

Ergodic BSDEs driven by Markov Chains

Samuel N. Cohen
University of Oxford

Ying Hu
Université de Rennes I

July 24, 2012

Abstract

We consider ergodic backward stochastic differential equations, in a setting where noise is generated by a countable state uniformly ergodic Markov chain. We show that for Lipschitz drivers such that a comparison theorem holds, these equations admit unique solutions. To obtain this result, we show by coupling and splitting techniques that uniform ergodicity estimates of Markov chains are robust to perturbations of the rate matrix, and that these perturbations correspond in a natural way to EBSDEs. We then consider applications of this theory to Markov decision problems with a risk-averse average reward criterion.

1 Introduction

Much work has gone into understanding optimal control with an average cost criterion, over an infinite horizon (for example, see the review paper Arapostathis et al. [1], or more recent work by Guo and Hernández-Lerma [14], and references therein). Provided an underlying controlled Markov process, this criterion provides a useful method for understanding those payoffs which value the future as much as the present, and hence are insensitive to short-term effects. Much of this work is constrained to look only at costs which depend on the current state of the process, and at the (linear) expectation of future costs; this is, in this sense, a linear theory (to be precise, the corresponding Hamiltonian operators are infima taken over linear functions in the costate variable, see Section 5.2). Therefore, these methods are unable to deal adequately with risk-averse optimization, which requires a *nonlinear* assessment of future costs.

Given the known connection between BSDEs and the theory of ‘nonlinear expectations’, as defined by Peng [20] (see Coquet et al. [9] and Cohen [3] for the precise details of this connection), it is not unreasonable to expect that ‘ergodic’ BSDEs would provide a natural framework for understanding these nonlinear settings.

In Fuhrman, Hu and Tessitore [12] (see also Richou [21], Debussche, Hu and Tessitore [10]), a class of value functions are considered which depend on the average cost, not only through the current state, but on the controlled stochastic dynamics, and can do so in a nonlinear way. These value functions are given by

Ergodic Backward Stochastic Differential Equations (EBSDEs), a generalisation of the Backward Stochastic Differential Equations developed by Pardoux and Peng [19]. In [12], the stochastic dynamics are given with reference to a general cylindrical Wiener process in a Hilbert space, and the ‘current state’ is described by a geometrically ergodic solution to a forward stochastic differential equation.

In this paper, we consider the ergodic BSDEs when noise is generated by a continuous-time discrete-state Markov chain. The basic BSDEs of this type, for finitely many states, were considered by one of the authors in [4, 5, 8]. We shall show that, with a discounted criterion, the infinite-horizon version of these BSDEs admits Markovian and time-invariant solutions. Assuming the underlying chain is uniformly ergodic, we show that discounted BSDEs preserve uniform ergodicity, in an appropriate sense. From here, we show that the ergodic BSDEs admit unique solutions.

The paper is organized as follows. In the next section, we deal with the theory of discounted BSDEs. In Section 3, we show the robustness of ergodicity estimates of Markov chains to perturbations of the rate matrix. We give a novel, but natural, partial ordering of rate matrices, and show that, given the perturbation is not too large as determined by this ordering, any ergodicity estimates are transferable. This result is itself of independent interest in the study of ergodic properties of Markov chains. Section 4 proves the existence and uniqueness of bounded Markovian solutions to EBSDEs. Finally, in Section 5, a numerical example and some applications of these equations in optimal ergodic control are discussed. We conclude with some thoughts on future extensions in Section 6.

1.1 Introducing BSDEs and EBSDEs on Markov Chains

Consider a continuous-time countable-state process X in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where X will be our fundamental Markov chain. Without loss of generality, we shall represent X as taking values from the standard basis vectors e_i of \mathbb{R}^N (where $N \in \mathbb{N} \cup \{\infty\}$ is the number of states, and \mathbb{R}^∞ denotes the space of infinite real sequences). We write \mathcal{X} for this set of basis vectors. For notational simplicity, we will think of all vectors as column vectors, and denote by z^* the transpose of z (so that z^*y is the Euclidean or ℓ_2 inner product). An element $\omega \in \Omega$ can be thought of as describing a path of the chain X .

Now let $\{\mathcal{F}_t\}$ be the completion of the filtration generated by X , that is,

$$\mathcal{F}_t = \sigma(\{X_s\}_{s \leq t}) \vee \{B \in \mathcal{F}_\infty : \mathbb{P}(B) = 0\}.$$

As X is a right-continuous pure jump process which does not jump at time 0, this filtration is right-continuous, and we assume $\mathcal{F} = \mathcal{F}_\infty = \bigvee_{t < \infty} \mathcal{F}_t$. We shall assume that X is a Markov chain under \mathbb{P} , in the filtration $\{\mathcal{F}_t\}$. For basic theory of continuous-time countable-state Markov chains, see for example Rogers and Williams [22, Vol. 1, p228ff], for the approach taken here, see Elliott, Aggoun and Moore [11, Part III]).

Let A denote the possibly infinite rate matrix¹ of the chain X . Note that $(A_t)_{ij} \geq 0$ for $i \neq j$ and $\sum_i (A_t)_{ij} = 0$ for all j (the columns of A all sum to 0).

¹In our notation, as in [11], A is the matrix with entries A_{ij} , where A_{ij} is the rate of jumping from state j to state i . Depending on the convention used, this is either the rate matrix or its transpose. In our notation A^* , the transpose of A , is the generator of the Markov chain.

We assume, for simplicity, that the entries in A_t are uniformly bounded, and so the chain is regular.

From the Doob–Meyer decomposition (see [11, Appendix B]), we write our chain in the following way

$$X_t = X_0 + \int_{]0,t]} A_u X_{u-} du + M_t, \quad (1)$$

where M is a locally-finite-variation pure-jump martingale in \mathbb{R}^N , and the chain starts in state $X_0 \in \mathbb{R}^N$. Our aim is to study EBSDEs, that is, infinite-horizon equations of the form

$$Y_t = Y_T + \int_t^T [f(X_{u-}, Z_u) - \lambda] du - \int_t^T Z_u^* dM_u, \quad 0 \leq t \leq T < \infty, \quad (2)$$

where $f : \mathcal{X} \times \mathbb{R}^N \rightarrow \mathbb{R}$ is a given function, Y is a real-valued càdlàg stochastic process, Z is a predictable process in \mathbb{R}^N such that

$$\int_0^t Z_u^* dM_u := \sum_i \int_0^t (Z_u)^i d(M)_u^i$$

is a martingale, square integrable up to finite times, (here $(\cdot)^i$ denotes the i th component of the vector), and $\lambda \in \mathbb{R}$. The key points distinguishing these equations from ‘classical’ BSDEs on Markov chains (as considered in [5]) are that this must hold both for all t and for all T , and that the value λ is a part of the solution.

Our key method will be based on considering the limits of the following ‘discounted’ BSDEs,

$$Y_t^\alpha = Y_T^\alpha + \int_t^T [f(X_{u-}, Z_u^\alpha) - \alpha Y_{u-}^\alpha] du - \int_t^T (Z_u^\alpha)^* dM_u, \quad 0 \leq t \leq T < \infty, \quad (3)$$

the existence of solutions to which we shall first establish.

Remark 1. We note that the use of left limits for X and Y in the driver terms of (2), (3) initially seems unconventional, for those used to the theory of BSDEs in a Brownian setting. However, it is the natural approach when the driver term can itself jump (see [6]), it allows us to perform Girsanov transformations without constantly having to switch between the left and right limits of the processes, and as the integral is with respect to Lebesgue measure and our processes have at most countably many jumps, in this case the equation is unchanged whether the left limits are included or not.

Of importance will be the following process and the associated spaces.

Definition 1. *Let*

$$\psi_t^x := \text{diag}(A_t x) - A_t \text{diag}(x) - \text{diag}(x) A_t^*,$$

for x a basis vector of \mathbb{R}^N . Write ψ_t for the process $\psi_t^{X_{t-}}$. Then ψ is a predictable process taking values in the symmetric, positive semidefinite matrices in $\mathbb{R}^{N \times N}$, with the property that

$$E \left[\left(\int_{]0,t]} Z_u^* dM_u \right)^2 \right] = \int_{]0,t]} E[Z_u^* \psi_u Z_u] du$$

for any t and any predictable processes Z of correct dimension (see [4]). For simplicity, we write

$$\|z\|_{M_t}^2 := z^* \psi_t z,$$

and note that this is a stochastic seminorm.

We define the following spaces of processes.

- $Y \in S^2$ if $E[\sup_t Y_t^2] < \infty$ and Y is càdlàg,
- $Z \in H_M^2$ if $E[\int_{]0,\infty[} \|Z_t\|_{M_t}^2 dt] < \infty$ and Z is predictable,
- $Z \in H_{M,t\text{-loc}}^2$ if $E[\int_{]0,t[} \|Z_s\|_{M_s}^2 ds] < \infty$ for all $t < \infty$ and Z is predictable (note that this is not the usual space of processes locally in H_M^2 , as this must hold for every deterministic t , rather than for a specific sequence of stopping times),
- $Z \sim_M Z'$ if $\|Z_t - Z'_t\|_{M_t} = 0$ for almost all t .

2 Discounted BSDE

We begin with the following result.

Theorem 1. *Let T be a finite deterministic time, and $f : \Omega \times [0, T] \times \mathbb{R} \times \mathbb{R}^N \rightarrow \mathbb{R}$ be a predictable function. If f is uniformly Lipschitz in y and z , that is, there exists a constant $c > 0$ such that*

$$|f(\omega, t, y, z) - f(\omega, t, y', z')| \leq c(|y - y'| + \|z - z'\|_{M_t}),$$

and

$$E\left[\int_{]0,T[} |f(\omega, t, 0, 0)|^2 dt\right] < \infty$$

then for any $\xi \in L^2(\mathcal{F}_T)$, there exists a unique solution $(Y, Z) \in S^2 \times H_M^2$ to the BSDE

$$\xi = Y_t - \int_{]t,T[} f(\omega, u, Y_u, Z_u) du + \int_{]t,T[} Z_u^* dM_u.$$

Proof. For the finite state case, this result is given in [4]. For the infinite state case, we use the martingale representation result established in [4], which naturally extends to general spaces, coupled with the existence result for BSDEs in general spaces established in [6].

We note that, unlike in [4, 6], we have not assumed that \mathcal{F}_0 is trivial, that is, that X_0 is deterministic. Hence Y_0 is also a random variable. This poses no problems for the theory of BSDEs, however it will be useful for us to note that, on the set $\{X_0 = x\}$, we will obtain a deterministic value Y_0^x , as \mathcal{F}_0 is nothing but the completion of $\sigma(X_0)$. \square

A key result in the analysis of BSDEs is the comparison theorem. In the case of BSDEs with Markov Chain noise, and in general for BSDEs with jumps, a further condition is required to ensure that the result holds. In [7, 3] a general condition under which the comparison theorem holds is presented, and in [5] a condition specific to Markov chain BSDEs was also given.

Definition 2. *For a driver f , we say that*

- f is balanced if

$$\frac{f(\omega, t, y, z) - f(\omega, t, y, z')}{\|z - z'\|_{M_t}^2} (z - z')^* \Delta M_t > -1,$$

- f is weakly balanced if

$$\left(\frac{f(\omega, t, y, z) - f(\omega, t, y, z')}{\|z - z'\|_{M_t}^2} \wedge 0 \right) (z - z')^* \Delta M_t > -1,$$

- f is strictly balanced if

$$\frac{f(\omega, t, y, z) - f(\omega, t, y, z')}{\|z - z'\|_{M_t}^2} (z - z')^* \Delta M_t > -1 + \gamma$$

for some $\gamma > 0$,

where all inequalities must hold for any $y \in \mathbb{R}$, any $z, z' \in \mathbb{R}^N$, and up to indistinguishability.

Clearly strictly balanced \Rightarrow balanced \Rightarrow weakly balanced. Our main attention is on those cases where f is (strictly) balanced, however, the comparison theorem holds assuming only that f is weakly balanced.

Theorem 2 (Finite-time comparison theorem). *Let (Y, Z) and (Y', Z') be the solutions to two BSDEs with drivers f and f' . Suppose f is weakly balanced and $f(\omega, t, y, z) \geq f'(\omega, t, y, z)$ for all (y, z) , $dt \times d\mathbb{P}$ -a.s. Then $Y_T \geq Y'_T$ a.s. implies $Y_t \geq Y'_t$ a.s. up to indistinguishability.*

Proof. This is simply Theorem 3 of [3], where a trivial modification of the last line of the proof of Lemma 1 of [3] is needed to exploit the definition of a weakly balanced driver. \square

In fact, in the balanced case, the proof of the comparison theorem is easy to deduce from the following lemma.

Lemma 1. *If f is Lipschitz and balanced, then for any predictable processes $Z, Z' \in H_{M, t\text{-loc}}^2$, any process $Y \in S^2$, any $T < \infty$, the measure \mathbb{Q}^T defined by*

$$\frac{d\mathbb{Q}^T}{d\mathbb{P}} = \mathcal{E} \left(\int_{]0, \cdot \wedge T]} \frac{f(\omega, t, Y_{t-}, Z_t) - f(\omega, t, Y_{t-}, Z'_t)}{\|Z_t - Z'_t\|_{M_t}^2} (Z_t - Z'_t)^* dM_t \right),$$

is a probability measure (where \mathcal{E} denotes the Doléans-Dade exponential), and

$$\tilde{M}_t = \int_{]0, t \wedge T]} (f(\omega, s, Y_{s-}, Z_s) - f(\omega, s, Y_{s-}, Z'_s)) ds + \int_{]0, t \wedge T]} (Z_s - Z'_s)^* dM_s$$

is a \mathbb{Q}^T -martingale.

Proof. First note that

$$L_t := \int_{]0, t \wedge T]} \frac{f(\omega, s, Y_{s-}, Z_s) - f(\omega, s, Y_{s-}, Z'_s)}{\|Z_s - Z'_s\|_{M_t}^2} (Z_s - Z'_s)^* dM_s$$

is a local martingale with quadratic variation

$$\langle L \rangle_t = \int_{]0, t \wedge T]} \frac{|f(\omega, s, Y_{s-}, Z_s) - f(\omega, s, Y_{s-}, Z'_s)|^2}{\|Z_s - Z'_s\|_{M_t}^4} \|Z_s - Z'_s\|_{M_t}^2 dt \leq cT,$$

for c a Lipschitz constant of f . Hence we know \tilde{M}_t^T is a true (BMO-)martingale with all moments finite (see [3, Lemma 2]). Then, as f is balanced, we see that $\Delta L_t > -1$, so \mathbb{Q}^T is a probability measure.

To show that the desired process is then an application of Girsanov's theorem. By Hölder's inequality we have, for any stopping time $\tau \leq T$, any $\epsilon < 1$,

$$E_{\mathbb{Q}^T}[\tilde{M}_\tau^{2-\epsilon}] \leq E_{\mathbb{P}}[L_T^{2/\epsilon}]^{\epsilon/2} E_{\mathbb{P}}[\tilde{M}_\tau^2]^{1-\epsilon/2}$$

which is uniformly bounded, so \tilde{M} is a true \mathbb{Q}^T -martingale. \square

The connection between these definitions of 'balanced' drivers and the condition given in [5] is given by the following lemma, which is presented for completeness

Lemma 2. *The following conditions are equivalent.*

(i) f is weakly balanced.

(ii) For any $z, z' \in \mathbb{R}^N$, up to indistinguishability, on the set where

$$(e_i^* A_t X_{t-})[(z - z')^*(e_i - X_{t-})] \geq 0$$

for all i , we have

$$(f(\omega, t, y, z) - f(\omega, t, y, z')) \wedge 0 \geq -(z - z')^* A_t X_{t-} \quad a.s.,$$

with equality only when $\|z - z'\|_{M_t} = 0$.

Proof. Clearly the conditions are trivial on the set $f(\omega, t, y, z) - f(\omega, t, y, z') = 0$, so we exclude this from consideration.

(ii) \Rightarrow (i). First note that the condition in Definition 2 is equivalent to

$$[(f(\omega, t, y, z) - f(\omega, t, y, z')) \wedge 0](z - z')^*(X_t - X_{t-}) > -(z - z')^* \psi_M(z - z') \quad a.s. \quad (4)$$

For fixed values of t, y, z, X_{t-} , suppose there is i with $(e_i^* A_t X_{t-})[(z - z')^*(e_i - X_{t-})] < 0$. Then let

$$h = \left(\frac{f(\omega, t, y, z) - f(\omega, t, y, z')}{(z - z')^*(e_i - X_{t-})} \vee 0 \right) (e_i - X_{t-}).$$

If there is no such i , then note that $(z - z')^* A_t X_{t-} < 0$ (as it is the compensator of a nondecreasing process), and let

$$h = \left(\frac{f(\omega, t, y, z) - f(\omega, t, y, z')}{(z - z')^* A_t X_{t-}} \vee 0 \right) A_t X_{t-}$$

In either case, $h^*(z - z') = (f(\omega, t, y, z) - f(\omega, t, y, z')) \wedge 0$, $\mathbf{1}^* h = 0$ and $h^* e_i \geq 0$ for all $e_i \neq X_{t-}$. In the latter case, note also that $\mathbf{1}^*(A_t X_{t-} - h) = 0$ and $e_i^*(A_t X_{t-} - h) > 0$ for all $e_i \neq X_{t-}$.

By considering all possible jumps, we then have that (4) simplifies to

$$(z - z')^*[h(e_j - X_{t-})^* + \psi_t](z - z') > 0$$

for all e_j with $e_j^* A_t X_{t-} > 0$. This is equivalent to

$$(z - z')^*[h(e_j - X_{t-})^* + (e_j - X_{t-})h^* + 2\psi_t](z - z') > 0.$$

As $z - z'$ is at most defined up to the addition of a constant, for each e_j , we can suppose without loss of generality that $(z - z')^* e_j = 0$. Hence it is enough to show that the symmetric matrix $[2\psi_t - hX_{t-}^* - X_{t-}h^*]$ is positive definite. However, as this matrix is of the same form as ψ_t , this is straightforward.

(i) \Rightarrow (ii). We know that (4) holds. Suppose that $(e_i^* A_t X_{t-})[(z - z')^*(e_i - X_{t-})] \geq 0$ for all i . Without loss of generality, we select a representation of $z - z'$ such that $(z - z')^* X_{t-} = 0$ and $(z - z')^* e_j = 0$ for all j with $e_j^* A_t X_{t-} = 0$. Note that this then implies $z - z'$ is componentwise nonnegative. Then (4) reduces to

$$(f(\omega, t, y, z) - f(\omega, t, y, z')) \wedge 0 > -\frac{(z - z')^* \psi_M(z - z')}{(z - z')^*(X_t)} = -\frac{\sum_i a_i (z - z')_i^2}{(z - z')^*(X_t)}$$

with the convention $0/0 = \infty$, where $a_i = e_i^* A_t X_{t-}$. This inequality must hold almost surely, so it must hold in the case $X_t = e_j$, for j maximizing $(z - z')^* e_j = (z - z')_j$. As $(z - z')_i / (z - z')_j < 1$, we have

$$\begin{aligned} (f(\omega, t, y, z) - f(\omega, t, y, z')) \wedge 0 &> -\frac{\sum_i a_i (z - z')_i^2}{(z - z')_j} \\ &> -\sum_i a_i (z - z')_i = -(z - z')^* A_t X_{t-} \end{aligned}$$

as desired. \square

Corollary 1. *If f is balanced, then we also have that for any $z, z' \in \mathbb{R}^N$, up to indistinguishability, if*

$$(e_i^* A_t X_{t-})[(z - z')^*(e_i - X_{t-})] \geq 0$$

for all i , then

$$f(\omega, t, y, z) - f(\omega, t, y, z') \geq -(z - z')^* A_t X_{t-} \quad a.s.,$$

with equality only when $\|z - z'\|_{M_t} = 0$.

Proof. The proof of (i) \Rightarrow (ii) in Lemma 2 is sufficient to prove this result, where the term ' $\wedge 0$ ' is removed throughout. \square

The following version of Tanaka's formula will be useful. We adopt the convention that $\text{sign}(x) = x/|x|$ for $x \neq 0$ and $\text{sign}(0) = 0$.

Lemma 3. *Let Y be a càdlàg process of finite variation. Then the dynamics of $|Y|$ are given by*

$$d|Y|_t = \text{sign}(Y_{t-})dY_t + \Delta L_t^Y$$

where ΔL_t^Y is the 'local time' jump process, with

$$\Delta L_t^Y = |Y_t|(1 - \text{sign}(Y_t Y_{t-})) \geq 0.$$

Proof. The dynamics of $|Y|$ are clear except when $Y_{s-} = 0$ or when Y jumps over zero, that is, when Y_{t-} and Y_t are of opposite sign.

If $Y_{s-} = 0$, then either $|Y_s| > 0$, in which case $\Delta|Y_s| = \Delta L_t^Y$, or $Y_s = 0$, in which case there is a version of the derivative of Y which is zero at s (as the measure induced by Y does not charge the point s). In either case, the dynamics are satisfied.

If Y_{t-} and Y_t are of opposite sign, we have

$$\begin{aligned}\Delta|Y|_t &= \text{sign}(Y_{t-})\Delta Y_t + (\Delta|Y|_t - \text{sign}(Y_{t-})\Delta Y_t) \\ &= \text{sign}(Y_{t-})\Delta Y_t + (|Y_t| - |Y_{t-}| - \text{sign}(Y_{t-})(Y_t - Y_{t-})) \\ &= \text{sign}(Y_{t-})\Delta Y_t + (|Y_t| - \text{sign}(Y_{t-})Y_t) \\ &= \text{sign}(Y_{t-})\Delta Y_t + |Y_t|(1 - \text{sign}(Y_{t-}Y_t))\end{aligned}$$

and so in both cases the dynamics hold. \square

We now seek to show that there exists a bounded solution to the infinite-horizon BSDE with discounting. The approach used to obtain this result is based on Briand and Hu [2] and Royer [23]. However, in our setting, the following result greatly simplifies the analysis.

Lemma 4. *Let Y satisfy the dynamics*

$$dY_t = \beta_t dt + Z_t^* dM_t$$

for some arbitrary predictable process β . If $|Y| \leq k$ for some k , then $|e_i^* Z| \leq 2k$ for all i , for at least one representative in $H_{M,t\text{-loc}}^2$. That is, Z is bounded componentwise by twice a bound on Y .

Proof. First note that as the jumps of M are totally inaccessible, if an inequality holds for every jump time, then it must hold almost everywhere on every set where jumps occur with positive probability, that is, up to a representative in $H_{M,t\text{-loc}}^2$. Now note that at any jump,

$$Y_t - Y_{t-} = \Delta Y_t = Z_t^* \Delta M_t = Z_t^*(X_t - X_{t-}).$$

Therefore,

$$|Z_t^*(X_t - X_{t-})| \leq 2k \quad a.s.$$

We now take a representative Z such that $Z^* X_{t-} \equiv 0$, which can be done as Z is only ever defined up to addition of a constant. Therefore $|Z_t^* X_t| \leq 2k$ for every X_t which can be reached with positive probability in a single jump. Taking $Z_t = 0$ for all X_t which cannot be reached in a single jump (which can be done up to equality \sim_M), we see that $|e_i^* Z| \leq 2k$. \square

Theorem 3. *Let $\alpha > 0$ and $f : \Omega \times \mathbb{R}^+ \times \mathbb{R}^N \rightarrow \mathbb{R}$ be*

- *uniformly Lipschitz (in its third component) with respect to the $\|\cdot\|_{M_t}$ -norm $d\mathbb{P} \times dt$ -a.e.*
- *balanced, in the sense of Definition 2 (omitting the y variable), and*
- *such that $|f(\omega, t, 0)|$ is uniformly bounded by $C \in \mathbb{R}$.*

Then there exists an adapted solution (Y, Z) , with Y càdlàg and $Z \in H_{M, t\text{-loc}}^2$, to the equation

$$Y_T = Y_t - \int_{]t, T]} (-\alpha Y_{u-} + f(\omega, u, Z_u)) du + \int_{]t, T]} Z_u^* dM_u, \quad 0 \leq t \leq T < \infty \quad (5)$$

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

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

$$0 = Y_t^T - \int_{]t, T]} (-\alpha Y_{u-}^T + f(\omega, u, Z_u^T)) du + \int_{]t, T]} (Z_u^T)^* dM_u \quad (6)$$

then $\lim_{T \rightarrow \infty} Y_t^T = Y_t$ a.s., uniformly on compact sets in t .

Proof. Uniqueness. We first show that, if a bounded solution exists, it is unique. Suppose we have two bounded solutions (Y, Z) and (Y', Z') to (5). For simplicity, write $\delta Y = Y - Y'$ and $\delta Z = Z - Z'$.

For $T > 0$, let \mathbb{Q}_1^T be the measure with density

$$\frac{d\mathbb{Q}_1^T}{d\mathbb{P}} = \mathcal{E} \left(\int_{]0, \wedge T]} \frac{(f(\omega, u, Z_u) - f(\omega, u, Z'_u)) (\delta Z_u)^* dM_u}{\|\delta Z_u\|_{M_u}^2} \right), \quad (7)$$

where \mathcal{E} denotes the Doléans-Dade exponential. As f is balanced, we can see that \mathbb{Q}_1^T is a probability measure. By Lemma 1, it follows that

$$\begin{aligned} & - \int_{]0, t]} \text{sign}(\delta Y_{u-}) (f(\omega, u, Z_u) - f(\omega, u, Z'_u)) du \\ & + \int_{]0, t]} \text{sign}(\delta Y_{u-}) \delta Z_u^* dM_u + \sum_{u \leq t} \Delta L_u^{(\delta Y_u)} \end{aligned}$$

is a \mathbb{Q}_1^T submartingale on $[0, T]$. (Note that the inclusion of $\text{sign}(\delta Y_{u-})$ simply exchanges Z and Z' , and does not affect the quadratic variation.)

By Tanaka's formula and Itô's formula, we have, for all $s \leq t \leq T$,

$$E_{\mathbb{Q}_1^T} [e^{-\alpha t} |\delta Y_t| - e^{-\alpha s} |\delta Y_s| | \mathcal{F}_s] \geq 0,$$

hence,

$$|\delta Y_s| \leq e^{-\alpha(t-s)} E_{\mathbb{Q}_1^T} [|\delta Y_t| | \mathcal{F}_s] \leq e^{-\alpha(t-s)} C,$$

for C a bound on $|\delta Y_t|$. This bound is independent of T , and collapses as $t \rightarrow \infty$. Hence $|\delta Y_s| = 0$, from which we see $Y_s = Y'_s$ a.s. for every s , and hence $Y = Y'$ up to indistinguishability as Y and Y' are càdlàg.

Existence. We now show that a bounded solution exists. Let (Y^T, Z^T) denote the solution to the time T -horizon BSDE, as defined in (6).

First, we show that Y^T is bounded. Again, we invoke Lemma 1 and let \mathbb{Q}_2^T denote the probability measure with density

$$\frac{d\mathbb{Q}_2^T}{d\mathbb{P}} = \mathcal{E} \left(\int_{]0, \wedge T]} \frac{f(\omega, u, Z_u^T) - f(\omega, u, 0)}{\|Z_u\|_{M_u}^2} (Z_u^T)^* dM_u \right).$$

Applying Tanaka's formula and Itô's formula to $e^{-\alpha t}|Y_t^T|$, we see that

$$|Y_t^T| \leq e^{\alpha t} E_{\mathbb{Q}_2^T} \left[\int_t^T e^{-\alpha u} |g(\omega, u, 0)| du \middle| \mathcal{F}_t \right] \leq C/\alpha \quad (8)$$

for C a bound on $|f(\omega, u, 0)|$. Hence Y^T is uniformly bounded.

Second, we show that Y^T forms a Cauchy sequence in T . For any $T' > T$, we use Lemma 1 to take the probability measure

$$\frac{d\mathbb{Q}_3^{T,T'}}{d\mathbb{P}} = \mathcal{E} \left(\int_{]0, \cdot \wedge T]} \frac{f(\omega, u, Z_u^T) - f(\omega, u, Z_u^{T'})}{\|Z_u^T - Z_u^{T'}\|_{M_u}^2} (Z_u^T - Z_u^{T'})^* dM_u \right).$$

Again applying Tanaka's formula, Itô's formula and the inequality (8), for $t < T$ we have

$$|Y_t^T - Y_t^{T'}| \leq e^{-\alpha(T-t)} E_{\mathbb{Q}_3^{T,T'}} [|Y_T^{T'}| | \mathcal{F}_t] \leq C e^{-\alpha(T-t)} / \alpha \quad a.s.$$

Hence we see that Y_t^T is a Cauchy sequence in T , and so a limit exists, and we denote it Y_t . The desired convergence uniformly on compacts is clear, and the bound established in (8) also holds for Y_t .

Finally, as Y_t^T is uniformly bounded, so is Z_t^T , by Lemma 4. As Y_t^T converges a.s. uniformly on compacts in t , its jumps converge a.s. uniformly on compacts, however this implies that the Z_t^T also converge a.s. uniformly on compacts, at least up to equivalence in $\|\cdot\|_{M_t}$. Taking Z as the limit of Z^T , we have our desired solution (Y, Z) . \square

To finish this section, we finally state a result on the existence of 'Markovian' solutions to these BSDEs, that is, when the BSDE solution (Y, Z) can be written as a function of the underlying state process X .

Theorem 4. *Suppose our Markov chain X is irreducible at every point of time. Consider either*

- *the situation of Theorem 1, when the terminal value of the BSDE is of the form $Q = \phi(X_T)$ for some deterministic function $\phi : \mathbb{R}^N \rightarrow \mathbb{R}$, or*
- *the situation of Theorem 3,*

and suppose that the driver f factors through X_{t-} , as a function of ω , that is, \tilde{f} can be written in the form $f(\omega, t, \dots) = \tilde{f}(X_{t-}(\omega), t, \dots)$ for some function \tilde{f} . Then there exists a function $v : [0, \infty[\times \mathcal{X} \rightarrow \mathbb{R}$ such that $Y_t = v(t, X_t)$ and $e_i^ Z_t = v(t, e_i)$, and this function satisfies the coupled ODE system*

$$\frac{dv(t, e_i)}{dt} = -\tilde{f}(e_i, t, v(t, e_i), v(t, \cdot)) - e_i^* A_t^* v(t, \cdot)$$

in the setting of Theorem 1, respectively

$$\frac{dv(t, e_i)}{dt} = -\alpha v(t, e_i) - \tilde{f}(e_i, t, v(t, \cdot)) - e_i^* A_t^* v(t, \cdot),$$

in the setting of Theorem 3, where $v(t, \cdot)$ refers to the vector in \mathbb{R}^N with entries $\{v(t, e_j)\}_{j=1}^N$.

Proof. For N finite, in the first case, this result is directly from [8]. It is easy to verify that the same arguments as in [8] will hold for N infinite. For the situation of Theorem 3, note that the finite-time approximations Y_t^T constructed in the proof of that theorem are all examples of the first case, with $\phi \equiv 0$. Hence we can find functions $v^T(\cdot, \cdot)$ satisfying the desired statements, and the proven convergence allows us to take $T \rightarrow \infty$. The only difficulty with this is the stated dynamics on $v(t, e_i)$, which can easily be verified directly, as in [8]. \square

Corollary 2. *Suppose that X is time-homogeneous (so $A = A_t$ is constant), and we are in the situation of Theorem 3, where our driver factors through X_{t-} , and does not depend on time. Then v does not depend on time, and we have the equation in \mathbb{R}^N*

$$\alpha v = -\tilde{f}(\cdot, v) - A^*v$$

where $v = v(t, \cdot)$ is a vector in \mathbb{R}^N , and $\tilde{f}(\cdot, v)$ refers to the vector with entries $\{\tilde{f}(e_i, v)\}_{i=1}^N$. If \tilde{f} is balanced, using the natural modification of Definition 2 (and correspondingly Lemma 2) this equation admits a unique solution.

Proof. Under the conditions of the corollary, as X is a Markov chain, our infinite-horizon BSDE (5) does not vary in t given the state X_t . It is then clear that the unique solution is a function purely of X_t . The dynamics are then the natural rewriting of those in Theorem 4, and the existence and uniqueness of the solution is the result of Theorem 3. \square

Remark 2. It is worth noting that this is a nontrivial algebraic statement, due to the nonlinearity of the function \tilde{f} . If we knew that $\tilde{f}(v) + A^*v$ was a strictly monotone function, in the sense that

$$\langle \tilde{f}(v) - \tilde{f}(v') + A^*(v - v'), v - v' \rangle \leq -\epsilon \|v - v'\|^2$$

for some $\epsilon > 0$ and all v, v' , then it would be possible to construct an existence result using standard techniques (see, for example [24, p.565]). Initially, this would appear to be true, at least for chains with finitely many states, using the fact that the nonlinear equation generates measures under which X is a Markov chain (see the proof of Lemma 12), and the Perron–Frobenius theorem or Krein–Rutman theorem to bound eigenvalues below zero. However, such an argument depends on the diagonalizability of the derivative of $\tilde{f}(v) + A^*v$ which is a non-trivial assumption. For example, one might consider the situation on four states with $\tilde{f} = 0$, and

$$A^* = \begin{bmatrix} -3 & 1 & 2 & 0 \\ 1 & -3 & 2 & 0 \\ 0 & 2 & -3 & 1 \\ 0 & 2 & 1 & -3 \end{bmatrix} \quad (9)$$

which is a very well behaved Markov chain generator, however is defective as a matrix and so it is easy to find vectors v with $v'Av > 0$. Hence this argument fails in general.

3 Uniformly ergodic chains

We aim to construct solutions to the ergodic BSDE (2), using the solutions to the discounted BSDEs (5). As our aim is to contemplate the long-run behaviour, we

shall now make the following two assumptions, for the remainder of the paper. First, the Markov chain X is time homogeneous (so $A_t = A$ for all t) and irreducible, and second, the BSDE drivers we consider will all factor through X , so our BSDE solutions will be functions of the current state of the Markov chain, and, in particular, Corollary 2 holds.

The key to using discounted BSDEs to approximate the ergodic BSDE is to study the ergodic behaviour of the Markov chain itself. We shall therefore make moderately restrictive assumptions on the Markov chain, sufficient to give explicit bounds on its convergence to its ergodic distribution, and then show that these bounds carry over to the solution of the relevant BSDEs.

Definition 3. *Let \mathcal{M} denote the set of probability measures on \mathcal{X} , with the topology inherited from considering them as a convex subset of $\ell_1(\mathcal{X})$ (the total variation topology, with norm $\|f\|_{TV} = \sum_x |f(x)|$). We write $P_t\mu$ for the law of X_t given $X_0 \sim \mu$.*

We say the Markov chain X is uniformly ergodic if there exists a measure π on \mathcal{X} , and constants $R, \rho > 0$, such that

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

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

We now prove a few standard properties of such chains. The following lemma is simply a variant of [18, Theorem 16.2.2(iv)]. For simplicity, we will write $E_x[\cdot]$ for $E[\cdot | X_0 = x]$.

Lemma 5. *Let X be a uniformly ergodic Markov chain, and let x_C be an arbitrary state. Let τ_C be the first hitting time of x_C . Then for some $\beta > 0$ (and hence for all β sufficiently small),*

$$G_{x_C}(\beta) := \sup_{x \in \mathcal{X}} E_x[e^{\beta\tau_C}] < \infty.$$

The following general continuity lemma allows us to take a bound, such as that established in Lemma 5, and force it to converge for small β .

Lemma 6. *Let T be a random variable, and consider $G(\beta) = \sup_{\nu} E_{\nu}[e^{\beta T}]$, where ν is a parameterization of measures under which the expectation is taken (in our case, the family of initial states of the Markov chain). If there exists $\beta^* > 0$ such that $G(\beta^*) \leq K$, for some $K < \infty$, then for any $\epsilon > 0$,*

$$G(\beta) \leq 1 + \epsilon \quad \text{for all } \beta \in \left[0, \frac{\beta^*}{2} \left(\frac{\epsilon}{K} \wedge 1\right)\right].$$

Note in particular that the bound does not depend on the underlying measures, given K and β^ .*

Proof. Consider $G^{\nu}(\beta) = E_{\nu}[e^{\beta T}]$. For $\delta > 0$ we have

$$E_{\nu}[Te^{\beta T}] \leq E_{\nu}[(\delta^{-1}e^{\delta T})e^{\beta T}] = \frac{1}{\delta} E_{\nu}[e^{(\delta+\beta)T}] = \frac{1}{\delta} G^{\nu}(\delta + \beta) < \infty,$$

so using the dominated convergence theorem to exchange the order of integration and differentiation, in the region $\delta + \beta \leq \beta^*$ we have

$$\frac{d}{d\beta} G^{\nu}(\beta) = \frac{d}{d\beta} E_{\nu}[e^{\beta T}] = E_{\nu}[Te^{\beta T}] \leq \frac{1}{\delta} G^{\nu}(\delta + \beta).$$

Hence, fixing $\delta = \beta^*/2$, we have a (uniform) bound, on the derivative $\frac{d}{d\beta}G^\nu(\beta)$ for all $\beta \leq \beta^*/2$. Therefore, as $G^\nu(0) = 1$, we know

$$G^\nu(\beta) \leq 1 + \beta \left(\frac{2}{\beta^*} G(\beta^*) \right) \leq 1 + \beta \left(\frac{2}{\beta^*} K \right)$$

which implies

$$G^\nu(\beta) \leq 1 + \epsilon \quad \text{for all } \beta \leq \frac{\beta^*}{2} \left(\frac{\epsilon}{K} \wedge 1 \right).$$

Taking a supremum, we obtain the uniform bound $G(\beta) \leq 1 + \epsilon$. \square

We now return to general properties of uniformly integrable Markov chains.

Lemma 7. *Let X be a uniformly ergodic Markov chain, and let $|a_{xx}| := -x^*Ax$ denote the rate of transitions out of state x . Then $\{|a_{xx}|\}_{x \in \mathcal{X}}$ is uniformly bounded above 0.*

Proof. We know that the probability of being in the same state after a fixed timestep δ is given by $\exp(-\delta|a_{xx}|)$. Taking the δ skeleton of the chain, from [18, Theorem 16.2.2(iii)] we see that the expected time until a transition from state x is uniformly bounded with respect to x . By a geometric distribution argument this expectation is given by $(1 - \exp(-\delta|a_{xx}|))^{-1}$, hence $|a_{xx}|$ is bounded away from zero. \square

Lemma 8. *Let X be a uniformly ergodic Markov chain on countable states, with bounded transition rates. Then the discrete time chain obtained by observing X at each of its jump times is also uniformly ergodic.*

Proof. Let x_C be an arbitrary state, which is petite as our chain has countably many states. Let N be the number of jumps up to hitting x_C , and let the i th jump be denoted T_i , with $T_0 = 0$. The uniform ergodicity of a Markov chain is then equivalent to stating that $\sup_x E_x[T_N] < \infty$, where E_x denotes the expectation conditional on $X_0 = x$. If α is an upper bound on the transition rates, then

$$\begin{aligned} E_x[T_N] &= E_x \left[\sum_{i=1}^N (T_i - T_{i-1}) \right] = E_x \left[\sum_{i=1}^{\infty} I_{\{T_{i-1} < T_N\}} E[T_i - T_{i-1} | \mathcal{F}_{T_{i-1}}] \right] \\ &\geq E_x \left[\sum_{i=1}^{\infty} I_{\{T_{i-1} < T_N\}} \alpha^{-1} \right] = \alpha^{-1} E_x[N] \end{aligned}$$

Hence $\sup_x E_x[N] < \infty$, so the Markov chain generated by observations at jump times is also uniformly ergodic. \square

To prove ergodicity properties, a standard technique is based on coupling copies of the Markov chain (see Lindvall [16] for an overview of these methods). We here shall work with the simplest form of coupling, when we take two independent copies X and Y of our chain, which have distinct initial states, and we study their first meeting time. We note that throughout this section, Y will be used to refer to a copy of a Markov chain, rather than to the solution of a BSDE.

Lemma 9. *Let X and Y be two independent copies of the uniformly ergodic chain with rate matrix A . Let $T = \inf\{t : X_t = Y_t\}$. Then there exists $\beta > 0$ such that*

$$G^*(\beta) := \sup_{x,y} E_{xy}[e^{\beta T}] < \infty,$$

where E_{xy} is the expectation conditional on $X_0 = x$ and $Y_0 = y$.

Proof. Let x_C be an arbitrary state, with ergodic probability $\pi(x_C)$. As the chain is uniformly ergodic, we know

$$\sup_y |P(Y_t = x_C | Y_0 = y) - \pi(x_C)| \rightarrow 0 \text{ as } t \rightarrow \infty,$$

so we can find $t^* < \infty$ such that $\sup_y |P(Y_{t^*} = x_C | Y_0 = y) - \pi(x_C)| < \min\{\pi(x_C), 1 - \pi(x_C)\}/2$, in particular, $P(Y_t = x_C) > p$ for all $t \geq t^*$, for some fixed $p > 0$.

Define the following sequence of stopping times,

$$\tau_0 = \inf\{t \geq t^* : X_t = x_C\}, \quad \tau_i = \inf\{t \geq \tau_{i-1} + t^* : X_t = x_C\}.$$

We shall show that $T_C = \inf\{\tau_i : X_{\tau_i} = Y_{\tau_i}\}$ has exponential moments, the result will then follow as $T \leq T_C$. As the two paths are independent, the definition of p implies that for each i , $P(Y_{\tau_i} = x_C | \mathcal{F}_{\tau_{i-1}}) \geq p$. Hence a geometric trials argument gives

$$\begin{aligned} E_{xy}[e^{\beta T_C}] &\leq E_{xy}\left[\sum_{i=0}^{\infty} e^{\beta \tau_i} (1-p)^i\right] \\ &= \sum_{i=0}^{\infty} E_{xy}\left[\prod_{j \leq i} \left(e^{\beta t^*} (1-p) E_{xy}[e^{\beta(\tau_j - \tau_{j-1} - t^*)} | \mathcal{F}_{\tau_{j-1} + t^*}]\right)\right] \\ &\leq \sum_{i=0}^{\infty} E_{xy}\left[\prod_{j \leq i} \left(e^{\beta t^*} (1-p) G_{x_C}(\beta)\right)\right] \\ &= \sum_{i=0}^{\infty} (e^{\beta t^*} (1-p) G_{x_C}(\beta))^i \\ &= \left(1 - (e^{\beta t^*} (1-p) G_{x_C}(\beta))\right)^{-1}. \end{aligned}$$

Hence, by taking β small enough that $e^{\beta t^*} G_{x_C}(\beta) < (1-p)^{-1/2}$, which can be done by Lemma 6, we have $G^*(\beta) < \infty$. \square

Corollary 3. *If X and Y are two independent copies of a uniformly ergodic Markov chain with bounded transition rates, and if N is the total number of jumps (of both copies) until they first meet, it follows that $\sup_{x,y} E_{xy}[e^{\beta N}] < \infty$ for some $\beta > 0$.*

Proof. By Lemma 8, the discrete time chain obtained by observing only when one of the chains jumps is also uniformly ergodic. Hence simply apply a discrete-time version of Lemma 9 to this setting, as the argument is identical. \square

The following lemma provides a ‘lower’ exponential moment.

Lemma 10. *Let X and Y be two independent copies of a uniformly ergodic chain with bounded transition rates, and let N be the total number of jumps (of both copies) until they first meet. For any fixed $\gamma \in]0, 1[$ we know,*

$$q_\gamma := \inf_{x,y} E_{xy}[\gamma^N] > 0.$$

Proof. By Lemma 8 the Markov chain generated by observations at jump times is also uniformly ergodic. Let $q(x, y; n) = P(N = n | X_0 = x, Y_0 = y)$. By Corollary 3 we know that there exists $\beta > 0$ such that $\sum_n e^{\beta n} q(x, y; n)$ is uniformly bounded for all x, y . Therefore, for some constant c , we know $0 \leq q(x, y; n) < ce^{-\beta n}$ for all x, y and n , hence $\sum_{n=m}^\infty q(x, y; n) < \frac{c}{1-e^\beta} e^{-\beta m+1}$. As we also know $\sum_n q(x, y; n) = 1$, this implies that, taking $m > \log(\frac{1-e^\beta}{c})/\beta - 1$, there is a positive lower bound on $\sum_{n=0}^{m-1} q(x, y; n)$ which is uniform in x and y , and hence $q_\gamma > 0$. \square

3.1 Perturbed rates

We shall see later that the solutions to an infinite horizon BSDE can be viewed as expectations under a perturbation of the rate matrix of the underlying Markov chain. Hence, we wish to show that if the underlying chain is uniformly ergodic then this remains the case under the perturbed measure. To do so, we use a variant of the Nummelin splitting (see [18, Chapter 5]), which is commonly used to prove ergodicity properties.

The class of perturbations under consideration are neatly expressed through the following definition.

Definition 4. *Consider A and B (possibly infinite) rate matrices, that is, matrices with $A_{ij} \geq 0$ for $i \neq j$ and $\sum_i A_{ij} = 0$ for all j , and similarly for B . We shall write $B \succeq A$ whenever $B - A$ is also a rate matrix.*

We shall say that B is controlled by A whenever there exists some $\gamma > 0$ such that $\gamma A \preceq B$. We shall say that B is strictly controlled by A whenever we also know that $B - \gamma A$ has diagonal entries bounded above by $-\gamma$.

Note that without loss of generality γ can be taken to be arbitrarily small, in particular, we assume it is less than or equal to a positive lower bound on $|a_{xx}|$, which exists by Lemma 7. Note also that the relation \succeq defined in this way is a valid partial ordering of rate matrices, which has not, to our knowledge, been previously explored.

Theorem 5. *Let A and B be rate matrices, and suppose B is strictly controlled by A with constant γ . If, under the measure induced by A , the process X is uniformly ergodic, then X is also uniformly ergodic under the measure induced by B . Furthermore, the constant R and the rate ρ of convergence to the ergodic distribution, in Definition 3, can be taken as functions only of A and γ , and the constant R can be made arbitrarily close to 1 (with a corresponding decrease in the rate ρ).*

The proof of this theorem is the main purpose of this section. We begin by defining a splitting of the Markov chain, à la [18]. We assume the conditions of the theorem throughout the remainder of the section, and fix a constant γ satisfying Definition 4.

Definition 5. We define the split space of X to be $\check{X} = \mathcal{X} \times \{0, 1\}$. For notational simplicity, we write $\mathcal{X}_0 := \mathcal{X} \times \{0\}$ and $\mathcal{X}_1 := \mathcal{X} \times \{1\}$, so that $\check{X} = \mathcal{X}_0 \cup \mathcal{X}_1$.

To split a measure ν on \mathcal{X} into a measure $\check{\nu}$ on \check{X} , for $\check{y} \in \check{X}$, if y is the projection of \check{y} , then

$$\check{\nu}(\check{y}) = \begin{cases} (1 - \gamma)\nu(y) & \text{for } \check{y} \in \mathcal{X}_0, \\ \gamma\nu(y) & \text{for } \check{y} \in \mathcal{X}_1. \end{cases}$$

If $\check{\nu}$ is a signed measure with a single negative component (e.g. a column of a rate matrix), then the negative component is not split, but is assigned to one of the two corresponding states in its entirety, yielding the two splittings $\check{\nu}^{(0)}, \check{\nu}^{(1)}$.

Finally we define the split rate matrix \mathcal{B} by

$$\mathcal{B}\check{x} = \begin{cases} \check{g}_x^{(0)} & \check{x} = x \times \{0\} \in \mathcal{X}_0, \\ \check{a}_x^{(1)} & \check{x} = x \times \{1\} \in \mathcal{X}_1, \end{cases}$$

where $g_x = (1 - \gamma)^{-1}(B - \gamma A)x$ and $a_x = Ax$.

Intuitively, transitions occur from $\check{x} \in \mathcal{X}_1$ following the vector Ax , and from $\check{x} \in \mathcal{X}_0$ following the vector $(1 - \gamma)^{-1}(B - \gamma A)x$, except that the result is randomly split between $(\mathcal{X}_0, \mathcal{X}_1)$ with probabilities $(1 - \gamma, \gamma)$. Note that, as in the classical Nummelin splitting, this split chain has marginal transition matrix B , that is, for any measure q on \mathcal{X} , we know $\int \mathcal{B}d\check{q} = \int Bdq$.

Our procedure is now to consider two independent copies of our Markov chain on the split space, and to show that the first meeting time of these chains admits exponential moments. This will allow us to prove the desired uniform ergodicity estimates. We prove this in the following extended lemma.

Lemma 11. Suppose B is strictly controlled by A , with constant $\gamma \leq \inf_x |a_{xx}|$, and let \check{X} and \check{Y} be copies of the chain on the split space, as described in Definition 5.

Let $T = \inf\{t : \check{X}_t = \check{Y}_t\}$ and

$$H^*(\beta) = \sup_{\check{x}, \check{y}} E[e^{\beta T} | \check{X}_0 = \check{x}, \check{Y}_0 = \check{y}].$$

For any $\epsilon > 0$, there exists a constant $\tilde{\beta}_{\gamma\epsilon}$ such that $H^*(\beta) < 1 + \epsilon$ for all $\beta \leq \tilde{\beta}_{\gamma\epsilon}$. Furthermore, $\tilde{\beta}_{\gamma\epsilon}$ is dependent only on γ and A , in particular, it does not depend on B except through γ .

Proof. We will use a renewal approach, first considering the times when a jump occurs that results in both \check{X} and \check{Y} being in \mathcal{X}_1 , when they weren't both in \mathcal{X}_1 previously. As the two chains are independent, this will occur when one chain is already in \mathcal{X}_1 , and the other jumps into \mathcal{X}_1 . Let K_t denote the number of times \check{X} or \check{Y} enters \mathcal{X}_1 up to time t , when the other is already in \mathcal{X}_1 , and t_k denote the time of the k th such transition.

We can bound $H^*(\beta)$ by assuming that \check{X} and \check{Y} first meet in \mathcal{X}_1 (that is, we ignore any prior meetings in \mathcal{X}_0). If we knew $K_T = k$, then we know that after t_k , neither copy leaves \mathcal{X}_1 before their final meeting. As transitions out of \mathcal{X}_1 are independent of where in \mathcal{X}_1 the chains move, their marginal transitions

on this set will follow the rate matrix γA . We write E^{*j} for the expectation conditioned on not leaving \mathcal{X}_j . Therefore, conditioning on $\check{X}_{t_k}, \check{Y}_{t_k}$ and $K_T = k$ and rescaling time, we see that

$$E[e^{\beta(T-t_k)} | \check{X}_{t_k}, \check{Y}_{t_k}, K_T = k] = E^{*1}[e^{\beta(T-t_k)} | \check{X}_{t_k}, \check{Y}_{t_k}] \leq G^*(\gamma^{-1}\beta),$$

where $G^*(\beta)$ be the function defined in Lemma 11, that is, the supremum over starting states of the moment generating function of the first hitting time *on the basic (unsplit) state space \mathcal{X} with rate matrix A* .

Write $E_{\check{x}\check{y}}$ for the expectation conditional on $\check{X}_0 = \check{x}, \check{Y}_0 = \check{y}$, and similarly $P_{\check{x}\check{y}}$ for the conditional probability. We have the bound

$$\begin{aligned} E_{\check{x}\check{y}}[e^{\beta T}] &= E_{\check{x}\check{y}}\left[E_{\check{x}\check{y}}[e^{\beta T} | \check{X}_{t_k}, \check{Y}_{t_k}, K_T]\right] \\ &\leq E_{\check{x}\check{y}}\left[\sum_{k=0}^{\infty} E[e^{\beta T} | \check{X}_{t_k}, \check{Y}_{t_k}, K_T = k] P_{\check{x}\check{y}}(K_T = k | \check{X}_{t_k}, \check{Y}_{t_k})\right] \\ &\leq \sum_{k=0}^{\infty} E_{\check{x}\check{y}}[e^{\beta t_k} E[e^{\beta(T-t_k)} | \check{X}_{t_k}, \check{Y}_{t_k}, K_T = k] P_{\check{x}\check{y}}(K_T = k | \check{X}_{t_k}, \check{Y}_{t_k})] \\ &\leq G^*(\gamma^{-1}\beta) \sum_{k=0}^{\infty} E_{\check{x}\check{y}}[e^{\beta t_k} P_{\check{x}\check{y}}(K_T = k | \check{X}_{t_k}, \check{Y}_{t_k})] \\ &\leq G^*(\gamma^{-1}\beta) \sum_{k=0}^{\infty} (E_{\check{x}\check{y}}[e^{2\beta t_k}] E_{\check{x}\check{y}}[P_{\check{x}\check{y}}(K_T = k | \check{X}_{t_k}, \check{Y}_{t_k})^2])^{1/2} \\ &\leq G^*(\gamma^{-1}\beta) \sum_{k=0}^{\infty} E_{\check{x}\check{y}}[e^{2\beta t_k}]^{1/2} P_{\check{x}\check{y}}(K_T = k)^{1/2}. \end{aligned} \tag{10}$$

We will seek to bound the components of the above final sum.

Probability component First consider $P_{\check{x}\check{y}}(K_T = k)$. Exactly as in the proof of Lemma 10, from points $x, y \in \mathcal{X}$, let $q(x, y; n)$ denote the probability of a pair of independent basic (unsplit) chains meeting in precisely n jumps, under the rate matrix γA , given starting values $X_0 = x, Y_0 = y$. Then a geometric trials argument yields, for $\check{x} = x \times \{1\}, \check{y} = y \times \{1\} \in \mathcal{X}_1$,

$$P_{\check{x}\check{y}}(K_T = 0) = \sum_n q(x, y; n) \gamma^n.$$

By another geometric trials argument, we see that

$$\begin{aligned} P_{\check{x}\check{y}}(K_T = k) &= E_{\check{x}\check{y}}\left[\left(\sum_n q(\check{X}_{t_k}, \check{Y}_{t_k}; n) \gamma^n\right) \prod_{0 \leq i < k} \left(1 - \sum_n q(\check{X}_{t_i}, \check{Y}_{t_i}; n) \gamma^n\right)\right] \\ &\leq (1 - \gamma)(1 - q_\gamma)^k, \end{aligned} \tag{11}$$

where $q_\gamma := \inf_{x,y} (\sum_n q(x, y; n) \gamma^n) > 0$ by Lemma 10.

Expectation component Now consider $E_{\check{x}\check{y}}[e^{\beta t_k}]$. To begin, consider

$$Q(\beta) := \sup_{\check{x}\check{y}} E_{\check{x}\check{y}}[e^{\beta t_1}],$$

the mgf of the first time \check{X} and \check{Y} will both be in \mathcal{X}_1 , following at least one of them entering \mathcal{X}_0 , maximised over starting conditions. We call a jump which results in one of \check{X}, \check{Y} changing between \mathcal{X}_0 and \mathcal{X}_1 a ‘layer shift’. Note that the two chains jump at the same time with probability zero, and so the process which counts how many of them are in \mathcal{X}_0 is skip free.

Now to bound $Q(\beta)$, we need to bound two components – the time σ_1^0 at which the first of the two chains enters \mathcal{X}_0 , and the time $t_1 - \sigma_1^0$ until both chains have returned to \mathcal{X}_1 . To obtain such a bound, we must also consider the possibility that, before t_1 , both chains will enter \mathcal{X}_0 . Define the stopping times

$$\begin{aligned}\sigma_2^i &= \inf\{t > \sigma_1^{i-1} : \check{X}_t, \check{Y}_t \in \mathcal{X}_0\}, \\ \sigma_1^i &= \inf\{t > \sigma_2^i : \check{X}_t \in \mathcal{X}_1 \text{ or } \check{Y}_t \in \mathcal{X}_1\},\end{aligned}$$

so that σ_i^j refers to the first time that precisely i of the chains are in \mathcal{X}_0 , given that they have both previously entered \mathcal{X}_0 precisely j times.

Let

$$\begin{aligned}J_0(\beta) &:= \sup_{\check{x}, \check{y} \in \mathcal{X}_1} E_{\check{x}\check{y}}[e^{\beta\sigma_1^0}], \\ J_1(\beta) &:= \sup_{\check{x} \in \mathcal{X}_0, \check{y} \in \mathcal{X}_1} E_{\check{x}\check{y}}[e^{\beta\sigma_2^1} | \sigma_2^1 < t_1], \\ J_2(\beta) &:= \sup_{\check{x}, \check{y} \in \mathcal{X}_0} E_{\check{x}\check{y}}[e^{\beta\sigma_1^1}], \\ J_3(\beta) &:= \sup_{\check{x} \in \mathcal{X}_0, \check{y} \in \mathcal{X}_1} E_{\check{x}\check{y}}[e^{\beta t_1} | t_1 < \sigma_2^1],\end{aligned}$$

and note that the choice of whether \check{x} or \check{y} is in \mathcal{X}_0 or \mathcal{X}_1 is arbitrary in the definition of J_1 and J_3 . This gives J_0 as the maximal mgf of the first time one chain enters \mathcal{X}_0 ; J_1 as the maximal mgf of the first time both chains are in \mathcal{X}_0 given one of them is now, and given this occurs before they both enter \mathcal{X}_1 ; J_2 as the first time one chain enters \mathcal{X}_1 given they are both in \mathcal{X}_0 ; and J_3 as the maximal mgf of the first time both chains are in \mathcal{X}_1 given one of them is now, and given this occurs before they both enter \mathcal{X}_0 .

We can now construct another geometric trials argument. If one of \check{X}, \check{Y} is in \mathcal{X}_0 , then the next layer shift will result either in them both being in \mathcal{X}_0 or in \mathcal{X}_1 . The probability of both of them moving to \mathcal{X}_1 in the next layer shift depends on the relative rates of transition, however looking at the relative probabilities of transitions we find

$$\begin{aligned}P(\sigma_2^1 > t_1) &\geq \frac{\inf\{|b_{xx} - \gamma a_{xx}|\}\gamma}{\inf\{|b_{xx} - \gamma a_{xx}|\}\gamma + \sup\{|a_{xx}|\}(1 - \gamma)} \\ &\geq \frac{\gamma^2}{\gamma^2 + \sup\{|a_{xx}|\}(1 - \gamma)} := p\end{aligned}$$

and $p > 0$ as A is bounded. Note that p does not depend on B . Let D count the number of times both \check{X} and \check{Y} enter \mathcal{X}_0 before both returning to \mathcal{X}_1 . We have

$$\begin{aligned}\sup_{\check{x}, \check{y}} E_{\check{x}, \check{y}}[e^{\beta(t_1 - \sigma_1^0)}] &= \sup_{\check{x}, \check{y}} E_{\check{x}, \check{y}}[e^{\beta(t_1 - \sigma_1^D) + \beta \sum_{d=0}^{D-1} (\sigma_1^{d+1} - \sigma_1^d)}] \\ &= \sup_{\check{x}, \check{y}} E_{\check{x}, \check{y}} \left[E[e^{\beta(t_1 - \sigma_1^D)} | \mathcal{F}_{\sigma_1^D}, D] e^{\beta \sum_{d=0}^{D-1} (\sigma_1^{d+1} - \sigma_1^d)} \right] \quad (12) \\ &\leq J_3(\beta) \sup_{\check{x}, \check{y}} E_{\check{x}, \check{y}} \left[e^{\beta \sum_{d=0}^{D-1} (\sigma_1^{d+1} - \sigma_1^d)} \right].\end{aligned}$$

As σ_i^D is not a stopping time (due to the presence of D), we now use a renewal argument.

$$\begin{aligned}
F &:= \sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}, \tilde{y}} \left[e^{\beta \sum_{d=0}^{D-1} (\sigma_1^{d+1} - \sigma_1^d)} \right] \\
&= \sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}, \tilde{y}} \left[E[e^{\beta \sum_{d=1}^{D-1} (\sigma_1^{d+1} - \sigma_1^d)} | \mathcal{F}_{\sigma_1^1}, D > 0] I_{\{D > 0\}} e^{\beta(\sigma_1^1 - \sigma_1^0)} \right] \\
&\leq \sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}, \tilde{y}} \left[F I_{\{D > 0\}} e^{\beta(\sigma_1^1 - \sigma_1^0)} \right] \\
&= F \sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}, \tilde{y}} \left[I_{\{D > 0\}} E[e^{\beta(\sigma_1^1 - \sigma_2^1)} | \mathcal{F}_{\sigma_2^1}] e^{\beta(\sigma_2^1 - \sigma_1^0)} \right] \\
&\leq F J_2(\beta) \sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}, \tilde{y}} \left[I_{\{D > 0\}} E_{\tilde{x}, \tilde{y}} [e^{\beta(\sigma_2^1 - \sigma_1^0)} | D > 0] \right] \\
&\leq F J_2(\beta) J_1(\beta) \sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}, \tilde{y}} [I_{\{D > 0\}}] \\
&\leq F J_2(\beta) J_1(\beta) (1 - p).
\end{aligned}$$

Hence provided $J_2(\beta) J_1(\beta) (1 - p) < 1$, we have

$$\sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}, \tilde{y}} [e^{\beta(t_1 - \sigma_1^0)}] \leq J_3(\beta) F \leq \frac{J_3(\beta)}{1 - J_1(\beta) J_2(\beta) (1 - p)}. \quad (13)$$

By Lemma 6, if we can uniformly bound J_1 and J_2 , then by reducing β we have a uniform bound arbitrarily close to 1. Hence for β sufficiently small this estimate will hold, and we can combine (12) and (13) to give

$$\sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}, \tilde{y}} [e^{\beta t_1}] \leq \frac{J_0(\beta) J_3(\beta)}{1 - J_1(\beta) J_2(\beta) (1 - p)}. \quad (14)$$

Therefore, to uniformly bound the mgf of the first return time of both chains to \mathcal{X}_1 , it is enough to find uniform bounds on J_i for each i .

Bounding J_i . Let τ_i denote the time of the i th jump of the chain \tilde{X} , and for $j \in \{0, 1\}$ let

$$R_j(i, \beta) = E[e^{\beta \tau_i} | \tilde{X}_t \in \mathcal{X}_j \text{ for all } t \in [0, \tau_i]] =: E^{*j}[e^{\beta \tau_i}],$$

so $R_j(i, \cdot)$ is the mgf of the i th transition of the chain, given it has not left \mathcal{X}_j . Then for $i \geq 2$,

$$\begin{aligned}
R_j(i, \beta) &= E^{*j}[e^{\beta \tau_i}] = E^{*j}[e^{\tau_{i-1}} E^{*j}[e^{\beta \tau_i - \tau_{i-1}} | \mathcal{X}_{\tau_{i-1}}]] \\
&\leq R_j(i-1, \beta) V_j(\beta)
\end{aligned}$$

where $V_j(\beta) = \sup_{\tilde{x} \in \mathcal{X}_j} E^{*j}[e^{\beta \tau_1} | \tilde{X}_0 = \tilde{x}]$ is the maximum of the mgfs of the next transition. By definition, looking at the first jump we have $R_j(1, \beta) \leq V_j(\beta)$, and so by recursion, $R_j(i, \beta) \leq (V_j(\beta))^i$.

To bound J_0 , we can instead bound the first time that \tilde{X} jumps to \mathcal{X}_0 , ignoring any previous jumps of \tilde{Y} . By a geometric trials argument, we see that

$$J_0(\beta) = \sum_{i=1}^{\infty} \gamma (1 - \gamma)^{i-1} R_0(i, \beta) \leq \frac{\gamma V_0(\beta)}{1 - (1 - \gamma) V_0(\beta)}. \quad (15)$$

To bound J_1 , we need only bound the time until \tilde{Y} enters \mathcal{X}_0 , as the condition $\sigma_2^1 < t_1$ provides an upper bound on σ_2^1 , so the conditional expectation is less than the unconditional expectation. Hence, by the same geometric trials argument,

$$J_1(\beta) \leq \frac{\gamma V_0(\beta)}{1 - (1 - \gamma)V_0(\beta)}. \quad (16)$$

To bound J_2 , as with J_0 , we can instead bound the first time that \tilde{X} jumps to \mathcal{X}_1 , ignoring any previous jumps of \tilde{Y} . Hence

$$J_2(\beta) = \sum_{i=1}^{\infty} (1 - \gamma)\gamma^{i-1} R_0(i, \beta) \leq \frac{(1 - \gamma)V_1(\beta)}{1 - \gamma V_1(\beta)}.$$

To bound J_3 , as with J_1 , we see that the condition $t_1 < \sigma_2^1$ provides an upper bound on t_1 , and so it is enough to bound the time until \tilde{X} enters \mathcal{X}_1 . Hence

$$J_3(\beta) \leq \frac{(1 - \gamma)V_1(\beta)}{1 - \gamma V_1(\beta)}.$$

Therefore to find bounds on J_i for each i , we need only bound the right hand side of (15) and (16).

As jumps of Markov chains are exponentially distributed, and conditional on not leaving \mathcal{X}_1 the chains move following rate matrix γA , we can easily see that, as $\gamma \leq \inf_x |a_{xx}|$,

$$V_1(\beta) = \sup_x \left(1 - \frac{\beta}{\gamma |a_{xx}|}\right)^{-1} \leq \left(1 - \frac{\beta}{\gamma^2}\right)^{-1}.$$

As B is strictly controlled by A , and conditionally on not leaving \mathcal{X}_0 the chains move following rate matrix $B - \gamma A$, we also have

$$V_0(\beta) = \sup_x \left(1 - \frac{\beta}{|b_{xx} - \gamma a_{xx}|}\right)^{-1} \leq \left(1 - \frac{\beta}{\gamma}\right)^{-1}.$$

So the right hand side of (15) and (16) respectively can be bounded by

$$\frac{\gamma V_0(\beta)}{1 - (1 - \gamma)V_0(\beta)} \leq 1 + \frac{\beta}{\gamma^3 - \beta}, \quad \frac{\gamma V_0(\beta)}{1 - (1 - \gamma)V_0(\beta)} \leq 1 + \frac{\beta}{\gamma(1 - \gamma) - \beta},$$

which are uniformly bounded for $\beta < \frac{\gamma}{2}(\gamma^2 \wedge (1 - \gamma))$.

Therefore, for β sufficiently small, there is a uniform bound on (14), which depends only on γ . Therefore, by Lemma 6, for any $\epsilon > 0$, there exists $\beta > 0$ such that

$$Q(\beta) = \sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}\tilde{y}}[e^{\beta t_1}] < 1 + \epsilon, \quad (17)$$

and this β depends only on γ and A .

Putting the pieces together. Now it is easy to see that, with $t_0 := 0$,

$$E_{\tilde{x}\tilde{y}}[e^{2\beta t_k}]^{1/2} \leq E_{\tilde{x}\tilde{y}}[E[e^{2\beta \sum_{j=1}^k (t_j - t_{j-1})} | \mathcal{F}_{t_{j-1}}]]^{1/2} \leq E_{\tilde{x}\tilde{y}}[e^{2\beta t_1}]^{k/2}.$$

Hence, returning to (10), we have

$$\begin{aligned}
\sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}\tilde{y}}[e^{\beta T}] &\leq G^*(\gamma^{-1}\beta) \sum_{k=0}^{\infty} \sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}\tilde{y}}[e^{2\beta tk}]^{1/2} \sup_{\tilde{x}, \tilde{y}} P_{\tilde{x}\tilde{y}}(K_T = k)^{1/2} \\
&\leq G^*(\gamma^{-1}\beta) \sum_{k=0}^{\infty} (Q(2\beta))^{k/2} (1-\gamma)^{1/2} (1-q_\gamma)^{k/2} \\
&= G^*(\gamma^{-1}\beta) \frac{(1-\gamma)^{1/2}}{1 - \sqrt{(1-q_\gamma)Q(2\beta)}}.
\end{aligned}$$

This is bounded whenever $Q(2\beta) < (1-q_\gamma)^{-1/2}$ and $G^*(\gamma^{-1}\beta)$ is bounded. We have just shown in (17) that this can be achieved with a uniform choice of β . Hence, there exists a uniform choice of $\beta^* > 0$ such that $\sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}\tilde{y}}[e^{\beta T}] < K$ for all $\beta < \beta^*$. By Lemma 6, for any $\epsilon > 0$ we can hence find $\tilde{\beta}_{\gamma\epsilon}$ such that

$$\sup_{\tilde{x}, \tilde{y}} E_{\tilde{x}\tilde{y}}[e^{\beta T}] \leq 1 + \epsilon \quad \text{for all } \beta < \tilde{\beta}_{\gamma\epsilon}.$$

To conclude, we note that this choice of $\tilde{\beta}_{\gamma\epsilon}$ depends only on γ and A . \square

Corollary 4. *Let X, Y be two independent copies of the Markov chain on \mathcal{X} with rate matrix B . Let T be the first meeting time of these chains. Then for any $\epsilon > 0$, there exists $\beta > 0$ such that $\sup_{x, y} E[e^{\beta T} | X_0 = x, Y_0 = y] \leq 1 + \epsilon$.*

Proof. This follows from Lemma 11, as the chain X is simply the marginal chain of \tilde{X} ignoring the splitting, and similarly for Y . \square

Given this result, we can now prove Theorem 5.

Proof of Theorem 5. For any $\mu, \nu \in \mathcal{M}$, consider two copies X, Y of the chain with rate matrix B , where $X_0 \sim \mu, Y_0 \sim \nu$. Recall we denote by $P_t\mu$ the law of X_t , and $\|f\|_{TV} = \sum_x |f(x)|$ is a norm on \mathcal{M} .

Consider the coupled process

$$\tilde{Y}_t = \begin{cases} Y_t & t < T \\ X_t & t \geq T. \end{cases}$$

It is easy to see that \tilde{Y}_t and Y_t have the same law. However, this implies,

$$\|P_t\mu - P_t\nu\|_{TV} = \sum_z |P(X_t = z) - P(\tilde{Y}_t = z)| \leq P(T > t).$$

By Corollary 4 and Chernoff's inequality, for any $\epsilon > 0$ there exists $\tilde{\beta}_{\gamma\epsilon} > 0$ dependent only on γ and A such that

$$\begin{aligned}
P(T > t) &\leq E[e^{\tilde{\beta}_{\gamma\epsilon} T} | X_0 \sim \mu, Y_0 \sim \nu] e^{-\tilde{\beta}_{\gamma\epsilon} t} \\
&\leq \sup_{x, y} E[e^{\tilde{\beta}_{\gamma\epsilon} T} | X_0 = x, Y_0 = y] e^{-\tilde{\beta}_{\gamma\epsilon} t} \\
&\leq (1 + \epsilon) e^{-\tilde{\beta}_{\gamma\epsilon} t}.
\end{aligned}$$

Therefore,

$$\|P_t\mu - P_t\nu\|_{TV} \leq (1 + \epsilon) e^{-\tilde{\beta}_{\gamma\epsilon} t}. \quad (18)$$

Replacing ν with $P_{s-t}\mu$ in (11), we see that for any $s > t$,

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

so $P_t\mu$ is a Cauchy sequence in \mathcal{M} . Hence it has a limit, $P_t\mu \rightarrow \pi$, which is invariant under P_t . Taking $\nu = \pi$ in (18), we then see

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

and so the desired convergence can be guaranteed, is independent of μ , and occurs uniformly in B given γ and A . \square

4 Ergodic BSDEs

As we now have estimates for the uniform ergodicity of our Markov chain, we shall use these to prove the existence of solutions to Ergodic BSDEs. The techniques used here are modifications of those in [10] and [12]. For ease of reference, we collect the various assumptions needed in one place.

Assumption 1. *Let X be a uniformly ergodic, time homogeneous Markov chain with bounded transition rates. Let $f : \mathcal{X} \times \mathbb{R}^N \rightarrow \mathbb{R}$ be a function satisfying the conditions of Theorem 3 (uniformly Lipschitz in the second component under $\|\cdot\|_{M_t}$, balanced and $\psi(\cdot, 0)$ uniformly bounded) which is also strictly balanced.*

The following lemma allows us to simply connect our analysis of uniform ergodicity to the theory of BSDEs, by means of Girsanov transformations.

Lemma 12. *Let $Z, Z' \in \mathbb{R}^N$ be any two vectors, defined up to equivalence \sim_M , and let f be our strictly balanced function. Let \mathbb{P}^x denote the measure under which X is a Markov chain with rate matrix A and initial state x . For each $T > 0$, define the measure with density*

$$\frac{d\mathbb{Q}^{x,T}}{d\mathbb{P}^x} = \mathcal{E}\left(\int_{]0, \cdot \wedge T]} \frac{f(X_{u-}, Z) - f(X_{u-}, Z')}{\|Z - Z'\|_{M_u}^2} (Z - Z')^* dM_u\right).$$

Then, on the interval $[0, T]$, X is a time-homogeneous Markov chain under $\mathbb{Q}^{x,T}$, and the corresponding rate matrix B does not depend on T . Furthermore, B is such that the Markov chain with rate matrix B is uniformly ergodic, with constants R and ρ which do not depend on Z or Z' .

Proof. By Lemma 1, $\mathbb{Q}^{x,T}$ is a probability measure. To show that X is a time homogeneous Markov chain under $\mathbb{Q}^{x,T}$, it is enough to show that an equation of the form of (1) must hold, that is, for $t < T$ we have a rate matrix B such that

$$X_t = X_0 + \int_{]0, t]} BX_{u-} du + \mathbb{Q}^{x,T}\text{-martingale} \quad (19)$$

However, we know that for $t < T$,

$$\begin{aligned} X_t &= X_0 + \int_{]0, t]} AX_{u-} du + M_t \\ &= X_0 + \int_{]0, t]} AX_{u-} du + \int_{]0, t \wedge T]} \left(\frac{f(X_{u-}, Z) - f(X_{u-}, Z')}{\|Z - Z'\|_{M_u}^2} \psi_u(Z - Z') \right) du + \tilde{M}_t^T \end{aligned}$$

and so, by uniqueness of the Doob–Meyer decomposition, we can uniquely define a matrix B which satisfies (19), through the equation

$$B\tilde{x} = A\tilde{x} + \frac{f(\tilde{x}, Z) - f(\tilde{x}, Z')}{\|Z - Z'\|_{M_u}^2} \psi_u^{\tilde{x}}(Z - Z') \quad \text{for } \tilde{x} \in \mathcal{X}$$

(recalling that $\tilde{x} \in \mathcal{X}$ is a basis vector of \mathbb{R}^N). We know that

$$\mathbf{1}^* B\tilde{x} = \mathbf{1}^* A\tilde{x} + \frac{f(\tilde{x}, Z) - f(\tilde{x}, Z')}{\|Z - Z'\|_{M_u}^2} \mathbf{1}^* \psi_u^{\tilde{x}}(Z - Z') = 0 + 0,$$

so to check that B is a valid rate matrix, we need to check that the only negative element can occur on the diagonal. As f is balanced, we know that for any i with $e_i^* A\tilde{x} > 0$,

$$\frac{f(\tilde{x}, Z) - f(\tilde{x}, Z')}{\|Z - Z'\|_{M_u}^2} ((Z - Z')^*(e_i - \tilde{x})) > -1$$

as Z and Z' are only defined up to equivalence \sim_M , hence at most up to the addition of a constant, we can assume without loss of generality that $(Z - Z')^* e_i = 0$, and so

$$\frac{f(\tilde{x}, Z) - f(\tilde{x}, Z')}{\|Z - Z'\|_{M_u}^2} (\tilde{x}^*(Z - Z')) < 1.$$

On the other hand, for this choice of representative $Z - Z'$, we then have

$$e_i^* \psi_u^{\tilde{x}}(Z - Z') = -(e_i^* A\tilde{x})(\tilde{x}^*(Z - Z'))$$

and so,

$$e_i^* B\tilde{x} = (e_i^* A\tilde{x}) \left(1 - \frac{f(\tilde{x}, Z) - f(\tilde{x}, Z')}{\|Z - Z'\|_{M_u}^2} (\tilde{x}^*(Z - Z')) \right) > 0.$$

The fact that B and A also have the same pattern of zeros can be checked in a similar way, as f is balanced. Therefore, X is a Markov chain under $\mathbb{Q}^{x, T}$ with rate matrix B , for all T .

As we also know f is strictly balanced, we have exactly

$$1 - \frac{f(\tilde{x}, Z) - f(\tilde{x}, Z')}{\|Z - Z'\|_{M_u}^2} (\tilde{x}^*(Z - Z')) \geq \gamma$$

for some $\gamma > 0$, and hence $e_i^*(B - \gamma A)\tilde{x} \geq 0$ for all $e_i \neq \tilde{x}$. Hence (as the columns all sum to zero) $B - \gamma A$ is also a rate matrix, and hence B is controlled by A . We also have

$$\tilde{x}^* B\tilde{x} = - \sum_{i: e_i \neq \tilde{x}} e_i^* B X_{u-} \leq -\gamma \sum_{i: e_i \neq \tilde{x}} e_i^* A\tilde{x} = \gamma \tilde{x}^* A\tilde{x},$$

so

$$(\tilde{x}^* B\tilde{x}) - \frac{\gamma}{2} (\tilde{x}^* A\tilde{x}) \leq \frac{\gamma}{2} (\tilde{x}^* A\tilde{x}).$$

As X is uniformly ergodic under the initial measure, $\tilde{x}^* A\tilde{x}$ is bounded away from zero, by Lemma 7. Therefore, for $\gamma^* < \gamma \left(\frac{\inf_{\tilde{x}} |\tilde{x}^* A\tilde{x}| \wedge 1}{2} \right)$, we see that B is strictly controlled by A with constant γ^* , independent of Z and Z' .

As we know that B is strictly controlled by A with a fixed constant γ^* , we can apply Theorem 5 to show that there exist constants R, ρ , independent of Z and Z' , such that X is uniformly ergodic under $\mathbb{Q}^{x,T}$ with constants R and ρ . \square

We now show that, for some sequence α_n , the solutions to the discounted BSDEs converge in an appropriate sense.

Lemma 13. *Let (Y^α, Z^α) be the unique bounded adapted solution to the discounted BSDE*

$$Y_T^\alpha = Y_t^\alpha - \int_{]t,T]} (-\alpha Y_{u-}^\alpha + f(X_{u-}, Z_u^\alpha)) du + \int_{]t,T]} (Z_u^\alpha)^* dM_u, \quad t \leq T < \infty,$$

and let $x_0 \in \mathcal{X}$ be an arbitrary state. By Corollary 2, we know there exists a function $v^\alpha : \mathcal{X} \rightarrow \mathbb{R}$ such that $Y_t^\alpha = v^\alpha(X_t)$. Then there exists a bound $C' < \infty$ such that

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

uniformly in x and α . Hence there exists a sequence $\alpha_n \rightarrow 0$ 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 bounded function $v : \mathcal{X} \rightarrow \mathbb{R}$ and some $\lambda \in \mathbb{R}$.

Proof. Let C be a bound on $|f(\cdot, 0)|$. Then we see from Theorem 3 that $|v^\alpha(\cdot)| \leq C/\alpha$. We first need to convert this into a bound which is uniform in α . To do so, we will view the process Y^α under a perturbed measure, and use our ergodicity result to control its growth.

We first note that, by Corollary 2, the solution process Z^α is constant, up to equivalence \sim_M .

For each $T > 0$, define the measure with density

$$\frac{d\mathbb{Q}^{x,\alpha,T}}{d\mathbb{P}^x} = \mathcal{E} \left(\int_{]0,\wedge T]} \frac{f(X_{u-}, Z_u^\alpha) - f(X_{u-}, 0)}{\|Z_u^\alpha\|_{M_u}^2} (Z_u^\alpha)^* dM_u \right).$$

By Lemma 12, under this measure X is still a uniformly ergodic Markov chain, with constants R, ρ which do not depend on α (as α only affects Z^α).

Using Girsanov's theorem in the form of Lemma 1, we can verify that for any x, α, T ,

$$v^\alpha(x) = E_{\mathbb{Q}^{x,\alpha,T}} \left[e^{-\alpha T} v^\alpha(X_T) + \int_{]0,T]} e^{-\alpha u} f(X_{u-}, 0) du \right].$$

As $|v(x)| \leq C/\alpha$, letting $T \rightarrow \infty$ we obtain

$$v^\alpha(x) = \lim_{T \rightarrow \infty} E_{\mathbb{Q}^{x,\alpha,T}} \left[\int_{]0,T]} e^{-\alpha u} f(X_{u-}, 0) du \right].$$

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

$$\begin{aligned}
|v^\alpha(x) - v^\alpha(x')| &= \left| \lim_{T \rightarrow \infty} E_{\mathbb{Q}^{x, \alpha, T}} \left[\int_{]0, T]} e^{-\alpha u} f(X_{u-}, 0) du \right] \right. \\
&\quad \left. - \lim_{T \rightarrow \infty} E_{\mathbb{Q}^{x', \alpha, T}} \left[\int_{]0, T]} e^{-\alpha u} f(X_{u-}, 0) du \right] \right| \\
&= \left| \lim_{T \rightarrow \infty} \int_{]0, T]} e^{-\alpha u} \int_{\mathcal{X}} f(X_{u-}, 0) (dP_u \delta_x - dP_u \delta_{x'}) du \right| \\
&\leq C \left| \lim_{T \rightarrow \infty} \int_{]0, T]} e^{-\alpha u} \|P_u \delta_x - P_u \delta_{x'}\|_{TV} du \right| \\
&\leq CR \left| \lim_{T \rightarrow \infty} \int_{]0, T]} e^{-\alpha u} e^{-\rho u} du \right| \\
&= CR(\alpha + \rho)^{-1}
\end{aligned}$$

And so we have a uniform bound $|v^\alpha(x) - v^\alpha(x')| \leq CR\rho^{-1}$. In addition, we recall that $|\alpha v^\alpha(x)| \leq C$.

Given this, we can use a diagonal procedure to construct a sequence $\alpha_n \downarrow 0$ such that

$$\begin{aligned}
\alpha_n v^{\alpha_n}(x_0) &\rightarrow \lambda \\
v^{\alpha_n}(x) - v^{\alpha_n}(x_0) &\rightarrow v(x) \quad \text{for all } x \in \mathcal{X}
\end{aligned}$$

for a function $v : \mathcal{X} \rightarrow \mathbb{R}$ (recall that \mathcal{X} is at most countable). Note that, for $C' = CR\rho^{-1}$, we have the bounds $|\lambda| \leq C'$ and $|v(x)| \leq C'$.

Finally, we notice that, for an arbitrary $x \in \mathcal{X}$,

$$\alpha_n v^{\alpha_n}(x) = \alpha_n v^{\alpha_n}(x_0) + \alpha_n (v^{\alpha_n}(x) - v^{\alpha_n}(x_0)) \rightarrow \lambda + 0$$

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

Theorem 6. *Let v, λ be as constructed in Lemma 13. The triple (Y, Z, λ) , where*

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

is the unique bounded Markovian solution, with $v(x_0) = 0$, to the Ergodic BSDE

$$Y_t = Y_T + \int_{]t, T]} [f(X_{u-}, Z_u) - \lambda] du + \int_{]t, T]} Z_u^* dM_u.$$

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

Proof. Let α_n be the sequence constructed in Lemma 13. As Z^{α_n} is only defined up to constant shifts, we know that $Y_t^{\alpha_n} := v^{\alpha_n}(X_t)$, $e_i^* Z_t^{\alpha_n} := v^{\alpha_n}(e_i) - v^{\alpha_n}(x_0)$ solves the discounted BSDE

$$Y_t^{\alpha_n} = Y_T^{\alpha_n} + \int_{]t, T]} (f(X_{u-}, Z_u^{\alpha_n}) - \alpha_n Y_{u-}^{\alpha_n}) du + \int_{]t, T]} (Z_u^{\alpha_n})^* dM_u.$$

Note that $Z_t^{\alpha_n}$ is constant in t , and by the bound established in Lemma 13, $|e_i^* Z_t^{\alpha_n}|$ is uniformly bounded. Hence, as the transition rates in A are bounded,

we also know $\|Z^{\alpha_n}\|_{M_t}^2$ is uniformly bounded, and dominated convergence yields $\lim(\int Z^{\alpha_n} dM) = \int(\lim Z^{\alpha_n})dM$. Therefore

$$\begin{aligned}
v(X_t) &= \lim_n (v^{\alpha_n}(X_t) - v^{\alpha_n}(x_0)) \\
&= \lim_n (v^{\alpha_n}(X_T) - v^{\alpha_n}(x_0)) + \lim_n \int_{]t, T]} (f(X_{u-}, Z^{\alpha_n}) - \alpha_n v^{\alpha_n}(X_{u-})) du \\
&\quad + \lim_n \int_{]t, T]} (Z^{\alpha_n})^* dM_u \\
&= v(X_T) + \int_{]t, T]} \lim_n (f(X_{u-}, Z^{\alpha_n}) - v^{\alpha_n}(x_0)) - \lim_n \alpha_n v^{\alpha_n}(X_{u-}) du \\
&\quad + \int_{]t, T]} \lim_n (Z^{\alpha_n} - v^{\alpha_n}(x_0))^* dM_u \\
&= v(X_T) + \int_{]t, T]} (f(X_{u-}, Z) - \lambda) du + \int_{]t, T]} Z^* dM_u,
\end{aligned}$$

that is, $(v(X_t), v(\cdot), \lambda)$ is a solution to the EBSDE.

To show the solution is unique, suppose (Y', Z', λ') is another bounded solution. Let $\tilde{Y} = Y - Y'$, $\tilde{Z} = Z - Z'$ and $\tilde{\lambda} = \lambda - \lambda'$. Then

$$\tilde{Y}_t = \tilde{Y}_T + \int_{]t, T]} [f(X_{u-}, Z) - f(X_{u-}, Z') - \tilde{\lambda}] du + \int_{]t, T]} \tilde{Z}_u^* dM_u$$

and defining the measure \mathbb{Q}_1^T as in (7)

$$\frac{d\mathbb{Q}_1^T}{d\mathbb{P}^x} = \mathcal{E} \left(\int_{]0, \wedge T]} \frac{f(X_{u-}, Z) - f(X_{u-}, Z')}{\|\tilde{Z}\|_{M_u}^2} (\tilde{Z}_u)^* dM_u \right),$$

we see that

$$\tilde{\lambda} = T^{-1} E_{\mathbb{Q}_1^T} [\tilde{Y}_T - \tilde{Y}_0].$$

As \tilde{Y} is uniformly bounded, taking $T \rightarrow \infty$ we see $\tilde{\lambda} = 0$, that is, $\lambda = \lambda'$. Substituting back, we see that $\tilde{Y}_0 = E_{\mathbb{Q}_1^T} [\tilde{Y}_T]$ for all T .

From Lemma 12, if Y' is a bounded Markovian solution with $Y'_t = v'(X_t)$, we see that X is still a time homogenous uniformly ergodic Markov chain under \mathbb{Q}_1^T , on the interval $[0, T]$. The family of measures \mathbb{Q}_1^T is then consistent, in that $\mathbb{Q}^T|_{\mathcal{F}_T} = \mathbb{Q}^{T'}|_{\mathcal{F}_T}$ for all $T \leq T'$. Therefore, conditioning on $X_0 = x$, we can extend them to a single measure \mathbb{Q}^x under which X is a time homogenous uniformly ergodic Markov chain for all times, and $\mathbb{Q}^x|_{\mathcal{F}_T} = \mathbb{Q}_1^T|_{\mathcal{F}_T, X_0=x}$ for all T . If $\tilde{\pi}$ is the ergodic measure of X under \mathbb{Q}^x , we then have

$$v(x) - v'(x) = \tilde{Y}_0 = \lim_T E_{\mathbb{Q}^x} [\tilde{Y}_T] = \int_{\mathcal{X}} (v(y) - v'(y)) d\tilde{\pi}(y),$$

and we see that the right hand side is independent of x , that is, $v'(x) = v(x) + c$ for some constant $c \in \mathbb{R}$, for all $x \in \mathcal{X}$. The equivalence $\|Z - Z'\|_{M_t}^2 = 0$ then follows from Corollary 2. \square

Corollary 5. *The sequences $\alpha_n v^{\alpha_n}$ and $v^{\alpha_n}(x) - v^{\alpha_n}(x_0)$ constructed in Lemma 13 converge for any choice of sequence $\{\alpha_n \downarrow 0\}$.*

Proof. Suppose this were not the case. Then we could construct two sequences with distinct limits, both of which would yield bounded Markovian solutions to the EBSDE. However we have shown that there is only one such solution, which is a contradiction. \square

Corollary 6. *The value λ in the EBSDE solution $(Y, Z, \lambda) = (v(X_t), v, \lambda)$ satisfies*

$$\lambda = \int_{\mathcal{X}} f(y, v) d\pi(y),$$

and hence, if we define a signed measure μ^x on \mathcal{X} by $\mu^x(y) = \int_{]0, \infty[} (\mathbb{P}_t^x(y) - \pi(y)) dt$, (which is well defined due to the uniform ergodicity of X under \mathbb{P}) we have

$$v(x) = c + \int_{\mathcal{X}} f(y, v) d\mu^x(y)$$

for some $c = \int_{\mathcal{X}} v(x) d\pi \in \mathbb{R}$.

Proof. The invariance of the ergodic distribution π implies that, for any fixed time T , any $g : \mathcal{X} \rightarrow \mathbb{R}$,

$$\int_{\mathcal{X}} E_{\mathbb{P}^x} [g(X_T)] d\pi(x) = \int_{\mathcal{X}} E_{\mathbb{P}^x} [g(X_{T-})] d\pi(x) = \int_{\mathcal{X}} g(x) d\pi(x).$$

We write

$$v(x) = E_{\mathbb{P}^x} \left[v(X_T) + \int_{]0, T]} (f(X_{u-}, v) - \lambda) du \right].$$

so by the stated invariance property,

$$\begin{aligned} \int_{\mathcal{X}} v(x) d\pi(x) &= \int_{\mathcal{X}} E_{\mathbb{P}^x} [v(X_T)] d\pi(x) + \int_{\mathcal{X}} E_{\mathbb{P}^x} \left[\int_{]t, T]} (f(X_{u-}, v) - \lambda) du \right] d\pi(x) \\ &= \int_{\mathcal{X}} v(x) d\pi(x) + \int_{]t, T]} \left[\int_{\mathcal{X}} E_{\mathbb{P}^x} [f(X_{u-}, v)] d\pi(x) \right] du - T\lambda \\ &= \int_{\mathcal{X}} v(x) d\pi(x) + \int_{]t, T]} \left[\int_{\mathcal{X}} f(x, v) d\pi(x) \right] du - T\lambda \\ &= \int_{\mathcal{X}} v(x) d\pi(x) + T \int_{\mathcal{X}} f(x, v) d\pi(x) - T\lambda \end{aligned}$$

and rearrangement yields the first result.

Consequently, we also have

$$\begin{aligned} v(x) &= \lim_{T \rightarrow \infty} E_{\mathbb{P}^x} \left[v(X_T) + \int_{]0, T]} (f(X_{u-}, v) - \lambda) du \right] \\ &= \int_{\mathcal{X}} v(y) d\pi(y) + \lim_{T \rightarrow \infty} \int_{]0, T]} \left(\int_{\mathcal{X}} f(y, v) d\mathbb{P}_u^x(y) - \int_{\mathcal{X}} f(y, v) d\pi(y) \right) du \\ &= \int_{\mathcal{X}} v(y) d\pi(y) + \int_{]0, \infty[} \left(\int_{\mathcal{X}} f(y, v) d(\mathbb{P}_u^x(y) - \pi(y)) \right) du \\ &= \int_{\mathcal{X}} v(y) d\pi(y) + \int_{\mathcal{X}} f(y, v) d\mu^x(y). \end{aligned}$$

\square

Remark 3. If the Lipschitz constant of f were sufficiently small, when compared with the coefficients in the ergodicity of X , the above corollary would provide a simple direct proof of existence for EBSDEs. Furthermore, while we do not have a comparison theorem for the Y component of BSDEs, from this result, if c is fixed, we can differentiate $v(x)$ in terms of $f(y, v(\cdot))$, and (again given a small enough Lipschitz constant) solve the resultant implicit equation. From this analysis, we can see that, if f is sufficiently small, then $v(x)$ is monotone increasing in $f(x, \cdot)$.

Remark 4. This representation of v in fact provides an intuitive meaning for v . We see here that, up to the addition of a constant, $v(x)$ is given by the cost f , integrated through time with respect to the ‘deviation from ergodicity’ measure μ^x . Therefore, it is natural to think of v as giving a ‘short-run additional expected cost’ term, while λ gives the long-run ergodic cost.

We now give a variant of the comparison theorem for EBSDEs, considering the λ component of the solution.

Theorem 7. *Let f and f' be two strictly balanced drivers, and (Y, Z, λ) and (Y', Z', λ') the corresponding EBSDE solutions. Then if $f(x, z) \geq f'(x, z)$ for all $x \in \mathcal{X}$, $Z \in \mathbb{R}^N$, then $\lambda \geq \lambda'$.*

Proof. Take the measure with density

$$\frac{d\mathbb{Q}^{x,T}}{d\mathbb{P}^x} = \mathcal{E}\left(\int_{]0, \cdot \wedge T]} \frac{f'(X_{u-}, Z) - f'(X_{u-}, Z')}{\|Z - Z'\|_{M_u}^2} (Z - Z')^* dM_u\right),$$

and the corresponding rate matrix B and resultant ergodic measure π_B . Then, if $Y_t = v(X_t)$ and $Y'_t = v'(X_t)$, integrating the difference of the EBSDEs yields

$$\begin{aligned} & \int v(x) - v'(x) d\pi_B \\ &= \int_{\mathcal{X}} E_{\mathbb{Q}^{x,T}} \left[Y_T - Y'_T + \int_{]0, T]} (f(X_{u-}, Z) - f'(X_{u-}, Z)) dt - T(\lambda - \lambda') \right] d\pi_B \\ &= \int_{\mathcal{X}} (v(x) - v'(x)) d\pi_B + \int_{]0, T]} \left(\int_{\mathcal{X}} f(x, Z) - f'(x, Z) d\pi_B \right) dt - T(\lambda - \lambda') \end{aligned}$$

and so

$$(\lambda - \lambda') = \int_{\mathcal{X}} (f(x, Z) - f'(x, Z)) d\pi_B \geq 0.$$

□

Remark 5. We note that we have also shown that, given that f is strictly balanced and A is the generator of a uniformly ergodic chain, the vector equation

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

admits a unique solution $(v(\cdot), \lambda) \in \mathbb{R}^N \times \mathbb{R}$, and that this solution is approximated by the equations considered in Corollary 2. This is without any further reference to the monotonicity of f . This can be compared with the classical equation for ergodic cost, which is simply

$$0 = f(\cdot) + A^*v - \lambda \mathbf{1}.$$

Remark 6. From a numerical perspective, as we expect that $f(\cdot, v) + A^*v$ will typically be ‘approximately monotone’ in v , solving this equation iteratively, for example through the equations

$$\begin{aligned}\lambda_n &= N^{-1}\mathbf{1}^*\left(f(\cdot, v_n) + A^*v_n\right) \\ v_{n+1} &= (A^*)^+\left(\lambda_n\mathbf{1} - f(\cdot, v_n)\right)\end{aligned}$$

where $(A^*)^+$ denotes the Moore-Penrose pseudoinverse of A^* , will often provide a simple numerical scheme in the finite state case.

In the infinite state case, a Monte-Carlo approach can be used to approximate the equations given by Corollary 6, however the convergence of such a scheme lies beyond the scope of this paper.

5 Applications and Examples

5.1 Rate uncertainty

As a first example, we consider the rate matrix A constructed (9), and the driver

$$f(x, v) = I_{x \in \zeta} + \min_{r \in [1/\beta, \beta]} \{(r-1)v^*Ax\},$$

where $\zeta \subseteq \mathcal{X}$ and $\beta > 1$. Essentially, this driver attempts to determine the ergodic probability of being in the set ζ , however it introduces an uncertainty about the overall transition rate of the chain, by scaling the rates up or down with the parameter β , so as to minimise the probability of being in ζ . A related BSDE over finite horizon, for a different choice of A , was considered in [8].

For $\beta = 2$, we solve the EBSDE numerically using the simple algorithm suggested in Remark 6. We list all possibilities for ζ up to symmetries, and choose v such that its values in each state sum to zero (or, equivalently in this case, its ergodic mean value is zero). For comparison, we also list the ergodic probability associated with the set ζ .

ζ	$v(e_1)$	$v(e_2)$	$v(e_3)$	$v(e_4)$	λ	$\pi(\zeta)$
\emptyset	0	0	0	0	0	0
$\{e_1\}$	0.1207	-0.0172	-0.0517	-0.0517	0.0345	0.125
$\{e_2\}$	-0.0652	0.1087	-0.0217	-0.0217	0.1304	0.375
$\{e_1, e_2\}$	0.1000	0.1000	-0.1000	-0.1000	0.2000	0.500
$\{e_1, e_3\}$	0.1500	-0.0500	0.0500	-0.1500	0.2000	0.500
$\{e_1, e_4\}$	0.0769	-0.0769	-0.0769	0.0769	0.0769	0.250
$\{e_2, e_3\}$	-0.1429	0.1429	0.1429	-0.1429	0.4286	0.750
$\{e_1, e_2, e_3\}$	0.1364	0.1364	0.0455	-0.3182	0.6364	0.875
$\{e_1, e_2, e_4\}$	0.0294	0.0294	-0.1471	0.0882	0.2941	0.625
$\{e_1, e_2, e_3, e_4\}$	0	0	0	0	1	1

Note that, in the classical case without the rate uncertainty term, we would have $\lambda_\zeta = \pi(\zeta)$. This is clearly not the case here, due to the nonlinearity introduced by the rate uncertainty. In fact, we can see here that λ defines, in some sense, an ‘ergodic capacity’ for the chain (however one can verify that the EBSDE solution is generally not given by the Choquet integral with respect to

this capacity). Note further that while the ergodic probability does not reveal the asymmetry between the states of the chain, the EBSDE is affected by this in a nontrivial way.

5.2 Classical Optimal Ergodic Control

We now give a more abstract example, indicating how an ergodic control problem can be seen in this framework. Problems of this sort have been extensively considered, see for example the classical paper of Kakumanu [15], or more recent work by Guo and Hernández-Lerma [14].

Consider the problem of minimizing

$$J(x, U) = \limsup_{T \rightarrow \infty} T^{-1} E_x^U \left[\int_{]0, T]} L(X_{u-}, U_t) dt \right]$$

where

- \mathcal{U} is a space of controls, which is a separable metric space,
- U is a \mathcal{U} valued predictable process,
- $L : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ is a bounded measurable cost function,
- E_x^U is the expectation under which at time t , for the path ω , X jumps from state e_i to state e_j at a rate $e_j^* A^{U_t(\omega)} e_i$, for some measurable matrix valued function $A^{(\cdot)} : \mathcal{U} \rightarrow$ rate matrices, and $X_0 = x$,
- for some $\gamma > 0$, for all $u \in \mathcal{U}$, the matrices A^u are uniformly bounded and strictly dominated by A with constant γ (in the sense of Definition 4), for some reference rate matrix A under which X is a uniformly ergodic Markov chain.

We shall write E for the expectation under which X is a uniformly ergodic Markov chain with rate matrix A .

We define the Hamiltonian

$$f(x, z) = \inf_{u \in \mathcal{U}} \{L(x, u) + z^*(A^u - A)x\}. \quad (20)$$

We notice that $f(\cdot, 0)$ is bounded, and by the assumption that the A^u are strictly controlled by A , it is easy to show that f is a Lipschitz function in z under the $\|\cdot\|_{M_t}$ norm, and that f is strictly balanced (by testing with z equal to each basis vector). Therefore, the EBSDE with driver $f(x, z)$ admits a unique bounded Markovian solution $(\bar{Y}_t, \bar{Z}_t, \bar{\lambda}) = (v(X_t), v, \bar{\lambda})$, where v is the vector with components $v(e_i)$.

If this infimum is attained, then there exists (assuming the continuum hypothesis, by McShane and Warfield [17]) a measurable function $\kappa : \mathcal{X} \times \mathbb{R}^N \rightarrow \mathcal{U}$ such that

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

We then have the following theorem.

Theorem 8. *In the setting described above, let (Y, Z, λ) be any (possibly non-Markovian) bounded solution to the EBSDE (2) with driver f . Then the following hold.*

(i) For an arbitrary control U we have $J(x, U) \geq \lambda = \bar{\lambda}$, and equality holds if and only if

$$L(X_{t-}, U_t) + Z_t^*(A^{U_t} - A)X_{t-} = f(X_{t-}, Z_t) \quad d\mathbb{P} \times dt - a.e.$$

(ii) If the infimum is attained in (20), then the control $\bar{U}_t = \kappa(X_t, Z_t)$ verifies $J(x, \bar{U}) = \bar{\lambda}$.

In particular, for the bounded Markovian solution to the EBSDE, in terms of the vector v the following hold.

(iii) For arbitrary controls U we have $J(x, U) = \bar{\lambda}$ if and only if

$$L(X_{t-}, U_t) + v^*(A^{U_t} - A)X_{t-} = f(X_{t-}, v) \quad d\mathbb{P} \times dt - a.e.$$

(iv) If the infimum is attained in (20), then the control $\bar{U}_t = \kappa(X_{t-}, v)$ verifies $J(x, \bar{U}) = \bar{\lambda}$, that is, we have an optimal feedback control.

Proof. That $\lambda = \bar{\lambda}$ is a consequence of Theorem 6, and points (iii) and (iv) follow directly from (i) and (ii).

To show (i), for $T > 0$ define the measures,

$$\frac{d\mathbb{Q}^{x,T}}{d\mathbb{P}^x} = \mathcal{E} \left(\int_{]0, \cdot \wedge T]} \frac{Z_s^*(A^{U_s} - A)X_{s-}}{\|Z_s\|_{M_s}^2} Z_s^* dM_s \right).$$

As $(Y, Z, \bar{\lambda})$ is a solution to the EBSDE, we have

$$\begin{aligned} \bar{\lambda} &= \frac{1}{T} E_{\mathbb{Q}^{x,T}} [Y_T - Y_0] + \frac{1}{T} E_{\mathbb{Q}^{x,T}} \left[\int_{]0, T]} L(X_{s-}, U_s) ds \right] \\ &\quad + \frac{1}{T} E_{\mathbb{Q}^{x,T}} \left[\int_{]0, T]} f(X_{s-}, Z_s) - Z_s^*(A^{U_s} - A)X_{s-} - L(X_{s-}, U_s) ds \right], \end{aligned}$$

and hence, as f is an infimum over the controls,

$$\bar{\lambda} + \frac{1}{T} E_{\mathbb{Q}^{x,T}} [Y_0 - Y_T] \leq \frac{1}{T} E_{\mathbb{Q}^{x,T}} \left[\int_{]0, T]} L(X_{s-}, U_s) ds \right].$$

As $Y_0 - Y_T$ is uniformly bounded, taking a limit we see that

$$\bar{\lambda} \leq \limsup_{T \rightarrow \infty} \frac{1}{T} E_{\mathbb{Q}^{x,T}} \left[\int_{]0, T]} L(X_{s-}, U_s) ds \right] = J(x, U).$$

Conversely, if the infimum is attained, we have $f(X_{s-}, Z_s) = L(X_{s-}, \bar{U}_s) + Z_s^*(A^{\bar{U}_s} - A)X_{s-}$ for some $\bar{U}_s = \kappa(X_{s-}, Z_s)$. Then equality holds throughout and we see that

$$\bar{\lambda} = \limsup_{T \rightarrow \infty} \frac{1}{T} E_{\mathbb{Q}^{x,T}} \left[\int_{]0, T]} L(X_{s-}, \bar{U}_s) ds \right] = J(x, \bar{U}).$$

□

Remark 7. From this analysis, we also see that one could equally define the cost functional $J(x, U)$ as the liminf, rather than the limsup, and the same conclusions would hold, with the same optimal cost $\bar{\lambda}$. This is simply because the function under consideration converges for an optimal policy, as given by the EBSDE solution.

5.3 Risk averse control

The strength of the EBSDE approach is, however, not made manifest by the classical linear setting. Through the use of EBSDEs, one can happily consider nonlinear examples, in particular, when the expectation E_x^U is replaced by a dynamically consistent nonlinear expectation, in the sense of Peng [20].

For example, one could consider a form of risk averse ergodic control under uncertainty. In this case, we are in a similar setting to above, but the control U_s does not yield a unique rate matrix A^{U_s} , rather a family of matrices $\{A^{U_s, w}\}_{w \in W_{U_s}}$, where W_{U_s} is an index set, and the $A^{U_s, w}$ are still all uniformly strictly controlled by a single rate matrix A , as in the previous section.

To ensure dynamic consistency, we assume that W_u is such that if w and w' are both in W_{U_s} , then there is an element $w'' \in W_{U_s}$ such that the matrix obtained by interchanging arbitrary columns of $A^{u, w}$ and $A^{u, w'}$ is equal to $A^{u, w''}$.

In this case, we could attempt to minimise the maximum over the relevant measures (or some other nonlinear functional) corresponding to the value functional

$$J(x, u) = \limsup_{T \rightarrow \infty} T^{-1} \sup_{w \in W_{U_s}} E_x^{U, w} \left[\int_{]0, T]} L(X_t, U_t) dt \right],$$

(or equivalently with a \liminf). This can be directly treated by our setting, simply by taking the non-concave Hamiltonian

$$f(u; x, z) = \inf_{u \in \mathcal{U}} \left\{ L(x, u) + \sup_{w \in W_u} \{z^*(A^{u, w} - A)x\} \right\}$$

and solving the corresponding nonlinear EBSDE. As above, one can verify that $J(x, u)$ is then the λ component of the solution to the EBSDE with driver $f(u, \cdot, \cdot)$. Extending Theorem 8 to this setting is then a straightforward task.

Remark 8. In [3], conditions are given under which a dynamically consistent nonlinear expectation in a general setting can be represented by means of a BSDE with a balanced driver. Suppose each choice of policy yielded a nonlinear expectation. For each policy which makes X a homogeneous Markov process (in the sense that the conditional nonlinear expectation of $\phi(X_t)$ given \mathcal{F}_s is a function of X_s and $t - s$, for any measurable ϕ and any $s \leq t$), the drivers of these BSDEs can be written in the form $g_u : \mathcal{X} \times \mathbb{R}^N \rightarrow \mathbb{R}$, and assume that these drivers are strictly balanced. Given this representation, the Hamiltonian for our cost minimization problem has the generic form

$$f(u; x, z) = \inf_{u \in \mathcal{U}} \{L(x, u) + g_u(x, z)\}.$$

6 Conclusions

We have seen that, for uniformly ergodic Markov chains, EBSDEs with strictly balanced drivers admit unique bounded Markovian solutions. The methods used to determine this result, while based on previous work on EBSDEs, require a different approach to the ergodicity of the underlying process, due to the presence of jumps.

In deriving this, we have constructed a partial ordering of the rate matrices, and shown that any chain bounded below by a uniformly ergodic chain

must also be uniformly ergodic, and that the rate coefficients can be uniformly bounded. In some sense, this result is similar to recent work by Galtchouk and Pergamenshchikov [13], who study geometric ergodicity properties of general Markov processes, under an assumption of a uniform Lyapunov function. On the other hand, our method is better suited to the study of EBSDEs, where the perturbation of the rate matrix arises directly from the driver of the BSDE. Future work may allow a weakening of our assumption of uniform ergodicity to a form of geometric ergodicity, however we expect that this will require some restriction of the class of Markov chains (for example, to stochastically monotone chains). Such techniques have been used for ergodic costs from a classical control perspective, see for example Guo and Hernández-Lerma [14]. Similarly it may be possible to weaken the assumption of strictly balanced EBSDE drivers to either balanced or weakly balanced.

The applications of the theory of EBSDEs are still in development, and the explicit computability of solutions to these equations, in terms of solving a single nonlinear vector equation, is of some interest. The consequences for risk sensitive ergodic control, and properties of the convergence of these solutions to the diffusion case, remain to be explored.

References

- [1] Aristotle Arapostathis, Vivek S. Borkar, Emmanuel Fernández-Gaucherand, Mrinal K. Ghosh, and Steven I. Marcus. Discrete-time controlled Markov processes with average cost criterion: A survey. *SIAM Journal on Control & Optimization*, 31(2):282–344, 1993.
- [2] Philippe Briand and Ying Hu. Stability of BSDEs with random terminal time and homogenization of semilinear elliptic PDEs. *Journal of Functional Analysis*, 155:455–494, 1998.
- [3] Samuel N. Cohen. Representing filtration consistent nonlinear expectations as g -expectations in general probability spaces. *Stochastic Processes and their Applications*, 122(4):1601–1626, 2012.
- [4] Samuel N. Cohen and Robert J. Elliott. Solutions of backward stochastic differential equations on Markov chains. *Communications on Stochastic Analysis*, 2(2):251–262, August 2008.
- [5] Samuel N. Cohen and Robert J. Elliott. Comparisons for backward stochastic differential equations on Markov chains and related no-arbitrage conditions. *The Annals of Applied Probability*, 20(1):267–311, 2010.
- [6] Samuel N. Cohen and Robert J. Elliott. Existence, uniqueness and comparisons for BSDEs in general spaces. to appear in *Annals of Applied Probability*, 2010.
- [7] Samuel N. Cohen, Robert J. Elliott, and Charles E.M. Pearce. A general comparison theorem for backward stochastic differential equations. *Advances in Applied Probability*, 42(3):878–898, 2010.

- [8] Samuel N. Cohen and Lukasz Szpruch. On Markovian solutions to Markov chain BSDEs. *Numerical Algebra, Control and Optimization*, 2(2):257–269, 2012.
- [9] François Coquet, Ying Hu, Jean Mémin, and Shige Peng. Filtration consistent nonlinear expectations and related g -expectations. *Probability Theory and Related Fields*, 123(1):1–27, May 2002.
- [10] Arnaud Debussche, Ying Hu, and Gianmario Tessitore. Ergodic BSDEs under weak dissipative assumptions. *Stochastic Processes and their Applications*, 121:407–426, 2011.
- [11] Robert J. Elliott, L. Aggoun, and J.B. Moore. *Hidden Markov Models: Estimation and Control*. Springer-Verlag, Berlin-Heidelberg-New York, 1994.
- [12] Marco Fuhrman, Ying Hu, and Gianmario Tessitore. Ergodic BSDEs and optimal ergodic control in Banach spaces. *SIAM Journal on Control & Optimization*, 48(3):1542–1566, 2009.
- [13] L. Galtchouk and S. Pergamenschikov. Geometric ergodicity for classes of homogeneous Markov chains. arxiv:1002.2341v2.
- [14] Xianping Guo and Onésimo Hernández-Lerma. Drift and monotonicity conditions for continuous-time controlled Markov chains with an average criterion. *IEEE Transactions on Automatic Control*, 48:236–245, 2001.
- [15] Prasadarao Kakumanu. Nondiscounted continuous time Markovian decision process with countable state space. *SIAM Journal on Control*, 10(1):210–220, 1972.
- [16] Torgny Lindvall. *Lectures on the Coupling Method*. Dover, Mineola, New York, 2002.
- [17] E. J. McShane and R. B. Warfield, Jr. On Filippov’s implicit functions lemma. *Proceedings of the American Mathematical Society*, 18(1):41–47, 1967.
- [18] Sean Meyn and Richard L Tweedie. *Markov Chains and Stochastic Stability*. Cambridge University Press, 2nd edition, 2009.
- [19] E. Pardoux and Shige Peng. Adapted solution of a backward stochastic differential equation. *Systems & Control Letters*, 14:55–61, 1990.
- [20] Shige Peng. *Backward Stochastic Differential Equations*, chapter 9: Backward SDE and related g -expectation, pages 141–159. Pitman Research Notes in Mathematics. Longman, 1997.
- [21] Adrien Richou. Ergodic BSDEs and related PDEs with Neumann boundary conditions. *Stochastic Processes and their Applications*, 119:2945–2969, 2009.
- [22] L.C.G. Rogers and D. Williams. *Diffusions, Markov Processes and Martingales*. Cambridge University Press, 2nd edition, 2000.

- [23] Manuela Royer. BSDEs with a random terminal time driven by a monotone generator and their links with PDEs. *Stochastics and Stochastics Reports*, 76:281–307, 2004.
- [24] Eberhard Zeidler. *Nonlinear Functional Analysis and its Applications*, volume II/B. Springer, 1989.