

Ergodic Backward Stochastic Difference Equations

by

Andrew L Allan
University of Oxford

and

Samuel N Cohen
University of Oxford

EPSRC Centre for Doctoral Training in Partial Differential Equations:
Analysis and Applications
Mathematical Institute
University of Oxford
Andrew Wiles Building
ROQ, Woodstock Road
Oxford, UK
OX2 6GG

October 2016

Ergodic Backward Stochastic Difference Equations

Andrew L. Allan* and Samuel N. Cohen†
Mathematical Institute, University of Oxford
andrew.allan@maths.ox.ac.uk
samuel.cohen@maths.ox.ac.uk

September 26, 2016

This is an Accepted Manuscript of an article published by Taylor & Francis in Stochastics: An International Journal of Probability and Stochastic Processes on 2nd September 2016, available online:
<http://www.tandfonline.com/10.1080/17442508.2016.1224881>.

Abstract

We consider ergodic backward stochastic differential equations in a discrete time setting, where noise is generated by a finite state Markov chain. We show existence and uniqueness of solutions, along with a comparison theorem. To obtain this result, we use a Nummelin splitting argument to obtain ergodicity estimates for a discrete time Markov chain which hold uniformly under suitable perturbations of its transition matrix. We conclude with an application of this theory to a treatment of an ergodic control problem.

Keywords: Ergodic BSDE, Markov Chain, Uniform Ergodicity, Nummelin Splitting, Ergodic control
MSC: 60J10, 93E20, 60F99

1 Introducing Discrete Time BSDEs

1.1 Introduction

Over the last 25 years, Backward Stochastic Differential Equations (BSDEs) have been extensively researched, and established as a fundamental object in

*Research supported by the Engineering and Physical Sciences Research Council [EP/L015811/1].

†Research supported by the Oxford–Man Institute for Quantitative Finance.

mathematical finance and stochastic control. Here, one is typically interested in the solution process (Y, Z) of an equation of the form

$$Y_t - \int_t^T f(\omega, u, Y_u, Z_u) du + \int_t^T Z_u dW_u = \xi,$$

for some generator (or driver) function f and some \mathcal{F}_T -measurable terminal condition ξ , where W is an n -dimensional Brownian motion.

When the cost of a control problem is considered in an ergodic manner, that is, it is the long run behaviour which is important, the associated equations turn out to be examples of a variation of these equations, named ‘Ergodic BSDEs’, which typically take the form

$$Y_T = Y_t - \int_t^T (f(X_u, Z_u) - \lambda) du + \int_t^T Z_u dW_u.$$

Having been first introduced in Fuhrman, Hu and Tessitore [10], the theory of these equations and their link to ergodic control problems have been researched over the last few years; see Richou [14], Debussche, Hu and Tessitore [8] and Cohen and Fedyashov [5]. In Cohen and Hu [6], these equations were considered in the case where noise is generated by a continuous time, countable state Markov chain.

This paper considers the discrete time analogue of these equations, where noise is generated by a discrete time, finite state Markov chain. In Section 2 we shall extend the existing theory of discrete time BSDEs, as given in Cohen and Elliott [3], to the case with an infinite time horizon. Section 3 is a brief digression to the study of uniformly ergodic Markov chains in discrete time, in which we shall show how to obtain ergodicity estimates for a Markov chain which are robust under a suitable perturbation of its transition matrix. These estimates have not been previously obtained in a discrete time setting and are of independent interest in the study of the ergodic properties of discrete time Markov chains.

We shall then make use of these estimates in Section 4 to construct a suitable limiting procedure with which to prove the existence of bounded Markovian solutions to discrete time Ergodic BSDEs. We will also prove a comparison theorem, and observe the relationship between the Markovian solution and the ergodic measure of the underlying Markov chain. Finally, in Section 5 we shall show how our theory can be applied to a particular ergodic control problem.

1.2 Finite horizon BSDEs

We will begin by introducing the theory of finite horizon, discrete time BSDEs, as established in [3].

Consider an underlying discrete time, finite state stochastic process X . Without loss of generality, we may assume that X takes values in the standard basis of \mathbb{R}^N , for some $N \in \mathbb{N}$. That is, for each $t \in \{0, 1, 2, \dots\}$,

$$X_t \in \mathcal{X} := \{e_1, e_2, \dots, e_N\},$$

where $e_k = (0, 0, \dots, 0, 1, 0, \dots, 0)^* \in \mathbb{R}^N$, and $[\cdot]^*$ denotes vector (or matrix) transposition.

Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$ be a filtered probability space, where \mathcal{F}_t is the completion of the σ -algebra generated by the process X up to time t , and assume that $\mathcal{F} = \mathcal{F}_\infty := \bigvee_{t \geq 0} \mathcal{F}_t$.

Remark 1.1. We have not assumed that \mathcal{F} is complete, nor that \mathcal{F}_0 contains all the \mathbb{P} -null sets in \mathcal{F}_∞ . Later, we will consider our processes under a variety of different probability measures, and we do not wish for \mathbb{P} -null events in the tail σ -algebra to be known at finite times, as these events may occur with positive probability under a new measure.

Note that the atoms of \mathcal{F}_t are sets of positive measure of the form

$$\{\omega' \in \Omega : X_s(\omega') = X_s(\omega) \text{ for all } s \leq t\}$$

for some $\omega \in \Omega$, up to a null set in \mathcal{F}_t . We will find that much of our analysis can be done by restricting our attention to individual atoms.

For $t \geq 1$, let $M_t := X_t - \mathbb{E}[X_t | \mathcal{F}_{t-1}]$, which defines a martingale difference sequence M .

We now define a backward stochastic difference equation (BSDE) as an equation of the form

$$Y_t - \sum_{t \leq u < T} f(\omega, u, Y_u, Z_u) + \sum_{t \leq u < T} Z_u^* M_{u+1} = \xi, \quad (1)$$

where $T > 0$ is a finite deterministic terminal time, $f : \Omega \times \{0, \dots, T-1\} \times \mathbb{R} \times \mathbb{R}^N \rightarrow \mathbb{R}$ is an adapted map, and $\xi : \Omega \rightarrow \mathbb{R}$ is an \mathcal{F}_T -measurable random variable. By “ f is adapted” we mean that, for all $t \in \{0, \dots, T-1\}$, the map from $\Omega \times \mathbb{R} \times \mathbb{R}^N \rightarrow \mathbb{R}$ given by $(\omega, y, z) \mapsto f(\omega, t, y, z)$ is $\mathcal{F}_t \otimes \mathcal{B}(\mathbb{R}) \otimes \mathcal{B}(\mathbb{R}^N)$ -measurable. Since there are only finitely many possible paths for X up to time t , \mathcal{F}_t is sufficiently coarse that all \mathcal{F}_t -measurable real-valued functions belong to $L^1(\Omega, \mathcal{F}_t, \mathbb{P})$ and $L^\infty(\Omega, \mathcal{F}_t, \mathbb{P})$.

A solution to the BSDE (1) is a pair (Y, Z) of adapted processes, taking values in $\mathbb{R} \times \mathbb{R}^N$, such that

$$Y_t(\omega) - \sum_{t \leq u < T} f(\omega, u, Y_u(\omega), Z_u(\omega)) + \sum_{t \leq u < T} Z_u^*(\omega) M_{u+1}(\omega) = \xi(\omega)$$

holds for all $t \in \{0, \dots, T\}$ and almost all ω .

Definition 1.2. We will denote by $\|\cdot\|_{M_{t+1}}$ the stochastic seminorm on the space of \mathbb{R}^N -valued, \mathcal{F}_t -measurable random variables given by

$$\|Z_t\|_{M_{t+1}}^2 := Z_t^* \mathbb{E}[M_{t+1} M_{t+1}^* | \mathcal{F}_t] Z_t = \mathbb{E}[(Z_t^* M_{t+1})^2 | \mathcal{F}_t],$$

and write $Z_t \sim_{M_{t+1}} Z'_t$ whenever $\|Z_t - Z'_t\|_{M_{t+1}} = 0$ a.s. We also denote by $\|\cdot\|_M$ the seminorm on the space of adapted processes in \mathbb{R}^N given by

$$\|Z\|_M^2 := \mathbb{E} \left[\sum_{0 \leq u < T} \|Z_u\|_{M_{u+1}}^2 \right] = \sum_{0 \leq u < T} \mathbb{E}[(Z_u^* M_{u+1})^2],$$

and write $Z \sim_M Z'$ whenever $\|Z - Z'\|_M = 0$.

It is clear that $Z \sim_M Z'$ if and only if $Z_t \sim_{M_{t+1}} Z'_t$ for all $t \in \{0, \dots, T-1\}$. It is also easy to see that \sim_M and $\sim_{M_{t+1}}$ are both equivalence relations.

We will construct the solutions of our BSDE by making use of the following Martingale Representation Theorem from Elliott and Yang [9].

Theorem 1.3. *Let L be a real-valued $\{\mathcal{F}_t\}_{t \geq 0}$ -adapted martingale. Then there exists an adapted \mathbb{R}^N -valued process Z such that*

$$L_t = L_0 + \sum_{0 \leq u < t} Z_u^* M_{u+1}$$

for all t . Further, the process Z is unique up to equivalence \sim_M .

Proof. By the Doob–Dynkin Lemma (see for example Shiryaev [16, p.174]), for each t , there exists a function $g_t : \mathcal{X} \times \mathcal{X} \times \dots \times \mathcal{X} \rightarrow \mathbb{R}$ such that

$$L_{t+1} = L_t + g_t(X_0, X_1, \dots, X_{t+1}).$$

Define Z_t as the \mathcal{F}_t -measurable, \mathbb{R}^N -valued random variable with components given by $e_k^* Z_t = g_t(X_0, X_1, \dots, X_t, e_k)$ for each $e_k \in \mathcal{X}$, and notice that

$$Z_t^* X_{t+1} = g_t(X_0, X_1, \dots, X_t, X_{t+1}) = L_{t+1} - L_t.$$

As L is a martingale, we see that $Z_t^* M_{t+1} = Z_t^* X_{t+1}$. Hence, $L_{t+1} = L_t + Z_t^* M_{t+1}$, and the result follows. The uniqueness of Z follows easily from the definition of \sim_M . \square

Corollary 1.4. *Let W be a real-valued \mathcal{F}_{t+1} -measurable random variable such that $\mathbb{E}[W | \mathcal{F}_t] = 0$. Then there exists an \mathbb{R}^N -valued \mathcal{F}_t -measurable random variable Z_t such that*

$$W = Z_t^* M_{t+1} \quad \text{a.s.}$$

Further, this variable is unique up to equivalence $\sim_{M_{t+1}}$.

Proof. Simply consider the martingale L defined by $L_s = \mathbb{1}_{\{s > t\}} W$, and apply Theorem 1.3. \square

The following theorem provides sufficient conditions for the existence of a unique solution to our finite horizon BSDE. In general, these conditions are also necessary for a unique solution to exist for all terminal conditions ξ , as shown in [3, Corollary 2].

Theorem 1.5. *Suppose our driver f is such that the following two assumptions hold:*

- (i) For any Y , if $Z \sim_M Z'$, then $f(\omega, t, Y_t, Z_t) = f(\omega, t, Y_t, Z'_t)$ a.s. for all t .

(ii) For any $z \in \mathbb{R}^N$, any t and for almost all ω , the map

$$y \mapsto y - f(\omega, t, y, z)$$

is a bijection $\mathbb{R} \rightarrow \mathbb{R}$.

Then, for any \mathcal{F}_T -measurable terminal condition ξ , the BSDE (1) has an adapted solution (Y, Z) . Moreover, this solution is unique, up to indistinguishability for Y and \sim_M for Z .

Proof. We shall construct the solution using backward induction. Note that $Y_T = \xi$ is the unique solution for Y at time T . Suppose that we have found a solution for Y at time $t + 1$. The one step dynamics of (1) are given by

$$Y_t - f(\omega, t, Y_t, Z_t) + Z_t^* M_{t+1} = Y_{t+1}. \quad (2)$$

Taking an $\mathbb{E}[\cdot | \mathcal{F}_t]$ expectation gives

$$Y_t - f(\omega, t, Y_t, Z_t) = \mathbb{E}[Y_{t+1} | \mathcal{F}_t], \quad (3)$$

and substituting this back into (2), we obtain

$$Z_t^* M_{t+1} = Y_{t+1} - \mathbb{E}[Y_{t+1} | \mathcal{F}_t]. \quad (4)$$

It follows that any adapted solution (Y_t, Z_t) at time t must satisfy both (3) and (4). By Corollary 1.4, there exists an \mathbb{R}^N -valued \mathcal{F}_t -measurable random variable Z_t such that (4) holds, and this variable is unique up to equivalence $\sim_{M_{t+1}}$.

With this Z_t , it follows from assumption (i) that equation (3) is now uniquely determined as an equation in Y_t . By assumption (ii), for almost all ω , this equation has a unique solution $Y_t(\omega)$. The pair (Y_t, Z_t) is then the unique solution of the BSDE at time t , and we can therefore construct the full solution (Y, Z) by backward induction. Since for each t , Y_t is unique up to equality a.s. and Z_t is unique up to $\sim_{M_{t+1}}$, it follows that the solution (Y, Z) is unique up to indistinguishability for Y , and \sim_M for Z . \square

1.3 BSDEs on Markov chains

Let us now consider the case when our underlying process X is a discrete time, finite state Markov chain under \mathbb{P} . To keep our notation somewhat consistent with the continuous time theory in [6], we shall define the transition matrix $A_t = [a_{ij}]$ of the Markov chain X by $a_{ij} = \mathbb{P}(X_{t+1} = i | X_t = j)$, so that the columns of the matrix give the probability distributions associated with the transitions of the chain. Note that X is allowed to be time-inhomogeneous, so that its transition matrix A_t can vary (deterministically) through time.

Notice that now $\mathbb{E}[X_{t+1} | \mathcal{F}_t] = A_t X_t$, so that $M_{t+1} = X_{t+1} - A_t X_t$.

Definition 1.6. We shall say that a driver f is *Markovian* if at every time t , the value of f only depends on ω through the value of X_t . That is, f is Markovian if it can be written as $f(\omega, t, y, z) = \tilde{f}(X_t, t, y, z)$ for some function $\tilde{f} : \mathcal{X} \times \mathbb{N} \times \mathbb{R} \times \mathbb{R}^N \rightarrow \mathbb{R}$.

Lemma 1.7. *Assume that X is a Markov chain with transition matrix A_t , and let f be a Markovian driver which satisfies the conditions of Theorem 1.5. Suppose that $\xi = \varphi(X_T)$ for some deterministic function $\varphi : \mathcal{X} \rightarrow \mathbb{R}$, and let (Y, Z) be the solution of the corresponding BSDE (1). Then there exists a function $v : \{0, 1, \dots, T\} \times \mathcal{X} \rightarrow \mathbb{R}$ such that, for each t , $Y_t = v(t, X_t)$ and the k^{th} component of Z_t is given by $e_k^* Z_t = v(t+1, e_k)$ (cf. [7, Theorem 3.2]).*

In particular, the solution (Y, Z) is also Markovian in the sense that Y_t and Z_t are both deterministic functions of X_t (and in fact Z_t is completely deterministic). Moreover, writing \mathbf{v}_{t+1} for the vector in \mathbb{R}^N with entries $e_k^ \mathbf{v}_{t+1} = v(t+1, e_k)$, the function v satisfies*

$$v(t, e_k) - \tilde{f}(e_k, t, v(t, e_k), \mathbf{v}_{t+1}) - \mathbf{v}_{t+1}^* A_t e_k = 0 \quad (5)$$

for all $t \in \{0, \dots, T-1\}$ and all $e_k \in \mathcal{X}$ such that $\mathbb{P}(X_t = e_k) > 0$.

Proof. We shall prove the existence of the function v by backward induction in t . Let $v(T, \cdot) := \varphi(\cdot)$, so that $Y_T = v(T, X_T)$.

Suppose that for some t we have found a function $v(t+1, \cdot) : \mathcal{X} \rightarrow \mathbb{R}$ such that $Y_{t+1} = v(t+1, X_{t+1})$. Recall from the proofs of Theorem 1.3 and Theorem 1.5, that the k^{th} component of Z_t is given by $e_k^* Z_t = g_t(X_0, X_1, \dots, X_t, e_k)$, where, in this case,

$$g_t(X_0, X_1, \dots, X_t, X_{t+1}) = v(t+1, X_{t+1}) - \sum_{i=1}^N v(t+1, e_i) e_i^* A_t X_t. \quad (6)$$

Note that the sum on the right-hand side of (6) does not depend on X_{t+1} , so that this term corresponds to the addition of a multiple of $\mathbf{1}$ to Z_t , where $\mathbf{1}$ denotes the vector in \mathbb{R}^N with all entries equal to 1. However, we infer from Definition 1.2 that $\|\mathbf{1}\|_{M_{t+1}} = 0$, so the addition of multiples of $\mathbf{1}$ does not change the value of Z_t up to equivalence $\sim_{M_{t+1}}$. We may therefore ignore this term, and deduce that $e_k^* Z_t = v(t+1, e_k)$ for all $e_k \in \mathcal{X}$.

By assumption, the driver f can be written as $f(\omega, t, y, z) = \tilde{f}(X_t, t, y, z)$ for some function \tilde{f} . From the proof of Theorem 1.5, Y_t is given by the (a.s. unique) solution of the equation:

$$Y_t - \tilde{f}(X_t, t, Y_t, Z_t) = \mathbb{E}[Y_{t+1} | \mathcal{F}_t].$$

Hence, since Z_t and $\mathbb{E}[Y_{t+1} | \mathcal{F}_t]$ are both deterministic functions of X_t , it follows that the same is true of Y_t , so there must exist a function $v(t, \cdot) : \mathcal{X} \rightarrow \mathbb{R}$ such that $Y_t = v(t, X_t)$. We therefore deduce the existence of the required function v by backward induction. The equation (5) follows from the one-step dynamics of (1). \square

2 Infinite Horizon BSDEs

2.1 Existence and uniqueness of solutions

We now consider infinite horizon BSDEs:

$$Y_T = Y_t - \sum_{t \leq u < T} f(\omega, u, Y_u, Z_u) + \sum_{t \leq u < T} Z_u^* M_{u+1}. \quad (7)$$

Here T is no longer a fixed terminal time. Instead, both t and T may take any finite values such that $0 \leq t < T$. A solution of this equation is an adapted process (Y, Z) such that (7) holds a.s. for every $0 \leq t < T$.

As a special case, we also introduce a ‘discounted’ BSDE as an infinite horizon equation of the form:

$$Y_T = Y_t - \sum_{t \leq u < T} (f(\omega, u, Z_u) - \alpha Y_u) + \sum_{t \leq u < T} Z_u^* M_{u+1}, \quad (8)$$

where $\alpha > 0$ is a fixed constant.

In this section we will show that infinite horizon BSDEs admit unique bounded solutions for a suitable class of drivers. Similar results are considered in a continuous time setting in Royer [15] and in [6], however the discrete time case has not been previously considered.

We will return to the case when X is a Markov chain in Section 2.2. However, the theory we present here for infinite horizon BSDEs holds much more generally.

Remark 2.1. Recall Definition 1.2. In light of our infinite horizon setting, we will henceforth write $Z \sim_M Z'$ if $Z_t \sim_{M_{t+1}} Z'_t$ for all $t \geq 0$.

Proposition 2.2. *Suppose that for any $t \geq 0$, $y \in \mathbb{R}$, $z, z' \in \mathbb{R}^N$ and any atom $F \in \mathcal{F}_t$, the driver f satisfies*

$$f(\omega, t, y, z) - f(\omega, t, y, z') > \min_{i \in \mathbb{J}_t^F} \{(z - z')^*(e_i - \mathbb{E}[X_{t+1} | \mathcal{F}_t])\} \quad (9)$$

on F , where $\mathbb{J}_t^F := \{i : \mathbb{P}(X_{t+1} = e_i | F) > 0\}$, unless

$$\min_{i \in \mathbb{J}_t^F} \{(z - z')^*(e_i - \mathbb{E}[X_{t+1} | \mathcal{F}_t])\} = 0,$$

in which case the inequality can be replaced by an equality. Then, for any adapted processes Y, Z, Z' , there exists a probability measure \mathbb{Q} on (Ω, \mathcal{F}) , which is equivalent to \mathbb{P} on \mathcal{F}_t for every finite t , and such that

$$\tilde{M}_t := - \sum_{0 \leq u < t} (f(\omega, u, Y_u, Z_u) - f(\omega, u, Y_u, Z'_u)) + \sum_{0 \leq u < t} (Z_u - Z'_u)^* M_{u+1}$$

is a martingale under \mathbb{Q} .

Proof. We will define \mathbb{Q} inductively in t . First let $\mathbb{Q}|_{\mathcal{F}_0} = \mathbb{P}|_{\mathcal{F}_0}$.

Next, suppose that we have defined \mathbb{Q} on \mathcal{F}_t for some $t \geq 0$. Let F be an atom of \mathcal{F}_t (so that $\mathbb{Q}(F) > 0$), and let

$$S_t^F := \{(Z_t - Z'_t)^*(e_i - \mathbb{E}_{\mathbb{P}}[X_{t+1} | \mathcal{F}_t]) : i \in \mathbb{J}_t^F\},$$

where here all variables are evaluated on F . Note that

$$\mathbb{E}_{\mathbb{P}}[(Z_t - Z'_t)^*(X_{t+1} - \mathbb{E}_{\mathbb{P}}[X_{t+1} | \mathcal{F}_t]) | F] = 0, \quad (10)$$

and hence that $\min(S_t^F) \leq 0$ and $\max(S_t^F) \geq 0$. Suppose that $\min(S_t^F) < 0$. By assumption, we can swap the roles of z and z' in (9), and deduce that

$$f(\omega, t, Y_t, Z_t) - f(\omega, t, Y_t, Z'_t) \in \text{conv}^\circ(S_t^F),$$

where here $\text{conv}^\circ(\cdot)$ denotes the interior of the convex hull of a subset of \mathbb{R} .

If $\min(S_t^F) = 0$, then it follows from (10) that $S_t^F = \{0\}$, and hence, by assumption, that

$$f(\omega, t, Y_t, Z_t) - f(\omega, t, Y_t, Z'_t) = 0.$$

It follows from the definition of convexity that, in either case, for each $i \in \mathbb{J}_t^F$, there exists $\mu_i^F > 0$ such that $\sum_{i \in \mathbb{J}_t^F} \mu_i^F = 1$, and

$$f(\omega, t, Y_t, Z_t) - f(\omega, t, Y_t, Z'_t) = \sum_{i \in \mathbb{J}_t^F} \mu_i^F (Z_t - Z'_t)^*(e_i - \mathbb{E}_{\mathbb{P}}[X_{t+1} | \mathcal{F}_t]). \quad (11)$$

Define $\mathbb{Q}(\{X_{t+1} = e_i\} \cap F) := \mu_i^F \mathbb{Q}(F)$ for $i \in \mathbb{J}_t^F$, and $\mathbb{Q}(\{X_{t+1} = e_i\} \cap F) := 0$ for $i \notin \mathbb{J}_t^F$.

The measure \mathbb{Q} is now defined on the whole of \mathcal{F}_{t+1} and therefore, by induction, on $\bigcup_{s \geq 0} \mathcal{F}_s$. By Carathéodory's extension theorem (see for example [16, p.152]), there exists a unique extension of \mathbb{Q} to a probability measure on $\mathcal{F} = \sigma(\bigcup_{s \geq 0} \mathcal{F}_s)$. By construction, \mathbb{Q} is equivalent to \mathbb{P} on \mathcal{F}_t , for any $t \geq 0$. Further, for any t, F and $i \in \mathbb{J}_t^F$ as above, $\mathbb{Q}(\{X_{t+1} = e_i\} | F) = \mu_i^F$.

Note that

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}}[\tilde{M}_{t+1} - \tilde{M}_t | \mathcal{F}_t] &= -(f(\omega, t, Y_t, Z_t) - f(\omega, t, Y_t, Z'_t)) \\ &\quad + (Z_t - Z'_t)^*(\mathbb{E}_{\mathbb{Q}}[X_{t+1} | \mathcal{F}_t] - \mathbb{E}_{\mathbb{P}}[X_{t+1} | \mathcal{F}_t]). \end{aligned}$$

It follows from (11) that the above expression is equal to zero on every atom F of \mathcal{F}_t . Hence, as required, \tilde{M} is a martingale under \mathbb{Q} . \square

Remark 2.3. Note that the measure \mathbb{Q} constructed in the proof of Proposition 2.2 is by no means uniquely defined, as in general there will be many ways of choosing the constants μ_i^F such that the required conditions hold. We also note that in general \mathbb{Q} will not be equivalent to \mathbb{P} on \mathcal{F} , which is why it was important that we did not assume that \mathcal{F}_0 contains all the \mathbb{P} -null sets in \mathcal{F} .

Definition 2.4. We shall say that a driver f is Lipschitz in z , if there exists $L > 0$ such that, for any t and any \mathcal{F}_t -measurable random variables Y_t, Z_t, Z'_t ,

$$|f(\omega, t, Y_t, Z_t) - f(\omega, t, Y_t, Z'_t)| \leq L \|Z_t - Z'_t\|_{M_{t+1}} \quad \text{a.s.}$$

We will now proceed to prove a result on the existence and uniqueness of bounded solutions to our infinite horizon BSDE.

Theorem 2.5. *Suppose that the driver f satisfies the following conditions:*

- for some $\alpha > 0$, $f(\omega, t, y, z) + \alpha y$ is continuous and decreasing in y ,
- f is Lipschitz in z ,
- f satisfies the conditions of Proposition 2.2, and
- $f(\omega, t, 0, 0)$ is uniformly bounded by some $C > 0$.

Then there exists an adapted solution (Y, Z) to the infinite horizon BSDE

$$Y_T = Y_t - \sum_{t \leq u < T} f(\omega, u, Y_u, Z_u) + \sum_{t \leq u < T} Z_u^* M_{u+1}, \quad (12)$$

satisfying $|Y_t| \leq C/\alpha$ for all t , and this solution is unique among bounded adapted solutions.

Furthermore, if (Y^T, Z^T) denotes the (unique) adapted solution to the finite horizon BSDE

$$0 = Y_t^T - \sum_{t \leq u < T} f(\omega, u, Y_u^T, Z_u^T) + \sum_{t \leq u < T} (Z_u^T)^* M_{u+1}, \quad (13)$$

then $Y_t^T \rightarrow Y_t$ in L^∞ as $T \rightarrow \infty$, uniformly on finite sets in t .

Proof. First we will show that if a bounded solution exists, then it is unique. Suppose we have two bounded solutions (Y, Z) and (Y', Z') to (12). Let $\delta Y = Y - Y'$ and $\delta Z = Z - Z'$, and consider the one-step dynamics, given by

$$\delta Y_{t+1} = \delta Y_t - (f(\omega, t, Y_t, Z_t) - f(\omega, t, Y'_t, Z'_t)) + \delta Z_t^* M_{t+1}. \quad (14)$$

By Proposition 2.2, there exists a probability measure \mathbb{Q}_1 such that the process defined by

$$(\omega, t) \mapsto - \sum_{0 \leq u < t} (f(\omega, u, Y'_u, Z_u) - f(\omega, u, Y'_u, Z'_u)) + \sum_{0 \leq u < t} \delta Z_u^* M_{u+1}$$

is a martingale under \mathbb{Q}_1 . It follows that

$$\delta Y_t - f(\omega, t, Y_t, Z_t) + f(\omega, t, Y'_t, Z'_t) = \mathbb{E}_{\mathbb{Q}_1}[\delta Y_{t+1} | \mathcal{F}_t].$$

By rearranging our first assumption, we know that

$$(1 + \alpha)|\delta Y_t| \leq |\delta Y_t - f(\omega, t, Y_t, Z_t) + f(\omega, t, Y'_t, Z'_t)| \quad (15)$$

and hence, by induction, that

$$|\delta Y_t| \leq (1 + \alpha)^{-(r-t)} \mathbb{E}_{\mathbb{Q}_1} [|\delta Y_r| \mid \mathcal{F}_t]$$

for all $r > t$. Note that the above holds \mathbb{Q}_1 -a.s., but since \mathbb{Q}_1 is equivalent to \mathbb{P} on \mathcal{F}_t for every finite t , there is no distinction here between \mathbb{Q}_1 -a.s. and \mathbb{P} -a.s. Since Y and Y' are both uniformly bounded, the same is true of δY . We then deduce, upon taking the limit as $r \rightarrow \infty$ in the above, that $\delta Y_t = 0$ a.s. It follows that $Y = Y'$, up to \mathbb{P} and \mathbb{Q}_1 -indistinguishability. Substituting back into (14), we deduce that $\delta Z_t^* M_{t+1} = 0$, and hence that $Z \sim_M Z'$.

We now proceed to prove existence. Our first assumption ensures that $y \mapsto y - f(\omega, t, y, z)$ is strictly increasing, and it follows from the Lipschitz condition that the driver of (13) satisfies the assumptions of Theorem 1.5, so with terminal condition $Y_T^T = 0$, the solution (Y^T, Z^T) exists uniquely. First we show that Y^T is bounded.

By Proposition 2.2, there exists a probability measure \mathbb{Q}_2^T such that

$$(\omega, t) \mapsto - \sum_{0 \leq u < t} (f(\omega, u, 0, Z_u^T) - f(\omega, u, 0, 0)) + \sum_{0 \leq u < t} (Z_u^T)^* M_{u+1}$$

defines a martingale under \mathbb{Q}_2^T , where we let $Z_t^T = 0$ for $t \geq T$. Using a similar rearrangement as in (15), it follows from the one-step dynamics of (13) that

$$\begin{aligned} (1 + \alpha) |Y_t^T| &\leq |Y_t^T - f(\omega, t, Y_t^T, Z_t^T) + f(\omega, t, 0, Z_t^T)| \\ &\leq \mathbb{E}_{\mathbb{Q}_2^T} [|Y_{t+1}^T| \mid \mathcal{F}_t] + |f(\omega, t, 0, 0)|. \end{aligned}$$

Using the above as the basis of an inductive argument yields

$$|Y_t^T| \leq \mathbb{E}_{\mathbb{Q}_2^T} \left[(1 + \alpha)^{-(r-t)} |Y_r^T| + \sum_{k=0}^{r-t-1} (1 + \alpha)^{-(k+1)} |f(\omega, t+k, 0, 0)| \mid \mathcal{F}_t \right]$$

for all $t < r \leq T$. Setting $r = T$ and using the fact that $Y_T^T = 0$, we deduce that Y^T is uniformly bounded by C/α .

Let $T' > T$. By Proposition 2.2, there exists a probability measure $\mathbb{Q}_3^{T, T'}$ such that

$$(\omega, t) \mapsto - \sum_{0 \leq u < t} (f(\omega, u, Y_u^{T'}, Z_u^T) - f(\omega, u, Y_u^{T'}, Z_u^{T'})) + \sum_{0 \leq u < t} (Z_u^T - Z_u^{T'})^* M_{u+1}$$

defines a martingale under $\mathbb{Q}_3^{T, T'}$. Using the same methods as before, it follows by induction that

$$|Y_t^T - Y_t^{T'}| \leq (1 + \alpha)^{-(r-t)} \mathbb{E}_{\mathbb{Q}_3^{T, T'}} [|Y_r^T - Y_r^{T'}| \mid \mathcal{F}_t]$$

for all $t < r \leq T$. Then, by the boundedness established above,

$$|Y_t^T - Y_t^{T'}| \leq (1 + \alpha)^{-(T-t)} \mathbb{E}_{\mathbb{Q}_3^{T, T'}} [|Y_T^{T'}| \mid \mathcal{F}_t] \leq \frac{C(1 + \alpha)^{-(T-t)}}{\alpha}.$$

Hence, we see that $\{Y_t^T\}_{T \geq t}$ is a Cauchy sequence in T with respect to the L^∞ norm. Therefore, for any $t \geq 0$, there exists an \mathcal{F}_t -measurable random variable Y_t , such that $Y_t^T \rightarrow Y_t$ in L^∞ as $T \rightarrow \infty$, and this convergence is clearly uniform on finite sets in t . Further, $Y = \{Y_t\}_{t \geq 0}$ is also uniformly bounded by C/α .

As $(Z_t^T)^* M_{t+1} = Y_{t+1}^T - \mathbb{E}_{\mathbb{P}}[Y_{t+1}^T | \mathcal{F}_t]$, we deduce that

$$|(Z_t^T - Z_t^{T'})^* M_{t+1}| \leq |Y_{t+1}^T - Y_{t+1}^{T'}| + \mathbb{E}_{\mathbb{P}}[|Y_{t+1}^T - Y_{t+1}^{T'}| | \mathcal{F}_t],$$

and it follows that $\{(Z_t^T)^* M_{t+1}\}_{T \geq t}$ is a Cauchy sequence in T . Hence, there exists an \mathcal{F}_{t+1} -measurable random variable U_{t+1} such that $(Z_t^T)^* M_{t+1} \rightarrow U_{t+1}$ in L^∞ as $T \rightarrow \infty$. By Corollary 1.4, there exists an \mathcal{F}_t -measurable random variable Z_t such that $U_{t+1} = Z_t^* M_{t+1}$.

It follows, from the assumed Lipschitz continuity in z and continuity in y , that $f(\omega, t, Y_t^T, Z_t^T) \rightarrow f(\omega, t, Y_t, Z_t)$ in L^∞ as $T \rightarrow \infty$. Finally, taking the limit as $T \rightarrow \infty$ in the one-step dynamics for Y^T , we deduce that

$$Y_{t+1} = Y_t - f(\omega, t, Y_t, Z_t) + Z_t^* M_{t+1},$$

and that this holds for every $t \geq 0$. Note that these are the one-step dynamics of (12). Hence, (Y, Z) is a bounded adapted solution of (12), which we know is unique among such solutions. \square

Corollary 2.6 (Comparison Theorem). *Let f, f' be two drivers which satisfy the conditions of Theorem 2.5, and let (Y, Z) and (Y', Z') be the bounded solutions of the associated infinite horizon BSDEs. If $f(\omega, t, y, z) \geq f'(\omega, t, y, z)$ for all t, y, z and almost all ω , then $Y \geq Y'$ up to indistinguishability.*

Furthermore, the comparison is strict. That is, if $\mathbb{1}_A Y_t = \mathbb{1}_A Y'_t$ for some t and $A \in \mathcal{F}_t$, then $Y = Y'$ on $A \times \{s \geq t\}$.

Proof. Consider the one-step dynamics of $Y - Y'$, given by

$$Y_{t+1} - Y'_{t+1} = Y_t - Y'_t - (f(\omega, t, Y_t, Z_t) - f'(\omega, t, Y'_t, Z'_t)) + (Z_t - Z'_t)^* M_{t+1}.$$

By Proposition 2.2, there exists a probability measure \mathbb{Q} such that

$$(\omega, t) \mapsto - \sum_{0 \leq u < t} (f'(\omega, u, Y'_u, Z'_u) - f(\omega, u, Y_u, Z_u)) + \sum_{0 \leq u < t} (Z_u - Z'_u)^* M_{u+1}$$

defines a martingale under \mathbb{Q} . It follows that

$$\begin{aligned} & Y_t - Y'_t - f(\omega, t, Y_t, Z_t) + f(\omega, t, Y'_t, Z'_t) \\ &= \mathbb{E}_{\mathbb{Q}}[Y_{t+1} - Y'_{t+1} | \mathcal{F}_t] + f(\omega, t, Y'_t, Z'_t) - f(\omega, t, Y_t, Z_t) \\ &\geq \mathbb{E}_{\mathbb{Q}}[Y_{t+1} - Y'_{t+1} | \mathcal{F}_t]. \end{aligned} \tag{16}$$

As $y \mapsto y - f(\omega, t, y, z)$ is strictly increasing, we see that $Y_t - Y'_t \geq 0$ whenever $\mathbb{E}_{\mathbb{Q}}[Y_{t+1} - Y'_{t+1} | \mathcal{F}_t] \geq 0$. (We can also see that $Y_t - Y'_t = 0$ only if $\mathbb{E}_{\mathbb{Q}}[Y_{t+1} - Y'_{t+1} | \mathcal{F}_t] = 0$, which will imply the strict comparison.)

By induction, we can apply this result to the finite horizon approximation with terminal condition $Y_T^T = (Y^T)'_T = 0$, and deduce that $Y_t^T \geq (Y^T)'_t$ for all $0 \leq t \leq T$. Theorem 2.5 shows that $Y_t^T - (Y^T)'_t \rightarrow Y_t - Y'_t$ for every t , and the result follows. Nonnegativity of the expectation then also guarantees the strict comparison holds. \square

2.2 γ -balanced drivers

We now return to the case when our underlying process X is a (possibly time-inhomogeneous) Markov chain under \mathbb{P} , with associated transition matrix A_t .

Definition 2.7. Let A and B be transition matrices and let $\gamma \in (0, 1)$. We shall say that B is γ -controlled by A (or B is controlled by A with constant γ), and write $B \succeq_\gamma A$, if $B \geq \gamma A$ componentwise. If $B \succeq_\gamma A$ and $A \succeq_\gamma B$ then we will write $A \sim_\gamma B$.

Remark 2.8. Writing $A = [a_{ij}]$ and $B = [b_{ij}]$, it is easy to check that the statement $B \succeq_\gamma A$ is equivalent to the condition

$$\gamma \leq \inf \left\{ \frac{b_{ij}}{a_{ij}} \mid i, j \text{ such that } a_{ij} \neq 0 \right\}.$$

We also note that $A \sim_\gamma B$ if and only if A and B have the same pattern of zero entries, and

$$\frac{b_{ij}}{a_{ij}} \in [\gamma, \gamma^{-1}]$$

for all i, j such that $a_{ij} \neq 0$.

Definition 2.9. We shall say that a driver f is γ -balanced if for some $\gamma \in (0, 1)$, there exists a random field $\psi : \Omega \times \mathbb{N} \times \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}^N$ which is adapted (in the same sense as f), and such that for any $t \geq 0$, $y \in \mathbb{R}$, $z, z' \in \mathbb{R}^N$ and for almost all ω , we have

(i) $f(\omega, t, y, z) - f(\omega, t, y, z') = (z - z')^* (\psi(\omega, t, z, z') - A_t X_t),$

(ii) for each $e_i \in \mathcal{X}$,

$$\frac{e_i^* \psi(\omega, t, z, z')}{e_i^* A_t X_t} \in [\gamma, \gamma^{-1}],$$

where here $0/0 := 1$, and

(iii) $\mathbf{1}^* \psi(\omega, t, z, z') = 1$, where we recall that $\mathbf{1}$ denotes the vector in \mathbb{R}^N with all entries equal to 1.

Lemma 2.10. *If a driver f is γ -balanced, then it is Lipschitz in z , with Lipschitz constant $1/\gamma$.*

Proof. Let $\psi = \psi(\omega, t, z, z')$ be the random field associated with f in Definition 2.9. Let $t \geq 0$ and let F be an atom of \mathcal{F}_t . Suppose that X_t takes the value e_j on F . Writing $A_t = [a_{ij}]$ and $\psi = [\psi_i]$, we note, from property (ii) of Definition 2.9, that $\psi_i \leq \gamma^{-1}a_{ij}$. Then, for any \mathcal{F}_t -measurable random variables Y_t, Z_t, Z'_t evaluated on F , we have

$$\begin{aligned} |f(\omega, t, Y_t, Z_t) - f(\omega, t, Y_t, Z'_t)| &= \left| \sum_{i=1}^N \psi_i (Z_t - Z'_t)^* (e_i - A_t e_j) \right| \\ &\leq \frac{1}{\gamma} \sum_{i=1}^N a_{ij} |(Z_t - Z'_t)^* (e_i - A_t e_j)| = \frac{1}{\gamma} \mathbb{E} \left[|(Z_t - Z'_t)^* M_{t+1}| \mid F \right]. \end{aligned}$$

This holds on every atom F of \mathcal{F}_t , and the result follows. \square

Lemma 2.11. *If f is a γ -balanced driver, then f satisfies the conditions of Proposition 2.2.*

Proof. Let $\psi = \psi(\omega, t, z, z')$ be the random field associated with f in Definition 2.9. Let $t \geq 0, y \in \mathbb{R}, z, z' \in \mathbb{R}^N$ and let F be any atom of \mathcal{F}_t . Recall the notation $\mathbb{J}_t^F := \{i : \mathbb{P}(X_{t+1} = e_i \mid F) > 0\}$. Writing $\psi = [\psi_i]$, we see, from the properties of ψ given in Definition 2.9, that $\psi_i > 0 \iff e_i^* A_t X_t > 0 \iff i \in \mathbb{J}_t^F$, and hence that $\sum_{i \in \mathbb{J}_t^F} \psi_i = 1$, and

$$f(\omega, t, y, z) - f(\omega, t, y, z') = \sum_{i \in \mathbb{J}_t^F} \psi_i (z - z')^* (e_i - \mathbb{E}[X_{t+1} \mid \mathcal{F}_t]). \quad (17)$$

Write $S_t^F := \{(z - z')^* (e_i - \mathbb{E}[X_{t+1} \mid \mathcal{F}_t]) : i \in \mathbb{J}_t^F\}$, where the conditional expectation is evaluated on F . Note that

$$\mathbb{E}[(z - z')^* (X_{t+1} - \mathbb{E}[X_{t+1} \mid \mathcal{F}_t]) \mid F] = 0, \quad (18)$$

and hence that $\min(S_t^F) \leq 0$ and $\max(S_t^F) \geq 0$. If $\min(S_t^F) < 0$, then (17) implies that

$$f(\omega, t, y, z) - f(\omega, t, y, z') > \min(S_t^F).$$

If $\min(S_t^F) = 0$, then (18) implies that $S_t^F = \{0\}$, and it clearly follows from (17) that

$$f(\omega, t, y, z) - f(\omega, t, y, z') = 0.$$

This establishes the conditions of Proposition 2.2. \square

Note that if $f : \Omega \times \mathbb{N} \times \mathbb{R}^N \rightarrow \mathbb{R}$ is γ -balanced then, provided that $f(\omega, t, 0)$ is bounded, it follows from the above that the driver of the discounted BSDE (8) satisfies the conditions of Theorem 2.5, so we have existence and uniqueness of bounded solutions.

The following lemma is a discrete time version of [2, Lemma 2].

Lemma 2.12. Fix $\gamma \in (0, 1)$. Let $\{f^u\}_{u \in \mathcal{U}}$ be a non-empty family of γ -balanced drivers that are independent of y , and for each $u \in \mathcal{U}$ let $\psi^u(\omega, t, z, z')$ be the random field associated with f^u . Let

$$g(\omega, t, z) := \inf_{u \in \mathcal{U}} \{f^u(\omega, t, z)\}.$$

Suppose $|g(\omega, t, z)| < \infty$ for all t, z and almost all ω . Then g is also γ -balanced, and similarly for $\sup_{u \in \mathcal{U}} \{f^u\}$.

Proof. By property (i) of Definition 2.9, the functions $\{f^u\}_{u \in \mathcal{U}}$ are continuous in z , uniformly in u . It follows that g is a continuous, and hence measurable function of z .

Fix $t \geq 0$. As \mathcal{F}_t can be made up from finitely many atoms, there exists a single null set $N_t \in \mathcal{F}_t$ such that $N' \subseteq N_t$ for every null set $N' \in \mathcal{F}_t$. For simplicity, define $g(\omega, t, z) = 0$ for $\omega \in N_t$.

For every $\omega \notin N_t$ and every $z, z' \in \mathbb{R}^N$, we can find a sequence $\{u_n\}_{n \geq 1}$ in \mathcal{U} (dependent on z, z') such that, omitting the arguments ω, t for clarity,

$$\lim_{n \rightarrow \infty} (z - z')^* \psi^{u_n}(z, z') = \inf_{u \in \mathcal{U}} \{(z - z')^* \psi^u(z, z')\}. \quad (19)$$

For each n , $\psi^{u_n}(z, z')$ is a stochastic vector, i.e. it belongs to the compact subset $\{\varphi \in \mathbb{R}^N : \varphi_i \geq 0, \sum_{i=1}^N \varphi_i = 1\}$. Hence, there exists a subsequence $\{u_{n_k}\}_{k \geq 1}$ such that

$$\lim_{k \rightarrow \infty} \psi^{u_{n_k}}(z, z') = \underline{\psi}(z, z'),$$

where $\underline{\psi}(z, z')$ is a stochastic vector which satisfies conditions (ii) and (iii) of Definition 2.9. In addition, $\underline{\psi}$ can be constructed in such a way that it is measurable in z, z' (see [4, Theorem A.10.5]). Further, by (19), we have

$$(z - z')^* \underline{\psi}(z, z') = \inf_{u \in \mathcal{U}} \{(z - z')^* \psi^u(z, z')\}.$$

By a similar argument, there exists a stochastic vector $\bar{\psi}(z, z')$, which also satisfies conditions (ii) and (iii) of Definition 2.9, is measurable in z, z' , and is such that

$$(z - z')^* \bar{\psi}(z, z') = \sup_{u \in \mathcal{U}} \{(z - z')^* \psi^u(z, z')\}.$$

We have, for each $u \in \mathcal{U}$,

$$f^u(z) - f^u(z') = (z - z')^* (\psi^u(z, z') - A_t X_t),$$

from which we deduce that

$$\begin{aligned} (z - z')^* \underline{\psi}(z, z') &\leq g(z) - g(z') + (z - z')^* A_t X_t \\ &\leq (z - z')^* \bar{\psi}(z, z'). \end{aligned}$$

With $0/0 := 1/2$, it follows that

$$\mu := \frac{(z - z')^* \bar{\psi}(z, z') - (g(z) - g(z')) - (z - z')^* A_t X_t}{(z - z')^* (\bar{\psi}(z, z') - \underline{\psi}(z, z'))} \in [0, 1].$$

Define $\hat{\psi}(z, z') := \mu \underline{\psi}(z, z') + (1 - \mu) \bar{\psi}(z, z')$. Then

$$g(z) - g(z') = (z - z')^* (\hat{\psi}(z, z') - A_t X_t),$$

and we see that the random field $\hat{\psi}(\omega, t, z, z')$ satisfies all the conditions of Definition 2.9 for the driver g , and hence that g is γ -balanced.

The proof for $\sup_{u \in \mathcal{U}} \{f^u\}$ is similar. \square

2.3 Markovian solutions

We showed in Lemma 1.7 that, if X , f and the terminal condition ξ are all Markovian, i.e. at time t they only depend on ω through the value of X_t , then so is the solution (Y, Z) of the corresponding finite horizon BSDE. We shall now extend this result to our infinite horizon equations. As previously, we assume throughout that X is a Markov chain with transition matrix A_t .

Lemma 2.13. *Let f be a Markovian driver satisfying the conditions of Theorem 2.5, and let (Y, Z) be the unique bounded solution of the associated infinite horizon BSDE. Then there exists a deterministic function $v : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}$ such that, for all t , $Y_t = v(t, X_t)$ and $e_k^* Z_t = v(t + 1, e_k)$ for all $e_k \in \mathcal{X}$.*

Moreover, writing \mathbf{v}_{t+1} for the vector in \mathbb{R}^N with entries $e_k^ \mathbf{v}_{t+1} = v(t + 1, e_k)$, the function v satisfies (5) for all t and all $e_k \in \mathcal{X}$ such that $\mathbb{P}(X_t = e_k) > 0$.*

Proof. Let (Y^T, Z^T) denote the (unique adapted) solution of the finite horizon BSDE (13). Note that we are in the situation of Lemma 1.7, with $\varphi \equiv 0$. Hence, for each $T \geq 0$, there exists a function $v^T : \{0, 1, \dots, T\} \times \mathcal{X} \rightarrow \mathbb{R}$ such that $Y_t^T = v^T(t, X_t)$ and $e_k^* Z_t^T = v^T(t + 1, e_k)$ for all $e_k \in \mathcal{X}$.

Recall, from Theorem 2.5 that $Y_t^T \rightarrow Y_t$ as $T \rightarrow \infty$ with respect to the L^∞ norm, uniformly on finite sets in t , from which we deduce the existence of a function $v : \mathbb{N} \times \mathcal{X} \rightarrow \mathbb{R}$ such that $Y_t = v(t, X_t)$ for all t .

Note that the one-step dynamics of the infinite horizon BSDE (12) are identical to the one-step dynamics of the corresponding finite horizon BSDE (13). Therefore, given the value of Y_{t+1} , we can construct the value of Z_t exactly as in the finite horizon case. In fact, since Y_{t+1} is Markovian, we can apply the same argument as in the proof of Lemma 1.7 to deduce that $e_k^* Z_t = v(t + 1, e_k)$ for all $e_k \in \mathcal{X}$. The equation (5) follows from the one-step dynamics of (12), just as in the finite horizon case. \square

Corollary 2.14. *Recall the conditions of Lemma 2.13. Suppose that X is time-homogeneous, so that $A = A_t$ does not depend on t , and that the driver $f = \tilde{f}(X_t, y, z)$ is also independent of t . Then the same is true of the function v given in Lemma 2.13.*

Proof. The BSDE (12) is given in this case by

$$Y_T = Y_t - \sum_{t \leq u < T} \tilde{f}(X_u, Y_u, Z_u) + \sum_{t \leq u < T} Z_u^* (X_{u+1} - A X_u).$$

Note that, given X_u , the summands are independent of u . By the uniqueness of solutions established in Theorem 2.5, the same must therefore be true of the solution (Y, Z) , and the result follows. \square

3 Uniformly Ergodic Markov Chains

The main aim of this section is to obtain ergodicity estimates for a discrete time Markov chain which hold uniformly for a suitable class of perturbations of its transition matrix. We will loosely follow the argument given in Section 3 of [6], by adapting the main steps into the discrete time setting. Henceforth, we will assume that all the Markov chains we consider are time-homogeneous.

3.1 Uniform ergodicity

The following lemma will turn out to be a useful tool for improving bounds on expectations, and improves slightly the estimate given in [6, Lemma 3.3].

Lemma 3.1. *Let T be a random variable, and consider $G(\beta) = \sup_{\nu} \mathbb{E}_{\nu}[e^{\beta T}]$, where ν is a parametrisation of probability measures under which the expectation is taken. Suppose there exist finite constants $\beta^*, K > 0$ such that $G(\beta^*) \leq K$. Then, for any $\epsilon > 0$,*

$$G(\beta) \leq 1 + \epsilon \quad \text{for all } \beta \in \left[0, \beta^* \left(\frac{\epsilon}{K} \wedge 1\right)\right].$$

Proof. Let $0 \leq c \leq \left(\frac{\epsilon}{K} \wedge 1\right)$, so that $e^{c\beta^* T} \leq 1 + ce^{\beta^* T}$ for any $T \geq 0$. Then, for any given measure ν in the class considered, we have

$$\mathbb{E}_{\nu}[e^{c\beta^* T}] \leq \mathbb{E}_{\nu}[\mathbb{1}_{\{T < 0\}} + \mathbb{1}_{\{T \geq 0\}}(1 + ce^{\beta^* T})] \leq 1 + c\mathbb{E}_{\nu}[e^{\beta^* T}] \leq 1 + \epsilon.$$

Taking the supremum over all measures ν , we deduce the result. \square

Definition 3.2. Let $\tilde{\mathcal{M}}(\mathcal{X})$ denote the vector space (over \mathbb{R}) of finite signed measures on $(\mathcal{X}, \mathcal{P}(\mathcal{X}))$. We will endow this space with the total variation norm, given by

$$\|\mu\|_{TV} := \frac{1}{2} \sum_{x \in \mathcal{X}} |\mu(\{x\})|,$$

for $\mu \in \tilde{\mathcal{M}}(\mathcal{X})$. By standard results, $\tilde{\mathcal{M}}(\mathcal{X})$ is a Banach space. Let \mathcal{M} denote the set of probability measures on $(\mathcal{X}, \mathcal{P}(\mathcal{X}))$. Note that \mathcal{M} is a closed subset of $\tilde{\mathcal{M}}(\mathcal{X})$.

Definition 3.3. Let X be a Markov chain on \mathcal{X} , and P_t its transition operator, so that $P_t \mu$ is the law of X_t given $X_0 \sim \mu$. We say that the chain X is *uniformly ergodic* if there exists a probability measure π on \mathcal{X} , and constants $R, \rho > 0$ such that

$$\sup_{\mu \in \mathcal{M}} \|P_t \mu - \pi\|_{TV} \leq R e^{-\rho t} \quad \text{for all } t \geq 0.$$

In this case π is the unique invariant measure for X .

Under our assumptions of discrete time chains on a finite state space, we have the following simple classification of uniformly ergodic chains.

Proposition 3.4. *Let X and Y be two independent copies of a Markov chain on the finite state space \mathcal{X} , and let $T = \inf\{t \geq 1 : X_t = Y_t\}$ be the first meeting time of the two chains. Then the following statements are equivalent:*

- (i) X is uniformly ergodic.
- (ii) X is irreducible up to transient states (i.e. X has only one closed communicating class) and is aperiodic.
- (iii) There exists $\beta > 0$ such that

$$G^*(\beta) := \sup_{x,y \in \mathcal{X}} \mathbb{E}_{xy} [e^{\beta T}] < \infty,$$

where \mathbb{E}_{xy} denotes expectation conditional on $X_0 = x$ and $Y_0 = y$.

Proof. The implications (i) \Rightarrow (ii) \Rightarrow (iii) are straightforward. The implication (iii) \Rightarrow (i) will be demonstrated in the course of the proof of Theorem 3.5. \square

Note that, by Lemma 3.1, we can make $G^*(\beta)$ arbitrarily close to 1 by choosing β sufficiently small.

3.2 Split chains

Recall from Definition 2.7, that we say B is γ -controlled by A , and write $B \succeq_\gamma A$, if $B \geq \gamma A$ componentwise, and we write $A \sim_\gamma B$ if both $B \succeq_\gamma A$ and $A \succeq_\gamma B$.

Theorem 3.5. *Let A and B be transition matrices such that B is γ -controlled by A for some $\gamma \in (0, 1)$. If, under the measure induced by A , the chain X is uniformly ergodic, then X is also uniformly ergodic under the measure induced by B . Furthermore, the constants R and ρ in Definition 3.3 depend only on A and γ , and R can be made arbitrarily close to 1 (with a corresponding decrease in ρ).*

In our discrete time, finite state setting, the fact that X is uniformly ergodic under the measure induced by B is trivial. The interesting part of Theorem 3.5 is the claim there exist constants of ergodicity R, ρ which hold for *all* transition matrices B such that $B \succeq_\gamma A$. We shall make use of this result in Section 4, in the proof of existence of bounded Markovian solutions to Ergodic BSDEs.

The proof of this theorem is the main purpose of this section. To do this we will use a Nummelin splitting argument, the broad theory of which is described in Meyn and Tweedie [13, Chapter 5]. We will assume the conditions of Theorem 3.5 throughout the remainder of the section.

Definition 3.6. Define the *split space* of \mathcal{X} to be $\tilde{\mathcal{X}} := \mathcal{X} \times \{0, 1\}$. We will also denote the *layers* of the splitting by $\mathcal{X}_0 := \mathcal{X} \times \{0\}$ and $\mathcal{X}_1 := \mathcal{X} \times \{1\}$, so that

$\tilde{\mathcal{X}} = \mathcal{X}_0 \cup \mathcal{X}_1$. Given $y \in \mathcal{X}$, denote the corresponding elements in \mathcal{X}_0 and \mathcal{X}_1 by \check{y}_0 and \check{y}_1 respectively.

Let $\nu \in \tilde{\mathcal{M}}(\mathcal{X})$. For each $y \in \mathcal{X}$, let

$$\check{\nu}(\{\check{y}_0\}) = (1 - \gamma)\nu(\{y\}), \quad \check{\nu}(\{\check{y}_1\}) = \gamma\nu(\{y\}).$$

This defines a split measure $\check{\nu}$ on $(\tilde{\mathcal{X}}, \mathcal{P}(\tilde{\mathcal{X}}))$. Similarly, we define the splitting of a column vector $\phi \in \mathbb{R}^N$ via the splitting map $\check{\Pi} : \mathbb{R}^N \rightarrow \mathbb{R}^{2N}$, defined by

$$\check{\Pi}(\phi) := (1 - \gamma) \begin{bmatrix} \phi \\ 0 \end{bmatrix} + \gamma \begin{bmatrix} 0 \\ \phi \end{bmatrix} = \begin{bmatrix} (1 - \gamma)\phi \\ \gamma\phi \end{bmatrix} \in \mathbb{R}^{2N},$$

where here 0 is the zero vector in \mathbb{R}^N .

Remark 3.7. Note that these splittings are consistent in the following way. Consider a measure ν as a vector in \mathbb{R}^N , so that $\nu(\{y\}) = y^*\nu$ for each $y \in \mathcal{X}$. Similarly, consider the split measure $\check{\nu}$ as a vector in \mathbb{R}^{2N} , so that

$$\check{\nu}(\{\check{y}_0\}) = \begin{bmatrix} y \\ 0 \end{bmatrix}^* \check{\nu}, \quad \check{\nu}(\{\check{y}_1\}) = \begin{bmatrix} 0 \\ y \end{bmatrix}^* \check{\nu}$$

for each $y \in \mathcal{X}$. Then we see that the split measure $\check{\nu}$ corresponds precisely to the split vector $\check{\Pi}(\nu)$.

Notice that if ν is a probability measure on \mathcal{X} , then $\check{\nu}$ is a probability measure on $\tilde{\mathcal{X}}$. Equivalently, if ϕ is a stochastic vector in \mathbb{R}^N , then $\check{\Pi}(\phi)$ is a stochastic vector in \mathbb{R}^{2N} .

Definition 3.8. We can define the splitting of a matrix by splitting its columns as in the previous definition. That is, given a matrix M , define

$$\check{M} := \begin{bmatrix} (1 - \gamma)M \\ \gamma M \end{bmatrix}.$$

Definition 3.9. Recall the assumptions of Theorem 3.5, in particular that A and B are transition matrices with $B \succeq_\gamma A$ for some $\gamma \in (0, 1)$. Let

$$C = (1 - \gamma)^{-1}(B - \gamma A),$$

and let $\mathcal{B} = [\check{C} | \check{A}]$, i.e. the augmented matrix formed from \check{C} and \check{A} . Note that C and \mathcal{B} are both transition matrices.

We can now define a new Markov chain \check{X} on the split space $\tilde{\mathcal{X}}$, which jumps according to transition matrix \mathcal{B} . Intuitively, transitions occur from a state $\check{x} \in \mathcal{X}_0$ according to the vector Cx , and from a state $\check{x} \in \mathcal{X}_1$ according to the vector Ax (where x is the projection of \check{x} in \mathcal{X}), except that the result is randomly split between the layers \mathcal{X}_0 and \mathcal{X}_1 with probabilities $1 - \gamma$ and γ respectively.

Lemma 3.10. *Let $\phi \in \mathbb{R}^N$. Then $\check{\Pi}(B\phi) = \mathcal{B}\check{\Pi}(\phi)$.*

Proof. By the definitions of $\check{\Pi}$, \mathcal{B} and C , we have

$$\begin{aligned}\mathcal{B}\check{\Pi}(\phi) &= (1 - \gamma)\check{C}\phi + \gamma\check{A}\phi \\ &= (1 - \gamma)\check{\Pi}(C\phi) + \gamma\check{\Pi}(A\phi) = \check{\Pi}(B\phi).\end{aligned}$$

□

Proposition 3.11. *Let X be a Markov chain on \mathcal{X} with initial distribution $X_0 \sim \nu$ and transition matrix B , and let \check{X} be a Markov chain on $\check{\mathcal{X}}$ with initial distribution $\check{X}_0 \sim \check{\nu}$ and transition matrix \mathcal{B} . Then X and \check{X} have the same marginal distributions, ignoring the splitting. That is, for any t and any $y \in \mathcal{X}$,*

$$y^*\mathbb{E}[X_t] = \mathbb{P}(X_t = y) = \mathbb{P}(\check{X}_t \in \{\check{y}_0, \check{y}_1\}).$$

Proof. Consider ν as a vector in \mathbb{R}^N and $\check{\nu}$ as a vector in \mathbb{R}^{2N} , as in Remark 3.7. By the Chapman–Kolmogorov equation, the distribution of X_t is given by $\mathbb{E}[X_t] = B^t\nu$, and the distribution of \check{X}_t is given by $\mathbb{E}[\check{X}_t] = \mathcal{B}^t\check{\nu}$. By applying Lemma 3.10 t times, we obtain $\check{\Pi}(B^t\nu) = \mathcal{B}^t\check{\Pi}(\nu)$. Recall from Remark 3.7 that $\check{\Pi}(\nu) = \check{\nu}$. Hence, we have $\check{\Pi}(\mathbb{E}[X_t]) = \mathbb{E}[\check{X}_t]$, and the result follows. □

The following lemma follows easily from the definition of \mathcal{B} and the basic notion of conditional probability.

Lemma 3.12. *Let \check{X} be as in the above proposition. Write \mathbb{E}^* for the expectation conditioned on \check{X} never leaving the layer it starts in (i.e. never jumps from \mathcal{X}_0 to \mathcal{X}_1 , or from \mathcal{X}_1 to \mathcal{X}_0). Then, under \mathbb{E}^* , \check{X} jumps according to the transition matrix*

$$\begin{bmatrix} C & 0 \\ 0 & A \end{bmatrix}.$$

Intuitively, under \mathbb{E}^* , \check{X} either jumps in \mathcal{X}_0 following transition matrix C , or in \mathcal{X}_1 following transition matrix A , depending on which layer it starts in.

3.3 Exponential moment bounds

We will now consider two independent copies of our Markov chain on the split space, and show that the first meeting time of these chains admits exponential moments.

Theorem 3.13. *Let \check{X} and \check{Y} be two independent Markov chains on the split space $\check{\mathcal{X}}$, each following transition matrix \mathcal{B} , as defined in Definition 3.9. Let $\check{S} := \inf\{t \geq 0 : \check{X}_t = \check{Y}_t\}$. Then, for any $\epsilon > 0$, there exists $\tilde{\beta} > 0$ such that*

$$H^*(\beta) := \sup_{\check{x}, \check{y} \in \check{\mathcal{X}}} \mathbb{E}_{\check{x}\check{y}}[e^{\beta\check{S}}] \leq 1 + \epsilon \quad \text{for all } \beta \in [0, \tilde{\beta}],$$

where $\mathbb{E}_{\check{x}\check{y}}$ is the expectation conditional on $\check{X}_0 = \check{x}$ and $\check{Y}_0 = \check{y}$. Furthermore, $\tilde{\beta}$ does not depend on B except through A and γ .

The proof of this theorem will be done in a number of steps. First, we will introduce some more notation.

Let K_t denote the number of jumps, up to time t , that result in both \check{X} and \check{Y} being in \mathcal{X}_1 when they were not both in \mathcal{X}_1 previously, and denote by t_k the time of the k^{th} such transition. That is, let $t_0 = 0$ and

$$t_k = \inf\{t > t_{k-1} : (\check{X}_t, \check{Y}_t \in \mathcal{X}_1) \setminus (\check{X}_{t-1}, \check{Y}_{t-1} \in \mathcal{X}_1)\}$$

for $k \geq 1$. Finally, let

$$\check{T} = \inf\{t \geq 1 : (\check{X}_t = \check{Y}_t \in \mathcal{X}_1) \cap (\check{X}_{t-1}, \check{Y}_{t-1} \in \mathcal{X}_1)\},$$

i.e. the first time the chains meet in \mathcal{X}_1 , both having jumped from a state in \mathcal{X}_1 .

Lemma 3.14. *Let $G^*(\beta)$ be the function defined in Proposition 3.4, that is, the supremum over starting states of the moment generating function of the first meeting time on the basic (unsplit) state space \mathcal{X} of two independent Markov chains with transition matrix A . Then, for any $\check{x}, \check{y} \in \check{\mathcal{X}}$,*

$$\mathbb{E}_{\check{x}\check{y}}[e^{\beta\check{T}}] \leq G^*(2\beta)^{\frac{1}{2}} \sum_{k=0}^{\infty} \mathbb{E}_{\check{x}\check{y}}[e^{2\beta t_k}]^{\frac{1}{2}} \mathbb{P}_{\check{x}\check{y}}(K_{\check{T}} = k)^{\frac{1}{2}}, \quad (20)$$

where $\mathbb{P}_{\check{x}\check{y}}$ is the probability conditioned on $\check{X}_0 = \check{x}$ and $\check{Y}_0 = \check{y}$.

Proof. Note that, conditioned on $\{K_{\check{T}} = k\}$, \check{X} and \check{Y} do not leave \mathcal{X}_1 between t_k , the k^{th} time they arrive in \mathcal{X}_1 when they were not both in \mathcal{X}_1 previously, and \check{T} , the time at which they next meet in \mathcal{X}_1 . Write \mathbb{E}^* for the expectation conditioned on both \check{X} and \check{Y} not leaving the layers they start in. Then

$$\mathbb{E}_{\check{x}\check{y}}[e^{2\beta(\check{T}-t_k)} \mid K_{\check{T}} = k] \leq \sup_{x,y \in \mathcal{X}_1} \mathbb{E}_{xy}^*[e^{2\beta\check{T}}] = G^*(2\beta),$$

since, by Lemma 3.12, under \mathbb{E}^* , \check{X} and \check{Y} jump in \mathcal{X}_1 following transition matrix A . Then

$$\begin{aligned} \mathbb{E}_{\check{x}\check{y}}[e^{\beta\check{T}} \mathbb{1}_{\{K_{\check{T}}=k\}}] &= \mathbb{E}_{\check{x}\check{y}}\left[e^{\beta t_k} \mathbb{E}_{\check{x}\check{y}}\left[e^{\beta(\check{T}-t_k)} \mathbb{1}_{\{K_{\check{T}}=k\}} \mid t_k\right]\right] \\ &\leq \mathbb{E}_{\check{x}\check{y}}[e^{2\beta t_k}]^{\frac{1}{2}} \mathbb{E}_{\check{x}\check{y}}\left[\left(\mathbb{E}_{\check{x}\check{y}}\left[e^{\beta(\check{T}-t_k)} \mathbb{1}_{\{K_{\check{T}}=k\}} \mid t_k\right]\right)^2\right]^{\frac{1}{2}} \\ &\leq \mathbb{E}_{\check{x}\check{y}}[e^{2\beta t_k}]^{\frac{1}{2}} \mathbb{E}_{\check{x}\check{y}}\left[\mathbb{E}_{\check{x}\check{y}}\left[e^{2\beta(\check{T}-t_k)} \mathbb{1}_{\{K_{\check{T}}=k\}} \mid t_k\right]\right]^{\frac{1}{2}} \\ &= \mathbb{E}_{\check{x}\check{y}}[e^{2\beta t_k}]^{\frac{1}{2}} \left(\mathbb{E}_{\check{x}\check{y}}\left[e^{2\beta(\check{T}-t_k)} \mid K_{\check{T}} = k\right] \mathbb{P}_{\check{x}\check{y}}(K_{\check{T}} = k)\right)^{\frac{1}{2}} \\ &\leq \mathbb{E}_{\check{x}\check{y}}[e^{2\beta t_k}]^{\frac{1}{2}} (G^*(2\beta) \mathbb{P}_{\check{x}\check{y}}(K_{\check{T}} = k))^{\frac{1}{2}}, \end{aligned}$$

and the result follows. \square

We shall now seek to bound the components of the sum in Lemma 3.14.

Lemma 3.15. *For states $x, y \in \mathcal{X}$, consider a pair of independent basic (un-split) chains X, Y following transition matrix A , with starting values $X_0 = x$ and $Y_0 = y$. For $n \geq 1$, let $q(x, y; n)$ denote the probability that the first positive time the chains meet is at time n . Let*

$$q_\gamma := \inf_{x, y \in \mathcal{X}} \sum_{n=1}^{\infty} q(x, y; n) \gamma^{2n},$$

where we recall that $\gamma \in (0, 1)$. Then $q_\gamma \in (0, 1)$ and, for any $\check{x}, \check{y} \in \check{\mathcal{X}}$ and any $k \geq 0$,

$$\mathbb{P}_{\check{x}\check{y}}(K_{\check{T}} = k) \leq \gamma^2 (1 - q_\gamma)^{k-1}.$$

Proof. The fact that $q_\gamma \in (0, 1)$ is clear, as we are working with a finite state space, and the independent chains almost surely meet after a finite time.

Now, suppose first that $\check{x}, \check{y} \in \mathcal{X}_1$. Then, under $\mathbb{P}_{\check{x}\check{y}}$, $\{K_{\check{T}} = 0\}$ is the event that the split chains \check{X} and \check{Y} meet before either of them leaves \mathcal{X}_1 . As they are independent, at each transition the probability that both chains will jump to \mathcal{X}_1 is γ^2 . Further, by Lemma 3.12, for as long as the chains remain in \mathcal{X}_1 , they jump following transition matrix A . Hence,

$$\mathbb{P}_{\check{x}\check{y}}(K_{\check{T}} = 0) = \sum_{n=1}^{\infty} q(x, y; n) \gamma^{2n} \leq \gamma^2.$$

Now let $k \geq 1$. Each time the chains both arrive in \mathcal{X}_1 when they were not both in \mathcal{X}_1 previously, they will either meet before one of them leaves \mathcal{X}_1 , or they will not. Considering these as successes and failures, a geometric trials argument yields

$$\begin{aligned} \mathbb{P}_{\check{x}\check{y}}(K_{\check{T}} = k) &= \mathbb{E}_{\check{x}\check{y}} \left[\left(\sum_{n=1}^{\infty} q(\check{X}_{t_k}, \check{Y}_{t_k}; n) \gamma^{2n} \right) \prod_{i=0}^{k-1} \left(1 - \sum_{n=1}^{\infty} q(\check{X}_{t_i}, \check{Y}_{t_i}; n) \gamma^{2n} \right) \right] \\ &\leq \gamma^2 (1 - q_\gamma)^k, \end{aligned}$$

where, in the notation of Definition 3.6, we define $q(\check{x}_1, \check{y}_1; n) = q(x, y; n)$ for notational simplicity. Now suppose that either $\check{x} \in \mathcal{X}_0$ or $\check{y} \in \mathcal{X}_0$. Clearly the chains must first both arrive in \mathcal{X}_1 before they can meet in \mathcal{X}_1 , so $\mathbb{P}_{\check{x}\check{y}}(K_{\check{T}} = 0) = 0$. For $k \geq 1$, we have

$$\begin{aligned} \mathbb{P}_{\check{x}\check{y}}(K_{\check{T}} = k) &= \mathbb{E}_{\check{x}\check{y}} \left[\left(\sum_{n=1}^{\infty} q(\check{X}_{t_k}, \check{Y}_{t_k}; n) \gamma^{2n} \right) \prod_{i=1}^{k-1} \left(1 - \sum_{n=1}^{\infty} q(\check{X}_{t_i}, \check{Y}_{t_i}; n) \gamma^{2n} \right) \right] \\ &\leq \gamma^2 (1 - q_\gamma)^{k-1}. \end{aligned}$$

Putting this together, we deduce the result. \square

Lemma 3.16. *Let q_γ be as in the previous lemma. Then there exists $\beta > 0$ such that, for any $\tilde{x}, \tilde{y} \in \tilde{\mathcal{X}}$ and any $k \geq 0$,*

$$\mathbb{E}_{\tilde{x}\tilde{y}}[e^{2\beta t_k}] \leq \left(\frac{1}{2} \left(1 + \frac{1}{1 - q_\gamma}\right)\right)^k.$$

Further, this β only depends on γ and q_γ .

Proof. It will be useful to consider two independent geometric random variables

$$Q_1 \sim \text{Geom}(1 - \gamma^2) \quad \text{and} \quad Q_2 \sim \text{Geom}(\gamma^2),$$

where we adopt the definition of the probability mass function of $Q \sim \text{Geom}(p)$ as being given by $\mathbb{P}(Q = k) = (1 - p)^{k-1}p$ for $k \geq 1$.

If $\tilde{x}, \tilde{y} \in \mathcal{X}_1$ then t_1 is the time taken for either of the chains to leave \mathcal{X}_1 , plus the time taken for both chains to return to \mathcal{X}_1 . It follows from an application of the strong Markov property that these times are independent. We therefore have that $t_1 \stackrel{d}{=} Q_1 + Q_2$. If either $\tilde{x} \in \mathcal{X}_0$ or $\tilde{y} \in \mathcal{X}_0$ then t_1 is just the time taken for both chains to arrive in \mathcal{X}_1 , so in this case $t_1 \stackrel{d}{=} Q_2$. Similarly, for $k \geq 2$, $t_k - t_{k-1}$ is the time taken for either of the chains to leave \mathcal{X}_1 , plus the time taken for both chains to return to \mathcal{X}_1 , and these times are independent, so it follows that $t_k - t_{k-1} \stackrel{d}{=} Q_1 + Q_2$. Hence, for all $k \geq 1$,

$$\mathbb{E}_{\tilde{x}\tilde{y}}[e^{2\beta(t_k - t_{k-1})}] \leq \mathbb{E}[e^{2\beta(Q_1 + Q_2)}].$$

By Lemma 3.1, we can make a new choice of $\beta > 0$, which depends only on γ and q_γ , such that

$$\mathbb{E}[e^{2\beta(Q_1 + Q_2)}] \leq \frac{1}{2} \left(1 + \frac{1}{1 - q_\gamma}\right).$$

It follows from another application of the strong Markov property that $t_k - t_{k-1}$ is independent of t_1, \dots, t_{k-1} for each k . Then, for $k \geq 1$,

$$\mathbb{E}_{\tilde{x}\tilde{y}}[e^{2\beta t_k}] = \prod_{j=1}^k \mathbb{E}_{\tilde{x}\tilde{y}}[e^{2\beta(t_j - t_{j-1})}] \leq \mathbb{E}[e^{2\beta(Q_1 + Q_2)}]^k \leq \left(\frac{1}{2} \left(1 + \frac{1}{1 - q_\gamma}\right)\right)^k.$$

□

Proof of Theorem 3.13. Substituting the bounds given by Lemmas 3.15 and 3.16 into (20), we obtain

$$\begin{aligned} \mathbb{E}_{\tilde{x}\tilde{y}}[e^{\beta \tilde{T}}] &\leq G^*(2\beta)^{\frac{1}{2}} \sum_{k=0}^{\infty} \left[\left(\frac{1}{2} \left(1 + \frac{1}{1 - q_\gamma}\right)\right)^k \cdot \gamma^2 (1 - q_\gamma)^{k-1} \right]^{\frac{1}{2}} \\ &= \gamma \left(\frac{G^*(2\beta)}{1 - q_\gamma}\right)^{\frac{1}{2}} \left(1 - \left(1 - \frac{q_\gamma}{2}\right)^{\frac{1}{2}}\right)^{-1}. \end{aligned}$$

This is valid provided that β is chosen to be at least as small as the β given in Lemma 3.16. The expression above is finite provided that β is also sufficiently small to ensure that $G^*(2\beta)$ is finite, which can be guaranteed by Proposition 3.4. Note that this choice of β depends only on A , γ and q_γ .

With this β , we have

$$H^*(\beta) = \sup_{\tilde{x}, \tilde{y} \in \tilde{\mathcal{X}}} \mathbb{E}_{\tilde{x}\tilde{y}}[e^{\beta\tilde{S}}] \leq \sup_{\tilde{x}, \tilde{y} \in \tilde{\mathcal{X}}} \mathbb{E}_{\tilde{x}\tilde{y}}[e^{\beta\tilde{T}}] < \infty.$$

By Lemma 3.1, for any $\epsilon > 0$ there exists $\tilde{\beta} > 0$ such that $H^*(\beta) \leq 1 + \epsilon$ for all $\beta \in [0, \tilde{\beta}]$. Finally, we see that $\tilde{\beta}$ depends only on A , γ , ϵ , and on q_γ , which is itself a function of A and γ . In particular, $\tilde{\beta}$ does not depend on B except through A and γ . \square

Corollary 3.17. *Let X and Y be two independent copies of the Markov chain on \mathcal{X} with transition matrix B . Let $T = \inf\{t \geq 0 : X_t = Y_t\}$ be the first meeting time of these chains. Then, for any $\epsilon > 0$, there exists $\tilde{\beta} > 0$ such that*

$$\sup_{x, y \in \mathcal{X}} \mathbb{E}_{xy}[e^{\beta T}] \leq 1 + \epsilon \quad \text{for all } \beta \in [0, \tilde{\beta}],$$

where, as usual, \mathbb{E}_{xy} is the expectation conditional on $X_0 = x$ and $Y_0 = y$. Furthermore, $\tilde{\beta}$ does not depend on B except through A and γ .

Proof. By Proposition 3.11, X has the same marginal distribution as \check{X} , ignoring the splitting, and similarly for Y . Hence T , the first meeting time of X and Y , is less than or equal (in distribution) to \check{S} , the first meeting time of \check{X} and \check{Y} , which we showed has the required bound in Theorem 3.13. \square

Proof of Theorem 3.5. Let X and Y be two independent copies of the Markov chain on \mathcal{X} with transition matrix B . Let $T = \inf\{t \geq 0 : X_t = Y_t\}$ be the first meeting time of these chains, as in the previous Corollary. Let $\mu, \nu \in \mathcal{M}$ (recall Definition 3.2). Write P_t for the transition operator of X (and Y), so that $P_t\mu$ is the law of X_t given $X_0 \sim \mu$, and $P_t\nu$ is the law of Y_t given $Y_0 \sim \nu$. It is easy to see that, conditioned on $\{T \leq t\}$, X_t and Y_t have the same distribution. It follows that

$$\|P_t\mu - P_t\nu\|_{TV} \leq \mathbb{P}(T > t | X_0 \sim \mu, Y_0 \sim \nu).$$

By Corollary 3.17, for any $\epsilon > 0$, there exists $\tilde{\beta} > 0$, which depends only on A , γ and ϵ , such that

$$\sup_{x, y \in \mathcal{X}} \mathbb{E}_{xy}[e^{\tilde{\beta}T}] \leq 1 + \epsilon.$$

Then, by Markov's inequality,

$$\begin{aligned} \|P_t\mu - P_t\nu\|_{TV} &\leq \mathbb{P}(T > t | X_0 \sim \mu, Y_0 \sim \nu) \\ &\leq \mathbb{E}[e^{\tilde{\beta}T} | X_0 \sim \mu, Y_0 \sim \nu] e^{-\tilde{\beta}t} \leq (1 + \epsilon)e^{-\tilde{\beta}t}. \end{aligned} \quad (21)$$

Let $s > t$. Replacing ν by $P_{s-t}\mu$ in the above, we obtain

$$\|P_t\mu - P_s\mu\|_{TV} \leq (1 + \epsilon)e^{-\tilde{\beta}t}.$$

Since this holds for all $s > t$, we see that $\{P_t\mu\}_{t \geq 0}$ is a Cauchy sequence in \mathcal{M} , and hence converges to an element $\pi \in \mathcal{M}$, and from (21), π is independent of the initial measure μ . Taking the limit as $s \rightarrow \infty$ in the above, it follows that

$$\sup_{\mu \in \mathcal{M}} \|P_t\mu - \pi\|_{TV} \leq (1 + \epsilon)e^{-\tilde{\beta}t}.$$

Hence, a Markov chain under the measure induced by B is uniformly ergodic. Further, as noted above, the rate of convergence $\tilde{\beta}$ does not depend on B except through A and γ , and we may take ϵ arbitrarily small, with a corresponding decrease in $\tilde{\beta}$. \square

4 Ergodic BSDEs

We now introduce an ‘Ergodic BSDE’ (EBSDE) as an infinite horizon equation of the form

$$Y_T = Y_t - \sum_{t \leq u < T} (f(X_u, Z_u) - \lambda) + \sum_{t \leq u < T} Z_u^* M_{u+1}. \quad (22)$$

A solution of this equation is a triple (Y, Z, λ) such that (22) holds a.s. for all finite values of t and T such that $0 \leq t < T$, where, as usual Y and Z are adapted processes of appropriate dimension, and where $\lambda \in \mathbb{R}$ is a constant. Note that, unlike in our discounted BSDEs where the constant α is given, here λ is to be found as part of the solution.

We will follow the method given in Section 4 of [6], which is itself based on the work of [10] and [8]. Throughout this section we will assume that X is a uniformly ergodic, time-homogeneous Markov chain on \mathcal{X} with transition matrix A , and that the driver $f : \mathcal{X} \times \mathbb{R}^N \rightarrow \mathbb{R}$ is γ -balanced, Markovian, and independent of y and t . Under these assumptions, we shall prove a result on the existence and uniqueness of solutions to the EBSDE (22).

Lemma 4.1. *Under the assumptions stated above, let Z, Z' be any two deterministic, or Markovian, \mathbb{R}^N -valued processes which do not depend on t , defined up to equivalence \sim_M . Then there exists a probability measure \mathbb{Q} on (Ω, \mathcal{F}) such that the following conditions hold:*

- \mathbb{Q} is equivalent to \mathbb{P} on \mathcal{F}_t for every finite t ,
- $\tilde{M}_t := -\sum_{0 \leq u < t} (f(X_u, Z_u) - f(X_u, Z'_u)) + \sum_{0 \leq u < t} (Z_u - Z'_u)^* M_{u+1}$ is a martingale under \mathbb{Q} , and
- under \mathbb{Q} , X is a uniformly ergodic Markov chain, and the constants of ergodicity R, ρ depend only on A and γ .

Proof. Let $\psi(\omega, t, z, z')$ be the random field associated with f , as given in Definition 2.9. By our assumptions on f and the processes Z, Z' , there exists a (deterministic) matrix $\Psi^{Z, Z'}$ such that

$$\psi(\omega, t, Z_t, Z'_t) = \Psi^{Z, Z'} X_t$$

holds for all t . It follows from the properties of ψ given in Definition 2.9 that $\Psi^{Z, Z'}$ is a transition matrix, and that $\Psi^{Z, Z'} \sim_\gamma A$.

Let \mathbb{Q} be the probability measure under which X is a Markov chain with transition matrix $\Psi^{Z, Z'}$. Since $\Psi^{Z, Z'} \sim_\gamma A$, the matrices $\Psi^{Z, Z'}$ and A have the same pattern of zero entries, so that the possible jumps of X are the same under \mathbb{Q} and \mathbb{P} . As we are in a discrete time, finite state setting, it follows immediately that \mathbb{Q} is equivalent to \mathbb{P} on \mathcal{F}_t for every finite t . We also have that

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}}[(Z_t - Z'_t)^* M_{t+1} \mid \mathcal{F}_t] &= (Z_t - Z'_t)^* (\Psi^{Z, Z'} X_t - A X_t) \\ &= f(X_t, Z_t) - f(X_t, Z'_t), \end{aligned}$$

from which it follows that $\mathbb{E}_{\mathbb{Q}}[\tilde{M}_{t+1} - \tilde{M}_t \mid \mathcal{F}_t] = 0$, so that \tilde{M} is a martingale under \mathbb{Q} as required. The final statement is the result of Theorem 3.5. \square

Remark 4.2. As we saw in Lemma 2.13 and Corollary 2.14, the Z part of the solution of our discounted BSDE is constant (i.e. deterministic and independent of t), though the result above holds just as well for any Markovian processes Z, Z' , provided they are still independent of t . If they do depend on t then X will in general be a time-inhomogeneous Markov chain under \mathbb{Q} , in which case the notion of uniform ergodicity is no longer meaningful. Nevertheless, in an analogous context, [2] uses the continuous time version of the ergodicity estimates we obtained in Section 3 (as given in [6]) to prove the existence of solutions of continuous time BSDEs up to unbounded stopping times.

Lemma 4.3. *For $\alpha > 0$, let (Y^α, Z^α) be the unique bounded solution (as given by Theorem 2.5) of the discounted BSDE*

$$Y_T^\alpha = Y_t^\alpha - \sum_{t \leq u < T} (f(X_u, Z_u^\alpha) - \alpha Y_u^\alpha) + \sum_{t \leq u < T} (Z_u^\alpha)^* M_{u+1}, \quad (23)$$

and let $x_0 \in \mathcal{X}$ be an arbitrary state. By Lemma 2.13 and Corollary 2.14, there exists a function $v^\alpha : \mathcal{X} \rightarrow \mathbb{R}$ such that $Y_t^\alpha = v^\alpha(X_t)$ and $e_k^* Z_t^\alpha = v^\alpha(e_k)$. Then there exists a bound $C' > 0$ such that

$$|v^\alpha(x) - v^\alpha(x_0)| \leq C', \quad \alpha |v^\alpha(x)| \leq C' \quad (24)$$

uniformly in x and α , and hence there exists a sequence $\alpha_n \searrow 0$ such that

$$(v^{\alpha_n}(x) - v^{\alpha_n}(x_0)) \rightarrow v(x) \quad \text{and} \quad \alpha_n v^{\alpha_n}(x) \rightarrow \lambda \quad \text{for all } x \in \mathcal{X},$$

for some $\lambda \in \mathbb{R}$ and some function $v : \mathcal{X} \rightarrow \mathbb{R}$.

Proof. Let $C > 0$ be a bound on $|f(\cdot, 0)|$. It then follows from Theorem 2.5 that $|v^\alpha(\cdot)| \leq C/\alpha$. By Lemma 4.1, there exists a measure \mathbb{Q}^α such that

$$(\omega, t) \mapsto - \sum_{0 \leq u < t} (f(X_u, Z_u^\alpha) - f(X_u, 0)) + \sum_{0 \leq u < t} (Z_u^\alpha)^* M_{u+1}$$

defines a martingale under \mathbb{Q}^α . Moreover, X is uniformly ergodic under \mathbb{Q}^α , and the constants of ergodicity R, ρ do not depend on α . Writing $P_t^\alpha \delta_x$ for the law of X_t under \mathbb{Q}^α given $X_0 = x$, we see from the proof of Theorem 3.5 that

$$\|P_t^\alpha \delta_x - P_t^\alpha \delta_{x'}\|_{TV} \leq R e^{-\rho t}$$

for all $x, x' \in \mathcal{X}$. From the one-step dynamics of (23), we obtain

$$\mathbb{E}_{\mathbb{Q}^\alpha}[v^\alpha(X_{t+1}) | \mathcal{F}_t] = (1 + \alpha)v^\alpha(X_t) - f(X_t, 0).$$

It follows by induction that

$$v^\alpha(X_0) = \mathbb{E}_{\mathbb{Q}^\alpha} \left[(1 + \alpha)^{-T} v^\alpha(X_T) + \sum_{k=0}^{T-1} (1 + \alpha)^{-(k+1)} f(X_k, 0) \mid \mathcal{F}_0 \right]$$

for all $T \geq 1$. As $|v^\alpha(\cdot)| \leq C/\alpha$, letting $T \rightarrow \infty$, we deduce that

$$v^\alpha(x) = \lim_{T \rightarrow \infty} \mathbb{E}_{\mathbb{Q}^\alpha} \left[\sum_{k=0}^{T-1} (1 + \alpha)^{-(k+1)} f(X_k, 0) \mid X_0 = x \right].$$

Then, for any $x, x' \in \mathcal{X}$, we have

$$\begin{aligned} |v^\alpha(x) - v^\alpha(x')| &= \left| \lim_{T \rightarrow \infty} \mathbb{E}_{\mathbb{Q}^\alpha} \left[\sum_{k=0}^{T-1} (1 + \alpha)^{-(k+1)} f(X_k, 0) \mid X_0 = x \right] \right. \\ &\quad \left. - \lim_{T \rightarrow \infty} \mathbb{E}_{\mathbb{Q}^\alpha} \left[\sum_{k=0}^{T-1} (1 + \alpha)^{-(k+1)} f(X_k, 0) \mid X_0 = x' \right] \right| \\ &= \lim_{T \rightarrow \infty} \left| \sum_{k=0}^{T-1} (1 + \alpha)^{-(k+1)} \int_{\mathcal{X}} f(X_k, 0) (d(P_k^\alpha \delta_x) - d(P_k^\alpha \delta_{x'})) \right| \\ &\leq \lim_{T \rightarrow \infty} 2C \sum_{k=0}^{T-1} (1 + \alpha)^{-(k+1)} \|P_k^\alpha \delta_x - P_k^\alpha \delta_{x'}\|_{TV} \\ &\leq \lim_{T \rightarrow \infty} 2CR \sum_{k=0}^{T-1} (1 + \alpha)^{-(k+1)} e^{-\rho k} = \frac{2CR}{1 + \alpha - e^{-\rho}}. \end{aligned}$$

We therefore have the bound

$$|v^\alpha(x) - v^\alpha(x')| \leq \frac{2CR}{1 - e^{-\rho}},$$

which holds for all $x, x' \in \mathcal{X}$ and all $\alpha > 0$. Since $\alpha|v^\alpha(\cdot)| \leq C$ (and we may certainly assume that $R > 1$), we see that (24) is satisfied with $C' = 2CR(1 - e^{-\rho})^{-1}$.

By the Bolzano–Weierstrass Theorem, there exists a sequence $\alpha_n \searrow 0$ such that $\alpha_n v^{\alpha_n}(x_0) \rightarrow \lambda$ and

$$(v^{\alpha_n}(x) - v^{\alpha_n}(x_0)) \rightarrow v(x)$$

for all $x \in \mathcal{X}$, for some $\lambda \in \mathbb{R}$ and some function $v : \mathcal{X} \rightarrow \mathbb{R}$. Note that $|\lambda|$ and $|v(\cdot)|$ are both bounded by C' . Finally, we notice that, for any $x \in \mathcal{X}$,

$$\alpha_n v^{\alpha_n}(x) = \alpha_n v^{\alpha_n}(x_0) + \alpha_n (v^{\alpha_n}(x) - v^{\alpha_n}(x_0)) \rightarrow \lambda,$$

so the convergence of this sequence to λ holds for all x . \square

We are now in a position to prove existence of solutions to our Ergodic BSDE.

Theorem 4.4. *Let v and λ be as constructed in Lemma 4.3. The triple (Y, Z, λ) , where*

$$Y_t := v(X_t), \quad e_k^* Z_t := v(e_k),$$

is the unique bounded, stationary (i.e. does not depend on t), Markovian solution, with $v(x_0) = 0$, to the Ergodic BSDE

$$Y_T = Y_t - \sum_{t \leq u < T} (f(X_u, Z_u) - \lambda) + \sum_{t \leq u < T} Z_u^* M_{u+1}. \quad (25)$$

Any other bounded solution (Y', Z', λ') satisfies $\lambda = \lambda'$, and any other bounded, stationary, Markovian solution (Y', Z', λ') satisfies $Y_t = Y'_t + c$ for some $c \in \mathbb{R}$, and $Z \sim_M Z'$.

Proof. Let $\{\alpha_n\}_{n \geq 1}$ be the sequence constructed in Lemma 4.3. We have that $Y_t^{\alpha_n} = v^{\alpha_n}(X_t)$ and $e_k^* Z_t^{\alpha_n} = v^{\alpha_n}(e_k)$ solve the discounted BSDE

$$Y_T^{\alpha_n} = Y_t^{\alpha_n} - \sum_{t \leq u < T} (f(X_u, Z_u^{\alpha_n}) - \alpha_n Y_u^{\alpha_n}) + \sum_{t \leq u < T} (Z_u^{\alpha_n})^* M_{u+1}.$$

However, since $\|\mathbf{1}\|_{M_{t+1}} = 0$, and f does not distinguish between values of Z_u up to equivalence $\sim_{M_{u+1}}$, we can equally write $Z_u^{\alpha_n} - v^{\alpha_n}(x_0)\mathbf{1}$ in the place of $Z_u^{\alpha_n}$ in the above. Note that $e_k^* Z_u^{\alpha_n} - v^{\alpha_n}(x_0) \rightarrow e_k^* Z_u$ as $n \rightarrow \infty$ for each $e_k \in \mathcal{X}$, and that, by the bound established in Lemma 4.3, $|e_k^* Z_u^{\alpha_n} - v^{\alpha_n}(x_0)|$ is uniformly bounded. Since f is Lipschitz in z , we deduce that

$$f(X_u, Z_u^{\alpha_n} - v^{\alpha_n}(x_0)\mathbf{1}) \rightarrow f(X_u, Z_u) \quad \text{as } n \rightarrow \infty \quad \text{a.s.}$$

It follows that

$$\begin{aligned}
Y_T &= \lim_{n \rightarrow \infty} (v^{\alpha_n}(X_T) - v^{\alpha_n}(x_0)) \\
&= \lim_{n \rightarrow \infty} (v^{\alpha_n}(X_t) - v^{\alpha_n}(x_0)) - \lim_{n \rightarrow \infty} \sum_{t \leq u < T} (f(X_u, Z_u^{\alpha_n}) - \alpha_n v^{\alpha_n}(X_u)) \\
&\quad + \lim_{n \rightarrow \infty} \sum_{t \leq u < T} (Z_u^{\alpha_n})^* M_{u+1} \\
&= v(X_t) - \lim_{n \rightarrow \infty} \sum_{t \leq u < T} (f(X_u, Z_u^{\alpha_n} - v^{\alpha_n}(x_0)\mathbf{1}) - \alpha_n v^{\alpha_n}(X_u)) \\
&\quad + \lim_{n \rightarrow \infty} \sum_{t \leq u < T} (Z_u^{\alpha_n} - v^{\alpha_n}(x_0)\mathbf{1})^* M_{u+1} \\
&= Y_t - \sum_{t \leq u < T} (f(X_u, Z_u) - \lambda) + \sum_{t \leq u < T} Z_u^* M_{u+1},
\end{aligned}$$

and we see that (Y, Z, λ) is indeed a solution of the EBSDE (25).

Suppose that (Y', Z', λ') is another bounded solution. Let $\tilde{Y} = Y - Y'$, $\tilde{Z} = Z - Z'$ and $\tilde{\lambda} = \lambda - \lambda'$. Then

$$\tilde{Y}_T = \tilde{Y}_0 - \sum_{0 \leq u < T} (f(X_u, Z_u) - f(X_u, Z'_u) - \tilde{\lambda}) + \sum_{0 \leq u < T} \tilde{Z}_u^* M_{u+1}. \quad (26)$$

By Proposition 2.2, there exists a measure \mathbb{Q} such that

$$(\omega, t) \mapsto - \sum_{0 \leq u < t} (f(X_u, Z_u) - f(X_u, Z'_u)) + \sum_{0 \leq u < t} \tilde{Z}_u^* M_{u+1}$$

defines a martingale under \mathbb{Q} . Taking an $\mathbb{E}_{\mathbb{Q}}$ expectation in (26), we obtain

$$\tilde{\lambda} = T^{-1} \mathbb{E}_{\mathbb{Q}}[\tilde{Y}_T - \tilde{Y}_0].$$

Since \tilde{Y} is uniformly bounded, taking the limit $T \rightarrow \infty$ gives $\tilde{\lambda} = 0$, so that $\lambda = \lambda'$. Substituting back into (26) and taking an $\mathbb{E}_{\mathbb{Q}}[\cdot | \mathcal{F}_0]$ expectation gives

$$\mathbb{E}_{\mathbb{Q}}[\tilde{Y}_T | \mathcal{F}_0] = \tilde{Y}_0. \quad (27)$$

Suppose further that Y' and Z' are Markovian and do not depend on t , so that in particular there exists a function $v' : \mathcal{X} \rightarrow \mathbb{R}$ such that $Y'_t = v'(X_t)$. Then the measure \mathbb{Q} may be taken to be the measure given by Lemma 4.1, so that X is still a uniformly ergodic Markov chain under \mathbb{Q} . Writing $\tilde{\pi}$ for the ergodic measure of X under \mathbb{Q} , it follows from (27) that, for any $x \in \mathcal{X}$,

$$v(x) - v'(x) = \lim_{T \rightarrow \infty} \mathbb{E}_{\mathbb{Q}}[\tilde{Y}_T | X_0 = x] = \int_{\mathcal{X}} (v(y) - v'(y)) d\tilde{\pi}(y).$$

Since the right hand side is independent of x , we see that $v(x) = v'(x) + c$ for all x , and hence that $Y_t = Y'_t + c$, for some $c \in \mathbb{R}$. In particular, if $v'(x_0) = 0$, then $c = 0$, and hence $Y = Y'$ up to indistinguishability.

With $\tilde{\lambda} = 0$ and $\tilde{Y}_t = c = \tilde{Y}_{t+1}$, we deduce from the one-step dynamics of \tilde{Y} that $\tilde{Z}_t^* M_{t+1} = 0$, and hence that $Z \sim_M Z'$. \square

Corollary 4.5. *The sequences $\{\alpha_n v^{\alpha_n}(x)\}_{n \geq 1}$ and $\{v^{\alpha_n}(x) - v^{\alpha_n}(x_0)\}_{n \geq 1}$ constructed in Lemma 4.3 converge to λ and $v(x)$ respectively for any choice of sequence $\alpha_n \searrow 0$.*

Proof. Suppose this were not the case. Then we could construct two sequences with distinct limits, both of which would yield bounded, stationary, Markovian solutions to the EBSDE, contradicting the uniqueness established in Theorem 4.4. \square

Corollary 4.6. *Let π denote the ergodic measure of X under \mathbb{P} . Writing \mathbf{v} for the vector in \mathbb{R}^N with entries $e_k^* \mathbf{v} = v(e_k)$, the value λ in the EBSDE solution $(Y, Z, \lambda) = (v(X_t), \mathbf{v}, \lambda)$ is given by*

$$\lambda = \int_{\mathcal{X}} f(x, \mathbf{v}) d\pi(x) = \sum_{x \in \mathcal{X}} f(x, \mathbf{v}) \pi(\{x\}).$$

Furthermore, there exists a probability measure $\pi_{\mathbf{v}}$ on \mathcal{X} such that

$$\lambda = \int_{\mathcal{X}} f(x, 0) d\pi_{\mathbf{v}}(x).$$

Proof. The invariance of the ergodic measure π implies that for any fixed time $t \geq 0$ and any function $g : \mathcal{X} \rightarrow \mathbb{R}$,

$$\int_{\mathcal{X}} \mathbb{E}[g(X_t) | X_0 = x] d\pi(x) = \int_{\mathcal{X}} g(x) d\pi(x). \quad (28)$$

For any $T > 0$, we have

$$v(X_0) = v(X_T) + \sum_{0 \leq u < T} (f(X_u, \mathbf{v}) - \lambda) - \sum_{0 \leq u < T} \mathbf{v}^* M_{u+1}. \quad (29)$$

Then, by the invariance property (28),

$$\begin{aligned} \int_{\mathcal{X}} v(x) d\pi(x) &= \int_{\mathcal{X}} \mathbb{E} \left[v(X_T) + \sum_{0 \leq u < T} (f(X_u, \mathbf{v}) - \lambda) \middle| X_0 = x \right] d\pi(x) \\ &= \int_{\mathcal{X}} v(x) d\pi(x) + T \int_{\mathcal{X}} f(x, \mathbf{v}) d\pi(x) - T\lambda, \end{aligned}$$

and rearranging gives the first result.

By Lemma 4.1, there exists a probability measure \mathbb{Q} such that, under \mathbb{Q} , the process defined by

$$(\omega, t) \mapsto - \sum_{0 \leq u < t} (f(X_u, \mathbf{v}) - f(X_u, 0)) + \sum_{0 \leq u < t} \mathbf{v}^* M_{u+1}$$

is a martingale, and X is still a uniformly ergodic Markov chain. Let $\pi_{\mathbf{v}}$ denote the ergodic measure of X under \mathbb{Q} . The second result then follows from (29), and the invariance property (28) applied with $\mathbb{E}_{\mathbb{Q}}$ and $\pi_{\mathbf{v}}$. \square

Remark 4.7. Under the additional assumption that X has no transient states, it follows from the one-step dynamics of (25), that the EBSDE solution $(v(X_t), \mathbf{v}, \lambda)$ is also a solution of the vector equation

$$\mathbf{v} - \mathbf{f}(\mathbf{v}) + \lambda \mathbf{1} - A^* \mathbf{v} = 0, \quad (30)$$

where $\mathbf{f}(\mathbf{v})$ denotes the vector with entries $e_k^* \mathbf{f}(\mathbf{v}) = f(e_k, \mathbf{v})$. Further, this solution is unique up to equality in λ and a constant shift in \mathbf{v} .

This equation also demonstrates why in general EBSDE drivers cannot have dependence on Y . Consider the case where $N = 2$ and

$$f(x, y, z) = x^* \mathbf{k} + y + z^* (B - A)x,$$

where $B \sim_\gamma A$ and $\mathbf{k} \in \mathbb{R}^2$ is a constant vector. Then f is γ -balanced, and equation (30) implies that $-\mathbf{k} + \lambda \mathbf{1} - B^* \mathbf{v} = 0$, which has no solutions when, for example,

$$B = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \quad \text{and} \quad e_1^* \mathbf{k} \neq e_2^* \mathbf{k}.$$

We now provide a comparison theorem for Ergodic BSDEs, concerning the λ part of the solution.

Theorem 4.8. *Let f and f' be two γ -balanced Markovian drivers, and let (Y, Z, λ) and (Y', Z', λ') be any corresponding bounded EBSDE solutions. If $f(x, z) \geq f'(x, z)$ for all $x \in \mathcal{X}$ and $z \in \mathbb{R}^N$, then $\lambda \geq \lambda'$.*

Proof. Write $(\bar{Y}, \bar{Z}, \lambda) = (v(X_t), \mathbf{v}, \lambda)$ and $(\bar{Y}', \bar{Z}', \lambda') = (v'(X_t), \mathbf{v}', \lambda')$ for the corresponding bounded, stationary, Markovian solutions given in Theorem 4.4. By Lemma 4.1, there exists a probability measure \mathbb{Q} such that, under \mathbb{Q} , the process defined by

$$(\omega, t) \mapsto - \sum_{0 \leq u < t} (f'(X_u, \mathbf{v}) - f'(X_u, \mathbf{v}')) + \sum_{0 \leq u < t} (\mathbf{v} - \mathbf{v}')^* M_{u+1}$$

is a martingale and X is still a uniformly ergodic Markov chain. Write $\hat{\pi}$ for the ergodic measure of X under \mathbb{Q} . Taking the difference of the EBSDEs, and applying the invariance property (28), we have

$$\begin{aligned} & \int_{\mathcal{X}} (v(x) - v'(x)) d\hat{\pi}(x) \\ &= \int_{\mathcal{X}} \mathbb{E}_{\mathbb{Q}} [v(X_T) - v'(X_T) \mid X_0 = x] d\hat{\pi}(x) \\ & \quad + \int_{\mathcal{X}} \mathbb{E}_{\mathbb{Q}} \left[\sum_{0 \leq u < T} (f(X_u, \mathbf{v}) - f'(X_u, \mathbf{v}) - (\lambda - \lambda')) \mid X_0 = x \right] d\hat{\pi}(x) \\ &= \int_{\mathcal{X}} (v(x) - v'(x)) d\hat{\pi}(x) + T \int_{\mathcal{X}} (f(x, \mathbf{v}) - f'(x, \mathbf{v})) d\hat{\pi}(x) - T(\lambda - \lambda'), \end{aligned}$$

and hence, by rearrangement,

$$\lambda - \lambda' = \int_{\mathcal{X}} (f(x, \mathbf{v}) - f'(x, \mathbf{v})) d\hat{\pi}(x) \geq 0.$$

□

Adapting a proof from Hu, Madec and Richou [11], we have the following additional interpretation.

Theorem 4.9. *Let f be a γ -balanced Markovian driver. Let (Y, Z, λ) be the corresponding (bounded) EBSDE solution, and for each T , let (Y^T, Z^T) be the solution of the finite horizon BSDE with driver f and \mathcal{F}_T -measurable terminal value $Y_T^T = \xi^T$, where $\sup_T \|\xi^T\|_{L^\infty} < \infty$. Then*

$$\left| \frac{Y_0^T}{T} - \lambda \right| \leq \frac{C}{T}$$

for some $C > 0$. In particular, $\frac{Y_0^T}{T} \rightarrow \lambda$ as $T \rightarrow \infty$.

Proof. We have

$$\begin{aligned} & Y_0^T - \xi^T + Y_T - Y_0 - \lambda T \\ &= - \sum_{0 \leq u < T} (f(X_u, Z_u) - f(X_u, Z_u^T)) + \sum_{0 \leq u < T} (Z_u - Z_u^T)^* M_{u+1}. \end{aligned}$$

By Proposition 2.2, there exists a probability measure \mathbb{Q}^T such that

$$(\omega, t) \mapsto - \sum_{0 \leq u < t} (f(X_u, Z_u) - f(X_u, Z_u^T)) + \sum_{0 \leq u < t} (Z_u - Z_u^T)^* M_{u+1}$$

defines a martingale under \mathbb{Q}^T , and it follows that

$$Y_0^T - \lambda T = \mathbb{E}_{\mathbb{Q}^T} [\xi^T - Y_T + Y_0 | \mathcal{F}_0].$$

The term on the right is bounded (uniformly in T) by some $C > 0$, and hence

$$\left| \frac{Y_0^T}{T} - \lambda \right| = \left| \frac{Y_0^T - \lambda T}{T} \right| \leq \frac{C}{T}.$$

□

Remark 4.10. It is clear from the proof that this result will hold assuming only that $\|\xi^T\|_{L^\infty}/T \rightarrow 0$ as $T \rightarrow \infty$, with a correspondingly slower rate of convergence.

5 Optimal Ergodic Control

As an application of the theory we have just developed for discrete time Ergodic BSDEs, we shall now present a treatment of an ergodic control problem. We will follow the method given in Section 5.2 of [6], which is itself based on the work of [10] and [8].

Suppose that

- Under \mathbb{P} , X is a uniformly ergodic Markov chain on \mathcal{X} with transition matrix A ,
- \mathcal{U} is a space of ‘controls’, which we assume to be a topological space, equal to a countable union of compact metrizable subsets of itself,
- $B^{(\cdot)}$ is a continuous function which maps each element $u \in \mathcal{U}$ to a transition matrix B^u such that $B^u \sim_\gamma A$, for some fixed $\gamma \in (0, 1)$,
- $L : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ is a cost function, which we assume to be bounded, measurable in x and continuous in u .

Define the ergodic cost

$$J(x, U) = \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_x^U \left[\sum_{0 \leq s < T} L(X_s, U_s) \right],$$

where

- U is a \mathcal{U} -valued adapted process, which we shall also refer to as a ‘control’,
- \mathbb{E}_x^U is the expectation under which $X_0 = x$, and for any given pair (ω, t) , X jumps according to transition matrix $B^{U_t(\omega)}$.

We wish to minimize $J(x, U)$ over all controls U . Our approach is to work with the probability measure under which X jumps at time t according to the transition matrix B^{U_t} . As we have made no Markov or time-homogeneity assumptions on U , X will in general not be a Markov chain under our new measure.

We define the Hamiltonian

$$f(x, z) = \inf_{u \in \mathcal{U}} \{L(x, u) + z^*(B^u - A)x\}. \quad (31)$$

Note that f is finite valued and that $f(\cdot, 0)$ is bounded. By Lemma 2.12, $f(X_t, z)$ is a γ -balanced Markovian driver. Hence, by Theorem 4.4, the EB-SDE with driver $f(X_t, z)$ admits a bounded, stationary, Markovian solution $(\bar{Y}, \bar{Z}, \bar{\lambda}) = (v(X_t), \mathbf{v}, \bar{\lambda})$, where as usual \mathbf{v} is the vector in \mathbb{R}^N with components $v(e_k)$.

If the infimum in (31) is attained then, by Filippov’s implicit function theorem (see either McShane and Warfield [12] or Beneš [1]), there exists a measurable function $\kappa : \mathcal{X} \times \mathbb{R}^N \rightarrow \mathcal{U}$ such that

$$f(x, z) = L(x, \kappa(x, z)) + z^*(B^{\kappa(x, z)} - A)x. \quad (32)$$

If the infimum in (31) is not attained, then applying Theorem 21.3.4 from [4] to the function $G(x, z, u) = L(x, u) + z^*(B^u - A)x$, we deduce, for any $\epsilon > 0$, the existence of a measurable function $\tau^\epsilon : \mathcal{X} \times \mathbb{R}^N \rightarrow \mathcal{U}$ such that

$$L(x, \tau^\epsilon(x, z)) + z^*(B^{\tau^\epsilon(x, z)} - A)x < f(x, z) + \epsilon. \quad (33)$$

Theorem 5.1. *In the setting described above, let (Y, Z, λ) be any (possibly non-Markovian) bounded solution to the EBSDE (22) with driver f . Then the following hold:*

- For any control U , we have $J(x, U) \geq \lambda = \bar{\lambda}$, with equality if

$$f(X_s, Z_s) = L(X_s, U_s) + Z_s^*(B^{U_s} - A)X_s \text{ for all } s \geq 0. \quad (34)$$

- If the infimum in (31) is attained, then the control $\bar{U}_t = \kappa(X_t, Z_t)$ satisfies $J(x, \bar{U}) = \bar{\lambda}$.
- Even if the infimum in (31) is not attained, there still exists a control \hat{U} such that $J(x, \hat{U}) = \bar{\lambda}$.

Proof. The fact that $\lambda = \bar{\lambda}$ follows immediately from Theorem 4.4.

Note that the process defined by

$$(\omega, t) \mapsto - \sum_{0 \leq s < t} Z_s^*(B^{U_s} - A)X_s + \sum_{0 \leq s < t} Z_s^*M_{s+1}$$

is a martingale under \mathbb{E}_x^U . Since $(Y, Z, \bar{\lambda})$ is a solution of the EBSDE, we have

$$Y_T = Y_0 - \sum_{0 \leq s < T} (f(X_s, Z_s) - \bar{\lambda}) + \sum_{0 \leq s < T} Z_s^*M_{s+1}$$

for any $T > 0$. Then

$$\begin{aligned} \bar{\lambda} &= \frac{1}{T} \mathbb{E}_x^U \left[Y_T - Y_0 + \sum_{0 \leq s < T} L(X_s, U_s) \right] \\ &\quad + \frac{1}{T} \mathbb{E}_x^U \left[\sum_{0 \leq s < T} (f(X_s, Z_s) - L(X_s, U_s) - Z_s^*(B^{U_s} - A)X_s) \right]. \end{aligned} \quad (35)$$

Since f is the infimum over all controls, we have

$$\bar{\lambda} \leq \frac{1}{T} \mathbb{E}_x^U \left[Y_T - Y_0 + \sum_{0 \leq s < T} L(X_s, U_s) \right].$$

As $Y_T - Y_0$ is uniformly bounded, it follows that

$$\bar{\lambda} \leq \limsup_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_x^U \left[\sum_{0 \leq s < T} L(X_s, U_s) \right] = J(x, U).$$

If (34) holds, then equality holds throughout the above, and we see that $\bar{\lambda} = J(x, U)$.

If the infimum in (31) is attained, then we know that (32) holds for some measurable function $\kappa : \mathcal{X} \times \mathbb{R}^N \rightarrow \mathcal{U}$. Setting $\bar{U}_s = \kappa(X_s, Z_s)$, we have that (34) holds, so that $\bar{\lambda} = J(x, \bar{U})$.

Even if the infimum in (31) is not attained, by (33), there exists a measurable function $\tau : \mathbb{N} \times \mathcal{X} \times \mathbb{R}^N \rightarrow \mathcal{U}$ such that

$$L(x, \tau(s, x, z)) + z^*(B^{\tau(s, x, z)} - A)x < f(x, z) + 2^{-s}.$$

Setting $\hat{U}_s = \tau(s, X_s, Z_s)$, (35) reduces to

$$\bar{\lambda} \geq \frac{1}{T} \mathbb{E}_x^{\hat{U}} \left[Y_T - Y_0 + \sum_{0 \leq s < T} L(X_s, \hat{U}_s) \right] - \frac{1}{T} \sum_{0 \leq s < T} 2^{-s},$$

and it follows that $\bar{\lambda} = J(x, \hat{U})$. \square

In particular, the controls $\bar{U}_s = \kappa(X_s, \mathbf{v})$ (assuming the function κ exists) and $\hat{U}_s = \tau(s, X_s, \mathbf{v})$ are optimal feedback controls.

Remark 5.2. We note that \hat{U}_s is time-dependent, but that for any $\epsilon > 0$, a time-homogenous feedback control satisfying $J(x, \hat{U}) \leq \bar{\lambda} + \epsilon$ can also be attained by setting $\hat{U}_s = \tau^\epsilon(X_s, \mathbf{v})$, where τ^ϵ satisfies (33).

Applying Theorem 4.9, we can see that the finite horizon cost function converges to the ergodic cost function, as indicated in the following corollary. Note that this does not guarantee convergence of the (not necessarily unique) optimal controls.

Corollary 5.3. *Let $J(x, U, T)$ denote the cumulative cost over time $[0, T]$, that is,*

$$J(x, U, T) = \mathbb{E}_x^U \left[\xi^T + \sum_{0 \leq s < T} L(X_s, U_s) \right],$$

where we have also included an \mathcal{F}_T -measurable terminal cost ξ^T , which we assume is bounded uniformly in T . Then

$$\frac{\inf_U J(x, U, T)}{T} \rightarrow \inf_U J(x, U) = \lambda \quad \text{as } T \rightarrow \infty,$$

where λ is the optimal ergodic cost, as in Theorem 5.1.

Proof. Let (Y^T, Z^T) be the solution of the finite horizon BSDE with driver f and terminal value ξ^T at time T . By a similar argument to the proof of Theorem 5.1, we see that $\mathbb{E}_x^U [Y_0^T] \leq J(x, U, T)$ for any control U , and hence, since $\mathbb{E}_x^U [Y_0^T]$ is independent of U , that $\inf_U J(x, U, T) = \mathbb{E}_x [Y_0^T]$. The result then follows from Theorem 4.9. \square

It is worth noticing that the λ we have obtained is optimal in the class of all strategies, not only among those of feedback type. If we had approached our problem purely through the equations obtained from (30) and (32), this optimality would require separate analysis.

As noted above, (34) is a sufficient condition to guarantee that $J(x, U) = \bar{\lambda}$. In [10] and [6], among others, in analogous contexts, it is stated that this is also a necessary condition. However, this is incorrect, as modifying the value of an optimal control over any finite time horizon does not affect the associated ergodic cost. Indeed, in our setting we have shown that we can always construct an optimal control such that (34) does not necessarily hold at any time s .

Remark 5.4. We have assumed that the Markov chain X and the cost function L are time-homogeneous. However, by considering the cyclic classes of X , it is not hard to extend the result of Theorem 5.1 to the case when both X and L are periodic in time.

References

- [1] V.E. Beneš, Existence of optimal strategies based on specified information, for a class of stochastic decision problems, *SIAM Journal on Control and Optimization*, 8(2) (1970) 179–188.
- [2] S.N. Cohen, Undiscounted Markov chain BSDEs to stopping times, *Journal of Applied Probability*, 51(1) (2014) 262–281.
- [3] S.N. Cohen and R.J. Elliott, A general theory of finite state Backward Stochastic Difference Equations, *Stochastic Processes and their Applications*, 120(4) (2010) 442–466.
- [4] S.N. Cohen and R.J. Elliott, *Stochastic Calculus and Applications*, Birkhäuser Basel-New York, 2nd ed., 2015.
- [5] S.N. Cohen and V. Fedyashov, Ergodic BSDEs with jumps and time dependence, arXiv:1406.4329.
- [6] S.N. Cohen and Y. Hu, Ergodic BSDEs Driven by Markov Chains, *SIAM Journal on Control and Optimization*, 51(5) (2013) 4138–4168.
- [7] S.N. Cohen and L. Szpruch, On Markovian Solutions to Markov chain BSDEs *Numerical Algebra, Control and Optimization*, 2(2) (2012) 257–269.
- [8] A. Debussche, Y. Hu and G. Tessitore, Ergodic BSDEs under weak dissipative assumptions, *Stochastic Processes and their Applications*, 121(3) (2011) 407–426.
- [9] R.J. Elliott and H. Yang, How to count and guess well: Discrete adaptive filters, *Applied Mathematics and Optimization*, 30(1) (1994) 51–78.

- [10] M. Fuhrman, Y. Hu and G. Tessitore, Ergodic BSDEs and optimal ergodic control in Banach spaces, *SIAM Journal on Control and Optimization*, 48(3) (2009) 1542–1566.
- [11] Y. Hu, P.-Y. Madec and A. Richou, A probabilistic approach to large time behaviour of mild solutions of HJB equations in infinite dimension *SIAM Journal on Control and Optimization*, 53(1) (2015) 378–398.
- [12] E.J. McShane and R.B. Warfield, On Filippov’s implicit functions lemma, *Proceedings of the American Mathematical Society*, 18(1) (1967) 41–47.
- [13] S. Meyn and R.L. Tweedie, *Markov Chains and Stochastic Stability*, Cambridge University Press, 2nd ed., 2009.
- [14] A. Richou, Ergodic BSDEs and related PDEs with Neumann boundary conditions, *Stochastic Processes and their Applications*, 119(9) (2009) 2945–2969.
- [15] M. Royer, BSDEs with a random terminal time driven by a monotone generator and their links with PDEs, *Stochastics and Stochastics Reports*, 76(4) (2004) 281–307.
- [16] A.N. Shiryaev, *Probability*, Springer, 2nd ed., 1996.