



Information Geometry of Reversible Markov Chains

Geoffrey Wolfer¹ · Shun Watanabe¹

Received: 11 June 2021 / Revised: 1 November 2021 / Accepted: 2 November 2021 /
Published online: 22 November 2021
© The Author(s) 2021

Abstract

We analyze the information geometric structure of time reversibility for parametric families of irreducible transition kernels of Markov chains. We define and characterize reversible exponential families of Markov kernels, and show that irreducible and reversible Markov kernels form both a mixture family and, perhaps surprisingly, an exponential family in the set of all stochastic kernels. We propose a parametrization of the entire manifold of reversible kernels, and inspect reversible geodesics. We define information projections onto the reversible manifold, and derive closed-form expressions for the e-projection and m-projection, along with Pythagorean identities with respect to information divergence, leading to some new notion of reversibilization of Markov kernels. We show the family of edge measures pertaining to irreducible and reversible kernels also forms an exponential family among distributions over pairs. We further explore geometric properties of the reversible family, by comparing them with other remarkable families of stochastic matrices. Finally, we show that reversible kernels are, in a sense we define, the minimal exponential family generated by the m-family of symmetric kernels, and the smallest mixture family that comprises the e-family of memoryless kernels.

Keywords Irreducible Markov chain · Reversible Markov chain · Exponential family · Mixture family · Information projection

The first author is supported by the Japan Society for the Promotion of Science (International Research Fellow, and under KAKENHI Grant 21F20378). The second author is supported in part by Japan Society for the Promotion of Science KAKENHI under Grant 20H02144.

✉ Geoffrey Wolfer
geo-wolfer@m2.tuat.ac.jp
Shun Watanabe
shunwata@cc.tuat.ac.jp

¹ Department of Computer and Information Sciences, Tokyo University of Agriculture and Technology, Tokyo, Japan

1 Introduction

Time reversibility is a fundamental property of many statistical laws of nature. Inspired by Schrödinger [1], Kolmogorov was the first [2], in his celebrated work [3,4], to investigate this notion in the context of Markov chains and diffusion processes. Reversible chains also find numerous applications in computer science, for instance in queuing networks [5] or Markov Chain Monte Carlo sampling algorithms [6]. For instance, a random walk over a weighted network corresponds to a reversible Markov chains [7, Section 3.2].

Reversible Markov operators enjoy a considerably richer mathematical structure than their non-reversible counterparts, enabling a wide range of analytical tools and techniques. Indeed, the significance of reversibility spans across surprisingly many areas of mathematics, from spectral theory [8, Chapter 12] to abstract algebra [9]. For instance, the mixing time of a reversible Markov chain, i.e. the time to guarantee closeness to stationarity, is controlled up to logarithmic factors by its absolute spectral gap (the difference of its two largest eigenvalues in magnitude). The diversity of the existing tools and analyses prompts our first question of whether reversibility can also be treated from an information geometry perspective.

Through the lens of information geometry, the manifold of all irreducible Markov kernels forms both an exponential family (e-family) and a mixture family (m-family). Our natural second question is whether we can find subfamilies of irreducible kernels that enjoy similar geometric properties, or in other words, can we find submanifolds that are autoparallel with respect to affine connections of interest? For instance, the set of doubly-stochastic matrices is known to form an m-family [10], while a tree model is an e-family of Markov kernels, if and only if it is an *FSMX model* [11].

In this article, we will answer these two questions, see that reversible irreducible Markov chains enjoy the structure of both exponential and mixture families, and explore their geometric properties.

1.1 Related work

The concept of exponential tilting of stochastic matrices using Perron-Frobenius (PF) theory can be traced back to the work of Miller [12]. The large deviation theory for Markov chains, whose crown achievement is showing that the convex conjugate of the log-PF root of the tilted kernel essentially controls the large deviation rate was further developed by Donsker and Varadhan [13], Gärtner [14], Dembo and Zeitouni [15]. Csiszár et al. [16] seem to be the first to recognize the exponential structure of the set of irreducible Markov kernels, in the context of information projections. Independently, Ito and Amari [17] implicitly introduced the notion of *asymptotic exponential families*, and exhibited irreducible Markov kernels as an example. Takeuchi and Barron [18] later formalized this definition (see also Takeuchi and Kawabata [19]), and Takeuchi and Nagaoka [20] subsequently proved that exponential families and their asymptotic counterparts are equivalent. Nakagawa and Kanaya [21] formally defined the exponential family of irreducible Markov chains and Nagaoka [22] later gave a full treatment in the language of information geometry, proving its dually flat struc-

ture. A notable collection of works has also explored the implications of this geometric structure for problems related to parameter estimation [10], hypothesis testing [21,23], large deviation theory [24], and hidden Markov models [25,26].

We refer the reader to Levin et al. [8] and Amari and Nagaoka [27] for thorough treatments of the theory of Markov chains and information geometry.

1.2 Outline and main results

In Section 2, we begin with a primer on reversible Markov chains, define exponential and mixture families, and briefly discuss the importance of affine structures for our analysis of exponential families. In Section 3, we define a time-reversal operation on parametric families, and show in Proposition 1 that both m -families and e -families are closed under this transformation. In Section 4 we introduce the concept of a reversible e -family, and provide a characterization (Theorem 2) of such family in terms of its carrier kernel and set of generator functions. Adapting the Kolmogorov criterion, we show that the necessary and sufficient conditions can be verified in a time that depends polynomially on the number of states. In Section 5, we prove that the set of all reversible and irreducible transition kernels is both an m -family, and an e -family (Theorem 3), construct a basis (Theorem 4), and derive a parametrization (Theorem 5) of the entire set of reversible kernels. In Section 6, we investigate information projections of an irreducible Markov chain onto its reversible submanifold. We show that the projections verify Pythagorean identities, and obtain closed-form expressions (Theorem 7). Additionally, we prove that the projections are always equidistant from an irreducible Markov kernel and its time-reversal (bisection property, Proposition 2). In Section 7, we show that reversible edge measures also form an e -family in distributions over pairs (Theorem 8). In Section 8, we briefly compare the geometric properties of reversible chains with several other natural families of Markov kernels. Finally, in Section 9, we characterize the reversible family as both the smallest exponential family that comprises symmetric kernels (Theorem 9), and the smallest mixture family that contains memoryless Markov kernels (Theorem 10).

2 Preliminaries

For $m \in \mathbb{N}$ we write $[m] = \{1, 2, \dots, m\}$. Let \mathcal{X} be a set such that $|\mathcal{X}| = m < \infty$, identified with $[m]$, where to avoid trivialities, we also assume that $m > 1$. We denote $\mathcal{P}(\mathcal{X})$ the probability simplex over \mathcal{X} , and $\mathcal{P}_+(\mathcal{X}) = \{\mu \in \mathcal{P}(\mathcal{X}) : \forall x \in \mathcal{X}, \mu(x) > 0\}$. All vectors will be written as row-vectors, unless otherwise stated. For some real matrices A and B , $\rho(A)$ is the spectral radius of A , $f[A]$ for $f: \mathbb{R} \rightarrow \mathbb{R}$ is the entry-wise application of f to A ; $A \circ B$ is the Hadamard product of A and B , $A > 0$ (resp. $A \geq 0$) means that A is an entry-wise positive (resp. non-negative) matrix. We will routinely identify a function $f: \mathcal{X}^2 \rightarrow \mathbb{R}$ with the linear operator $f: \mathbb{R}^{\mathcal{X}} \rightarrow \mathbb{R}^{\mathcal{X}}$.

2.1 Irreducible Markov chains

We let $(\mathcal{X}, \mathcal{E})$ be a strongly connected directed graph, where \mathcal{X} is the set of vertices, and $\mathcal{E} \subset \mathcal{X}^2$ the set of edges. Let $\mathcal{F}(\mathcal{X}, \mathcal{E})$ be the set of all real functions over the set \mathcal{E} , identified with the totality of functions over \mathcal{X}^2 that are null outside of \mathcal{E} , and let $\mathcal{F}_+(\mathcal{X}, \mathcal{E}) \subset \mathcal{F}(\mathcal{X}, \mathcal{E})$ be the subset of positive functions over \mathcal{E} . Similarly, we define $\mathcal{P}(\mathcal{E}) = \mathcal{P}(\mathcal{X}^2) \cap \mathcal{F}_+(\mathcal{X}, \mathcal{E})$, the set of distributions whose mass is concentrated on the edge set \mathcal{E} . We write $\mathcal{W}(\mathcal{X})$ for the set of row-stochastic transition kernels over the state space \mathcal{X} , and $\mathcal{W}(\mathcal{X}, \mathcal{E})$ for the subset of irreducible kernels whose support is \mathcal{E} , i.e.

$$\mathcal{W}(\mathcal{X}) \triangleq \left\{ P \in \mathbb{R}^{\mathcal{X}^2} : P \geq 0, \forall x \in \mathcal{X}, \sum_{x' \in \mathcal{X}} P(x, x') = 1 \right\},$$

$$\mathcal{W}(\mathcal{X}, \mathcal{E}) \triangleq \mathcal{F}_+(\mathcal{X}, \mathcal{E}) \cap \mathcal{W}(\mathcal{X}),$$

and where $P(x, x')$ corresponds to the transition probability from state x to state x' ¹. For $P \in \mathcal{W}(\mathcal{X}, \mathcal{E})$, there exists a unique $\pi \in \mathcal{P}_+(\mathcal{X})$, such that $\pi P = \pi$ [8, Corollary 1.17], which we call the stationary distribution of P . When $\mathcal{E} = \mathcal{X}^2$ and if there is no ambiguity about the space under consideration, we may write more simply \mathcal{F} , \mathcal{F}_+ instead of $\mathcal{F}(\mathcal{X}, \mathcal{X}^2)$, $\mathcal{F}_+(\mathcal{X}, \mathcal{X}^2)$ (a similar notation will apply to all subsequently defined spaces).

2.2 Reversibility

For an irreducible kernel P , we write $Q = \text{diag}(\pi)P$ for the edge measure matrix, [8, (7.5)], which corresponds to stationary pair-probabilities of P , i.e. $Q(x, x') = \mathbb{P}_\pi(X_t = x, X_{t+1} = x')$, and denote the set of irreducible edge measures by

$$\mathcal{Q}(\mathcal{X}, \mathcal{E}) \triangleq \{\text{diag}(\pi)P : P \in \mathcal{W}(\mathcal{X}, \mathcal{E}), \pi P = \pi\} \subset \mathcal{P}(\mathcal{E}).$$

Note that this definition is equivalent to

$$\mathcal{Q}(\mathcal{X}, \mathcal{E}) = \left\{ Q \in \mathcal{P}(\mathcal{E}) : \sum_{x' \in \mathcal{X}} Q(x, x') = \sum_{x' \in \mathcal{X}} Q(x', x) \right\}. \quad (1)$$

We further denote P^* for the uniquely defined *time-reversal* of P , that verifies $P^*(x, x') = \pi(x')P(x', x)/\pi(x)$, and write $Q^* = Q^\top$ for its corresponding edge measure, where \top denotes matrix transposition. When Q is symmetric (i.e. $Q^* = Q$), the chain verifies the *detailed balance equation*,

$$\pi(x)P(x, x') = \pi(x')P(x', x),$$

¹ We note that in information theory, $P(x, x')$ is often denoted by $P(x'|x)$. Our choice follows the applied probability literature (see e.g. Levin et al. [8]), and allows us to extend the notation seamlessly to general functions over \mathcal{X}^2 .

i.e. $P^* = P$, and we say that the Markov chain is *reversible*. Observe that in this case, for P irreducible over \mathcal{E} , the edge set must also be symmetric ($\mathcal{E} = \mathcal{E}^*$, where $\mathcal{E}^* \triangleq \{(x, x') \in \mathcal{X}^2 : (x', x) \in \mathcal{E}\}$). We write $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ for the set of all reversible kernels that are irreducible over $(\mathcal{X}, \mathcal{E})$. For $f, g \in \mathbb{R}^{\mathcal{X}}$, $\langle f, g \rangle_{\pi} \triangleq \sum_{x \in \mathcal{X}} f(x)g(x)\pi(x)$ defines an inner product. We call $\ell_2(\pi)$ the corresponding Hilbert space. The time-reversal is the adjoint operator of P in $\ell_2(\pi)$, i.e. the unique linear operator that verifies $\langle Pf, g \rangle_{\pi} = \langle f, P^*g \rangle_{\pi}, \forall f, g \in \mathbb{R}^{\mathcal{X}}$ (represented here as column vectors). As a consequence, when P is reversible, it is also self-adjoint in $\ell_2(\pi)$, and the spectrum of P is real.

2.3 Mixture family and exponential family

For later convenience we consider the following three equivalent definitions of a mixture family.

Definition 1 (*m-family of transition kernels*) We say that a family of irreducible transition kernels \mathcal{V}_m is a mixture family (m-family) of irreducible transition kernels on $(\mathcal{X}, \mathcal{E})$ when one of the following (equivalent) statements (i), (ii), (iii) holds.

(i) [28] There exist affinely independent $Q_0, Q_1, \dots, Q_d \in \mathcal{Q}(\mathcal{X}, \mathcal{E})$ such that

$$\mathcal{V}_m = \left\{ P_{\xi} \in \mathcal{W}(\mathcal{X}, \mathcal{E}) : Q_{\xi} = \sum_{i=1}^d \xi^i Q_i + \left(1 - \sum_{i=1}^d \xi^i\right) Q_0, \xi \in \mathcal{E} \right\},$$

where $\mathcal{E} = \{\xi \in \mathbb{R}^d : Q_{\xi}(x, x') > 0, \forall (x, x') \in \mathcal{E}\}$, and Q_{ξ} is the edge measure that pertains to P_{ξ} .

(ii) [27, 2.35] There exists $C, F_1, \dots, F_d \in \mathcal{F}(\mathcal{X}, \mathcal{E})$, such that $C, C + F_1, \dots, C + F_d$ are affinely independent,

$$\sum_{x, x'} C(x, x') = 1, \quad \sum_{x, x'} F_i(x, x') = 0, \forall i \in [d],$$

and

$$\mathcal{V}_m = \left\{ P_{\xi} \in \mathcal{W}(\mathcal{X}, \mathcal{E}) : Q_{\xi} = C + \sum_{i=1}^d \xi^i F_i, \xi \in \mathcal{E} \right\}$$

where $\mathcal{E} = \{\xi \in \mathbb{R}^d : Q_{\xi}(x, x') > 0, \forall (x, x') \in \mathcal{E}\}$, and Q_{ξ} is the edge measure that pertains to P_{ξ} .

(iii) [10, Section 4.2] There exist $k \in \mathbb{N}, g_1, \dots, g_k \in \mathcal{F}(\mathcal{X}, \mathcal{E})$ and $c_1, \dots, c_k \in \mathbb{R}$, such that

$$\mathcal{V}_m = \left\{ P \in \mathcal{W}(\mathcal{X}, \mathcal{E}) : \sum_{x, x'} Q(x, x') g_i(x, x') = c_i, \forall i \in [k] \right\}.$$

Note that \mathcal{E} is an open set, ξ is called the mixture parameter and d is the dimension of the family \mathcal{V}_m .

Definition 2 (*e-family of transition kernels*) Let $\Theta \subset \mathbb{R}^d$, be some connected parameter space that contains an open ball centered at 0. We say that the parametric family of irreducible transition kernels

$$\mathcal{V}_e = \left\{ P_\theta : \theta = (\theta^1, \dots, \theta^d) \in \Theta \right\}$$

is an exponential family (e-family) of transition kernels on $(\mathcal{X}, \mathcal{E})$ with natural parameter θ , whenever

- (i) For all $\theta \in \Theta$, $P_\theta \in \mathcal{W}(\mathcal{X}, \mathcal{E})$.
- (ii) There exist functions

$$\begin{aligned} K &: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}, \\ R &: \Theta \times \mathcal{X} \rightarrow \mathbb{R}, \\ g_1, \dots, g_d &: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}, \\ \psi &: \Theta \rightarrow \mathbb{R}, \end{aligned}$$

such that $\forall (x, x', \theta) \in \mathcal{X}^2 \times \Theta$,

$$\log P_\theta(x, x') = K(x, x') + \sum_{i=1}^d \theta^i g_i(x, x') + R(\theta, x') - R(\theta, x) - \psi(\theta), \quad (2)$$

when $(x, x') \in \mathcal{E}$, and $P_\theta(x, x') = 0$ otherwise.

When fixing some $\theta \in \Theta$, we may later write for convenience ψ_θ for $\psi(\theta)$ and R_θ for $R(\theta, \cdot) \in \mathbb{R}^{\mathcal{X}}$. The carrier kernel K , the collection of generator functions g_1, \dots, g_d and the parameter range Θ define the family entirely. The remaining functions R_θ and ψ_θ will be determined uniquely by PF theory, from the constraint of P_θ being row-stochastic (see for example the proof of Proposition 1). In fact, we can define the mapping \mathfrak{s} that constructs a proper irreducible stochastic matrix from any linear operator defined by an irreducible matrix over $(\mathcal{X}, \mathcal{E})$.

$$\begin{aligned} \mathfrak{s} &: \mathcal{F}_+(\mathcal{X}, \mathcal{E}) \rightarrow \mathcal{W}(\mathcal{X}, \mathcal{E}) \\ \tilde{P}(x, x') &\mapsto P(x, x') = \frac{\tilde{P}(x, x')v(x')}{\rho(\tilde{P})v(x)}, \end{aligned} \quad (3)$$

where $\rho(\tilde{P})$ and v are respectively the PF root and right PF eigenvector of \tilde{P} .

Remark 1 In Feigin et al. [29], Küchler and Sørensen [30], Hudson [31], Stefanov [32], Küchler and Sørensen [33], Sørensen [34], an exponential family of transition

kernels has the form

$$\log P_\theta(x, x') = K(x, x') + \sum_{i=1}^d \theta^i g_i(x, x') - \phi(\theta, x'),$$

for some function $\phi: \Theta \times \mathcal{X} \rightarrow \mathbb{R}$. Our Definition 2 however follows the one of Nagaoka [22], Hayashi and Watanabe [10], Watanabe and Hayashi [23], that is endowed with a more compelling geometrical structure [10, Remark 3].

Following the information geometry philosophy [27], we view the e-families or m-families that we defined, as d -dimensional submanifolds of $\mathcal{W}(\mathcal{X}, \mathcal{E})$ with corresponding chart maps $\theta, \xi: \mathcal{W}(\mathcal{X}, \mathcal{E}) \rightarrow \mathbb{R}^d$. We can give more geometrical, parametrization-free definitions of e-families and m-families of irreducible transition kernel over $(\mathcal{X}, \mathcal{E})$, as autoparallel submanifolds of $\mathcal{W}(\mathcal{X}, \mathcal{E})$ with respect to the e-connection and m-connection [22, Section 6]. We will prefer, however, to mostly cast our analysis in the language of linear algebra, and defer analysis of the relationship with differential geometry concepts to Section 5.3. This choice is motivated by the existence of a known correspondence between affine functions over \mathcal{E} and the manifold $\mathcal{W}(\mathcal{X}, \mathcal{E})$ [22] that we now describe. Denote,

$$\mathcal{N}(\mathcal{X}, \mathcal{E}) \triangleq \left\{ h \in \mathcal{F}(\mathcal{X}, \mathcal{E}) : \exists (c, f) \in (\mathbb{R}, \mathbb{R}^{\mathcal{X}}), \right. \\ \left. \forall (x, x') \in \mathcal{E}, h(x, x') = f(x') - f(x) + c \right\}. \tag{4}$$

Then $\mathcal{F}(\mathcal{X}, \mathcal{E})$ defines a $|\mathcal{E}|$ -dimension vector space, while $\mathcal{N}(\mathcal{X}, \mathcal{E})$ is an $|\mathcal{X}|$ -dimensional vector space [22, Section 3]. Introducing the mapping,

$$\Delta: \mathcal{F}(\mathcal{X}, \mathcal{E}) \rightarrow \mathcal{W}(\mathcal{X}, \mathcal{E}) \\ f \mapsto \Delta(f) = \mathfrak{s}(\exp \circ f), \tag{5}$$

we see from the expression at (3) that Δ gives a diffeomorphism from the quotient linear space

$$\mathcal{G}(\mathcal{X}, \mathcal{E}) \triangleq \mathcal{F}(\mathcal{X}, \mathcal{E}) / \mathcal{N}(\mathcal{X}, \mathcal{E})$$

to $\mathcal{W}(\mathcal{X}, \mathcal{E})$ and a subset \mathcal{V} of $\mathcal{W}(\mathcal{X}, \mathcal{E})$ is an e-family if and only if there exists an affine subspace \mathcal{A} of the quotient space $\mathcal{G}(\mathcal{X}, \mathcal{E})$ such that $\mathcal{V} = \Delta(\mathcal{A})$ (we identify a coset with a representative function in that coset). In this case, the correspondence is one-to-one, and the dimension of the affine space and the submanifold coincide [22, Theorem 2]. In particular, this entails that $\dim \mathcal{W}(\mathcal{X}, \mathcal{E}) = |\mathcal{E}| - |\mathcal{X}|$ Nagaoka [22, Corollary 1].

Remark 2 For Definition 2, unless stated otherwise, we will henceforth assume that the g_i form an independent family in $\mathcal{G}(\mathcal{X}, \mathcal{E})$. This will ensure that the family is well-behaved in the sense of Hayashi and Watanabe [10, Lemma 4.1].

3 Time-reversal of parametric families

We begin by extending the definition of a time-reversal to families of Markov chains.

Definition 3 (*Time-reversal family*) We say that the family of irreducible transition kernels \mathcal{V}^* is the time-reversal of the family of irreducible transition kernels \mathcal{V} when $\mathcal{V}^* = \{P^* : P \in \mathcal{V}\}$, where P^* denotes the time-reversal of P .

We now state the fundamental fact that the quality of being an e-family or an m-family of transition kernels is closed under this time-reversal operation.

Proposition 1 *The following statements hold.*

Time reversal of m-family: Let \mathcal{V}_m be an m-family over $(\mathcal{X}, \mathcal{E})$, then \mathcal{V}_m^ is an m-family over $(\mathcal{X}, \mathcal{E}^*)$. Furthermore, if \mathcal{V}_m is the m-family generated by $Q_1, \dots, Q_d \in \mathcal{Q}(\mathcal{X}, \mathcal{E})$ (following the notation at Definition 1-(i)), then the time-reversal m-family is given by*

$$\mathcal{V}_m^* = \left\{ P_\xi \in \mathcal{W}(\mathcal{X}, \mathcal{E}^*) : Q_\xi = \sum_{i=1}^d \xi^i Q_i^* + \left(1 - \sum_{i=1}^d \xi_i\right) Q_0^*, \xi \in \mathcal{E}^* \right\},$$

where Q_ξ pertains to P_ξ and with

$$\mathcal{E}^* = \left\{ \xi \in \mathbb{R}^d : Q_\xi(x, x') > 0, \forall (x, x') \in \mathcal{E}^* \right\} = \mathcal{E}.$$

Time reversal of e-family: Let \mathcal{V}_e be an e-family over $(\mathcal{X}, \mathcal{E})$, then \mathcal{V}_e^ is an e-family over $(\mathcal{X}, \mathcal{E}^*)$. Furthermore, if \mathcal{V}_e is the e-family generated by K and g_1, \dots, g_d (following the notation at Definition 2), then the time-reversal e-family is given by $\mathcal{V}^* = \{P_\theta^* : \theta \in \Theta\}$ such that*

$$\log P_\theta^*(x, x') = K(x', x) + \sum_{i=1}^d \theta^i g_i(x', x) + L_\theta(x') - L_\theta(x) - \psi_\theta,$$

when $(x, x') \in \mathcal{E}^*$, $P_\theta^*(x, x') = 0$ otherwise, and where L_θ is the left PF eigenvector of the non-negative irreducible matrix

$$\tilde{P}_\theta^*(x, x') = \exp \left(K(x', x) + \sum_{i=1}^d \theta^i g_i(x', x) \right).$$

Proof Since the edge measure Q_ξ^* of the time-reversal P_ξ^* is the transpose of Q_ξ corresponding to P_ξ , it is easy to obtain the expression of the time-reversal, and to see that \mathcal{V}_m^* is a mixture family. It remains to show that this also holds true for e-families. From the definition of an exponential family (2), and the requirement that

P_θ be row-stochastic, it must be that for any $x \in \mathcal{X}$,

$$\sum_{x' \in \mathcal{X}} \exp(K(x, x') + \sum_{i=1}^d \theta^i g_i(x, x')) e^{R_\theta(x')} = e^{\psi_\theta} e^{R_\theta(x)},$$

or more concisely, writing $\tilde{P}_\theta(x, x') = \exp(K(x, x') + \sum_{i=1}^d \theta^i g_i(x, x'))$ for $x, x' \in \mathcal{E}$ and $\tilde{P}_\theta(x, x') = 0$ otherwise, $\tilde{P}_\theta \exp[R_\theta] = e^{\psi_\theta} \exp[R_\theta]$. By positivity of the exponential function, the vector $\exp[R_\theta] \in \mathbb{R}^{\mathcal{X}}$ is positive. Thus, from the PF theorem, e^{ψ_θ} corresponds to the spectral radius of \tilde{P}_θ , and $\exp[R_\theta]$ its (right) associated eigenvector. There must therefore also exist a left positive eigenvector, which we denote by $\exp[L_\theta]$, such that

$$\exp[L_\theta] \tilde{P}_\theta = e^{\psi_\theta} \exp[L_\theta].$$

Defining the positive normalized measure

$$\pi_\theta(x) \triangleq \frac{\exp(L_\theta(x) + R_\theta(x))}{\sum_{x'' \in \mathcal{X}} \exp(L_\theta(x'') + R_\theta(x''))}, \tag{6}$$

it is easily verified that π_θ is the stationary distribution of P_θ . Notice that θ, K and g_i determine uniquely $L_\theta, R_\theta, \psi_\theta$ and π_θ by the PF theorem. Recall that the adjoint of a transition kernel P can be written $P^*(x, x') = \pi(x')P(x', x)/\pi(x)$, thus we can compute the time-reversal as

$$P_\theta(x', x) \frac{\pi_\theta(x')}{\pi_\theta(x)} = \exp \left(K(x', x) + L_\theta(x') - L_\theta(x) + \sum_{i=1}^d \theta^i g_i(x', x) - \psi_\theta \right),$$

when $(x', x) \in \mathcal{E}$, and 0 for $(x', x) \notin \mathcal{E}$. The requirements of Definition 2 for an e-family are all fulfilled, which concludes the proof. □

Remark 3 Recall that for a distribution $\mu \in \mathcal{P}(\mathcal{X})$, we can by *exponential change of measure* – also known as *exponential tilting* – construct the natural exponential family of μ :

$$\mu_\theta(x) = \mu(x) \exp(\theta x - A(\theta)),$$

where $A(\theta)$ is a normalization function that ensures $\mu_\theta \in \mathcal{P}(\mathcal{X})$ for all $\theta \in \mathbb{R}$. The idea of exponential change of measure for distributions can be traced back to Chernoff [35], and was later termed tilting [36,37]. Similarly, given some function $g : \mathcal{X}^2 \rightarrow \mathbb{R}$ we can tilt an irreducible kernel P (e.g. Miller [12]), by first constructing $\tilde{P}_\theta(x, x') = P(x, x')e^{\theta g(x, x')}$, and then rescaling the newly obtained irreducible matrix ² with the mapping s . When $\theta = 0$, notice that we recover the original P . But while

² Interestingly, the large deviation rate of g is given by the convex conjugate of the log-PF root of \tilde{P}_θ [15, Chapter 3].

in our definition,

$$P_\theta = \mathfrak{s}(P \circ \exp[\theta g]) = \frac{1}{\rho(P \circ \exp[\theta g])} \text{diag } v_\theta^{-1}(P \circ \exp[\theta g]) \text{diag } v_\theta$$

denotes the kernel tilted involving the right PF eigenvector v_θ , we could alternatively define the Markov kernel P'_θ by tilting P with the left PF eigenvector u_θ :

$$P'_\theta = \frac{1}{\rho(P \circ \exp[\theta g])} \text{diag } u_\theta^{-1}(P \circ \exp[\theta g])^\top \text{diag } u_\theta.$$

Observe that the right and left tilted versions of P with identical θ share the same stationary distribution $\pi_\theta \propto u_\theta \circ v_\theta$ (6) and that they are in fact each other’s time-reversal ($P'_\theta = P_\theta^*$), i.e. they form a pair of adjoint linear operators over the space $\ell_2(\pi_\theta)$.

4 Reversible exponential families

The previous section extended the time-reversal operation to parametric families of transition kernels. It seems then natural to investigate fixed points, i.e. parametric families that remain invariant under this transformation. We say that an irreducible e-family $\mathcal{V}(\mathcal{X}, \mathcal{E})$ is reversible when P is reversible $\forall P \in \mathcal{V}(\mathcal{X}, \mathcal{E})$. In this case, \mathcal{E} coincides with \mathcal{E}^* and $\mathcal{V}^*(\mathcal{X}, \mathcal{E}^*)$ with $\mathcal{V}(\mathcal{X}, \mathcal{E})$. Observe first that an e-family obtained from tilting a reversible P is not generally reversible, making it clear that the reversible nature of the family cannot be determined solely by the properties of the carrier kernel K . It is however easy to see that an e-family is reversible when K and all the generator functions g_1, \dots, g_d are symmetric. Moreover, for a state space \mathcal{X} of size 2, any exponential family would be reversible regardless of symmetry, showing that this condition is not always necessary. In this section, we give a complete characterization of this invariant set. Additionally, we explore the algorithmic cost of checking whether this property is verified from the description of the carrier kernel and generators of a given e-family. Before diving into the general theory of reversible e-families, let us consider the following simple examples.

Example 1 (*Lazy random walk on the m-cycle*) For $\mathcal{X} = [m]$, and

$$\mathcal{E}_{cy} = \left\{ (i, j) \in [m]^2 : |i - j| \bmod m - 1 \in \{0, 1\} \right\},$$

let

$$P_\theta(x, x') = \exp \left(\theta_1 \sum_{i=1}^m \delta_i(x) \delta_i(x') + \theta_2 \sum_{i=1}^m (\delta_i(x) \delta_{i+1}(x') - \delta_{i+1}(x) \delta_i(x')) - \psi(\theta) \right),$$

where $\psi(\theta) = \log(e^{\theta_1} + e^{\theta_2} + e^{-\theta_2})$ and $\delta_{m+1} = \delta_1$. This e-family $\{P_\theta\}_{\theta \in \mathbb{R}^2}$ corresponds to the set of biased lazy random walks on the m -cycle given by

$$\begin{aligned} P_\theta(x, x) &= \frac{e^{\theta_1}}{e^{\theta_1} + e^{\theta_2} + e^{-\theta_2}}, \\ P_\theta(x, x + 1) &= \frac{e^{\theta_2}}{e^{\theta_1} + e^{\theta_2} + e^{-\theta_2}}, \\ P_\theta(x + 1, x) &= \frac{e^{-\theta_2}}{e^{\theta_1} + e^{\theta_2} + e^{-\theta_2}}. \end{aligned}$$

Observe that $P_{(\theta_1, \theta_2)}^* = P_{(\theta_1, -\theta_2)}$, and thus P_θ is not a reversible e-family. The subfamily $\{P_\theta : \theta_1 \in \mathbb{R}, \theta_2 = 0\}$, however, i.e. unbiased lazy random walks on the m -cycle, form a reversible e-family.

Example 2 (Birth-and-death chains) For $\mathcal{X} = [m]$ and $\mathcal{E}_{\text{bd}} = \{(i, j) : |i - j| \leq 1\}$, a Markov kernel having its support on \mathcal{E}_{bd} is referred to as a birth-and-death chain. Since every birth-and-death chain is reversible [8, Section 2.5], $\mathcal{W}(\mathcal{X}, \mathcal{E}_{\text{bd}})$ is a reversible e-family.

We first recall Kolmogorov’s characterization of reversibility, which will be instrumental in our argument. For $\mathcal{E} \subset \mathcal{X}^2$ such that $(\mathcal{X}, \mathcal{E})$ is a strongly connected directed graph, we write $\Gamma(\mathcal{X}, \mathcal{E})$ for the set of *finite directed closed paths* in the graph $(\mathcal{X}, \mathcal{E})$. Formally, we treat γ as a map $[n] \rightarrow \mathcal{E}$ such that $\gamma(t) = (x_t, x_{t+1})$ with $x_{n+1} = x_1$ and we write $|\gamma| = n$ for the length of the path. For each $\gamma \in \Gamma(\mathcal{X}, \mathcal{E})$, we also introduce the reverse closed path $\gamma^* \in \Gamma(\mathcal{X}, \mathcal{E}^*)$ given by $\gamma^*(t) = (x_{t+1}, x_t)$. Namely, if $\gamma \in \Gamma(\mathcal{X}, \mathcal{E})$, we can write γ informally as a succession of edges such that the starting and finishing states agree (i.e. as an element of \mathcal{E}^n).

$$\gamma = ((x_1, x_2), (x_2, x_3), \dots, (x_{n-1}, x_n), (x_n, x_1)), x_i \in \mathcal{X}, \forall i \in [n].$$

Note that γ is not necessarily a *cycle*, i.e. in our definition, multiple occurrences of the same point of the space are allowed.

Theorem 1 (Kolmogorov’s criterion [3]) *Let P irreducible over $(\mathcal{X}, \mathcal{E})$. P is reversible if and only if for all $\gamma \in \Gamma(\mathcal{X}, \mathcal{E})$,*

$$\prod_{t=1}^{|\gamma|} P(\gamma(t)) = \prod_{t=1}^{|\gamma|} P(\gamma^*(t)).$$

Example 3 When $|\mathcal{X}| = 2$, all chains are reversible. For $|\mathcal{X}| = 3$, only one equation needs to be verified for P to be reversible:

$$P(1, 2)P(2, 3)P(3, 1) = P(1, 3)P(3, 2)P(2, 1).$$

We now extend the definition of reversibility to arbitrary irreducible functions $\mathcal{F}_+(\mathcal{X}, \mathcal{E})$ (non-negative on \mathcal{X}^2 and positive exactly on \mathcal{E}) based on Kolmogorov’s

criterion, and further introduce the concept of log-reversibility for $\mathcal{F}(\mathcal{X}, \mathcal{E})$, that considers sums instead of products.

Definition 4 (*Reversible and log-reversible functions*) Let $\mathcal{E} \subset \mathcal{X}^2$ such that $\mathcal{E} = \mathcal{E}^*$. reversible: A function $h \in \mathcal{F}_+(\mathcal{X}, \mathcal{E})$ is reversible whenever it satisfies that,

$$\prod_{t=1}^{|\gamma|} h(\gamma(t)) = \prod_{t=1}^{|\gamma|} h(\gamma^*(t)).$$

for all finite directed closed paths $\gamma \in \Gamma(\mathcal{X}, \mathcal{E})$.

log-reversible: A function $h \in \mathcal{F}(\mathcal{X}, \mathcal{E})$ is log-reversible whenever it satisfies that,

$$\sum_{t=1}^{|\gamma|} h(\gamma(t)) = \sum_{t=1}^{|\gamma|} h(\gamma^*(t)).$$

for all finite directed closed paths $\gamma \in \Gamma(\mathcal{X}, \mathcal{E})$.

Remark These definitions do not rely on connectedness properties of \mathcal{E} per se, but we will assume irreducibility nonetheless. Observe that when h is represented by an irreducible row-stochastic matrix, the definition of reversibility of h as a function and as a Markov operator coincide by Kolmogorov’s criterion (Theorem 1). Clearly, for $h \in \mathcal{F}(\mathcal{X}, \mathcal{E})$, $\exp[h]$ being reversible is equivalent to h being log-reversible. We could endow the set of positive reversible functions on \mathcal{E} with a group structure by considering the standard multiplicative operation on functions. We will choose however (Lemma 4), to rather construct and focus on the vector space of log-reversible functions.

Lemma 1 Let $h \in \mathcal{F}_+(\mathcal{X}, \mathcal{E})$ such that $\text{rank}(h) = 1$. Then h is a reversible function.

Proof Consider a closed path $\gamma \in \Gamma(\mathcal{X}, \mathcal{E})$. Writing $h = u_h^\top v_h$, we successively have that

$$\prod_{t=1}^{|\gamma|} h(\gamma(t)) = \prod_{t=1}^{|\gamma|} u_h(x_t) v_h(x_{t+1}) = \prod_{t=1}^{|\gamma|} u_h(x_t) \prod_{t=1}^{|\gamma|} v_h(x_t) = \prod_{t=1}^{|\gamma|} h(\gamma^*(t)).$$

□

Theorem 2 (*Characterization of reversible e-family*) Let $\mathcal{V}(\mathcal{X}, \mathcal{E})$ be an irreducible e-family of Markov chains, with natural parametrization θ , generated by K and $(g_i)_{i \in [d]}$. The following two statements are equivalent.

- (i) $\mathcal{V}(\mathcal{X}, \mathcal{E})$ is reversible.
- (ii) $\mathcal{E} = \mathcal{E}^*$ and $\mathcal{V}(\mathcal{X}, \mathcal{E})$ is such that the carrier kernel K and generator functions $g_i, \forall i \in [d]$ are all log-reversible functions.

Proof We apply Kolmogorov’s criterion to some arbitrary family member. Let γ be some finite closed path in $(\mathcal{X}, \mathcal{E})$,

$$\prod_{t=1}^{|\gamma|} P_\theta(\gamma(t)) = \prod_{t=1}^{|\gamma|} P_\theta(\gamma^*(t)).$$

Rewriting the left-hand side,

$$\begin{aligned} & \prod_{t=1}^{|\gamma|} P_\theta(\gamma(t)) \\ &= \prod_{t=1}^{|\gamma|} \exp\left(K(x_t, x_{t+1}) + R(\theta, x_{t+1}) - R(\theta, x_t) + \sum_{i=1}^d \theta^i g_i(x_t, x_{t+1}) - \psi(\theta)\right) \\ &= \exp(-|\gamma| \psi(\theta)) \exp\left(\sum_{t=1}^{|\gamma|} \left[K(x_t, x_{t+1}) + \sum_{i=1}^d \theta^i g_i(x_t, x_{t+1})\right]\right). \end{aligned}$$

Proceeding in a similar way with the right-hand side, we obtain

$$\begin{aligned} & \sum_{t=1}^{|\gamma|} K(x_t, x_{t+1}) + \sum_{i=1}^d \theta^i \sum_{t=1}^{|\gamma|} g_i(x_t, x_{t+1}) \\ &= \sum_{t=1}^{|\gamma|} K(x_{t+1}, x_t) + \sum_{i=1}^d \theta^i \sum_{t=1}^{|\gamma|} g_i(x_{t+1}, x_t). \end{aligned}$$

When K and the g_i are log-reversible, this equality is verified for any closed path, and every member of the family is therefore reversible. Taking $\theta = 0$ yields the reversibility requirement for K . Further taking $\theta^i = \delta_i(j)$ for $j \in [d]$ similarly yields the requirement for g_j . □

This path checking approach, although mathematically convenient, is not algorithmically efficient. In order to determine whether a full-support kernel –or function– is reversible, the number of distinct Kolmogorov equations that must be checked is

$$\sum_{k=3}^{|\mathcal{X}|} \binom{|\mathcal{X}|}{k} \frac{(k-1)!}{2}, \quad [38, \text{Proposition 2.1}]$$

which corresponds to the maximal number of cycles (i.e. closed paths such that the only repeated vertices are the first and last one) in a complete graph over $|\mathcal{X}|$ nodes. Such testing algorithm becomes rapidly intractable as $|\mathcal{X}|$ increases. However for Markov kernels, we know that this is equivalent to verifying the detailed balance equation, which can be achieved in (at most) polynomial time $\mathcal{O}(|\mathcal{X}|^3)$, by solving a linear system in order to find π . We show that this idea naturally extends to verifying reversibility of functions, enabling us to design an algorithm of time complexity $\mathcal{O}(|\mathcal{X}|^3)$.

Lemma 2 Let $h \in \mathcal{F}_+(\mathcal{X}, \mathcal{E})$ irreducible. h is reversible if and only if $\Pi_h \circ h^\top$ is a symmetric matrix, with $\Pi_h = v_h^\top u_h$ the PF projection of h , where u_h and v_h are respectively the left and right PF eigenvectors of h , normalized such that $u_h v_h^\top = 1$.

Proof Treat h as the linear operator $h: \mathbb{R}^{\mathcal{X}} \rightarrow \mathbb{R}^{\mathcal{X}}$. Suppose first that h is reversible. We apply PF theory, which guarantees that the following Cesàro averages converge [39, Example 8.3.2] to some positive projection,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^n \left(\frac{h}{\rho(h)} \right)^k = \Pi_h = v_h^\top u_h, \tag{7}$$

Fix $(x, x') \in \mathcal{X}^2$ such that $h(x, x') \neq 0$. For $k \in \mathbb{N}$, we write $\Gamma_k(x, x') \subset \Gamma(\mathcal{X}, \mathcal{E})$ the set of all directed closed paths γ with $(x_1, x_2, \dots, x_k) \in \mathcal{X}^k$ such that

$$\gamma = ((x, x_1), (x_1, x_2), \dots, (x_{k-1}, x_k), (x_k, x'), (x', x)).$$

For any such cycle, it holds (perhaps vacuously if $\Gamma_k(x, x') = \emptyset$) that

$$h(x, x_1) \cdots h(x_k, x')h(x', x) = h(x, x')h(x', x_k) \cdots h(x_1, x).$$

Summing this equality over all possible paths in $\Gamma_k(x, x')$ (i.e. summing over all $(x_1, \dots, x_k) \in \mathcal{X}^k$, with the assumption that $h(x, x') = 0$ whenever $(x, x') \notin \mathcal{E}$), we obtain

$$h^k(x, x')h(x', x) = h(x, x')h^k(x', x).$$

In the case where $h(x, x') = 0$, the above equation holds by symmetry of \mathcal{E} . For $n \in \mathbb{N}$, appropriately rescaling on both sides with the PF root, summing over all $k \in \{0, \dots, n\}$ and taking the limit at $n \rightarrow \infty$, (7) yields detailed balance equations with respect to the projection Π_h ,

$$\Pi_h(x, x')h(x', x) = h(x, x')\Pi_h(x', x),$$

or in other words, reversibility of h implies symmetry of $\Pi_h \circ h^\top$.

To prove necessity, we suppose now that this symmetry holds, with Π_h the PF projection of h . We know that $\text{rank}(\Pi_h) = \text{rank}(v_h^\top u_h) = 1$, and that Π_h is positive. Consider some finite directed closed path γ . Rearranging products yields

$$\prod_{t=1}^{|\gamma|} h(\gamma(t)) = \left(\prod_{t=1}^{|\gamma|} \frac{\Pi_h(\gamma(t))}{\Pi_h(\gamma^*(t))} \right) \prod_{t=1}^{|\gamma|} h(\gamma^*(t)),$$

but the first factor on the right-hand side vanishes, from the fact that rank one functions are always reversible (Lemma 1). This concludes the proof of the lemma. \square

Notice that we can define $\pi_h(x) \triangleq u_h(x)/v_h(x)$ the positive entry-wise ratio of the PF eigenvectors. We can then restate Lemma 2 in terms of the familiar detailed balance equation $\pi_h(x)h(x, x') = \pi_h(x')h(x', x)$.

Corollary 1 *Let $h \in \mathcal{F}(\mathcal{X}, \mathcal{E})$ some irreducible function. h is log-reversible if and only if there exists $f \in \mathbb{R}^{\mathcal{X}}$ such that $\forall x, x' \in \mathcal{X}, h(x, x') = h(x', x) + f(x') - f(x)$.*

Remark: when h is known to be reversible, one can compute π_h in $\mathcal{O}(|\mathcal{X}|)$, by adapting the technique of [40]; unfortunately, it is not possible to check for reversibility using this method. If the space becomes large, the reader can consider iterative (power) methods to compute the PF projector, potentially further reducing the verification time cost. We end this section with a technical lemma that will allow us in later sections to swiftly compute expectations of functions under certain reversibility or skew-symmetry properties.

Lemma 3 *Let P irreducible with associated edge measure matrix Q . For a function $g: \mathcal{X}^2 \rightarrow \mathbb{R}$, we write $Q[g] = \sum_{x, x' \in \mathcal{X}} Q(x, x')g(x, x')$.*

- (i) *If g is log-reversible, $Q[g] = Q^*[g]$.*
- (ii) *If g is skew-symmetric and P is reversible, $Q[g] = 0$.*
- (iii) *If there exists $f \in \mathbb{R}^{\mathcal{X}}$ such that for all $x, x' \in \mathcal{X}, g(x, x') = f(x') - f(x)$, $Q[g] = 0$ (regardless of P being reversible).*

Proof Claim (iii) follows by property of edge measure Q .

$$\begin{aligned} Q[g] &= \sum_{x, x' \in \mathcal{X}} Q(x, x')(f(x') - f(x)) \\ &= \sum_{x' \in \mathcal{X}} f(x') \sum_{x \in \mathcal{X}} Q(x, x') - \sum_{x \in \mathcal{X}} f(x) \sum_{x' \in \mathcal{X}} Q(x, x') \\ &= \sum_{x' \in \mathcal{X}} f(x')\pi(x') - \sum_{x \in \mathcal{X}} f(x)\pi(x) = 0. \end{aligned}$$

From Corollary 1, claim (iii), and re-indexing,

$$Q[g] = Q[g^T] + \sum_{x, x' \in \mathcal{X}} Q(x, x')(f(x') - f(x)) = Q^*[g],$$

which yields (i). To prove (ii), consider g such that $g(x', x) = -g(x, x')$. Then by re-indexing and symmetry of Q ,

$$Q[g] = \sum_{x', x \in \mathcal{X}} Q(x', x)g(x', x) = - \sum_{x', x \in \mathcal{X}} Q(x, x')g(x, x') = -Q[g].$$

□

5 The e-family of reversible Markov kernels

In Section 5.1, we begin by analyzing the affine structure of the space of log-reversible functions, derive its dimension, construct a basis, and deduce that the manifold of all irreducible reversible Markov kernels forms an exponential family. The dimension of this family confirms the well-known fact that the number of free parameters for a reversible kernel is only about half of what is required for the general case, hence that reversible chains serve in a sense as a “natural intermediate” [41, Section 5] in terms of model complexity. In Section 5.2, we proceed to derive a systematic parametrization of the manifold $\mathcal{W}(\mathcal{X}, \mathcal{E})$, similar in spirit to the one given in Ito and Amari [17], and in Nagaoka [22, Example 1]. In Section 5.3, we connect our results to general differential geometry, and point out that reversible kernels $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ form a doubly autoparallel submanifold in $\mathcal{W}(\mathcal{X}, \mathcal{E})$. Finally, we conclude with a brief discussion on reversible geodesics (Section 5.4).

5.1 Affine structures

Identifying \mathcal{X} with $[m]$, we can endow the set with the natural order induced from \mathbb{N} . In this section, we will henceforth assume that \mathcal{E} is symmetric, and consider the following subsets of \mathcal{E} ,

$$\begin{aligned} T_-(\mathcal{E}) &\triangleq \{(x, x') \in \mathcal{E} : x' > x\}, \\ T_+(\mathcal{E}) &\triangleq \{(x, x') \in \mathcal{E} : x' < x\}, \\ T_0(\mathcal{E}) &\triangleq \{(x, x') \in \mathcal{E} : x' = x\}, \end{aligned}$$

and

$$T(\mathcal{E}) \triangleq \{(x, x') \in \mathcal{E} : x' \leq x, (x, x') \neq (m, x_\star), x_\star = \arg \min_{x \in \mathcal{X}} \{(m, x) \in \mathcal{E}\}\}.$$

We immediately observe that the following cardinality relations hold

$$\begin{aligned} |T_+(\mathcal{E})| &= |T_-(\mathcal{E})|, \\ |T(\mathcal{E})| &= |T_+(\mathcal{E})| + |T_0(\mathcal{E})| - 1 = \frac{|\mathcal{E}| + |T_0(\mathcal{E})|}{2} - 1, \end{aligned} \quad (8)$$

and that from irreducibility, $x^\star \neq m$. The last expression in (8) highlights the fact that $|T(\mathcal{E})|$ is independent of any ordering of elements of \mathcal{X} . Note also that the element (m, x_\star) in the definition of $T(\mathcal{E})$ plays no special role, and could be replaced with any other element of $T_0(\mathcal{E}) \cup T_+(\mathcal{E})$. We define the sets of symmetric and log-reversible functions (Definition 4) over the graph $(\mathcal{X}, \mathcal{E})$, respectively by

$$\begin{aligned} \mathcal{F}_{\text{sym}}(\mathcal{X}, \mathcal{E}) &\triangleq \{h \in \mathcal{F}(\mathcal{X}, \mathcal{E}) : \forall x, x' \in \mathcal{X}, h(x, x') = h(x', x)\}, \\ \mathcal{F}_{\text{rev}}(\mathcal{X}, \mathcal{E}) &\triangleq \{h \in \mathcal{F}(\mathcal{X}, \mathcal{E}) : h \text{ is log-reversible}\}. \end{aligned}$$

We note that $\mathcal{F}_{\text{sym}}(\mathcal{X}, \mathcal{E})$ is isomorphic to the vector space of symmetric matrices whose entries are null outside of \mathcal{E} , thus $\dim \mathcal{F}_{\text{sym}}(\mathcal{X}, \mathcal{E}) = |T_+(\mathcal{E})| + |T_0(\mathcal{E})|$. We now show that $\mathcal{F}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ is also a vector space, and that it contains $\mathcal{N}(\mathcal{X}, \mathcal{E})$ defined at (4).

Lemma 4 *The following vector subspace inclusions hold:*

$$\mathcal{N}(\mathcal{X}, \mathcal{E}) \stackrel{(i)}{\subset} \mathcal{F}_{\text{rev}}(\mathcal{X}, \mathcal{E}) \stackrel{(ii)}{\subset} \mathcal{F}(\mathcal{X}, \mathcal{E}).$$

Proof To verify (ii), we argue that $\mathcal{F}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ is closed by linear combinations from properties of the sum. The fact that the null function is trivially reversible concludes this claim. For (i), consider an element $h \in \mathcal{N}(\mathcal{X}, \mathcal{E})$, such that $h(x, x') = f(x') - f(x) + c$. Then $h(x, x') = h(x', x) + 2f(x') - 2f(x)$, and from Corollary 1, $h \in \mathcal{F}_{\text{rev}}(\mathcal{X}, \mathcal{E})$, thus the inclusion holds. The set is closed by linear combinations by properties of sums again, and taking $f = 0, c = 0$ is allowed, whence claim (i). \square

Remark 4 In fact, defining

$$\mathcal{N}_0(\mathcal{X}, \mathcal{E}) \triangleq \left\{ h \in \mathcal{F}(\mathcal{X}, \mathcal{E}) : \exists f \in \mathbb{R}^{\mathcal{X}}, \right. \\ \left. \forall (x, x') \in \mathcal{E}, h(x, x') = f(x') - f(x) \right\}, \tag{9}$$

Corollary 1 implies that $\mathcal{F}_{\text{rev}}(\mathcal{X}, \mathcal{E}) = \mathcal{F}_{\text{sym}}(\mathcal{X}, \mathcal{E}) \oplus \mathcal{N}_0(\mathcal{X}, \mathcal{E})$.

It is then possible to further define the quotient space of reversible generator functions

$$\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E}) \triangleq \mathcal{F}_{\text{rev}}(\mathcal{X}, \mathcal{E}) / \mathcal{N}(\mathcal{X}, \mathcal{E}).$$

Theorem 3 *The following statements hold.*

- (i) *The set of reversible generators $\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ can be endowed with a $|T(\mathcal{E})|$ -dimensional vector space structure.*
- (ii) *The set $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ of irreducible and reversible Markov kernels over $(\mathcal{X}, \mathcal{E})$ forms an e -family of dimension $\dim \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E}) = |T(\mathcal{E})|$.*

Proof Let g be a log-reversible function over $(\mathcal{X}, \mathcal{E})$. From Corollary 1, there exists $f \in \mathbb{R}^{\mathcal{X}}$ such that $g(x, x') = g(x', x) + f(x') - f(x)$, or writing $h(x, x') = g(x, x') + \tilde{f}(x) - \tilde{f}(x')$ with $\tilde{f} = f/2$ (i.e. $h = (g + g^T)/2$), it holds that $h(x, x') = h(x', x)$, i.e. h is symmetric. $\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ thus also corresponds to the alternative quotient space

$$\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E}) \cong \mathcal{F}_{\text{sym}}(\mathcal{X}, \mathcal{E}) / \mathbb{R}, \tag{10}$$

and as a consequence $\dim \mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E}) = |T_+(\mathcal{E})| + |T_0(\mathcal{E})| - 1 = |T(\mathcal{E})|$. This concludes the proof of (i). Let $g \in \mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})$, and recall the definition (5) of the diffeomorphism Δ . By Theorem 2, $\Delta(\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})) \subset \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$. Conversely, let $P \in \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$. Then by the Kolmogorov criterion (Theorem 1), $\log[P] \in$

$\mathcal{F}_{\text{rev}}(\mathcal{X}, \mathcal{E})$, and there exist $(g, f, c) \in \mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E}) \times \mathbb{R}^{\mathcal{X}} \times \mathbb{R}$ such that for any $x, x' \in \mathcal{X}$, $\log P(x, x') = g(x, x') + f(x') - f(x) + c$ (where c is unique, f is unique up to an additive constant, and both can be recovered from PF theory). In other words, there exists $g \in \mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})$, with $P = s \circ \exp[g]$, hence $P \in \Delta(\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E}))$, proving that

$$\Delta(\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})) = \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E}).$$

Claim (ii) then follows from Nagaoka [22, Theorem 2], as discussed at the end of Section 2. □

Corollary 2 *For the set of positive Markov kernel, $|T_0(\mathcal{E})| = |\mathcal{X}|$ and $|\mathcal{E}| = |\mathcal{X}|^2$, thus $\dim \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{X}^2) = |\mathcal{X}|(|\mathcal{X}| + 1)/2 - 1$. This is in line with the known number of degrees of freedom of reversible Markov chains [9,41].*

Theorem 4 *The family of functions $g_{ij} = \delta_i^\top \delta_j + \delta_j^\top \delta_i$, for $(i, j) \in T(\mathcal{E})$, forms a basis of $\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})$.*

Proof We begin by proving the independence of the family in the quotient space $\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})$. Since g_{ij} is symmetric in the sense that $g_{ij} = g_{ji}^\top$, it trivially verifies the log-reversibility property, thus belongs to $\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})$. Let now $g \in \mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ be such that

$$g = \sum_{(i,j) \in T(\mathcal{E})} \alpha_{ij} g_{ij},$$

with $\alpha_{ij} \in \mathbb{R}$, for any $(i, j) \in T(\mathcal{E})$, and suppose that $g = 0_{\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})}$. Our first step is to observe that necessarily $g = 0_{\mathcal{F}(\mathcal{X}, \mathcal{E})}$, i.e. g must be the null vector in the ambient space. Let us suppose for contradiction that there exist $(f, c) \in (\mathbb{R}^{\mathcal{X}}, \mathbb{R})$ such that $g(x, x') = f(x') - f(x) + c$ and either $c \neq 0$ or f is not constant over \mathcal{X} . Since by definition, $(m, x_\star), (x_\star, m) \notin T(\mathcal{E})$,

$$\begin{aligned} 0 &= g(m, x_\star) = f(m) - f(x_\star) + c, \\ 0 &= g(x_\star, m) = f(x_\star) - f(m) + c, \end{aligned}$$

therefore summing the latter equalities yields $c = 0$, thus f cannot be constant. But then, g is both symmetric and skew-symmetric, which leads to a contradiction, and $g = 0_{\mathcal{F}(\mathcal{X}, \mathcal{E})}$. Since the family $\{g_{ij} : (i, j) \in T(\mathcal{E})\}$ is independent in the ambient space $\mathcal{F}(\mathcal{X}, \mathcal{E})$, the coefficients $\alpha_{ij}, (i, j) \in T(\mathcal{E})$ must be null, and as result, the family is also linearly independent in $\mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})$. Finally, since from Theorem 3, $|T(\mathcal{E})| = \dim \mathcal{G}_{\text{rev}}(\mathcal{X}, \mathcal{E})$, the family is maximally independent, hence constitutes a basis of the quotient vector space. □

Remark 5 An alternative way of showing the linear independence of the family $\{g_{ij} : (i, j) \in T(\mathcal{E})\}$ in Theorem 4 consists in verifying that (i) the family is independent in \mathcal{F}_{sym} , (ii) $\mathbb{R} \not\subset \text{span} \{g_{ij} : (i, j) \in T(\mathcal{E})\}$, and then invoking (10).

5.2 Parametrization of the manifold of reversible kernels

Recall that from [22, Example 1], in the complete graph case ($\mathcal{E} = \mathcal{X}^2$), we can find an explicit parametrization for $\mathcal{W}(\mathcal{X}, \mathcal{X}^2)$. Indeed, picking any $x_\star \in \mathcal{X}$, we can easily verify that for the two cases where $x' = x_\star$ and $x' \neq x_\star$,

$$\begin{aligned} \log P(x, x') &= \sum_{i=1}^{|\mathcal{X}|} \sum_{\substack{j=1 \\ j \neq x_\star}}^{|\mathcal{X}|} \log \frac{P(i, j)P(j, x_\star)}{P(i, x_\star)P(x_\star, x_\star)} \delta_i(x)\delta_j(x') \\ &\quad + \log P(x, x_\star) - \log P(x', x_\star) + \log P(x_\star, x_\star). \end{aligned}$$

In the remainder of this section, we show how to derive a similar parametrization for $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$. We start by recalling the definition of the expectation parameter of an exponential family of kernels. For an e-family \mathcal{V}_e , following the notation of Definition 2, we define

$$\eta_i(\theta) \triangleq Q_\theta[g_i] = \sum_{x, x' \in \mathcal{X}} Q_\theta(x, x')g_i(x, x'),$$

and call $\eta = (\eta_1, \dots, \eta_d)$ the expectation parameter of the family. We will first derive η and later convert to the natural parameter θ using the following lemma.

Lemma 5 *For a given exponential family, we can express the chart transition maps between the expectation and natural parameters $\theta \circ \eta^{-1}$ and $\eta \circ \theta^{-1}$. Extending the notation at Lemma 3,*

(i)

$$\eta_i(\theta) = Q_\theta[g_i] = \sum_{x, x' \in \mathcal{X}} Q_\theta(x, x')g_i(x, x').$$

(ii)

$$\begin{aligned} \theta^i(\eta) &= \left(\frac{\partial}{\partial \eta_i} Q_\eta \right) [\log P_\eta - K] \\ &= \sum_{x, x' \in \mathcal{X}} \left(\frac{\partial}{\partial \eta_i} Q_\eta(x, x') \right) (\log P_\eta(x, x') - K(x, x')). \end{aligned}$$

In particular, when the carrier kernel verifies $K = 0$, we more simply have

$$\theta^i(\eta) = \left(\frac{\partial}{\partial \eta_i} Q_\eta \right) [\log P_\eta].$$

Proof It is well-known that $\eta_i(\theta) = \frac{\partial}{\partial \theta^i} \psi_\theta = Q_\theta[g_i]$ [10, Lemma 5.1], [22, Theorem 4], [21, (28)], therefore we only need to show (ii). Let g_1, g_2, \dots, g_d be a

collection of independent functions of $\mathcal{G}(\mathcal{X}, \mathcal{E})$. Consider the exponential family as in Definition 2. Recall that for two transition kernels P_1, P_2 respectively irreducible over $(\mathcal{X}, \mathcal{E}_1)$ and $(\mathcal{X}, \mathcal{E}_2)$, and with stationary distributions π_1 and π_2 , the information divergence of P_1 from P_2 is given by

$$D(P_1 || P_2) = \begin{cases} \sum_{(x,x') \in \mathcal{E}_1} \pi_1(x) P_1(x, x') \log \frac{P_1(x, x')}{P_2(x, x')}, & \text{when } \mathcal{E}_1 \subset \mathcal{E}_2, \\ \infty & \text{otherwise.} \end{cases} \tag{11}$$

Writing P_0 for P_θ when $\theta = 0$,

$$\begin{aligned} D(P_\theta || P_0) &= \sum_{x, x' \in \mathcal{X}} Q_\theta(x, x') \log \frac{P_\theta(x, x')}{P_0(x, x')} \\ &= \sum_{x, x' \in \mathcal{X}} Q_\theta(x, x') \left[\sum_{i=1}^d \theta^i g_i(x, x') + R_\theta(x') - R_\theta(x) - \psi_\theta - R_0(x') + R_0(x) + \psi_0 \right] \\ &= \sum_{i=1}^d \theta^i \eta_i - \psi_\theta + \psi_0, \end{aligned}$$

where for the last equality we used (i) of the present lemma and Lemma 3-(iii). Moreover, by a direct computation,

$$Q_\theta [-\log P_0] = \psi_0 - Q_\theta[K].$$

Thus, the potential function is given by

$$\varphi(\eta) \triangleq \sum_{i=1}^d \theta^i \eta_i - \psi_\theta = Q_\theta [\log P_\theta] - Q_\theta[K] = Q_\eta [\log P_\eta] - Q_\eta[K]. \tag{12}$$

By taking the derivative, we recover $\frac{\partial}{\partial \eta_i} \varphi(\eta) = \theta^i(\eta)$ [22, (17)]. Moreover, from (12), we have that

$$\begin{aligned} \frac{\partial}{\partial \eta_i} \varphi(\eta) &= \frac{\partial}{\partial \eta_i} (Q_\eta [\log P_\eta] - Q_\eta[K]) \\ &= \left(\frac{\partial}{\partial \eta_i} Q_\eta \right) [\log P_\eta - K] + Q_\eta \left[\frac{\partial \log P_\eta}{\partial \eta_i} \right] - Q_\eta \left[\frac{\partial K}{\partial \eta_i} \right] \\ &= \left(\frac{\partial}{\partial \eta_i} Q_\eta \right) [\log P_\eta - K], \end{aligned}$$

where for the last equality, we used the fact that $Q_\eta[\partial K/\partial \eta_i] = 0$, and that from P_η being stochastic,

$$\sum_{x,x' \in \mathcal{X}} Q_\eta(x, x') \frac{\partial}{\partial \eta_i} \log P_\eta(x, x') = \sum_{x,x' \in \mathcal{X}} \pi_\eta(x) \frac{\partial}{\partial \eta_i} P_\eta(x, x') = 0.$$

This finishes proving (ii) of the lemma. □

Theorem 5 *Let $P \in \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$, with stationary distribution π . Using the basis $g_{ij} = \delta_i^\top \delta_j + \delta_j^\top \delta_i$, we can write Q , the edge measure matrix associated with P , as a member of the m -family of reversible kernels,*

$$Q = \frac{g_\star}{2} + \sum_{(i,j) \in T(\mathcal{E})} (g_{ij} - g_\star) \frac{Q(i, j)}{1 + \delta_i(j)},$$

where $g_\star = \delta_m^\top \delta_{x_\star} + \delta_{x_\star}^\top \delta_m$, and we can write P as a member of the e -family,

$$\begin{aligned} \log P(x, x') &= \sum_{(i,j) \in T(\mathcal{E})} \frac{1}{2} \log \frac{P(i, j)P(j, i)}{P(m, x_\star)P(x_\star, m)} \left(\frac{g_{ij}(x, x')}{1 + \delta_i(j)} \right) \\ &\quad + \frac{1}{2} \log \pi(x') - \frac{1}{2} \log \pi(x) + \frac{1}{2} \log P(m, x_\star)P(x_\star, m), \end{aligned}$$

when $(x, x') \in \mathcal{E}$, $P(x, x') = 0$ otherwise, and where $x_\star = \arg \min_{x \in \mathcal{X}} \{(m, x) \in \mathcal{E}\}$.

Proof Let us consider the basis

$$g_{ij} = \delta_i^\top \delta_j + \delta_j^\top \delta_i,$$

and taking $K = 0$, we are looking for a parametrization of the type

$$\begin{aligned} \tilde{P}_\theta(x, x') &= \exp \left(\sum_{(i,j) \in T(\mathcal{E})} \theta^{ij} g_{ij}(x, x') \right), \\ P_\theta(x, x') &= \mathfrak{s}(\tilde{P}_\theta)(x, x') = \tilde{P}_\theta(x, x') \exp(R_\theta(x') - R_\theta(x) - \psi_\theta), \end{aligned}$$

where $\exp \psi_\theta$ and $\exp[R_\theta]$ are respectively the PF root and right PF eigenvector of \tilde{P}_θ . We first derive a parametrization of the edge measure Q_η as a member of an m -family (following Definition 1-(ii) with respect to the expectation parameter η). For $(i, j) \in \mathcal{X}^2$, by Lemma 5-(i),

$$\eta_{ij} = Q_\eta[g_{ij}] = Q_\eta(i, j) + Q_\eta(j, i) = 2Q_\eta(i, j),$$

and thus, from symmetry of Q_η and since $Q_\eta \in \mathcal{P}(\mathcal{X}^2)$,

$$Q_\eta(i, j) = \begin{cases} 0 & \text{when } (i, j) \notin \mathcal{E} \\ \eta_{ij}/2 & \text{when } (i, j) \in T(\mathcal{E}), i \neq j \\ \eta_{ji}/2 & \text{when } (j, i) \in T(\mathcal{E}), i \neq j \\ \eta_{ii}/2 & \text{when } (i, i) \in T_0(\mathcal{E}) \\ \frac{1}{2} \left(1 - \sum_{(i,j) \in T(\mathcal{E})} \frac{\eta_{ij}}{1 + \delta_i(j)} \right) & \text{when } (i, j) \in \{(m, x_\star), (x_\star, m)\}, \end{cases}$$

and more compactly, for $(x, x') \in \mathcal{X}^2$,

$$Q_\eta = \frac{g_\star}{2} + \sum_{(i,j) \in T(\mathcal{E})} \frac{g_{ij} - g_\star}{2(1 + \delta_i(j))} \eta_{ij},$$

where g_\star is defined as in the statement of the theorem. We differentiate by η_{ij} for $(i, j) \in T(\mathcal{E})$, to obtain

$$\frac{\partial Q_\eta}{\partial \eta_{ij}} = \frac{g_{ij} - g_\star}{2(1 + \delta_i(j))}.$$

Invoking (ii) of Lemma 5, we convert the expectation parametrization to a natural one,

$$\theta^{ij}(\eta) = \sum_{(x,x') \in \mathcal{X}} \left(\frac{\partial}{\partial \eta_{ij}} Q_\eta(x, x') \right) \log P(x, x'),$$

so that

$$\begin{aligned} \theta^{ij}(\eta) &= \frac{1}{1 + \delta_i(j)} \sum_{(x,x') \in \mathcal{X}} \left(\frac{g_{ij}(x, x') - g_\star(x, x')}{2} \right) \log P(x, x') \\ &= \frac{1}{2(1 + \delta_i(j))} \log \frac{P(i, j)P(j, i)}{P(m, x_\star)P(x_\star, m)}. \end{aligned}$$

Notice that $\tilde{P}_\theta = \tilde{P}_\theta^\top$, hence the right and left PF eigenvector are identical, i.e. $R_\theta = L_\theta$ and as is known (see (6)), the stationary distribution is given by $\pi_\theta = \exp[2R_\theta] / \sum_{x \in \mathcal{X}} \exp(2R_\theta(x))$. In fact, we can easily verify that the right PF eigenvector is given by $\exp[R_\theta] = \sqrt{\pi}$, and that the PF root is

$$\exp \psi_\theta = (P(m, x_\star)P(x_\star, m))^{-1/2}.$$

Indeed, letting $x \in \mathcal{X}$, from detailed balance of P , we have

$$\sum_{x' \in \mathcal{X}} \tilde{P}_\theta(x, x') \sqrt{\pi(x')} = \sum_{x' \in \mathcal{X}} \sqrt{\frac{P(x, x')P(x', x)}{P(m, x_\star)P(x_\star, m)}} = \frac{\sqrt{\pi(x)}}{\sqrt{P(m, x_\star)P(x_\star, m)}}.$$

□

5.3 The doubly autoparallel submanifold of reversible kernels

Recall that we can view $\mathcal{W} = \mathcal{W}(\mathcal{X}, \mathcal{E})$ as a smooth manifold of dimension $d = \dim \mathcal{W} = |\mathcal{E}| - |\mathcal{X}|$. For each $P \in \mathcal{W}$, we can then consider the tangent plane T_P at P , endowed with a d -dimensional vector space structure. Together with the manifold, we define an information geometric structure consisting of a Riemannian metric, called the *Fisher information metric* \mathfrak{g} , and a pair of *torsion-free affine connections* $\nabla^{(e)}$ and $\nabla^{(m)}$ respectively called *e-connection* and *m-connection*, that are dual with respect to \mathfrak{g} , i.e. for any vector fields $X, Y, Z \in \Gamma(T\mathcal{W})$,

$$X\mathfrak{g}(Y, Z) = \mathfrak{g}(\nabla_X^{(e)} Y, Z) + \mathfrak{g}(Y, \nabla_X^{(m)} Z),$$

where $\Gamma(T\mathcal{W})$ is the set of all sections over the tangent bundle. We now review an explicit construction for $\mathfrak{g}, \nabla^{(m)}, \nabla^{(e)}$.

Construction in the natural chart map.

Consider a parametric family $\mathcal{V} = \{P_\theta : \theta \in \Theta\}$ with Θ open subset of \mathbb{R}^d . For any $n \in \mathbb{N}$, we define the path measure $Q_\theta^{(n)} \in \mathcal{P}(\mathcal{X}^n)$ induced from the kernel P_θ .

$$Q_\theta^{(n)}(x_1, x_2, \dots, x_n) = \pi_\theta(x_1) \prod_{t=1}^{n-1} P_\theta(x_t, x_{t+1}).$$

Nagaoka [22] defines the Fisher metric as

$$\begin{aligned} \mathfrak{g}_{ij}(\theta) &\triangleq \sum_{(x,x') \in \mathcal{E}} Q_\theta(x, x') \partial_i \log P_\theta(x, x') \partial_j \log P_\theta(x, x'), \\ &= \sum_{(x,x') \in \mathcal{E}} \partial_i \log P_\theta(x, x') \partial_j Q_\theta(x, x'), \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \mathfrak{g}_{ij}^n(\theta), \end{aligned}$$

and the dual affine e/m-connections of $\{P_\theta : \theta \in \Theta\}$ by their Christoffel symbols,

$$\begin{aligned} \Gamma_{ij,k}^{(e)}(\theta) &\triangleq \sum_{(x,x') \in \mathcal{E}} \partial_i \partial_j \log P_\theta(x, x') \partial_k Q_\theta(x, x') = \lim_{n \rightarrow \infty} \frac{1}{n} \Gamma_{ij,k}^{(e),n}(\theta), \\ \Gamma_{ij,k}^{(m)}(\theta) &\triangleq \sum_{(x,x') \in \mathcal{E}} \partial_i \partial_j Q_\theta(x, x') \partial_k \log P_\theta(x, x') = \lim_{n \rightarrow \infty} \frac{1}{n} \Gamma_{ij,k}^{(m),n}(\theta), \end{aligned}$$

where $\mathfrak{g}_{ij}^n(\theta), \Gamma_{ij,k}^{(e),n}(\theta), \Gamma_{ij,k}^{(m),n}(\theta)$ are the Fisher metric, and Christoffel symbols of the e/m-connections that pertain to the distribution family $\left\{ Q_\theta^{(n)} \right\}_{\theta \in \Theta}$.

Autoparallelity.

Connections allow us to talk about covariant derivatives and parallelity of vectors fields.

Definition 5 A submanifold \mathcal{V} is called autoparallel in \mathcal{W} with respect to a connection ∇ , when for any vector fields $\forall X, Y \in \Gamma(T\mathcal{V})$, it holds that

$$\nabla_X Y \in \Gamma(T\mathcal{V}).$$

A submanifold \mathcal{V} of \mathcal{W} is then an e-family (resp. m-family) if and only if it is autoparallel with respect to $\nabla^{(e)}$ (resp. $\nabla^{(m)}$) [22, Theorem 6]. As the manifold of reversible kernels is both an e-family and an m-family, it is called doubly autoparallel [42, Definition 1].

Theorem 6 *The manifold $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ of irreducible and reversible Markov chains over $(\mathcal{X}, \mathcal{E})$ is a doubly autoparallel submanifold in $\mathcal{W}(\mathcal{X}, \mathcal{E})$ with dimension*

$$\dim \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E}) = \frac{|\mathcal{E}| + |T_0(\mathcal{E})|}{2} - 1,$$

where $T_0(\mathcal{E}) = \{(x, x') \in \mathcal{E} : x = x'\}$.

Proof The set of reversible Markov chains is an e-family (Theorem 3), and an m-family (Theorem 5). □

5.4 Reversible geodesics

In this section, we let two irreducible reversible kernels P_0 and P_1 over $(\mathcal{X}, \mathcal{E})$, and discuss the geodesics that connect them with respect to $\nabla^{(e)}$ and $\nabla^{(m)}$. Although already guaranteed (see for example Ohara and Ishi [42, Proposition 1]), we offer alternative elementary proofs that any kernel lying on either e/m-geodesic is irreducible and reversible.

m-geodesics.

By irreducibility, there exist unique $Q_0, Q_1 \in \mathcal{Q}(\mathcal{X}, \mathcal{E})$ corresponding to P_0, P_1 . Moreover, by reversibility Q_0 and Q_1 are symmetric. We let

$$G_m(P_0, P_1) \triangleq \{P_\xi : Q_\xi = \xi Q_1 + (1 - \xi)Q_0 : \xi \in [0, 1]\}, \tag{13}$$

be the m-geodesic (auto-parallel curve with respect to the m-connection) connecting P_0 and P_1 . Then $G_m(P_0, P_1)$ forms an m-family of dimension 1. For any $\xi \in [0, 1]$, the matrix Q_ξ is symmetric as convex combination of two symmetric matrices. Q_ξ takes value 0 exactly when Q_0, Q_1 , i.e. P_0, P_1 take value 0. Furthermore, writing π_0 (resp. π_1) the unique stationary distribution of P_0 (resp. P_1),

$$\begin{aligned} \sum_{x'} Q_\xi(x, x') &= \xi \pi_1(x) + (1 - \xi)\pi_0(x), \\ \sum_x Q_\xi(x, x') &= \xi \pi_1(x') + (1 - \xi)\pi_0(x'), \end{aligned}$$

thus Q_ξ always defines a proper associated stochastic irreducible stochastic P_ξ . *e-geodesics.*

We consider the auto-parallel curve with respect to the e-connection that connect P_0 and P_1 ,

$$G_e(P_0, P_1) \triangleq \left\{ P_0(x, x') \exp\left(\theta \ln \frac{P_1(x, x')}{P_0(x, x')} + R_\theta(x') - R_\theta(x) - \psi_\theta\right) : \theta \in [0, 1] \right\}. \tag{14}$$

The set $G_e(P_0, P_1)$ forms an e-family of dimension 1. Indeed, from Theorem 2, and since P_0 and P_1 are reversible by hypothesis, it suffices to verify that $(x, x') \mapsto P_1(x, x')/P_0(x, x')$ is a reversible function over $(\mathcal{X}, \mathcal{E})$. This follows from a simple application of the Kolmogorov criterion (Theorem 1).

6 Reversible information projections

Reversible Markov kernels, as self-adjoint linear operators, enjoy a set of powerful yet brittle spectral properties. The eigenvalues are real, the second largest in magnitude controls the time to stationarity of the Markov process [8, Chapter 12], and all are stable under perturbation and estimation [43]. However, any deviation from reversibility carries steep consequences, as the spectrum can suddenly become complex, and partially loses control over the mixing time. Furthermore, eigenvalue perturbation results that were dimensionless [44, Corollary 4.10 (Weyl’s inequality)] now come at a cost possibly exponential in the dimension [44, Theorem 1.4 (Ostrowski-Elsner)]. For some irreducible P with stationary distribution π , it is therefore interesting to find the *closest* representative that is reversible, so as to enable Hilbert space techniques. Computing the closest reversible transition kernel with respect to a norm induced from an inner product was considered in Nielsen and Weber [45], who showed that the problem reduces to solving a convex minimization problem with a unique solution.

In this section, we examine this problem under a different notion of distance. We consider information projections onto the reversible family of transition kernels $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$, for some symmetric edge set \mathcal{E} . We define the m-projection and the e-projection of P onto the set of reversible transition kernels $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ respectively as

$$P_m \triangleq \arg \min_{\bar{P} \in \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})} D(P \parallel \bar{P}), \quad P_e \triangleq \arg \min_{\bar{P} \in \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})} D(\bar{P} \parallel P),$$

where $D(\cdot \parallel \cdot)$ is the informational divergence, that was defined at (11). These two generally distinct projections (D is not symmetric in its arguments) correspond to the closest reversible chains when considering information divergence as a measure of distance. Under a careful choice of the connection graph of the reversible family, we derive closed-form expressions for P_m and P_e , along with Pythagorean identities, as illustrated in Figure 1.




Fig. 1 Information projections P_e and P_m of P onto $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{X}^2)$ in the full support case ($\mathcal{E} = \mathcal{X}^2$) (Theorem 7), Pythagorean identities (Theorem 7), and the bisection property (Proposition 2)

Theorem 7 Let P be irreducible over $(\mathcal{X}, \mathcal{E})$.

m-projection.

The *m-projection* P_m of P onto $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E} \cup \mathcal{E}^*)$ is given by

$$P_m = \frac{P + P^*}{2}.$$

Moreover, for any $\bar{P} \in \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E} \cup \mathcal{E}^*)$, P_m satisfies the following Pythagorean identity.

$$D(P || \bar{P}) = D(P || P_m) + D(P_m || \bar{P}).$$

e-projection. When $\mathcal{E} \cap \mathcal{E}^*$ is a strongly connected directed graph, the *e-projection* P_e of P onto $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E} \cap \mathcal{E}^*)$ is given by

$$P_e = \mathfrak{s}(\tilde{P}_e), \quad \text{with } \tilde{P}_e(x, x') = \sqrt{P(x, x')P^*(x, x')},$$

and where \mathfrak{s} is the stochastic rescaling mapping defined at (3). Moreover, for any $\bar{P} \in \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E} \cap \mathcal{E}^*)$, P_e satisfies the following Pythagorean identity.

$$D(\bar{P} || P) = D(\bar{P} || P_e) + D(P_e || P).$$

Proof Our first order of business is to show that P_m and P_e belong respectively to $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E} \cup \mathcal{E}^*)$ and $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E} \cap \mathcal{E}^*)$. It is easy to see that $P_m(x, x') > 0$ exactly when (x, x') or (x', x) belongs to \mathcal{E} , hence $P_m \in \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E} \cup \mathcal{E}^*)$, and that $P_e(x, x') > 0$ whenever (x, x') belongs to both \mathcal{E} and \mathcal{E}^* . Moreover, since the time-reversal operation

preserves the stationary distribution of an irreducible chain, P_m has the same stationary distribution $\pi_m = \pi$, and a straightforward computation shows that P_m satisfies the detailed balance equation. To prove reversibility of P_e , we rewrite

$$\begin{aligned} \log P_e(x, x') &= \frac{1}{2} \log P(x, x')P(x', x) - \log \rho(\tilde{P}_e) \\ &\quad + \log \left(\sqrt{\pi(x')}v_e(x') \right) - \log \left(\sqrt{\pi(x)}v_e(x) \right). \end{aligned}$$

From Corollary 1, $\log[P_e] \in \mathcal{F}_{\text{rev}}(\mathcal{X}, \mathcal{E})$, thus $P_e \in \mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$.

To prove optimality of P_m , it suffices to verify the following Pythagorean identity

$$D(P || \bar{P}) = D(P || P_m) + D(P_m || \bar{P}).$$

Writing $Q_m = \text{diag}(\pi)P_m$, notice that $P_m = (P + P^*)/2$ is equivalent to $Q_m = (Q + Q^*)/2$. We then have

$$\begin{aligned} &D(P || P_m) + D(P_m || \bar{P}) - D(P || \bar{P}) \\ &= \sum_{x, x' \in \mathcal{X}} \left(Q(x, x') \log \frac{P(x, x')}{P_m(x, x')} + Q_m(x, x') \log \frac{P_m(x, x')}{\bar{P}(x, x')} \right. \\ &\quad \left. - Q(x, x') \log \frac{P(x, x')}{\bar{P}(x, x')} \right) \\ &= \sum_{x, x' \in \mathcal{X}} (Q_m(x, x') - Q(x, x')) \log \frac{P_m(x, x')}{\bar{P}(x, x')} \\ &= \sum_{x, x' \in \mathcal{X}} \left(\frac{Q^*(x, x') - Q(x, x')}{2} \right) \log \frac{P_m(x, x')}{\bar{P}(x, x')} \\ &= \frac{1}{2} Q^* [\log(P_m/\bar{P})] - \frac{1}{2} Q [\log(P_m/\bar{P})] = 0, \end{aligned}$$

where the last equality stems from (i) of Lemma 3 and reversibility of P_m and \bar{P} . Similarly, to prove optimality of P_e , it suffices to verify that

$$D(\bar{P} || P) = D(\bar{P} || P_e) + D(P_e || P).$$

By reorganizing terms

$$D(\bar{P} || P_e) + D(P_e || P) - D(\bar{P} || P) = \bar{Q} [\log(P/P_e)] - Q_e [\log(P/P_e)].$$

From the definition of $P_e(x, x')$,

$$\log \frac{P(x, x')}{P_e(x, x')} = \frac{1}{2} \log \frac{P(x, x')}{P(x', x)} + \frac{1}{2} \log \frac{\pi(x)}{\pi(x')} + \log \frac{v_e(x)}{v_e(x')} + \log \rho(\tilde{P}_e).$$

The first three terms being skew-symmetric, reversibility of \tilde{P} and (ii) of Lemma 3 yield that

$$\tilde{Q} [\log(P/P_e)] = \log \rho(\tilde{P}_e).$$

By a similar argument, $Q_e [\log(P/P_e)] = \log \rho(\tilde{P}_e)$, which concludes the proof. \square

In other words, the m-projection is given by the natural *additive reversibilization* [46, (2.4)] of P , while the e-projection is achieved by some newly defined *exponential reversibilization* of P .

The difference between the m-projection and the e-projection is illustrated in the following example.

Example 4 Let us consider the family of biased lazy random walks $P_\theta = P_{(\theta_1, \theta_2)}$, given in Example 1. Note that $\mathcal{E} = \mathcal{E}^*$. The m-projection P_m of P_θ onto $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ is the unbiased lazy random walk given by $P_m = P_{(\theta', 0)}$ with $\theta' = \theta_1 - \log \cosh \theta_2$, i.e.

$$P_m(x, x) = \frac{e^{\theta_1}}{e^{\theta_1} + e^{\theta_2} + e^{-\theta_2}},$$

$$P_m(x, x + 1) = P_m(x + 1, x) = \frac{e^{\theta_2} + e^{-\theta_2}}{2(e^{\theta_1} + e^{\theta_2} + e^{-\theta_2})}.$$

On the other hand, the e-projection P_e of P_θ onto $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ is the unbiased lazy random walk given by $P_e = P_{(\theta_1, 0)}$, i.e.

$$P_e(x, x) = \frac{e^{\theta_1}}{e^{\theta_1} + 2}, \quad P_e(x, x + 1) = P_e(x + 1, x) = \frac{1}{e^{\theta_1} + 2}.$$

Remark 6 We observe that, although the m-projection preserves the stationary distribution, this is not true for P_e , which exhibits a stationary distribution π_e generally different from π . Furthermore, while the solution for the m-projection is always properly defined by taking union of the edge sets, our expression for the e-projection requires additional constraints on the connection graph of P . Indeed, taking the intersection $\mathcal{E} \cap \mathcal{E}^*$, we always obtain a symmetric set, but can lose strong connectedness. We note but do not pursue the fact that reversibility can be defined for the less well-behaved set of reducible chains. In this case, π need not be unique, or could take null values, and the kernel could have a complex spectrum.

Finally, we show that for any irreducible P , both its reversible projections P_m and P_e are equidistant from P and its time-reversal P^* (see also Fig. 1).

Proposition 2 (Bisection property) *Let P irreducible, and let P_m (resp. P_e) the m-projection (resp. e-projection) of P onto $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$.*

$$D(P||P_m) = D(P^*||P_m), \quad D(P_e||P) = D(P_e||P^*).$$

Proof For P_1 irreducible over $(\mathcal{X}, \mathcal{E}_1)$ and P_2 irreducible over $(\mathcal{X}, \mathcal{E}_2)$, it is easy to see that

$$\mathcal{E}_1 \subset \mathcal{E}_2 \implies D(P_1 || P_2) = D(P_1^* || P_2^*).$$

Then take $P_2 = P_m$ for the first equality, and $P_1 = P_e$ for the second. □

7 The e-family of reversible edge measures

Recall that $\mathcal{P}(\mathcal{X}^2)$, the set of all distributions over \mathcal{X}^2 , forms an e-family [27, Example 2.8]. For some e-family of irreducible transition kernels $\mathcal{V}_e \subset \mathcal{W}(\mathcal{X}, \mathcal{X}^2)$, one may wonder whether the corresponding family of edge measures also forms an e-family of distributions in $\mathcal{P}(\mathcal{X}^2)$. We begin by illustrating that this holds in particular for the e-family obtained by tilting a memoryless Markov kernel.

Example 5 Consider the degenerate Markov kernel corresponding to an iid process $P(x, x') = \pi(x')$ for $\pi \in \mathcal{P}(\mathcal{X})$. For a given function $g: \mathcal{X} \rightarrow \mathbb{R}$, and $\theta \in \mathbb{R}$, construct $\tilde{P}_\theta(x, x') = P(x, x')e^{\theta g(x')} = \pi(x')e^{\theta g(x')}$. Then $v_\theta = \mathbf{1}$ is right eigenvector of \tilde{P}_θ with eigenvalue $\rho(\theta) = \sum_{x' \in \mathcal{X}} \pi(x')e^{\theta g(x')}$. Letting $\pi_\theta(x) = \pi(x)e^{\theta g(x)}/\rho(\theta)$, we see that π_θ is the left PF eigenvector of \tilde{P}_θ , and the stationary distribution of the rescaled P_θ . We can therefore write,

$$Q_\theta(x, x') = \exp(\log \pi(x)\pi(x') + \theta(g(x) + g(x')) - 2 \log \rho(\theta)),$$

thus $\{Q_\theta\}_{\theta \in \Theta}$ forms an exponential family of distributions over \mathcal{X}^2 . This fact can be further understood in the following manner. An e-family of distributions $\{\pi_\theta\}_\theta$ induces an e-family of memoryless Markov kernels $\{P_\theta\}_\theta$ with $P_\theta(x, x') = \pi_\theta(x')$ (see Lemma 7 for a proof of this fact for the set of all memoryless kernels), and thus with edge measures $Q_\theta(x, x') = \pi_\theta(x)\pi_\theta(x')$. Since the 2-iid extension $\{\pi_\theta(x)\pi_\theta(x')\}_\theta$ of the e-family $\{\pi_\theta\}_\theta$ is also an e-family, it follows that $\{Q_\theta(x, x')\}_\theta$ forms an e-family.

In the remainder of this section, we show that the subset of positive reversible edge measures $\mathcal{Q}_{\text{rev}} = \mathcal{Q}_{\text{rev}}(\mathcal{X}, \mathcal{X}^2)$, induced from the e-family of reversible positive kernels, forms a submanifold of $\mathcal{P}(\mathcal{X}^2)$ that is autoparallel with respect to the e-connection, i.e. \mathcal{Q}_{rev} is an e-family of distribution of over pairs. Our proof will rely on the definition of a *Markov map*.

Definition 6 (e.g. Nagaoka [47]) We say that $M: \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{Y})$ is a Markov map, when there exists a transition kernel P_M from \mathcal{X} to \mathcal{Y} (also called a channel) such that for any $\mu \in \mathcal{P}(\mathcal{X})$,

$$M(\mu) = \sum_{x \in \mathcal{X}} P_M(x, \cdot)\mu(x).$$

Let \mathcal{U} and \mathcal{V} be smooth submanifolds (statistical models) of $\mathcal{P}(\mathcal{X})$ and $\mathcal{P}(\mathcal{Y})$ respectively. When there exists a pair of Markov maps $M: \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{Y})$, $N: \mathcal{P}(\mathcal{Y}) \rightarrow$

$\mathcal{P}(\mathcal{X})$ such that their restrictions $M|_{\mathcal{U}}, N|_{\mathcal{V}}$ are bijections between \mathcal{U} and \mathcal{V} , and are the inverse mappings of each other, we say that \mathcal{U} and \mathcal{V} are Markov equivalent, and write $\mathcal{U} \cong \mathcal{V}$.

Lemma 6 *It holds that*

$$\mathcal{Q}_{\text{rev}}(\mathcal{X}, \mathcal{X}^2) \cong \mathcal{P} \left(\left[\frac{|\mathcal{X}| (|\mathcal{X}| + 1)}{2} \right] \right).$$

Proof Identify $\mathcal{X} = [m]$, and consider $Q \in \mathcal{Q}_{\text{rev}}$ such that

$$Q = \begin{pmatrix} \eta_{11} & \frac{\eta_{12}}{2} & \frac{\eta_{13}}{2} & \cdots & \frac{\eta_{1m}}{2} \\ \frac{\eta_{12}}{2} & \eta_{22} & \frac{\eta_{23}}{2} & \cdots & \frac{\eta_{2m}}{2} \\ \vdots & & \ddots & & \vdots \\ \vdots & & & \eta_{(m-1)(m-1)} & \frac{\eta_{(m-1)m}}{2} \\ \frac{\eta_{1m}}{2} & \cdots & \frac{\eta_{23}}{2} & \cdots & \eta_{mm} \end{pmatrix}$$

and where $\eta_{mm} = 1 - \sum_{i \leq j, (i,j) \neq (m,m)} \eta_{ij}$. We flatten the definition of Q .

$$Q = \left(\eta_{11}, \eta_{22}, \dots, \eta_{mm}, \frac{\eta_{12}}{2}, \frac{\eta_{12}}{2}, \dots, \underbrace{\frac{\eta_{ij}}{2}, \frac{\eta_{ij}}{2}}_{(i,j): i < j}, \dots, \frac{\eta_{(m-1)m}}{2}, \frac{\eta_{(m-1)m}}{2} \right).$$

Let the matrix E with $m(m - 1)/2$ columns and $m(m - 1)$ rows be such that,

$$E^T = \begin{pmatrix} 1 & 1 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \dots & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & 0 & 1 & 1 \end{pmatrix}.$$

Block matrix multiplication yields

$$Q \begin{pmatrix} I_m & 0 \\ 0 & E \end{pmatrix} = (\eta_{11}, \eta_{22}, \dots, \eta_{mm}, \eta_{12}, \dots, \eta_{ij}, \dots, \eta_{(m-1)m}) \in \mathcal{P} \left(\left[\frac{m(m + 1)}{2} \right] \right),$$

and further observing that for $F = \frac{1}{2} E^T$, it holds that $FE = I_{m(m-1)/2}$. Thus the mappings defined by $\begin{pmatrix} I_m & 0 \\ 0 & E \end{pmatrix}$ and $\begin{pmatrix} I_m & 0 \\ 0 & \frac{1}{2} E^T \end{pmatrix}$ are Markov maps and verify $\begin{pmatrix} I_m & 0 \\ 0 & \frac{1}{2} E^T \end{pmatrix} \begin{pmatrix} I_m & 0 \\ 0 & E \end{pmatrix} = I_{m(m+1)/2}$. This finishes proving the claim. □

Theorem 8 *The set \mathcal{Q}_{rev} forms an e-family and an m-family of $\mathcal{P}(\mathcal{X}^2)$ with dimension $|\mathcal{X}| (|\mathcal{X}| + 1)/2 - 1$. Moreover, \mathcal{Q} does not form an e-family in $\mathcal{P}(\mathcal{X}^2)$ (except when $|\mathcal{X}| = 2$).*

Proof Since $\mathcal{Q}_{\text{rev}} \subset \mathcal{P}(\mathcal{X}^2)$, the claim stems from the equivalence between (i) and (ii) of Nagaoka [47, Theorem 1], and application of Lemma 6, and the fact that $\dim \mathcal{P} \left(\left[\frac{|\mathcal{X}|(|\mathcal{X}|+1)}{2} \right] \right) = \frac{|\mathcal{X}|(|\mathcal{X}|+1)}{2} - 1$. In order to prove that \mathcal{Q} is not an e-family in $\mathcal{P}(\mathcal{X}^2)$, we first construct the following family of edge measures over three states.

$$\mathcal{Q}_0^{(3)} = \frac{1}{13} \begin{pmatrix} 1 & 1 & 2 \\ 2 & 1 & 2 \\ 1 & 3 & 1 \end{pmatrix}, \quad \mathcal{Q}_1^{(3)} = \frac{1}{13} \begin{pmatrix} 1 & 3 & 1 \\ 2 & 1 & 2 \\ 2 & 1 & 1 \end{pmatrix}.$$

Computing the point on the e-geodesic in $\mathcal{P}(\mathcal{X}^2)$ at parameter value 1/2, yields

$$\mathcal{Q}_{1/2}^{(3)} \propto \begin{pmatrix} 1 & \sqrt{3} & \sqrt{2} \\ 2 & 1 & 2 \\ \sqrt{2} & \sqrt{3} & 1 \end{pmatrix},$$

which does not belong to \mathcal{Q} . We can readily expand the above example to general state space size, $m > 3$, by considering the one-padded versions of the above $\mathcal{Q}_i^{(m)} \propto \begin{pmatrix} \mathcal{Q}_i^{(3)} & \mathbf{1}_3^\top \mathbf{1}_{m-3} \\ \mathbf{1}_{m-3}^\top \mathbf{1}_3 & \mathbf{1}_{m-3}^\top \mathbf{1}_{m-3} \end{pmatrix}$, for $i \in \{0, 1\}$. □

- Remark 7** (i) Nagaoka [47, Theorem 1-(iv)], actually proves the stronger result that \mathcal{Q}_{rev} forms an α -family in $\mathcal{P}(\mathcal{X}^2)$, for any $\alpha \in \mathbb{R}$ (see Amari and Nagaoka [27, Section 2.6] for a definition of α -families).
 (ii) We note but do not pursue here the fact that a more refined treatment over some irreducible edge set $\mathcal{E} \subsetneq \mathcal{X}^2$ is possible.

8 Comparison of remarkable families of Markov chains

We briefly compare the geometric properties of reversible kernels with that of several other remarkable families of Markov chains, and compile a summary in Table 1.

Family of all kernels irreducible over $(\mathcal{X}, \mathcal{E})$: $\mathcal{W}(\mathcal{X}, \mathcal{E})$. This family is known to form both an e-family and an m-family of dimension $|\mathcal{E}| - |\mathcal{X}|$ [22, Corollary 1].

Family of all reversible kernels irreducible over $(\mathcal{X}, \mathcal{E})$: $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$. We show in Theorem 3 and Theorem 6 that $\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$ is both an e-family and m-family of dimension $T(\mathcal{E})$, where

$$|T(\mathcal{E})| = \frac{|\mathcal{E}| + |T_0(\mathcal{E})|}{2} - 1,$$

with $T_0(\mathcal{E}) = \{(x, x') \in \mathcal{E} : x = x'\}$.

Family of positive memoryless (iid) kernels: $\mathcal{W}_{\text{iid}}(\mathcal{X}, \mathcal{X}^2)$. This family comprises degenerate irreducible kernels that correspond to iid processes, i.e. where all rows are equal to the stationary distribution. Notice that for $P \in \mathcal{W}_{\text{iid}}$, irreducibility forces P

Table 1 Summary of geometric properties of submanifolds of irreducible Markov kernels ($|\mathcal{X}| \geq 3$)

Manifold	m-family	e-family	Dimension	Parametrization
$\mathcal{W}(\mathcal{X}, \mathcal{E})$	✓	✓	$ \mathcal{E} - \mathcal{X} $	
$\mathcal{W}(\mathcal{X}, \mathcal{X}^2)$	✓	✓	$ \mathcal{X} (\mathcal{X} - 1)$	[22, Example 1]
$\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{E})$	✓	✓	$ T(\mathcal{E}) $	Theorem 5
$\mathcal{W}_{\text{rev}}(\mathcal{X}, \mathcal{X}^2)$	✓	✓	$ \mathcal{X} (\mathcal{X} + 1)/2 - 1$	Theorem 5
$\mathcal{W}_{\text{bis}}(\mathcal{X}, \mathcal{X}^2)$	✓	✗	$(\mathcal{X} - 1)^2$	[10, Example 4]
$\mathcal{W}_{\text{sym}}(\mathcal{X}, \mathcal{X}^2)$	✓	✗	$ \mathcal{X} (\mathcal{X} - 1)/2$	(16)
$\mathcal{W}_{\text{iid}}(\mathcal{X}, \mathcal{X}^2)$	✗	✓	$ \mathcal{X} - 1$	(15)
$G_m(P_0, P_1)$	✓	✗	1	(13)
$G_e(P_0, P_1)$	✗	✓	1	(14)

We also include, for completeness, the one dimensional manifolds defined by the e-geodesic $G_e(P_0, P_1)$ and m-geodesic $G_m(P_0, P_1)$ between two irreducible kernels P_0 and P_1 , as defined in Section 5.4. Note that for the binary case $|\mathcal{X}| = 2$, $\mathcal{W}_{\text{sym}} = \mathcal{W}_{\text{bis}}$ forms an e-family. Whenever available, the last column offers a reference to some explicit parametrization of the family

to be positive. We show that \mathcal{W}_{iid} is an e-family of dimension $|\mathcal{X}| - 1$ (Lemma 7), but not an m-family (Lemma 8).

Lemma 7 \mathcal{W}_{iid} forms an e-family of dimension $|\mathcal{X}| - 1$.

Proof For $\mathcal{X} = [m]$, let us consider the following parametrization proposed by Ito and Amari [17]:

$$\begin{aligned} \log P(x, x') &= \sum_{i=1}^{m-1} \log \frac{P(m, i)P(i, m)}{P(m, m)P(m, m)} \delta_i(x') \\ &\quad + \sum_{i=1}^{m-1} \sum_{j=1}^{m-1} \log \frac{P(i, j)P(m, m)}{P(m, j)P(i, m)} \delta_i(x) \delta_j(x') \\ &\quad + \log P(x, m) - \log P(x', m) + \log P(m, m). \end{aligned}$$

This corresponds to the basis

$$\begin{aligned} g_i &= \mathbf{1}^\top \delta_i, & i \in [m - 1], \\ g_{ij} &= \delta_i^\top \delta_j, & i, j \in [m - 1], \end{aligned}$$

with parameters

$$\theta^i = \log \frac{P(m, i)P(i, m)}{P(m, m)P(m, m)}, \quad \theta^{ij} = \log \frac{P(i, j)P(m, m)}{P(m, j)P(i, m)}.$$

Let P irreducible with stationary distribution π . Suppose first that P is memoryless, i.e. for all $x, x' \in \mathcal{X}$, $P(x, x') = \pi(x')$. In this case, for all $i, j \in [m - 1]$, the

coefficient θ^{ij} vanishes, and for all $i \in [m - 1]$, it holds that $\theta^i = \pi(i)/\pi(m)$, so that we can write more simply

$$\log P(x, x') = \sum_{i=1}^{m-1} \log \frac{\pi(i)}{\pi(m)} \delta_i(x') + \log \pi(m). \tag{15}$$

Conversely, now suppose that $\theta^{ij} = 0$ for any $i, j \in [m - 1]$. Then the matrix

$$\tilde{P}(x, x') = \exp \left(\sum_{i=1}^{m-1} \log \frac{P(m, i)P(i, m)}{P(m, m)P(m, m)} \delta_i(x') \right)$$

has rank one, the right PF eigenvector is constant, and P is memoryless. As a result, \mathcal{W}_{id} is an e-family of \mathcal{W} such that $\theta^{ij} = 0$ for every $i, j \in [m - 1]$. □

Lemma 8 \mathcal{W}_{id} does not form an m -family.

Proof We prove the case $|\mathcal{X}| = 2$ and $p \neq 1/2$,

$$P_0 = \begin{pmatrix} p & 1 - p \\ p & 1 - p \end{pmatrix}, \quad P_1 = \begin{pmatrix} 1 - p & p \\ 1 - p & p \end{pmatrix}.$$

Computing the corresponding edge measures,

$$Q_0 = \begin{pmatrix} p^2 & p(1 - p) \\ p(1 - p) & (1 - p)^2 \end{pmatrix}, \quad Q_1 = \begin{pmatrix} (1 - p)^2 & p(1 - p) \\ p(1 - p) & p^2 \end{pmatrix}.$$

But then if we let

$$Q_{1/2} = \frac{1}{2}Q_0 + \frac{1}{2}Q_1 = \begin{pmatrix} \frac{1}{2}(p^2 + (1 - p)^2) & p(1 - p) \\ p(1 - p) & \frac{1}{2}(p^2 + (1 - p)^2) \end{pmatrix},$$

we see that the stationary distribution is $\pi_{1/2} = \mathbf{1}/2$, and

$$P_{1/2} = \begin{pmatrix} p^2 + (1 - p)^2 & 2p(1 - p) \\ 2p(1 - p) & p^2 + (1 - p)^2 \end{pmatrix}.$$

But for $p \neq 0$, $P_{1/2}$ does not belong to \mathcal{W}_{id} , hence the family is not an m -family. The proof can be extended to the more general $\mathcal{X} = [m], m > 2$ by considering instead the two kernels defined by $\pi_p = (p, 1 - p, 1, \dots, 1)/(m - 1)$ and π_{1-p} for $p \in (0, 1), p \neq 1/2$. □

For simplicity, in the remainder of this section, we mostly consider the full support case.

Family of positive doubly-stochastic kernel: $\mathcal{W}_{\text{bis}}(\mathcal{X}, \mathcal{X}^2)$. Recall that a kernel P is said to be doubly-stochastic, or bi-stochastic, when P and P^T are both stochastic

matrices. In this case, the stationary distribution is always uniform. It is known that the set of doubly stochastic Markov chains forms an m -family of dimension $(|\mathcal{X}| - 1)^2$ [10, Example 4]. However, as a consequence of Lemma 10, it does not form an e -family (except when $|\mathcal{X}| = 2$).

Family of positive symmetric kernel: $\mathcal{W}_{\text{sym}}(\mathcal{X}, \mathcal{X}^2)$. A Markov kernel is symmetric, when $P = P^\top$, hence this family lies at the intersection between reversible and doubly-stochastic families of Markov kernels, which are both m -families. This implies that symmetric kernels also form an m -family. In fact, Lemma 9 shows that the dimension of this family is $|\mathcal{X}|(|\mathcal{X}| - 1)/2$. Lemma 10, however, shows that \mathcal{W}_{sym} only forms an e -family for $|\mathcal{X}| = 2$.

Lemma 9 \mathcal{W}_{sym} forms an m -family of dimension $|\mathcal{X}|(|\mathcal{X}| - 1)/2$.

Proof To prove the claim, we will rely on Definition 1-(ii) of a mixture family. Consider the functions $s_0: \mathcal{X}^2 \rightarrow \mathbb{R}$ and $s_{ij}: \mathcal{X}^2 \rightarrow \mathbb{R}$ for $i, j \in \mathcal{X}, i > j$ such that for any $x, x' \in \mathcal{X}, s_0(x, x') = \delta_x(x')/|\mathcal{X}|$ and $s_{ij} = \delta_i^\top \delta_j + \delta_j^\top \delta_i - 2\delta_i^\top \delta_i$. Let $Q \in \mathcal{Q}$, we verify that for any $x, x' \in \mathcal{X}$,

$$Q(x, x') = s_0(x, x') + \sum_{\substack{i, j \in \mathcal{X} \\ i > j}} s_{ij}(x, x')Q(i, j), \tag{16}$$

and moreover

$$\sum_{x, x' \in \mathcal{X}} s_0(x, x') = 1, \quad \sum_{x, x' \in \mathcal{X}} s_{ij}(x, x') = 0, \forall i, j \in \mathcal{X}, i > j.$$

It remains to show that the $s_0, s_0 + s_{ij}$, for $i > j$, are affinely independent, or equivalently, that the s_{ij} , for $i > j$, are linearly independent. Let $s = \sum_{i > j} \alpha_{ij}s_{ij}$ with $\alpha_{ij} \in \mathbb{R}$, for any $i > j$, be such that $s = 0$. For any $i > j$, taking $x = i, x' = j$ yields $\alpha_{ij} = 0$, thus the family is independent, hence constitutes a basis, and the dimension is $|\{i, j \in \mathcal{X}: i > j\}| = |\mathcal{X}|(|\mathcal{X}| - 1)/2$. \square

Lemma 10 For $|\mathcal{X}| \geq 2$,

- (i) The set \mathcal{W}_{sym} does not form an e -family, unless $|\mathcal{X}| = 2$.
- (ii) The set \mathcal{W}_{bis} does not form an e -family, unless $|\mathcal{X}| = 2$.

Proof We first treat the case $|\mathcal{X}| = 2$ for (i) and (ii). Notice that

$$P_\theta = \begin{pmatrix} e^\theta & 1 \\ \frac{1}{1+e^\theta} & \frac{1}{1+e^\theta} \end{pmatrix}$$

for $\theta \in \mathbb{R}$ satisfies $P_\theta \in \mathcal{W}_{\text{sym}}$, and that the latter expression exhausts all irreducible symmetric chains. We can therefore write

$$\mathcal{W}_{\text{sym}} = \{P_\theta: P_\theta(x, x') = \exp(\delta_x(x')\theta - \log(e^\theta + 1)), \theta \in \mathbb{R}\},$$

which follows the definition at (2) of an e-family with carrier kernel $K = 0$, generator $g(x, x') = \delta_x(x')$, natural parameter θ , $R_\theta = 0$ and potential function $\psi_\theta = \log(e^\theta + 1)$. Furthermore, for $|\mathcal{X}| = 2$, it is easy to see that symmetric and doubly-stochastic families coincide, hence \mathcal{W}_{bis} is also an e-family.

We now prove (i) for $|\mathcal{X}| = 3$. We will consider two positive symmetric Markov kernels P_0 and P_1 , and look at the e-geodesic

$$G_e(P_0, P_1) \triangleq \left\{ \mathfrak{s}(\tilde{P}_\theta) : \tilde{P}_\theta = P_0(x, x')^{1-\theta} P_1(x, x')^\theta, \theta \in [0, 1] \right\},$$

where the map \mathfrak{s} , defined in (3), enforces stochasticity. The matrix $P_\theta(x, x') = \mathfrak{s}(\tilde{P}_\theta)$ is symmetric, if and only if the right eigenvector of \tilde{P}_θ is constant. This, in turn, is equivalent to the rows of \tilde{P}_θ being all equal. Consider the two symmetric kernels

$$P_0 = \begin{pmatrix} \alpha & 2/3 - \alpha & 1/3 \\ 2/3 - \alpha & \alpha & 1/3 \\ 1/3 & 1/3 & 1/3 \end{pmatrix}, \quad P_1 = \frac{1}{3} \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix},$$

with free parameter $\alpha \neq 1/3$, and let us inspect the curve at parameter $\theta = 1/2$. For $P_{1/2}$ to be symmetric, it is necessary that

$$\sqrt{\alpha} + \sqrt{2/3 - \alpha} = 2\sqrt{1/3},$$

whose unique solution is precisely $\alpha = 1/3$. Invoking Nagaoka [22, Corollary 3] finishes proving (i) for $|\mathcal{X}| = 3$. We extend the proof to $|\mathcal{X}| \geq 4$ using the padding argument of Theorem 8, considering $\frac{1}{m} \begin{pmatrix} 3P_i & \mathbf{1}^\top \mathbf{1} \\ \mathbf{1}^\top \mathbf{1} & \mathbf{1}^\top \mathbf{1} \end{pmatrix}$. Suppose for contradiction that (ii) is false, i.e. bi-stochastic matrices form an e-family. Take then any e-geodesic between two arbitrary symmetric kernels. The latter operators being reversible, so is the geodesic. But then this curve must also be composed entirely of symmetric matrices, hence the geodesic is symmetric, which contradicts (i). \square

Remark 8 For $|\mathcal{X}| \geq 3$, the following hierarchies hold:

$$\begin{array}{l} \mathcal{W}_{\text{iid}} \stackrel{\text{e-family}}{\subsetneq} \mathcal{W}_{\text{rev}} \stackrel{\text{e-family}}{\subsetneq} \mathcal{W}, \\ \mathcal{W}_{\text{sym}} \stackrel{\text{m-family}}{\subsetneq} \mathcal{W}_{\text{rev}} \stackrel{\text{m-family}}{\subsetneq} \mathcal{W}, \\ \mathcal{W}_{\text{sym}} \stackrel{\text{m-family}}{\subsetneq} \mathcal{W}_{\text{bis}} \stackrel{\text{m-family}}{\subsetneq} \mathcal{W}. \end{array}$$

9 Generation of the reversible family

In this final section, we consider the family of positive Markov kernels $\mathcal{W} = \mathcal{W}(\mathcal{X}, \mathcal{X}^2)$ i.e. where the support $\mathcal{E} = \mathcal{X}^2$. We first show that \mathcal{W}_{rev} is in a sense the smallest exponential family that contains \mathcal{W}_{sym} , the family of symmetric Markov

kernels. Our notion of minimality relies on the following definition of the exponential hull of some submanifold of \mathcal{W} .

Definition 7 (*Exponential hull*) Let $\mathcal{V} \subset \mathcal{W}$.

$$\text{e-hull}(\mathcal{V}) = \left\{ \mathfrak{s}(\tilde{P}) : \log[\tilde{P}] = \sum_{i=1}^k \alpha_i \log[P_i], \right. \\ \left. k \in \mathbb{N}, \alpha_1, \dots, \alpha_k \in \mathbb{R}, \sum_{i=1}^k \alpha_i = 1, P_1, \dots, P_k \in \mathcal{V} \right\},$$

where \mathfrak{s} is defined in (3).

Remark: When $U = \frac{1}{|\mathcal{X}|} \mathbf{1}^\top \mathbf{1} \in \mathcal{V}$, the constraint $\sum_{i=1}^k \alpha_i = 1$ is redundant. Indeed, since U corresponds to the origin in e-coordinates, the linear hull and affine hull coincide in this case.

By (10), the reversible family is generated by symmetric functions. Even though not all symmetric functions correspond to symmetric kernels, the reversible family is spanned by the symmetric kernels as follows.

Theorem 9 For $|\mathcal{X}| \geq 3$, it holds that

$$\text{e-hull}(\mathcal{W}_{\text{sym}}) = \mathcal{W}_{\text{rev}}.$$

Proof We begin by proving the inclusion $\text{e-hull}(\mathcal{W}_{\text{sym}}) \subset \mathcal{W}_{\text{rev}}$. Let $P \in \text{e-hull}(\mathcal{W}_{\text{sym}})$, then there exist a positive $\tilde{P} \in \mathcal{F}_+$, and $k \in \mathbb{N}, \alpha_1, \dots, \alpha_k \in \mathbb{R}, P_1, \dots, P_k \in \mathcal{W}_{\text{sym}}$ such that $\log \tilde{P}(x, x') = \sum_{i=1}^k \alpha_i \log P_i(x, x')$. Observe that the function $\sum_{i=1}^k \alpha_i \log P_i(x, x')$ is symmetric in x and x' , thus $\log \tilde{P}(x, x')$ is log-reversible, and P is reversible.

We now prove the second inclusion $\mathcal{W}_{\text{rev}} \subset \text{e-hull}(\mathcal{W}_{\text{sym}})$. We let

$$\mathcal{H} = \text{span} \left(\{ \log[P] : P \in \mathcal{W}_{\text{sym}} \} \cup \mathcal{N} \right).$$

Recall from Theorem 4 that the functions $g_{ij} = \delta_i^\top \delta_j + \delta_j^\top \delta_i$, for $(i, j) \in T(\mathcal{X}^2)$, form a basis of the quotient space $\mathcal{G}_{\text{rev}} = \mathcal{F}_{\text{rev}}/\mathcal{N}$. It suffices therefore to show that $\{g_{ij} : (i, j) \in T(\mathcal{X}^2)\} \subset \mathcal{H}$. Introduce a free parameter $t \in (0, 1), t \neq 1/2$, and let us fix $(i, j) \in T_+(\mathcal{X})$. Consider $P_{ij,t} \in \mathcal{W}_{\text{sym}}$ defined as follows

$$P_{ij,t}(x, x') \triangleq \begin{cases} 2(1-t)/|\mathcal{X}| & \text{if } (x, x') \in \{(i, i), (j, j)\}, \\ 2t/|\mathcal{X}| & \text{if } (x, x') \in \{(i, j), (j, i)\}, \\ 1/|\mathcal{X}| & \text{otherwise,} \end{cases}$$

and the functions $\hat{h}_{ij}, \tilde{h}_{ij}$

$$\hat{h}_{ij} = \log |\mathcal{X}| + \log [P_{ij,t}] = a(\delta_i^\top \delta_i + \delta_j^\top \delta_j) + b(\delta_i^\top \delta_j + \delta_j^\top \delta_i), \\ \tilde{h}_{ij} = \log |\mathcal{X}| + \log [P_{ij,1-t}] = b(\delta_i^\top \delta_i + \delta_j^\top \delta_j) + a(\delta_i^\top \delta_j + \delta_j^\top \delta_i),$$

where for simplicity we wrote $a = \log 2(1 - t)$ and $b = \log 2t \neq a$. Since the function $((x, x') \mapsto \log |\mathcal{X}|) \in \mathcal{N}$, we have $\hat{h}_{ij}, \tilde{h}_{ij} \in \mathcal{H}$. Notice that we can write

$$g_{ij} = \frac{b\hat{h}_{ij} - a\tilde{h}_{ij}}{b^2 - a^2},$$

hence also $g_{ij} \in \mathcal{H}$. Introduce the function

$$h_{ij} = \frac{a\hat{h}_{ij} - b\tilde{h}_{ij}}{a^2 - b^2} = \delta_i^\top \delta_i + \delta_j^\top \delta_j \in \mathcal{H},$$

and observe that we can rewrite the identity $I = \mathbf{1}^\top \mathbf{1} - \sum_{(i,j) \in T_+(\mathcal{X}^2)} g_{ij}$ with $\mathbf{1}^\top \mathbf{1}$ being a constant function. It follows that $I \in \mathcal{H}$, and for any $j \in \mathcal{X}$, we can express

$$g_{jj} = \frac{2}{|\mathcal{X}| - 2} \left(\sum_{\substack{i \in \mathcal{X} \\ i > j}} h_{ij} + \sum_{\substack{i \in \mathcal{X} \\ i < j}} h_{ji} - I \right) \in \mathcal{H}.$$

As a result, $\{g_{ij} : (i, j) \in T(\mathcal{X}^2)\} \subset \mathcal{H}$, and the theorem follows. □

Remark 9 Observe that in the above proof, it is crucial that $|\mathcal{X}| \geq 3$. For $|\mathcal{X}| = 2$, we can only have $h_{21} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, and cannot construct $g_{11} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ nor $g_{22} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$. This is consistent with the observation that $\text{e-hull}(\mathcal{W}_{\text{sym}}) \neq \mathcal{W}_{\text{rev}}$ for $|\mathcal{X}| = 2$.

Secondly, we show that \mathcal{W}_{rev} is also the smallest mixture family that contains \mathcal{W}_{iid} , the family of Markov kernels that correspond to iid processes. For this, we define minimality in terms of a mixture hull.

Definition 8 (Mixture hull) Let $\mathcal{V} \subset \mathcal{W}$.

$$\text{m-hull}(\mathcal{V}) = \left\{ P : Q \in \mathcal{Q}, Q = \sum_{i=1}^k \alpha_i Q_i, \right. \\ \left. k \in \mathbb{N}, \alpha_1, \dots, \alpha_k \in \mathbb{R}, P_1, \dots, P_k \in \mathcal{V} \right\},$$

where Q (resp. Q_i) pertains to P (resp. P_i).

Theorem 10 *It holds that*

$$\text{m-hull}(\mathcal{W}_{\text{iid}}) = \mathcal{W}_{\text{rev}}.$$

Proof Let $P \in \text{m-hull}(\mathcal{W}_{\text{iid}})$, then the corresponding edge measure can be expressed as a linear combination $\sum_{i=1}^k \alpha_i Q_i$, with $k \in \mathbb{N}, \alpha_1, \dots, \alpha_k \in \mathbb{R}$, and where the Q_i pertain to some degenerate iid kernel $P_i = \mathbf{1}^\top \pi_i$. This implies that $Q_i(x, x') = \pi_i(x)\pi_i(x')$, hence Q_i is symmetric. In turn, Q is symmetric, i.e. P is reversible, and $\text{m-hull}(\mathcal{W}_{\text{iid}}) \subset \mathcal{W}_{\text{rev}}$.

For $(i, j) \in \mathcal{X}^2, i \geq j$, and $\varepsilon \in [0, 1]$, consider the mixture distribution

$$\pi_{ij,\varepsilon} = \frac{\varepsilon}{|\mathcal{X}|} \mathbf{1} + (1 - \varepsilon) \frac{\delta_i + \delta_j}{2} \in \mathcal{P}(\mathcal{X}).$$

A direct computation yields that the pair probabilities of the iid process can be written as

$$Q_{ij,\varepsilon}(x, x') = \frac{\varepsilon^2}{|\mathcal{X}|^2} + \frac{\varepsilon(1 - \varepsilon)}{2|\mathcal{X}|} \{ \delta_i(x) + \delta_j(x) + \delta_i(x') + \delta_j(x') \} + \frac{(1 - \varepsilon)^2}{4} \times \{ \delta_i(x)\delta_i(x') + \delta_i(x)\delta_j(x') + \delta_j(x)\delta_i(x') + \delta_j(x)\delta_j(x') \}.$$

We first show that $\{Q_{ij,0} : i \geq j\}$ forms a basis of \mathcal{F}_{sym} . Let $\{\alpha_{ij} \in \mathbb{R} : i \geq j\}$ be such that $\sum_{i \geq j} \alpha_{ij} Q_{ij,0} = 0$. Consider first $x, x' \in \mathcal{X}$ such that $x > x'$.

$$\sum_{i \geq j} \alpha_{ij} Q_{ij,0}(x, x') = \frac{1}{4} \sum_{i \geq j} \alpha_{ij} \delta_i(x) \delta_j(x') = \frac{1}{4} \alpha_{xx'} = 0 \text{ and } \alpha_{xx'} = 0.$$

By a similar argument for the case $x < x'$, we obtain that $\alpha_{xx'} = 0$ for any $x \neq x'$. Inspecting now the diagonal for $x \in \mathcal{X}$,

$$\sum_{i \geq j} \alpha_{ij} Q_{ij,0}(x, x) = \sum_{i \in \mathcal{X}} \alpha_{ii} Q_{ii,0}(x, x) = \sum_{i \in \mathcal{X}} \alpha_{ii} \delta_i(x) = \alpha_{xx} = 0.$$

This implies that the family $\{Q_{ij,0} : i \geq j\}$ is independent. Since $\dim \mathcal{F}_{\text{sym}} = |\mathcal{X}|(|\mathcal{X}| + 1)/2 = |\{Q_{ij,0} : i \geq j\}|$, it is maximally so, thus forms a basis. However, the basis elements are not in \mathcal{W}_{id} . We therefore examine the case $\varepsilon > 0$, and leverage the property that in normed vector spaces, finite linearly independent systems are stable under small perturbations (see Lemma 11 reported below for convenience) in order to show existence of a basis in \mathcal{W}_{id} .

Lemma 11 (Costara and Popa [48, p.9, Exercise 35]) *Let $n \in \mathbb{N}$, X is a normed vector space and x_1, \dots, x_n are n linearly independent elements in X . Then there exists $\eta > 0$ such that if y_1, y_2, \dots, y_n are such that $\|y_i\| < \eta$ for $i = 1, \dots, n$, then $x_1 + y_1, x_2 + y_2, \dots, x_n + y_n$ are also n linearly independent elements in X .*

Let us consider $(\mathcal{F}_{\text{sym}}, \|\cdot\|_{1,1})$, the space of real symmetric matrices equipped with the entry-wise ℓ_1 norm. For any $i \geq j$ and for any $\varepsilon \in (0, 1)$,

$$\begin{aligned} & \|Q_{ij,\varepsilon} - Q_{ij,0}\|_{1,1} \triangleq \sum_{x,x' \in \mathcal{X}} |Q_{ij,\varepsilon}(x, x') - Q_{ij,0}(x, x')| \\ & \leq |\mathcal{X}|^2 \left| \frac{\varepsilon^2}{|\mathcal{X}|^2} \right| + \left| \frac{\varepsilon(1 - \varepsilon)}{2|\mathcal{X}|} \right| \sum_{x,x'} |\delta_i(x) + \delta_j(x) + \delta_i(x') + \delta_j(x')| \end{aligned}$$

$$\begin{aligned}
& + \left| \frac{\varepsilon(2-\varepsilon)}{4} \right| \sum_{x, x' \in \mathcal{X}} |\delta_i(x)\delta_i(x') + \delta_i(x)\delta_j(x') + \delta_j(x)\delta_i(x') + \delta_j(x)\delta_j(x')| \\
& \leq \varepsilon^2 + 2\varepsilon |\mathcal{X}| + 2\varepsilon |\mathcal{X}|^2,
\end{aligned}$$

thus

$$\|Q_{ij,\varepsilon} - Q_{ij,0}\|_{1,1} \leq 5\varepsilon |\mathcal{X}|^2.$$

Let η as defined in Lemma 11, with respect to the basis $\{Q_{ij,0} : i \geq j\}$, and choose $0 < \varepsilon < \frac{\eta}{5|\mathcal{X}|^2}$. Then $\|Q_{ij,\varepsilon} - Q_{ij,0}\|_{1,1} < \eta$, thus the family $\{Q_{ij,\varepsilon} : i \geq j\}$ is a also basis for \mathcal{F}_{sym} that lies in \mathcal{W}_{id} , whence the theorem. \square

Acknowledgements Many thanks to the anonymous referees for the helpful comments, which helped us to improve the presentation of this manuscript. We also thank Hiroshi Nagaoka for an enlightening discussion, and constructive remarks, as well as Alexander Yu. Mitrophanov for commenting on a previous version of this manuscript.

Data Availability Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Schrödinger, E.: Über die umkehrung der naturgesetze. Sitzungsberichte der preussischen Akademie der Wissenschaften, physikalische mathematische Klasse **8**(N9), 144–153 (1931)
2. Dobrushin, R.L., Sukhov, Y.M., Fritz, J.: A.N. Kolmogorov - the founder of the theory of reversible Markov processes. Russian Math. Surv. **43**(6), 157–182 (1988)
3. Kolmogorov, A.: Zur theorie der Markoffschen ketten. Math. Ann. **112**(1), 155–160 (1936)
4. Kolmogorov, A.: Zur umkehrbarkeit der statistischen naturgesetze. Math. Ann. **113**(1), 766–772 (1937)
5. Kelly, F.P.: Reversibility and stochastic networks. Cambridge University Press, Cambridge (2011)
6. Brooks, S., Gelman, A., Jones, G., Meng, X.L.: Handbook of Markov Chain Monte Carlo. CRC Press, Boca Raton (2011)
7. Aldous, D., Fill, J.: Reversible Markov chains and random walks on graphs (2002)
8. Levin, D.A., Peres, Y., Wilmer, E.L.: Markov chains and mixing times, second edition. American Mathematical Soc (2009)
9. Pistone, G., Rogantin, M.P.: The algebra of reversible Markov chains. Ann. Inst. Stat. Math. **65**(2), 269–293 (2013)
10. Hayashi, M., Watanabe, S.: Information geometry approach to parameter estimation in Markov chains. Ann. Stat. **44**(4), 1495–1535 (2016). <https://doi.org/10.1214/15-AOS1420>

11. Takeuchi, J., Nagaoka, H.: Information geometry of the family of Markov kernels defined by a context tree. In: 2017 IEEE Information Theory Workshop (ITW), IEEE, pp 429–433 (2017)
12. Miller, H.: A convexity property in the theory of random variables defined on a finite Markov chain. *The Annals of mathematical statistics* pp 1260–1270 (1961)
13. Donsker, M.D., Varadhan, S.S.: Asymptotic evaluation of certain Markov process expectations for large time, i. *Commun. Pure Appl. Math.* **28**(1), 1–47 (1975)
14. Gärtner, J.: On large deviations from the invariant measure. *Theory Probab. Applic.* **22**(1), 24–39 (1977)
15. Dembo, A., Zeitouni, O.: *Large deviations techniques and applications*, vol 38 (1998)
16. Csiszár, I., Cover, T., Choi, B.S.: Conditional limit theorems under Markov conditioning. *IEEE Trans. Inform. Theory* **33**(6), 788–801 (1987)
17. Ito, H., Amari, S.i.: Geometry of information sources. In: *Proceedings of the 11th Symposium on Information Theory and Its Applications (SITA '88)*, pp 57–60 (1988)
18. Takeuchi, J.i., Barron, A.R.: Asymptotically minimax regret by Bayes mixtures. In: *Proceedings. 1998 IEEE International Symposium on Information Theory (Cat. No. 98CH36252)*, IEEE, p 318 (1998)
19. Takeuchi, J., Kawabata, T.: Exponential curvature of Markov models. In: *2007 IEEE International Symposium on Information Theory*, IEEE, pp 2891–2895 (2007)
20. Takeuchi, J., Nagaoka, H.: On asymptotic exponential family of Markov sources and exponential family of Markov kernels (2017)
21. Nakagawa, K., Kanaya, F.: On the converse theorem in statistical hypothesis testing for Markov chains. *IEEE Trans. Inform. Theory* **39**(2), 629–633 (1993)
22. Nagaoka, H.: The exponential family of Markov chains and its information geometry. In: *The proceedings of the Symposium on Information Theory and Its Applications*, vol 28, pp 601–604, (2005) (available on [arXiv:1701.06119](https://arxiv.org/abs/1701.06119))
23. Watanabe, S., Hayashi, M.: Finite-length analysis on tail probability for Markov chain and application to simple hypothesis testing. *Ann. Appl. Probab.* **27**(2), 811–845 (2017). <https://doi.org/10.1214/16-AAP1216>
24. Moulos, V., Anantharam, V.: Optimal Chernoff and Hoeffding bounds for finite state Markov chains. *arXiv preprint arXiv:1907.04467* (2019)
25. Hayashi, M.: Local equivalence problem in hidden Markov model. *Inform. Geom.* **2**(1), 1–42 (2019)
26. Hayashi, M.: Information geometry approach to parameter estimation in hidden Markov model. Bernoulli to appear (2021+)
27. Amari, Si., Nagaoka, H.: *Methods of information geometry*, vol 191. American Mathematical Soc (2007)
28. Fujiwara, A.: *Foundations of information geometry*. Makino Shoten (2015)
29. Feigin, P.D., et al.: Conditional exponential families and a representation theorem for asymptotic inference. *Ann. Statist.* **9**(3), 597–603 (1981)
30. Küchler, U., Sørensen, M.: On exponential families of Markov processes. *J. Statist. Plann. Infer* **66**(1), 3–19 (1998)
31. Hudson, I.L.: Large sample inference for Markovian exponential families with application to branching processes with immigration. *Australian J. Stat.* **24**(1), 98–112 (1982)
32. Stefanov, V.T.: Explicit limit results for minimal sufficient statistics and maximum likelihood estimators in some Markov processes: exponential families approach. *Ann. Stat.* 1073–1101 (1995)
33. Küchler, U., Sørensen, M.: Exponential families of stochastic processes: A unifying semimartingale approach. *International Statistical Review/Revue Internationale de Statistique* pp 123–144 (1989)
34. Sørensen, M.: On sequential maximum likelihood estimation for exponential families of stochastic processes. *International Statistical Review/Revue Internationale de Statistique* pp 191–210 (1986)
35. Chernoff, H.: A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Ann. Math. Stat.* pp 493–507 (1952)
36. Gallager, R.G.: *Information theory and reliable communication*, vol. 2. Springer, Berlin (1968)
37. Van Campenhout, J., Cover, T.: Maximum entropy and conditional probability. *IEEE Trans. Inform. Theory* **27**(4), 483–489 (1981)
38. Brill, P., Hlynka, M., Jiang, Q., et al.: Reversibility checking for Markov chains. *Commun. Stochastic Anal.* **12**(2), 2 (2018)
39. Meyer, C.D.: *Matrix analysis and applied linear algebra*, vol 71. Siam (2000)
40. Suomela, P.: Invariant measures of time-reversible Markov chains. *J. Appl. Prob.* pp 226–229 (1979)

41. Diaconis, P., Rolles, S.W., et al.: Bayesian analysis for reversible Markov chains. *Ann. Stat.* **34**(3), 1270–1292 (2006)
42. Ohara, A., Ishi, H.: Doubly autoparallel structure on the probability simplex. In: *Information Geometry and its Applications IV*, Springer pp 323–334 (2016)
43. Hsu, D., Kontorovich, A., Levin, D.A., Peres, Y., Szepesvári, C., Wolfer, G.: Mixing time estimation in reversible markov chains from a single sample path. *Ann. App. Probab.* **29**(4), 2439–2480 (2019). <https://doi.org/10.1214/18-AAP1457>
44. Stewart, G.W.: *Matrix perturbation theory*. Academic Press, Boston (1990)
45. Nielsen, A., Weber, M.: Computing the nearest reversible Markov chain. *Numer. Linear Algebra Appl.* **22**(3), 483–499 (2015)
46. Fill, J.A.: Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, with an application to the exclusion process. *The annals of applied probability* pp 62–87 (1991)
47. Nagaoka, H.: Information-geometrical characterization of statistical models which are statistically equivalent to probability simplexes. In: *2017 IEEE International Symposium on Information Theory (ISIT)*, IEEE, pp 1346–1350 (2017)
48. Costara, C., Popa, D.: *Exercises in functional analysis*, vol. 26. Springer Science & Business Media, Berlin (2013)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.