anonymous referees for their careful reading, many constructive remarks and suggestions. The referees' comments led to a significant improvement of the manuscript.

The first author was partially supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 502386356.

Conflict of interest

The authors declare there is no conflicts of interest.

References

1. M. G. Bellemare, W. Dabney, R. Munos, A distributional perspective on reinforcement learning, in *International Conference on Machine Learning*, (2017), 449–458.
2. M. G. Bellemare, W. Dabney, M. Rowland, *Distributional Reinforcement Learning*, MIT Press, 2023.
3. M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley & Sons, Inc., New York, 1994. <https://doi.org/10.1002/9780470316887>
4. M. Rowland, M. Bellemare, W. Dabney, R. Munos, Y. W. Teh, An analysis of categorical distributional reinforcement learning, in *International Conference on Artificial Intelligence and Statistics*, (2018), 29–37.
5. E. Krasheninnikova, J. García, R. Maestre, F. Fernández, Reinforcement learning for pricing strategy optimization in the insurance industry, *Eng. Appl. Artif. Intell.*, **80** (2019), 8–19. <https://doi.org/10.1016/j.engappai.2019.01.010>

6. P. N. Kolm, G. Ritter, Modern perspectives on reinforcement learning in finance, *J. Mach. Learn. Finance*, **1** (2019). <https://doi.org/10.2139/ssrn.3449401>
7. P. Embrechts, C. Klüppelberg, T. Mikosch, *Modelling Extremal Events: For Insurance and Finance*, Springer-Verlag, Berlin, 1997. <https://doi.org/10.1007/978-3-642-33483-2>
8. V. Zhuang, Y. Sui, No-regret reinforcement learning with heavy-tailed rewards, in *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, (2021), 3385–3393.
9. A. M. Medina, S. Yang, No-regret algorithms for heavy-tailed linear bandits, in *Proceedings of The 33rd International Conference on Machine Learning* (eds. M. F. Balcan and K. Q. Weinberger), PMLR, (2016), 1642–1650. <https://proceedings.mlr.press/v48/medina16.html>
10. X. Yu, H. Shao, M. R. Lyu, I. King, Pure exploration of multi-armed bandits with heavy-tailed payoffs, in *Conference on Uncertainty in Artificial Intelligence*, (2018), 937–946.
11. H. Shao, X. Yu, I. King, M. R. Lyu, Almost optimal algorithms for linear stochastic bandits with heavy-tailed payoffs, in *Advances in Neural Information Processing Systems* (eds. S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi and R. Garnett), Curran Associates, Inc., 2018.
12. S. Lu, G. Wang, Y. Hu, L. Zhang, Optimal algorithms for lipschitz bandits with heavy-tailed rewards, in *International Conference on Machine Learning*, PMLR, (2019), 4154–4163.
13. S. R. Chowdhury, A. Gopalan, Bayesian optimization under heavy-tailed payoffs, in *Advances in Neural Information Processing Systems* (eds. H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox and R. Garnett), Curran Associates, Inc., 2019.
14. A. Dubey, A. Pentland, Thompson sampling on symmetric alpha-stable bandits, in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, (2019), 5715–5721. <https://doi.org/10.24963/ijcai.2019/792>
15. G. Alsmeyer, F. Buckmann, Stability of perpetuities in Markovian environment, *J. Differ. Equations Appl.*, **23** (2017), 699–740. <https://doi.org/10.1080/10236198.2016.1271878>
16. D. Buraczewski, E. Damek, T. Mikosch, *Stochastic Models with Power-law Tails*, Springer, 2016, <https://doi.org/10.1007/978-3-319-29679-1>
17. T. Morimura, M. Sugiyama, H. Kashima, H. Hachiya, T. Tanaka, Nonparametric return distribution approximation for reinforcement learning, in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, (2010), 799–806.
18. K. J. Chung, M. J. Sobel, Discounted mdp's: Distribution functions and exponential utility maximization, *SIAM J. Control Optim.*, **25** (1987), 49–62. <https://doi.org/10.1137/0325004>
19. W. Vervaat, On a stochastic difference equation and a representation of nonnegative infinitely divisible random variables, *Adv. Appl. Probab.*, **11** (1979), 750–783. <https://doi.org/10.2307/1426858>
20. C. M. Goldie, R. Grübel, Perpetuities with thin tails, *Adv. Appl. Probab.*, **28** (1996), 463–480, <https://doi.org/10.2307/1428067>.
21. C. M. Goldie, R. A. Maller, Stability of perpetuities, *Ann. Probab.*, **28** (2000), 1195–1218. <https://doi.org/10.1214/aop/1019160331>

22. U. Rösler, A fixed point theorem for distributions, *Stochastic Process. Appl.*, **42** (1992), 195–214. [https://doi.org/10.1016/0304-4149\(92\)90035-O](https://doi.org/10.1016/0304-4149(92)90035-O)
23. N. H. Bingham, C. M. Goldie, J. L. Teugels, *Regular Variation*, Cambridge University Press, Cambridge, 1987. <https://doi.org/10.1017/CBO9780511721434>
24. D. Cline, Infinite series of random variables with regularly varying tails, *Inst. Appl. Math. Statist.*, (1983).
25. A. Brandt, The stochastic equation $Y_{n+1} = A_n Y_n + B_n$ with stationary coefficients, *Adv. Appl. Probab.*, **18** (1986), 211–220. <https://doi.org/10.2307/1427243>
26. T. Erhardsson, Conditions for convergence of random coefficient AR(1) processes and perpetuities in higher dimensions, *Bernoulli*, **20** (2014), 990–1005. <https://doi.org/10.3150/13-BEJ513>
27. O. Kallenberg, *Foundations of Modern Probability*, Springer, 2021. <https://doi.org/10.1007/978-3-030-61871-1>
28. P. Bougerol, N. Picard, Strict stationarity of generalized autoregressive processes, *Ann. Probab.*, **20** (1992), 1714–1730.
29. T. Mikosch, *Regular Variation, Subexponentiality and Their Applications in Probability Theory*, Eindhoven University of Technology Eindhoven, The Netherlands, 1999.



AIMS Press

© 2023 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>)