PORTUGALIAE MATHEMATICA
Vol. 64 Fasc. 3 – 2007
Nova Série

SIMPLICITY OF LYAPUNOV SPECTRA: A SUFFICIENT CRITERION

Artur Avila and Marcelo Viana

Abstract: We exhibit an explicit sufficient condition for the Lyapunov exponents of a linear cocycle over a Markov map to have multiplicity 1. This builds on work of Guivarc'h–Raugi and Gol'dsheid–Margulis, who considered products of random matrices, and of Bonatti–Viana, who dealt with the case when the base dynamics is a subshift of finite type. Here the Markov structure may have infinitely many symbols and the ambient space needs not be compact. As an application, in another paper we prove the Zorich–Kontsevich conjecture on the Lyapunov spectrum of the Teichmüller flow in the space of translation surfaces.

Contents

Received: August 15, 2006.
Work supported by the Brazil–France Agreement in Mathematics. M.V. is also supported by Pronex and Faperj. This research was partially conducted during the period A.A. served as a Clay Research Fellow.

1 – Introduction and statements

Let $\hat{f}: \hat{\Sigma} \to \hat{\Sigma}$ be an invertible measurable map and $\hat{A}: \hat{\Sigma} \to \mathrm{GL}(d, \mathbb{C})$ be a measurable function with values in the group of invertible $d \times d$ complex matrices. These data define a linear cocycle $\hat{F}_A$ over the map $\hat{f}$, through

$$\hat{F}_A: \hat{\Sigma} \times \mathbb{C}^d \to \hat{\Sigma} \times \mathbb{C}^d , \qquad \hat{F}_A(\hat{x}, v) = \big(\hat{f}(\hat{x}), \hat{A}(\hat{x})\, v\big) .$$

Note that $\hat{F}_A^n(x, v) = \big(\hat{f}^n(\hat{x}), \hat{A}^n(\hat{x})\big)$, where $\hat{A}^n(\hat{x}) = \hat{A}\big(\hat{f}^{n-1}(\hat{x})\big) \cdots \hat{A}\big(\hat{f}(\hat{x})\big)\, \hat{A}(\hat{x})$ and $\hat{A}^n(\hat{x})$ is the inverse of $\hat{A}^{-n}\big(\hat{f}^n(\hat{x})\big)$ if $n < 0$.

Let $\hat{\mu}$ be an $\hat{f}$-invariant probability measure on $\hat{\Sigma}$ relative to which the logarithms of the norms of $\hat{A}$ and its inverse are integrable. By the theorem of Oseledets [13], at μ-almost every $\hat{x} \in \hat{\Sigma}$ there exist numbers $\lambda_1(\hat{x}) > \lambda_2(\hat{x}) > \cdots > \lambda_k(\hat{x})$ and a decomposition $\mathbb{C}^d = E^1_{\hat{x}} \oplus E^2_{\hat{x}} \oplus \cdots \oplus E^k_{\hat{x}}$ into vector subspaces such that

$$\hat{A}(\hat{x})\, E^i_{\hat{x}} = E^i_{\hat{f}(\hat{x})} \qquad \text{and} \qquad \lambda_i(\hat{x}) = \lim_{|n| \to \infty} \frac{1}{n} \log \big\|\hat{A}^n(\hat{x})\, v\big\|$$

for every non-zero $v \in E^i_{\hat{x}}$ and $1 \le i \le k$. We call $\dim E^i_{\hat{x}}$ the *multiplicity* of $\lambda_i(\hat{x})$.

We assume that $\hat{\mu}$ is ergodic. Then the *Lyapunov exponents* $\lambda_i(\hat{x})$ are constant on a full measure subset of $\hat{\Sigma}$ and so are the dimensions of the Oseledets subspaces $E^i_{\hat{x}}$. The *Lyapunov spectrum* of $\hat{A}$ is the set of all Lyapunov exponents. We say that the Lyapunov spectrum is *simple* if it contains exactly d distinct values ($k = d$) or, equivalently, if every Lyapunov exponent λ_i has multiplicity 1. The main result in this paper, to be stated below, provides an explicit sufficient condition for the Lyapunov spectrum to be simple. We begin by describing the class of cocycles to which it applies. In Appendix A we discuss some extensions and applications.

1.1. Symbolic dynamics

We take $\hat{\Sigma} = \mathbb{N}^{\mathbb{Z}}$, the full shift space with countably many symbols, and $\hat{f}: \hat{\Sigma} \to \hat{\Sigma}$ to be the shift map:

$$\hat{f}\big((x_n)_{n \in \mathbb{Z}}\big) = (x_{n+1})_{n \in \mathbb{Z}} .$$

Let us call *cylinder* of $\hat{\Sigma}$ any subset of the form

$$\Big[\iota_m, ..., \iota_{-1};\, \iota_0;\, \iota_1, ..., \iota_n\Big] = \Big\{\hat{x}:\ x_j = \iota_j \ \text{ for } \ j = m, ..., n\Big\} .$$

Cylinders of $\Sigma^u = \mathbb{N}^{\{n\geq 0\}}$ and $\Sigma^s = \mathbb{N}^{\{n<0\}}$ are defined similarly, corresponding to the cases $m=0$ and $n=-1$, respectively, and they are represented as $[\iota_0, \iota_1, ..., \iota_n]^u$ and $[\iota_m, ..., \iota_{-1}]^s$, respectively. We endow $\hat{\Sigma}$, Σ^u, Σ^s with the topologies generated by the corresponding cylinders. Let $P^u\colon \hat{\Sigma} \to \Sigma^u$ and $P^s\colon \hat{\Sigma} \to \Sigma^s$ be the natural projections. We also consider the one-sided shift maps $f^u\colon \Sigma^u \to \Sigma^u$ and $f^s\colon \Sigma^s \to \Sigma^s$ defined by

$$f^u \circ P^u = P^u \circ \hat{f} \quad \text{and} \quad f^s \circ P^s = P^s \circ \hat{f}^{-1} .$$

For each $\hat{x} = (x_n)_{n\in\mathbb{Z}}$ in $\hat{\Sigma}$, we denote $x^u = P^u(\hat{x})$ and $x^s = P^s(\hat{x})$. Then $\hat{x} \mapsto (x^s, x^u)$ is a homeomorphism from $\hat{\Sigma}$ to the product $\Sigma^s \times \Sigma^u$. In what follows we often identify the two sets through this homeomorphism. When there is no risk of ambiguity, we also identify the *local stable set*

$$W^s_{\text{loc}}(x^u) = W^s_{\text{loc}}(\hat{x}) = \Big\{(y_n)_{n\in\mathbb{Z}}\colon\ x_n = y_n \text{ for all } n \geq 0\Big\} \quad \text{with } \Sigma^s$$

and the *local unstable set*

$$W^u_{\text{loc}}(x^s) = W^u_{\text{loc}}(\hat{x}) = \Big\{(y_n)_{n\in\mathbb{Z}}\colon\ x_n = y_n \text{ for all } n < 0\Big\} \quad \text{with } \Sigma^u ,$$

via the projections P^s and P^u.

In Section A.1 we shall discuss how more general situations may often be reduced to this one.

1.2. Product structure

Let $\mu^u = P^u_*\hat{\mu}$ and $\mu^s = P^s_*\hat{\mu}$ be the images of the ergodic $\hat{f}$-invariant probability measure $\hat{\mu}$ under the natural projections. It is easy to see that these are ergodic invariant probabilities for f^u and f^s, respectively. We take μ^s and μ^u to be positive on cylinders. Moreover, we assume $\hat{\mu}$ to be equivalent to their product, meaning there exists a measurable function $\rho\colon \hat{\Sigma} \to (0,\infty)$ such that

$$\hat{\mu} = \rho(\hat{x})\,(\mu^s \times \mu^u) , \qquad \hat{x} \in \hat{\Sigma} .$$

We assume that ρ is bounded from zero and infinity. For convenience of notation, we state this condition as follows: there exists some constant $K > 0$ such that

$$\frac{1}{K} \leq \frac{\rho(z^s, x^u)}{\rho(z^s, y^u)} \leq K \quad \text{and} \quad \frac{1}{K} \leq \frac{\rho(x^s, z^u)}{\rho(y^s, z^u)} \leq K \tag{1}$$

for all $x^s, y^s, z^s \in \Sigma^s$ and $x^u, y^u, z^u \in \Sigma^u$. Notice that $\{\hat{\mu}_{x^u} = \rho(\cdot, x^u)\mu^s : x^u \in \Sigma^u\}$ is a *disintegration* of $\hat{\mu}$ into conditional probabilities along local stable sets. By this we mean (see Rokhlin [15] or [2, Appendix C]) that $\hat{\mu}_{x^u}(W^s_{\mathrm{loc}}(x^u)) = 1$ for μ^u-almost every x^u and

$$\hat{\mu}(D) = \int \hat{\mu}_x\big(D \cap W^s_{\mathrm{loc}}(x^u)\big)\, d\mu^u(x^u)$$

for any measurable set $D \subset \hat{\Sigma}$. Analogously, $\{\hat{\mu}_{x^s} = \rho(x^s, \cdot)\mu^u : x^s \in \Sigma^s\}$ is a disintegration of $\hat{\mu}$ along local unstable sets. Since the density ρ is positive, the measures $\hat{\mu}_{x^u}$, $x^u \in \Sigma^u$ are all equivalent, and so are all $\hat{\mu}_{x^s}$, $x^s \in \Sigma^s$. Condition (1) just means that the Radon–Nikodym derivatives

$$\frac{d\hat{\mu}_{x^u}}{d\hat{\mu}_{y^u}} \text{ with } x^u, y^u \in \Sigma^u \quad \text{and} \quad \frac{d\hat{\mu}_{x^s}}{d\hat{\mu}_{y^s}} \text{ with } x^s, y^s \in \Sigma^s$$

are uniformly bounded from zero and infinity. This will be used to obtain the bounded distortion properties (6) and (14) below.

We also assume that the conditional probabilities $\hat{\mu}_{x^u}$ and $\hat{\mu}_{x^s}$ vary continuously with the base point, in the sense that the functions

$$\Sigma^u \ni x^u \mapsto \int \phi\, d\hat{\mu}_{x^u} \quad \text{and} \quad \Sigma^s \ni x^s \mapsto \int \psi\, d\hat{\mu}_{x^s} \tag{2}$$

are continuous, for any bounded measurable functions $\phi\colon \Sigma^s \to \mathbb{R}$ and $\psi\colon \Sigma^u \to \mathbb{R}$. Equivalently,

$$x^u \mapsto \hat{\mu}_{x^u}\big([\iota_m, \dots, \iota_{-1}]^s\big) \quad \text{and} \quad x^s \mapsto \hat{\mu}_{x^u}\big([\iota_0, \iota_1, \dots, \iota_n]^u\big)$$

are continuous for every choice of the ι_j's. This will be used to obtain (7) and Lemma 2.5.

In Section A.2 we show that these hypotheses hold, in particular, whenever the system satisfies a distortion summability condition. Indeed, in that case the density ρ may be taken continuous and bounded from zero and infinity. In general, the hypothesis (2) can probably be avoided: that is the case at least when the cocycle is locally constant; see the appendix of [1] and also Remark 4.6 below.

1.3. Invariant holonomies

Concerning the function $\hat{A}\colon \hat{\Sigma} \to \mathrm{GL}(d, \mathbb{C})$, we assume that it is continuous and admits stable and unstable holonomies:

Definition 1.1. We say $\hat{A}$ admits *stable holonomies* if the limit

$$H^s_{\hat{x},\hat{y}} = \lim_{n\to+\infty} \hat{A}^n(\hat{y})^{-1}\hat{A}^n(\hat{x})$$

exists for any pair of points $\hat{x}$ and $\hat{y}$ in the same local stable set, and depends continuously on $(\hat{x},\hat{y})$. *Unstable holonomies* $H^u_{\hat{x},\hat{y}}$ are defined in a similar way, with $n \to -\infty$ and $\hat{x}$ and $\hat{y}$ in the same local unstable set. □

Notice that stable holonomies $H^s_{\hat{x},\hat{y}} : \mathbb{C}^d \to \mathbb{C}^d$ are linear maps and they satisfy

(**a**) $H^s_{\hat{x},\hat{z}} = H^s_{\hat{z},\hat{y}} \cdot H^s_{\hat{x},\hat{z}}$ and $H^s_{\hat{x},\hat{x}} = \mathrm{id}$,

(**b**) $\hat{A}(\hat{y}) \cdot H^s_{\hat{x},\hat{y}} = H^s_{\hat{f}(\hat{x}),\hat{f}(\hat{y})} \cdot \hat{A}(\hat{x})$,

over all points for which the relations make sense. Similar remarks apply for the unstable holonomies.

For example, if $\hat{A}$ is *locally constant*, meaning that it is constant on each cylinder $[\iota]$, $\iota \in \mathbb{N}$, then $H^s_{\hat{x},\hat{y}} \equiv \mathrm{id}$ and $H^u_{\hat{x},\hat{y}} \equiv \mathrm{id}$. In Section A.3 we discuss other situations where these structures occur.

1.4. Statement of main result

Let $\hat{p} \in \hat{\Sigma}$ be a periodic point of $\hat{f}$ and $q \geq 1$ be its period. We call $\hat{z} \in \hat{\Sigma}$ a *homoclinic point* of $\hat{p}$ if $\hat{z} \in W^u_{\mathrm{loc}}(\hat{p})$ and there exists some multiple $l \geq 1$ of q such that $\hat{f}^l(\hat{z}) \in W^s_{\mathrm{loc}}(\hat{p})$. Then we define the *transition map*

$$\psi_{p,z} : \mathbb{C}^d \to \mathbb{C}^d \ , \qquad \psi_{p,z} = H^s_{\hat{f}^l(\hat{z}),\hat{p}} \cdot \hat{A}^l(\hat{z}) \cdot H^u_{\hat{p},\hat{z}} \ .$$

The following notion is our main criterion for simplicity of the Lyapunov spectrum. We refer to (p) as the *pinching property* and to (t) as the *twisting property*.

Definition 1.2. We say that $\hat{A} : \hat{\Sigma} \to \mathrm{GL}(d,\mathbb{C})$ is *simple* for $\hat{f}$ if there exists some periodic point $\hat{p} \in \hat{\Sigma}$ of $\hat{f}$ and some homoclinic point $\hat{z} \in \hat{\Sigma}$ of $\hat{p}$ such that

(**p**) All the eigenvalues of $\hat{A}^q(\hat{p})$ have distinct absolute values.

(**t**) For any invariant subspaces (sums of eigenspaces) E and F of $\hat{A}^q(\hat{p})$ with $\dim E + \dim F = d$, we have $\psi_{p,z}(E) \cap F = \{0\}$. □

Remark 1.3. Let θ_j, $j{=}1,\ldots,d$, represent the eigenspaces of $\hat{A}^q(\hat{p})$. For $d = 2$ the twisting condition means that $\psi_{p,z}(\theta_i) \neq \theta_j$ for all $1 \leq i,j \leq 2$. For $d = 3$ it means that $\psi_{p,z}(\theta_i)$ is outside the plane $\theta_j \oplus \theta_k$ and θ_i is outside the plane

$\psi_{p,z}(\theta_j \oplus \theta_k)$, for all choices of $1 \le i, j, k \le 3$. In general, this condition is equivalent to saying that the matrix of the transition map in a basis of eigenvectors of $\hat{A}^q(\hat{p})$ has all its algebraic minors different from zero. Indeed, it may be restated as saying that the determinant of the square matrix

$$\begin{pmatrix} B_{1,i_1} & \cdots & B_{1,i_r} & \delta_{1,j_1} & \cdots & \delta_{1,j_s} \\ \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\ B_{d,i_1} & \cdots & B_{d,i_r} & \delta_{d,j_1} & \cdots & \delta_{d,j_s} \end{pmatrix}$$

is non-zero for any $I = \{i_1, ..., i_s\}$ and $J = \{j_1, ..., j_r\}$ with $r + s = d$, where the $\delta_{i,j}$ are Dirac symbols and the $B_{i,j}$ are the entries of the matrix of $\psi_{p,z}$ in the basis of eigenvectors. Up to sign, this determinant is the algebraic minor $B[J^c \times I]$ corresponding to the lines $j \notin J$ and columns $i \in I$. □

Theorem A. *If $\hat{A} \colon \hat{\Sigma} \to \mathrm{GL}(d, \mathbb{C})$ is simple for $\hat{f}$ then all the Lyapunov exponents of the cocycle $\hat{F}_A$ for the measure $\hat{\mu}$ have multiplicity* 1.

Simplicity of the Lyapunov spectrum for independent random matrices was investigated in the eighties by Guivarc'h, Raugi [8], and Gol'dsheid, Margulis [7]. Theorem A also extends the main conclusions of Bonatti, Viana [4], who treated the case when the base dynamics f is a subshift of finite type.

The present extension has been carried out to include in the theory such examples as the Zorich cocycles, whose base dynamics are not of finite type. It has been conjectured by Zorich and Kontsevich [9, 19, 20] that the corresponding Lyapunov exponents have multiplicity 1. As an application of these ideas, in [1] we prove this conjecture. See also the comments in Appendix A to the present paper.

Let us point out that we improve [4] not only in that here we allow for infinite Markov structures and non-compact ambient spaces, but also because our criterion is sharper: whereas we only ask the cocycle to be simple, [4] needed a similar hypothesis on all exterior powers as well.

1.5. Outline of the proof

The starting point is the following observation. Let $\ell \in \{1, ..., d-1\}$ be fixed and assume the cocycle has ℓ Lyapunov exponents that are strictly larger than the remaining ones. Let $E(\hat{x})$ be the sum of the Oseledets subspaces associated to those largest exponents at a generic point $\hat{x} \in \hat{\Sigma}$. Then $\hat{x} \mapsto E(\hat{x})$ defines a measurable invariant section of the Grassmannian space of ℓ-dimensional subspaces of $\mathbb{C}^d$. This section is invariant along local unstable sets, meaning that

$$E(\hat{y}) = H^u_{\hat{x},\hat{y}} \cdot E(\hat{x}) \quad \text{for all } \hat{y} \in W^u_{\text{loc}}(\hat{x}) ,$$

because the hypotheses in Section 1.3 imply that

$$\hat{A}^n(\hat{y}) = H^u_{\hat{f}^n(\hat{x}),\hat{f}^n(\hat{x})} \cdot \hat{A}^n(\hat{x}) \cdot H^u_{\hat{y},\hat{x}} \quad \text{for all } n < 0 ,$$

and the norms of the unstable holonomies are bounded. Let $\hat{m}$ be the probability measure on $\hat{\Sigma} \times \text{Grass}(\ell, d)$ which projects down to $\hat{\mu}$ and has the Dirac measures $\delta_{E(\hat{x})}$ as conditional probabilities along the Grassmannian fibers. Then $\hat{m}$ is an invariant measure for the action of $\hat{A}$ on the Grassmannian bundle $\hat{\Sigma} \times \text{Grass}(\ell, d)$ and, typically, it is the unique one whose conditional probabilities are invariant under unstable holonomies.

To try and prove the theorem, we consider the space of all probability measures $\hat{m}$ on $\hat{\Sigma} \times \text{Grass}(\ell, d)$ that project down to $\hat{\mu}$, are invariant under the action of the cocycle, and whose conditional probabilities $\hat{m}_{\hat{x}}$ along the Grassmannian fibers are invariant under unstable holonomies. Proposition 4.2 ensures that such *invariant u-states* do exist. In Proposition 4.4 we prove that the projection m^u of any u-state $\hat{m}$ to $\Sigma^u \times \text{Grass}(\ell, d)$ admits conditional probabilities $m^u_{x^u}$ along the Grassmannian fibers that depend continuously on the base point x^u. This is very important for our arguments: continuity allows us to show that the kind of behavior the cocycle exhibits on the periodic point $\hat{p}$ in Definition 1.2 propagates to almost all orbits on the whole $\hat{\Sigma}$. Let us explain this.

Firstly, in Proposition 3.1, we use a simple martingale argument to show that the measure $\hat{m}$ may be recovered from m^u through

$$\hat{m}_{\hat{x}} = \lim_{n\to\infty} \hat{A}^n\big(\hat{f}^{-n}(\hat{x})\big)_* \, m^u_{P^u(\hat{f}^{-n}(\hat{x}))} \quad \hat{\mu}\text{-almost everywhere} . \tag{3}$$

The assumption that $\hat{A}^q(\hat{p})$ has ℓ largest eigenvalues implies that $\hat{A}^{qn}(\hat{p})_* \eta$ converges to the Dirac measure on the sum of the eigenspaces associated to the largest eigenvalues, for any probability measure η on $\text{Grass}(\ell, d)$ that gives zero weight to the hyperplane section defined by the other invariant subspaces. A crucial step, carried out in Section 6, is to prove that *the limit on the right hand side of* (3) *is a Dirac measure for almost every* $\hat{x}$. The proof has two main parts. In Proposition 5.1 we use the assumption that the cocycle is simple to show that the conditional probabilities of m give zero weight to hyperplane sections of the Grassmannian. Then, in Proposition 6.1, we use the continuity property in the previous paragraph, and the assumption that the cocycle is simple, to show that the behavior on the periodic point we just described does propagate to almost every orbit.

This proves that $\hat{m}_{\hat{x}} = \delta_{\xi(\hat{x})}$ almost everywhere, where $\xi(\hat{x})$ is some ℓ-subspace. In view of what we wrote before, $\xi(\hat{x})$ should correspond to the subspace $E(\hat{x})$ associated to the largest Lyapunov exponents. To prove that this is indeed so, we must also find the complementary invariant subspace. This is done by applying the previous theory to the adjoint (relative to some Hermitian form) cocycle $\hat{B} = \hat{A}^*$ over the inverse map $\hat{f}^{-1}$. Since our hypotheses are symmetric under time reversion, the same arguments as before yield an ℓ-dimensional section $\hat{x} \mapsto \xi^*(\hat{x})$ which is invariant under the action of $\hat{B}$ and under stable holonomies.

Let $\eta(\hat{x})$ be the orthogonal complement of $\xi^*(\hat{x})$. Then ξ and η are $\hat{A}$-invariant sections with complementary dimensions. Using the simplicity assumption once more, we check that $\xi(\hat{x})$ and $\eta(\hat{x})$ are transverse to each other at almost every point. The final step is to deduce from (3) that the Lyapunov exponents of $\hat{A}$ along ξ are strictly larger than those along η.

2 – Preliminary observations

Here we recall a few basic notions and prove a number of technical facts that will be useful in the sequel. The reader may be well advised to skip this section in a first reading, and then come back to it when a specific result or concept is needed.

2.1. Exterior powers and Grassmannians

Fix any $\ell \in \{1, ..., d-1\}$. The ℓth exterior power of $\mathbb{C}^d$, denoted by $\Lambda^\ell(\mathbb{C}^d)$, is the vector space of alternate ℓ-forms $\omega\colon (\mathbb{C}^d)^* \times \cdots \times (\mathbb{C}^d)^* \to \mathbb{C}$ on the dual space $(\mathbb{C}^d)^*$. It has

$$\dim \Lambda^\ell(\mathbb{C}^d) = \binom{d}{\ell}.$$

Every element of $\Lambda^\ell(\mathbb{C}^d)$ may be written as a sum of elements of the form $\omega_1 \wedge \cdots \wedge \omega_\ell$ with $\omega_i \in (\mathbb{C}^d)^{**}$. We represent by $\Lambda^\ell_v(\mathbb{C}^d)$ the subset of elements of this latter form, that we call *ℓ-vectors*. Any ℓ-vector may be written as $c\, w_1 \wedge \cdots \wedge w_\ell$, where $c \in \mathbb{C}$ and the w_i are orthogonal unit vectors (relative to any fixed Hermitian form). Hence, $\Lambda^\ell_v(\mathbb{C}^d)$ is a closed subset of $\Lambda^\ell(\mathbb{C}^d)$.

Since the bi-dual space is canonically isomorphic to $\mathbb{C}^d$, we may think of the ω_i as vectors in $\mathbb{C}^d$. Thus, there is a natural projection π_v from $\Lambda^\ell_v(\mathbb{C}^d) \backslash \{0\}$

to the Grassmannian $\mathrm{Grass}(\ell, d)$ of ℓ-dimensional subspaces of $\mathbb{C}^d$, associating to each non-zero ℓ-vector $\omega_1 \wedge \cdots \wedge \omega_\ell$ the subspace generated by $\{\omega_1, ..., \omega_\ell\}$. Two ℓ-vectors have the same image under π_v if and only if one is a multiple of the other. In other words, π_v induces a bijection between $\mathrm{Grass}(\ell, d)$ and the projective space $\mathbb{P}\Lambda_v^\ell(\mathbb{C}^d)$ of the space of ℓ-vectors.

The ℓth exterior power $\Lambda^\ell(B) \colon \Lambda^\ell(\mathbb{C}^d) \to \Lambda^\ell(\mathbb{C}^d)$ of an operator $B \colon \mathbb{C}^d \to \mathbb{C}^d$ is defined by

$$\Lambda^\ell(B)\,(\omega)\,(\phi_1, ..., \phi_\ell) = \omega(\phi_1 \circ B, ..., \phi_\ell \circ B) \ .$$

Notice that $\Lambda^\ell(B)(\omega_1 \wedge \cdots \wedge \omega_\ell) = B(\omega_1) \wedge \cdots \wedge B(\omega_\ell)$, and so $\Lambda^\ell(B)$ preserves the set $\Lambda_v^\ell(\mathbb{C}^d)$ of ℓ-vectors. Moreover, assuming B is invertible,

$$\pi_v \circ \Lambda^\ell(B) = B_\# \circ \pi_v \quad \text{on} \ \ \Lambda_v^\ell(\mathbb{C}^d) \ , \tag{4}$$

where $B_\#$ denotes the action of B on the Grassmannian.

Let H be a hyperplane, that is, a codimension 1 linear subspace of the vector space $\Lambda^\ell(\mathbb{C}^d)$. Then H may be written as

$$H = \Big\{\omega \in \Lambda^\ell(\mathbb{C}^d) \colon \ \omega \wedge \upsilon = 0\Big\}$$

for some non-zero $\upsilon \in \Lambda^{d-\ell}(\mathbb{C}^d)$. We call the hyperplane *geometric* if υ may be chosen a $(d-\ell)$-vector, that is, $\upsilon = \upsilon_{\ell+1} \wedge \cdots \wedge \upsilon_d$ for some choice of vectors υ_i in $\mathbb{C}^d = (\mathbb{C}^d)^{**}$. By definition, a *hyperplane section* of $\mathrm{Grass}(\ell, d)$ is the image under the projection π_v of the intersection of $\Lambda_v^\ell(\mathbb{C}^d)$ with some geometric hyperplane H of $\Lambda^\ell(\mathbb{C}^d)$. Note that, given any ℓ-vector $\omega = \omega_1 \wedge \cdots \wedge \omega_\ell$,

$$\omega \in H \quad \Leftrightarrow \quad \omega \wedge \upsilon = 0 \quad \Leftrightarrow \quad \pi_v(\omega) \cap \pi_v(\upsilon) \neq \{0\} \ .$$

Hence, the hyperplane section of $\mathrm{Grass}(\ell, d)$ associated to H contains precisely the ℓ-dimensional subspaces that have non-trivial intersection with the $(d-\ell)$-dimensional subspace generated by υ. The *orthogonal hyperplane section* to $V \in \mathrm{Grass}(\ell, d)$ is the hyperplane section associated to its orthogonal complement $V^\perp$.

To any Hermitian form on $\mathbb{C}^d$ there is a canonically associated one on $\Lambda^\ell(\mathbb{C}^d)$ such that the set of ℓ-vectors $e_{i_1} \wedge \cdots \wedge e_{i_\ell}$, $1 \leq i_1 < \cdots < i_\ell \leq d$ obtained from an arbitrary orthonormal basis $e_1, ..., e_d$ of the space E is an orthonormal basis of its exterior power. If B is a unitary operator then so is $\Lambda^\ell(B)$. Let $e_1, ..., e_d$ be an orthonormal basis of $\mathbb{C}^d$. We use the *polar decomposition* $B = K'DK$ of a linear isomorphism $B \colon \mathbb{C}^d \to \mathbb{C}^d$, where K and K' are unitary operators, and D is a diagonal operator (with respect to the chosen basis) with positive eigenvalues $a_1, ..., a_d$. The a_i are called *singular values* of B; we always take them to be numbered in non-increasing order.

2.2. Eccentricity of linear maps

Let $L\colon \mathbb{C}^d \to \mathbb{C}^d$ be a linear isomorphism and $1 \le \ell \le d$. The *ℓ-dimensional eccentricity* of L is defined by

$$\mathcal{E}(\ell, L) = \sup\left\{\frac{m(L \,|\, \xi)}{\|L \,|\, \xi^\perp\|} \colon\ \xi \in \operatorname{Grass}(\ell, d)\right\}, \qquad m(L \,|\, \xi) = \left\|(L \,|\, \xi)^{-1}\right\|^{-1}.$$

We call *most expanded ℓ-subspace* any $\xi \in \operatorname{Grass}(\ell, d)$ that realizes the supremum. These always exist, since the Grassmannian is compact and the expression depends continuously on ξ. These notions may be expressed in terms of the polar decomposition of L with respect to any orthonormal basis: denoting by $a_1, ..., a_d$ the eigenvalues of the diagonal operator D, in non-increasing order, then $\mathcal{E}(\ell, L) = a_\ell / a_{\ell+1}$. The supremum is realized by any subspace ξ whose image under K is a sum of ℓ eigenspaces of D such that the product of the eigenvalues is $a_1 \cdots a_\ell$. It follows that $\mathcal{E}(\ell, L) \ge 1$, and the most expanded ℓ-subspace is unique if and only if the eccentricity is larger than 1.

Let $e_1, ..., e_d$ be a basis of eigenvectors of D corresponding to the eigenvalues $a_1, ..., a_d$. For any $I \subset \{1, ..., d\}$ we represent $E_I = \bigoplus_{i \in I} e_i$. Given any $\eta \in \operatorname{Grass}(\ell, d)$ one may find a subset $I = \{i_1, ..., i_\ell\}$ of $\{1, ..., d\}$ such that η is the graph of a linear map

$$E_I \to E_J\ , \qquad e_i \mapsto \sum_{j \in J} \eta(i, j)\, e_j\ ,$$

where J is the complement of I. We say that $\eta' \in \operatorname{Grass}(\ell, d)$ is in the ε-neighborhood $B_\varepsilon(\eta)$ of η if (for some choice of I) it may also be written as the graph of a linear map from E_I to E_J such that all corresponding coefficients $\eta(i, j)$ and $\eta'(i, j)$ differ by less than ε. Given a hyperplane section H of $\operatorname{Grass}(\ell, d)$, defined by some $(d-\ell)$-vector υ, and given $\delta > 0$, we represent by H_δ the union of the hyperplane sections defined by all the $(d-\ell)$-vectors in the $B_\delta(\eta)$.

Lemma 2.1. *Given $C \ge 1$ and $\delta > 0$ there exists $\varepsilon > 0$ such that, for any $\eta \in \operatorname{Grass}(\ell, d)$ and any diagonal operator D with eccentricity $\mathcal{E}(\ell, D) \le C$, one may find a hyperplane section H of $\operatorname{Grass}(\ell, d)$ such that $D^{-1}(B_\varepsilon(\eta)) \subset H_\delta$.*

Proof: Choose $I = \{i_1, ..., i_\ell\}$ such that η is a graph over the subspace generated by $e_{i_1}, ..., e_{i_\ell}$. In other words, η admits a basis of the form

$$\left\{e_i + \sum_{j \in J} \eta(i, j) e_j \colon\ i \in I\right\},$$

where $J=\{j_1,...,j_{\ell-d}\}$ is the complement of I inside $\{1,...,d\}$. Let $a_1,...,a_d$ be the eigenvalues of D, in non-increasing order. Then

$$\left\{f_i = e_i + \sum_{j\in J} \frac{a_i}{a_j}\,\eta(i,j)\,e_j:\ i\in I\right\},$$

is a basis of $D^{-1}(\eta)$. We claim that there exist $\alpha\in I$ and $\beta\in J$ such that $a_\alpha/a_\beta\le K$: if $I=\{1,...,\ell\}$ it suffices to take $\alpha=\ell$ and $\beta=\ell+1$; otherwise, we may always choose $\beta\in\{1,...,\ell\}\backslash I$ and $\alpha\in I\backslash\{1,...,\ell\}$, and then we even have $a_\alpha/a_\beta\le 1$. This proves the claim. Now let

$$\upsilon = e_{j_1}\wedge\cdots\wedge e_{\alpha,\beta}\wedge\cdots\wedge e_{j_\ell}\ ,\qquad e_{\alpha,\beta} = e_\alpha \pm \frac{a_\alpha}{a_\beta}\,\eta(\alpha,\beta)\,e_\beta$$

be the $(d-\ell)$-vector given by the wedge products of all e_j, $j\in J$, except that e_β is replaced by $e_{\alpha,\beta}$. Notice that

$$\begin{aligned} D^{-1}(\eta)\wedge\upsilon &= f_{i_1}\wedge\cdots\wedge f_{i_\ell}\wedge e_{j_1}\wedge\cdots\wedge e_{\alpha,\beta}\wedge\cdots\wedge e_{j_\ell}\\ &= \Big[e_{i_1}\wedge\cdots\wedge e_{i_\ell}\Big]\wedge\Big[e_{j_1}\wedge\cdots\wedge\pm(a_\alpha/a_\beta)\,\eta(\alpha,\beta)\,e_\beta\wedge\cdots\wedge e_{j_\ell}\Big]\\ &\quad+\Big[e_{i_1}\wedge\cdots\wedge(a_\alpha/a_\beta)\,\eta(\alpha,\beta)\,e_\beta\wedge\cdots\wedge e_{i_\ell}\Big]\wedge\Big[e_{j_1}\wedge\cdots\wedge e_\alpha\wedge\cdots\wedge e_{j_\ell}\Big].\end{aligned}$$

Choosing the sign $\pm$ appropriately, the two terms cancel out and so $D^{-1}(\eta)\wedge\upsilon=0$. This means that $D^{-1}(\eta)$ belongs to the hyperplane section H defined by υ. In just the same way, given any η' in the ε-neighborhood of η we may find a $(d-\ell)$-vector

$$\upsilon' = e_{j_1}\wedge\cdots\wedge e'_{\alpha,\beta}\wedge\cdots\wedge e_{j_\ell}\ ,\qquad e'_{\alpha,\beta} = e_\alpha \pm \frac{a_\alpha}{a_\beta}\,\eta'(\alpha,\beta)\,e_\beta$$

such that $D^{-1}(\eta')$ belongs to the hyperplane section defined by υ'. Since $a_\alpha/a_\beta\le K$ and $|\eta(\alpha,\beta)-\eta'(\alpha,\beta)|<\varepsilon$, we have that $\upsilon'\in B_\delta(\upsilon)$ as long as ε is small enough. Then $D^{-1}(\eta')\in H_\delta$ for all η' in the ε-neighborhood of η, as claimed. ∎

Proposition 2.2. *Let $\mathcal{N}$ be a weak* compact family of probabilities on* Grass(ℓ,d) *such that all $\nu\in\mathcal{N}$ give zero weight to all hyperplane sections. Let $L_n:\mathbb{C}^d\to\mathbb{C}^d$ be linear isomorphisms such that $(L_n)_*\nu_n$ converges to a Dirac measure δ_ξ as $n\to\infty$, for some sequence ν_n in $\mathcal{N}$. Then the eccentricity $\mathcal{E}(\ell,L_n)$ goes to infinity and the image $L_n(\zeta_n^a)$ of the most expanding ℓ-subspace of L_n converges to ξ.*

Proof: Let $L_n : \mathbb{C}^d \to \mathbb{C}^d$, $\nu_n \in \mathcal{N}$, and $\xi \in \operatorname{Grass}(\ell, d)$ be as in the statement. Consider the polar decomposition $L_n = K_n' D_n K_n$, where D_n has eigenvalues $a_1, \ldots, a_d$, in non-increasing order.

We begin by reducing to the case $K_n = K_n' = \mathrm{id}$. Let $\mathcal{M} = \mathrm{U}(\ell, d)_* \mathcal{N}$, where $\mathrm{U}(\ell, d)$ is the group of transformations induced on $\operatorname{Grass}(\ell, d)$ by the unitary group. It is clear that all $\mu \in \mathcal{M}$ give zero weight to every hyperplane section of $\operatorname{Grass}(\ell, d)$. Notice also that $\mathcal{M}$ is weak* compact: given any sequence $\mu_j = (U_j)_* \nu_j$ with $\nu_j \in \mathcal{N}$ and $U_j \in \mathrm{U}(\ell, d)$, up to considering subsequences one may assume that ν_j converges to some $\nu \in \mathcal{N}$ in the weak* topology and U_j converges to some $U \in \mathrm{U}(\ell, d)$ uniformly on $\operatorname{Grass}(\ell, d)$, and then $(U_j)_* \nu_j$ converges to $U_* \nu \in \mathcal{M}$ in the weak* topology. Let $\mu_n = (K_n)_* \nu_n \in \mathcal{M}$. Then $(K_n' D_n)_* \mu_n$ converges to δ_ξ. In addition, up to considering a subsequence, we may assume that K_n' converges to some $K' \in \mathrm{U}(\ell, d)$ uniformly on $\operatorname{Grass}(\ell, d)$. Note that $\big((K')^{-1} K_n' D_n\big)_* \mu_n$ converges to δ_η, where $\eta = (K')^{-1}(\xi)$. Since $(K')^{-1} K_n'$ converges uniformly to the identity, this implies that $(D_n)_* \mu_n$ also converges to the Dirac measure at η.

Now, since $\mathcal{M}$ and the space of hyperplane sections of $\operatorname{Grass}(\ell, d)$ are compact, we may find $\delta > 0$ such that $\nu(H_\delta) < 1/2$ for every $\mu \in \mathcal{N}$ and every hyperplane section H of $\operatorname{Grass}(\ell, d)$. On the other hand, given any $\varepsilon > 0$ we have

$$\mu_n\big(D_n^{-1}(B_\varepsilon(\eta))\big) \;=\; (D_n)_* \mu_n(B_\varepsilon(\eta)) \;>\; 1/2$$

for every large n. Then $D_n^{-1}(B_\varepsilon(\eta))$ can not contained in H_δ, for any hyperplane section H. In view of Lemma 2.1, this implies that $\mathcal{E}(\ell, L_n) = \mathcal{E}(\ell, D_n)$ goes to infinity as $n \to \infty$, as claimed in the first part of the lemma.

The second part is a consequence, through similar arguments. Given any $\varepsilon > 0$, fix $\delta > 0$ small enough so that $\nu(H_\delta) < \varepsilon$ for any $\nu \in \mathcal{N}$ and any hyperplane section H of $\operatorname{Grass}(\ell, d)$. Let $H^n \subset \operatorname{Grass}(\ell, d)$ be the hyperplane section orthogonal to the most expanding direction ζ_n^a of L_n. By definition, the complement $\operatorname{Grass}(\ell, d) \setminus H_\delta^n$ of the δ-neighborhood of H^n consists of the elements of $\operatorname{Grass}(\ell, d)$ that avoid any $(d-\ell)$-subspace δ-close to $(\zeta_n^a)^\perp$. Since the eccentricity of L_n goes to infinity,

$$L_n\big(\operatorname{Grass}(\ell, d) \setminus H_\delta^n\big) \;\subset\; B_\varepsilon\big(L_n(\zeta_n^a)\big)$$

for every large n. Then, the $(L_n)_* \nu_n$-measure of $B_\varepsilon(L_n(\zeta_n^a))$ is larger than $1 - \varepsilon$. Since $(L_n)_* \nu_n$ converges to the Dirac measure at ξ, it follows that $\xi \in B_\varepsilon(L_n(\zeta_n^a))$ for every large n. As $\varepsilon > 0$ is arbitrary, this proves the second claim in the proposition. ∎

2.3. Quasi-projective maps

Let $v \mapsto [v]$ be the canonical projection from $\mathbb{C}^d$ minus the origin to the projective space $\mathbb{P}(\mathbb{C}^d)$. We call $P_\#\colon \mathbb{P}(\mathbb{C}^d) \to \mathbb{P}(\mathbb{C}^d)$ a *projective map* if there is some $P \in \mathrm{GL}(d,\mathbb{C})$ that induces $P_\#$ through $P_\#([v]) = [P(v)]$. It was pointed out by Furstenberg [6] that the space of projective maps has a natural compactification, the space of *quasi-projective maps*, defined as follows. The quasi-projective map $Q_\#$ induced by a non-zero, possibly non-invertible, linear map $Q\colon \mathbb{C}^d \to \mathbb{C}^d$ is given by $Q_\#([v]) = [Q(v_1)]$ where v_1 is any vector such that $v - v_1$ is in $\ker Q$. Observe that $Q_\#$ is defined and continuous on the complement of the projective subspace $\ker Q_\# = \big\{[v]\colon v \in \ker Q\big\}$. The space of quasi-projective maps inherits a topology from the space of non-zero linear maps, through the natural projection $Q \mapsto Q_\#$. Clearly, every quasi-projective map $Q_\#$ is induced by some linear map Q such that $\|Q\| = 1$. It follows that the space of quasi-projective maps in $\mathbb{P}(\mathbb{C}^d)$ is compact for this topology.

This notion has been extended to transformations on Grassmannian manifolds, by Gol'dsheid, Margulis [7]. Namely, one calls $P_\#\colon \mathrm{Grass}(\ell,d) \to \mathrm{Grass}(\ell,d)$ a *projective map* if there is $P \in \mathrm{GL}(d,\mathbb{C})$ that induces $P_\#$ through $P_\#(\xi) = P(\xi)$. Note that P may always be taken such that the map $\Lambda^\ell(P)$ it induces on $\Lambda^\ell(\mathbb{C}^d)$ has norm 1. Let $\mathcal{Q}$ be the closure of the set of all transformations $\Lambda^\ell(P)$ with P invertible. Since every $\Lambda^\ell(P)$ preserves the closed subset $\Lambda^\ell_v(\mathbb{C}^d)$, so does every $Q \in \mathcal{Q}$. The *quasi-projective* map $Q_\#$ induced on $\mathrm{Grass}(\ell,d)$ by a map $Q \in \mathcal{Q}$ is given by $Q_\#(\pi_v(\omega)) = \pi_v(Q(\omega))$ for any ℓ-vector ω in the complement of $\ker Q$. The space of all quasi-projective maps on $\mathrm{Grass}(\ell,d)$ inherits a topology from $\mathcal{Q}$, through the natural projection $Q \mapsto Q_\#$, and it is compact for this topology, since we may always take Q with norm equal to 1.

Lemma 2.3. *The kernel* $\ker Q_\# = \pi_v(\ker Q)$ *of any quasi-projective map is contained in some hyperplane section of* $\mathrm{Grass}(\ell,d)$.

Proof: We only have to check that $\ker Q$ is contained in a geometric hyperplane of $\Lambda^\ell(\mathbb{C}^d)$. Let P_n be any sequence of linear invertible maps such that every $\Lambda^\ell(P_n)$ has norm 1 and they converge to Q. Consider the polar decomposition $P_n = K_n' D_n K_n$ where $D_n = \mathrm{diag}[a_1^n,\dots,a_d^n]$ relative to some orthonormal basis $e_1,\dots,e_d$. Then $\Lambda^\ell(P_n) = \Lambda^\ell(K_n')\,\Lambda^\ell(D_n)\,\Lambda^\ell(K_n)$ is the polar decomposition of $\Lambda^\ell(P_n)$, where $\Lambda^\ell(D_n)$ is diagonal relative to the basis $e_{i_1}\wedge\cdots\wedge e_{i_\ell}$, $i_1<\cdots<i_\ell$ of $\Lambda^\ell(\mathbb{C}^d)$. Denote $e = e_1\wedge\cdots\wedge e_\ell$. Since the eigenvalues a_i^n, $i=1,\dots,d$, are in non-increasing order,

$$a_1^n\cdots a_\ell^n = \big\|\Lambda^\ell(D_n)(e)\big\| = \big\|\Lambda^\ell(D_n)\big\| = \big\|\Lambda^\ell(P_n)\big\| = 1\ .$$

Taking the limit over a convenient subsequence, we get that $Q=\Lambda^\ell(K')\,\mathcal{D}\Lambda^\ell(K)$ for some unitary operators K, K' and some norm 1 operator $\mathcal{D}$ diagonal with respect to the basis $e_{i_1}\wedge\cdots\wedge e_{i_\ell}$. Moreover, $\|\mathcal{D}(e)\|=1$ and the kernel of $\mathcal{D}$ is contained in the hyperplane section $H(e)$ orthogonal to e. Let $\omega=\Lambda^\ell(K)^{-1}(e)$ and $H=\Lambda^\ell(K)^{-1}(H(e))$ be the hyperplane section orthogonal to ω. Then

$$\eta\in\ker Q \;\Leftrightarrow\; \Lambda^\ell(K)\,\eta\in\ker\mathcal{D} \;\Rightarrow\; \Lambda^\ell(K)\,\eta\in H(e) \;\Leftrightarrow\; \eta\in H\ ,$$

and this proves the statement. ∎

The *weak* topology* in the space of probability measures on $\mathrm{Grass}(\ell,d)$ is characterized by the property that a sequence $(\nu_n)_n$ converges to a probability ν if and only if, given any continuous function $g\colon \mathrm{Grass}(\ell,d)\to\mathbb{R}$, the integrals $\int g\,d\nu_n$ converge to $\int g\,d\nu$. It is well-known that this topology is metrizable and compact, because the space of continuous functions on the Grassmannian contains countable dense subsets.

Lemma 2.4. *If $(P_n)_n$ is a sequence of projective maps converging to some quasi-projective map Q of* $\mathrm{Grass}(\ell,d)$*, and $(\nu_n)_n$ is a sequence of probability measures in* $\mathrm{Grass}(\ell,d)$ *converging weakly to some probability ν with $\nu(\ker Q)=0$, then $(P_n)_*\nu_n$ converges weakly to $Q_*\nu$.*

Proof: Let $(K_m)_m$ be a basis of neighborhoods of $\ker Q$ such that $\nu(\partial K_m)=0$ for all m. Given any continuous $g\colon \mathrm{Grass}(\ell,d)\to\mathbb{R}$, and given $\varepsilon>0$, fix $m\geq 1$ large enough so that $\nu(K_m)\leq\varepsilon$. Then fix $n_0\geq m$ so that $\nu_n(K_m)\leq\nu(K_m)+\varepsilon\leq 2\varepsilon$,

$$\left|\int_{K_m^c}(g\circ Q)\,d\nu_n-\int_{K_m^c}(g\circ Q)\,d\nu\right|\leq\varepsilon\quad\text{and}\quad\sup_{K_m^c}\left|g\circ P_n-g\circ Q\right|\leq\varepsilon$$

for all $n\geq n_0$. Then, splitting into integrals over K_m and over K_m^c,

$$\left|\int(g\circ P_n)\,d\nu_n-\int(g\circ Q)\,d\nu\right|\leq 2\varepsilon+3\varepsilon\sup|g|$$

for all $n\geq n_0$. This proves the lemma. ∎

For notational simplicity, in what follows we drop the subscript # and use the same symbol to represent a linear transformation and its action on any of the spaces $\mathrm{Grass}(\ell,d)$, $0<\ell<d$. In particular, we also denote by $\hat F_A$ the Grassmannian cocycles $\hat\Sigma\times\mathrm{Grass}(\ell,d)\to\hat\Sigma\times\mathrm{Grass}(\ell,d)$ defined by $\hat A$ over $\hat f$.

2.4. Bounded distortion

Let $k \geq 1$ be fixed. For each $I = (\iota_0, \ldots, \iota_{k-1})$ denote by $f_I^{u,k} : \Sigma^u \to [I]^u$ the inverse branch of $f^{u,k} = (f^u)^k$ with values in the cylinder $[I]^u = [\iota_0, \ldots, \iota_{k-1}]^u$. Moreover, define

$$J f_I^{u,k}(x^u) = \hat{\mu}_{x^u}([I]^s) \quad \text{for each } \; x^u \in \Sigma^u \;, \tag{5}$$

where $[I]^s = [\iota_0, \ldots \iota_{k-1}]^s$. The boundedness condition (1) gives

$$\frac{1}{K} \leq \frac{J f_I^{u,k}(x^u)}{J f_I^{u,k}(y^u)} \leq K \tag{6}$$

for every I and any pair of points x^u and y^u in Σ^u. This will be used in the proof of Lemma 2.6 and Corollary 4.7. Moreover, the continuity condition (2) implies that the function

$$x^u \mapsto J f_I^{u,k}(x^u) \tag{7}$$

is continuous on Σ^u, for every choice of I. In both cases, we also have dual objects and statements for inverse branches $f_I^{s,k}$ of the iterates of f^s. From (2) we also get the following fact, which will be used in the proof of Proposition 4.4.

Lemma 2.5. *Let $\Phi : \hat{\Sigma} \to \mathbb{R}$ be a bounded measurable function such that, for every fixed $x^s \in \Sigma^s$, the function $x^u \mapsto \Phi(x^s, x^u)$ is continuous at some $z^u \in \Sigma^u$. Then*

$$x^u \mapsto \int \Phi(x^s, x^u) \, d\hat{\mu}_{x^u}(x^s) \quad \text{is continuous at } \; z^u \; .$$

There is also a dual statement obtained by interchanging the roles of x^s and x^u.

Proof: Let $z^u \in \Sigma^u$ and $\varepsilon > 0$ be fixed. Define $\phi(x^s) = \Phi(x^s, z^u)$ for every $x^s \in \Sigma^s$. The continuity condition (2) gives that

$$\left| \int \phi(x^s) \, d\hat{\mu}_{x^u}(x^s) - \int \phi(x^s) \, d\hat{\mu}_{z^u}(x^s) \right| < \varepsilon \tag{8}$$

for any x^u in some neighborhood Z_0 of the point z^u. Let Z_n, $n \geq 0$ be a decreasing basis of neighborhoods of z^u. The assumption that Φ is continuous on the second variable means that for every x^s there exists some $n \geq 1$ such that

$$\left| \Phi(x^s, x^u) - \phi(x^s) \right| < \varepsilon \quad \text{for all } \; x^u \in Z_n \; .$$

Let $V(k,\varepsilon) \subset \Sigma^s$ be the set of points $x^s \in \Sigma^s$ for which we may take $n \le k$. Consider k large enough so that the $\hat\mu_{z^u}$-measure of $V(k,\varepsilon)^c$ is less than ε. Then, using condition (1),

$$\hat\mu_{x^u}\big(V(k,\varepsilon)^c\big) < K\varepsilon \quad \text{ for every } \; x^u \in \Sigma^u \;.$$

The difference $\big|\int \Phi(x^s,x^u)\, d\hat\mu_{x^u}(x^s) - \int \phi(x^s)\, d\hat\mu_{x^u}(x^s)\big|$ is bounded above by

$$\int_{V(k,\varepsilon)} \Big|\Phi(x^s,x^u) - \phi(x^s)\Big|\, d\hat\mu_{x^u}(x^s) \;+\; 2\sup|\Phi|\,\hat\mu_{x^u}\big(V(k,\varepsilon)^c\big)$$

and so, for any $x^u \in Z_k$,

$$\left|\int \Phi(x^s,x^u)\, d\hat\mu_{x^u}(x^s) - \int \phi(x^s)\, d\hat\mu_{x^u}(x^s)\right| < \varepsilon + 2K\varepsilon \sup|\Phi| \;. \tag{9}$$

Putting (8) and (9) together, we conclude that

$$\left|\int \Phi(x^s,x^u)\, d\hat\mu_{x^u}(x^s) - \int \Phi(x^s,z^u)\, d\hat\mu_{z^u}(x^s)\right| < 2\varepsilon + 2K\varepsilon \sup|\Phi|$$

for every x^u in the neighborhood Z_k of z^u. This proves the lemma. ∎

Given any measurable set $F \subset \Sigma^u$ and any $I = (\iota_0, \dots, \iota_{k-1})$, we have

$$\hat f^{-k}\big([I]^s \times F\big) \;=\; \Sigma^s \times f_I^{u,k}(F) \;=\; (P^u)^{-1}\big(f_I^{u,k}(F)\big)\,.$$

Consequently, since $\hat\mu$ is invariant under $\hat f$ and $\mu^u = P^u_*\hat\mu$,

$$\int_F Jf_I^k(x^u)\, d\mu^u(x^u) \;=\; \int_F \hat\mu_{x^u}\big([I]^s\big)\, d\mu^u(x^u) \;=\; \hat\mu\big([I]^s \times F\big) \;=\; \mu^u\big(f_I^{u,k}(F)\big)\,.$$

Thus, $Jf_I^{u,k}$ is a Radon–Nikodym derivative of the measure $F \mapsto \mu^u\big(f_I^{u,k}(F)\big)$ with respect to μ^u. An equivalent formulation is

$$\int (\psi \cdot Jf_I^{u,k})\, d\mu^u \;= \int_{[I]^u} (\psi \circ f^{u,k})\, d\mu^u \;.$$

for any bounded measurable function $\psi\colon \Sigma^u \to \mathbb{R}$, the previous equality corresponding to the case $\psi = \mathcal{X}_F$. Considering $F = \{x^u \in \Sigma^u \colon Jf_I^{u,k}(x^u) = 0\}$, we get that $Jf_I^{u,k}(f^{u,k}(z^u)) > 0$ for μ^u-almost every $z^u \in [I]^u$. Therefore,

$$Jf^{u,k}\colon \Sigma^u \to (0,+\infty)\;, \qquad Jf^{u,k}(z^u) = \frac{1}{Jf_I^{u,k}\big(f^{u,k}(z^u)\big)} \quad \text{when } z^u \in [I]^u \tag{10}$$

is well defined μ^u-almost everywhere. Moreover, given any bounded measurable function $\xi\colon [I]^u \to \mathbb{R}$ and denoting $\psi = (\xi \cdot Jf^{u,k}) \circ f_I^{u,k}$, we have that

$$\int (\xi \circ f_I^{u,k})\, d\mu^u = \int (\psi \cdot Jf_I^{u,k})\, d\mu^u = \int_{[I]^u} (\psi \circ f^{u,k})\, d\mu^u = \int (\xi \cdot Jf^{u,k})\, d\mu^u .$$

In particular, taking $\xi = \mathcal{X}_B$,

$$\mu\big(f^{u,k}(B)\big) = \int_B Jf^{u,k}\, d\mu^u \qquad \text{for every measurable } \; B \subset [I]^u .$$

In other words, $Jf^{u,k}$ is a *Jacobian* of μ^u for the kth iterate of f^u.

Lemma 2.6. *Given any $I = (\iota_0, ..., \iota_{k-1})$ and any $z^u \in [I]^u$,*

$$\hat{f}_*^k \hat{\mu}_{z^u} = Jf^{u,k}(z^u)\,\big(\hat{\mu}_{f^{u,k}(z^u)} \,|\, [I]^s\big) .$$

Moreover, a dual statement is true for $\hat{f}_*^{-k} \hat{\mu}_{z^s}$.

Proof: Let $x^u = f^{u,k}(z^u)$. Clearly, $z^u = f_I^{u,k}(x^u)$ and $\hat{f}^k$ maps $W^s_{\text{loc}}(z^u)$ bijectively to $[I]^s \times \{x^u\} \subset W^s_{\text{loc}}(x^u)$. Consider any $J = (\iota_l, ..., \iota_{-1})$, where $l < 0$, and denote $JI = (\iota_l, ..., \iota_0, ..., \iota_{k-1})$. By the definition (5),

$$\hat{\mu}_{x^u}\big([JI]^s\big) = Jf_{JI}^{u,k+l}(x^u) \quad \text{ and } \quad (\hat{f}_*^k \hat{\mu}_{z^u})\big([JI]^s\big) = \hat{\mu}_{z^u}\big([J]^s\big) = Jf_J^{u,l}(z^u) .$$

Since $f_{JI}^{u,k+l} = f_J^{u,l} \circ f_I^{u,k}$, we have that

$$Jf_{JI}^{u,k+l}(x^u) = Jf_I^{u,k}(x^u)\, Jf_J^{u,l}(z^u) \qquad \text{at } \; \mu^u\text{-almost every point} . \tag{11}$$

Using the continuity property (7), one concludes that the equality in (11) holds everywhere on $\operatorname{supp}\mu^u = \Sigma^u$. Replacing the previous pair of relations, we find that

$$\hat{\mu}_{x^u}\big([JI]^s\big) = Jf_I^{u,k}(x^u)\,(\hat{f}_*^k \hat{\mu}_{z^u})\big([JI]^s\big)$$

for every $z^u \in \Sigma^u$ and any choice of $J = (\iota_l, ..., \iota_{-1})$. This means that

$$\big(\hat{\mu}_{x^u} \,|\, [I]^s\big) = Jf_I^{u,k}(x^u)(\hat{f}_*^k \hat{\mu}_{z^u}) ,$$

which, in view of the definition (10), is just another way of writing the claim in the lemma. The dual statement is proved in just the same way. ∎

2.5. Backward averages

For each $x^u \in \Sigma^u$ and $k \geq 1$ let the *backward average* measure μ^u_{k,x^u} of the map f^u be defined on Σ^u by

$$\mu^u_{k,x^u} = \sum_{f^{u,k}(z^u)=x^u} \frac{1}{Jf^{u,k}(z^u)}\,\delta_{z^u} = \sum_I Jf_I^{u,k}(x^u)\,\delta_{f_I^{u,k}(x^u)}\,,$$

where the last sum is over all $I = (\iota_0, \ldots, \iota_{k-1})$. From (5) we get that

$$\sum_{f^{u,k}(z^u)=x^u} \frac{1}{Jf^{u,k}(z^u)} = \sum_I Jf_I^{u,k}(x^u) = \sum_I \hat{\mu}_{x^u}\big([I]^s\big) = 1 \tag{12}$$

for every $x^u \in \Sigma^u$. In other words, every μ^u_{k,x^u} is a probability measure. The definition also implies that

$$\int \mu^u_{k,x^u}(F)\,d\mu^u(x^u) = \sum_I \int_{f^{u,k}(F\cap[I]^u)} Jf_I^{u,k}\,d\mu^u = \sum_I \mu^u\big(F\cap[I]^u\big) = \mu^u(F)$$

for every measurable subset F of Σ^u. Thus,

$$\iint \psi(z^u)\,d\mu^u_{k,x^u}(z^u)\,d\mu^u(x^u) = \int \psi(x^u)\,d\mu^u(x^u) \tag{13}$$

for any bounded measurable function ψ on Σ^u. It is important to notice that the next result is stated for every (not just almost every) point x^u:

Lemma 2.7. *For every $x^u \in \Sigma^u$ and every cylinder $[J]^u \subset \Sigma^u$,*

$$K\mu^u\big([J]^u\big) \geq \limsup_n \frac{1}{n}\sum_{k=0}^{n-1}\mu^u_{k,x^u}\big([J]^u\big) \geq \liminf_n \frac{1}{n}\sum_{k=0}^{n-1}\mu^u_{k,x^u}\big([J]^u\big) \geq \frac{1}{K}\,\mu^u\big([J]^u\big)$$

Proof: Given any positive μ^u-measure set $X \subset \Sigma^u$, define

$$\mu^u_{k,X} = \frac{1}{\mu^u(X)}\int_X \mu^u_{k,z^u}\,d\mu^u(z^u)\,.$$

From the definition of the Jacobian one gets that

$$\mu^u_{k,X}(F) = \frac{1}{\mu^u(X)}\,\mu^u\big(F\cap(f^u)^{-k}(X)\big)$$

for every measurable set F and every $k \geq 1$. Since μ^u is ergodic, it follows that

$$\frac{1}{n}\sum_{k=0}^{n-1}\mu^u_{k,X}(F) \to \mu^u(F)\,. \tag{14}$$

Take $F=[J]^u$ and $X=\Sigma^u$. Assuming k is larger than the length of J, we have that $f_I^{u,k}(X)=[I]^u$ intersects $[J]^u$ if and only if it is contained in it. Then, $f_I^{u,k}(y^u)\in[J]^u$ if and only if $f_I^{u,k}(x^u)\in[J]^u$, for any $y^u\in X$. Together with (6), this implies that

$$\frac{1}{K}\le\frac{\mu^u_{k,y^u}\big([J]^u\big)}{\mu^u_{k,x^u}\big([J]^u\big)}\le K$$

for all $y^u\in X$, and so

$$\frac{1}{K}\le\frac{\mu^u_{k,X}\big([J]^u\big)}{\mu^u_{k,x^u}\big([J]^u\big)}\le K\ .$$

Combined with (14), this implies the statement of the lemma. ∎

As a direct consequence, for every cylinder $[J]^u\subset\Sigma^u$ and every $x^u\in\Sigma^u$,

$$\limsup_k\,\mu^u_{k,x^u}\big([J]^u\big)\;\ge\;K^{-1}\,\mu^u\big([J]^u\big)\,. \tag{15}$$

This fact will be used in the proof of Lemma 5.2.

2.6. Holonomy reduction

Fix an arbitrary point $x_-\in\Sigma^s$ and then, for each $\hat{x}\in\hat{\Sigma}$, denote by $\phi^u(\hat{x})$ the unique point in $W^u_{\mathrm{loc}}(x_-)\cap W^s_{\mathrm{loc}}(\hat{x})$. Using the stable holonomies in Definition 1.1, define $\hat{A}^u\colon\hat{\Sigma}\to\mathrm{GL}(d,\mathbb{C})$ by

$$\hat{A}^u(\hat{x})\;=\;H^s_{\hat{f}(\hat{x}),\phi^u(\hat{f}(\hat{x}))}\cdot\hat{A}(\hat{x})\cdot H^s_{\phi^u(\hat{x}),\hat{x}}\;=\;H^s_{\hat{f}(\phi^u(\hat{x})),\phi^u(\hat{f}(\hat{x}))}\cdot\hat{A}\big(\phi^u(\hat{x})\big)\,. \tag{16}$$

Equivalently, the cocycle $\hat{F}_{A^u}$ defined by $\hat{A}^u$ over f is conjugate to the cocycle $\hat{F}_A$ defined by $\hat{A}$ through the conjugacy

$$\Phi\colon\hat{\Sigma}\times\mathbb{C}^d\to\hat{\Sigma}\times\mathbb{C}^d\ ,\qquad\Phi(\hat{x},v)=\big(\hat{x},H^s_{\hat{x},\phi^u(\hat{x})}\big)\,.$$

Consequently, the two cocycles have the same Lyapunov exponents, and either one is simple if and only if the other one is. So, for the purpose of proving Theorem A one may replace $\hat{A}$ by either $\hat{A}^u$. On the other hand, the second equality in (16) implies that $\hat{A}^u$ is constant on every local stable set, and so

$$\hat{A}^u(\hat{x})=A^u(x^u)\qquad\text{for some}\quad A^u\colon\Sigma^u\to\mathrm{GL}(d,\mathbb{C})\ .$$

There is a dual construction, using unstable holonomies, where one finds a map $\hat{A}^s\colon\hat{\Sigma}\to\mathrm{GL}(d,\mathbb{C})$ that is constant on every local unstable set and such that the cocycle it defines over f is also conjugate to $\hat{F}_A$.

From now on, and until the end of Section 6, we consider $\hat{A}^u$ instead of $\hat{A}$. Notice that the corresponding stable holonomies are trivial

$$H^s_{\hat{x},\hat{y}} = \mathrm{id} \qquad \text{for all} \quad \hat{x} \text{ and } \hat{y} \ ,$$

because $\hat{A}^u$ is constant on local stable sets. For simplicity, we omit the superscripts u in the notations for $\hat{A}^u$, A^u, $\hat{F}_{A^u}$, Σ^u, P^u, f^u, x^u, μ^u, m^u, $H^u_{\hat{x},\hat{y}}$, $f_I^{u,k}$, etc, that is, we just represent these objects as $\hat{A}$, A, $\hat{F}_A$, Σ, P, f, x, μ, m, $H_{\hat{x},\hat{y}}$, f_I^k, etc.

3 – Convergence of conditional probabilities

Let $\hat{\pi}\colon \hat{\Sigma}\times\mathrm{Grass}(\ell,d) \to \hat{\Sigma}$ and $\pi\colon \Sigma^u\times\mathrm{Grass}(\ell,d) \to \Sigma^u$ be the natural projections. The value of $\ell \in \{1,...,d-1\}$ will be fixed till very near the end. Note that if $\hat{m}$ is an $\hat{F}_A$-invariant probability on $\hat{\Sigma}\times\mathrm{Grass}(\ell,d)$ then $m = (P\times\mathrm{id})_*\,\hat{m}$ is an F_A-invariant probability on $\Sigma\times\mathrm{Grass}(\ell,d)$. Moreover, if $\hat{\pi}_*\hat{m} = \hat{\mu}$ then $\pi_* m = \mu$. Given $\hat{x} \in \hat{\Sigma}$ we denote $x^n = P(\hat{f}^{-n}(\hat{x}))$ for $n \geq 0$.

Proposition 3.1. *Let $\hat{m}$ be any $\hat{F}_A$-invariant probability on $\hat{\Sigma}\times\mathrm{Grass}(\ell,d)$ such that $\hat{\pi}_*\hat{m} = \hat{\mu}$. Let $\{m_x\colon x \in \Sigma\}$ be a disintegration of the measure $m = (P\times\mathrm{id})_*\,\hat{m}$ along the Grassmannian fibers. Then the sequence of probability measures*

$$A^n(x^n)_*\, m_{x^n}$$

on $\mathrm{Grass}(\ell,d)$ converges in the weak* topology as $n \to \infty$, for $\hat{\mu}$-almost every $\hat{x} \in \hat{\Sigma}$.

Starting the proof, let $\mathcal{B}$ be the Borel σ-algebra of Σ. Consider the sequence $(\mathcal{B}_n)_n$ of σ-algebras of $\hat{\Sigma}$ defined by $\mathcal{B}_0 = P^{-1}(\mathcal{B})$ and $\mathcal{B}_n = \hat{f}(\mathcal{B}_{n-1})$ for $n \geq 1$. In other words, $\mathcal{B}_n$ is the σ-algebra generated by all cylinders $[\iota_{-n},...;\iota_0;...,\iota_m]$ with $m \geq 0$ and $\iota_j \in \mathbb{N}$. Fix any continuous function $g\colon \mathrm{Grass}(\ell,d) \to \mathbb{R}$. For $\hat{x} \in \hat{\Sigma}$ and $n \geq 0$, define

$$\hat{I}_n(\hat{x}) \;=\; \hat{I}_n(g,\hat{x}) \;=\; \int g\, d\big(A^n(x^n)_*\, m_{x^n}\big) \;=\; \int \big(g \circ A^n(x^n)\big)\, dm_{x^n} \ .$$

Notice that $\hat{I}_n$ is $\mathcal{B}_n$-measurable: it can be written as $\hat{I}_n = I_n \circ P \circ \hat{f}^{-n}$, where I_n is the $\mathcal{B}$-measurable function

$$I_n(x) \;=\; I_n(g,x) \;=\; \int \big(g \circ A^n(x)\big)\, dm_x \ .$$

Lemma 3.2. *For μ-almost every $x \in \Sigma$ and any $n \geq 0$ and $k \geq 1$,*

$$I_n(x) = \sum_{z \in f^{-k}(x)} \frac{1}{Jf^k(z)} I_{n+k}(z) = \int I_{n+k}(z)\, d\mu_{k,x}(z) .$$

Proof: Since the measure m is invariant under F_A^k, its disintegration must satisfy

$$m_x = \sum_{z \in f^{-k}(x)} \frac{1}{Jf^k(z)} A^k(z)_* m_z = \int \big(A^k(z)_* m_z\big)\, d\mu_{k,x}(z) \tag{17}$$

for μ-almost every $x \in \Sigma$. Then,

$$\begin{aligned} I_n(x) &= \int \big(g \circ A^n(x)\big)\, dm_x = \int \big(g \circ A^n(x)\big)\, d\left(\sum_{z \in f^{-k}(x)} \frac{1}{Jf^k(z)} A^k(z)_* m_z\right) \\ &= \sum_{z \in f^{-k}(x)} \frac{1}{Jf^k(z)} \int \big(g \circ A^{n+k}(z)\big)\, dm_z = \sum_{z \in f^{-k}(x)} \frac{1}{Jf^k(z)} I_{n+k}(z) , \end{aligned}$$

for μ-almost every $x \in \Sigma$, as claimed. ∎

The next lemma means that each $\hat{I}_n$ is the *conditional expectation* of $\hat{I}_{n+k}$ with respect to the σ-algebra $\mathcal{B}_n$ for all $k \geq 1$, and so the sequence $(\hat{I}_n, \mathcal{B}_n)_n$ is a martingale.

Lemma 3.3. *For any $n \geq 0$ and $k \geq 1$ and any $\mathcal{B}_n$-measurable function $\psi\colon \hat{\Sigma} \to \mathbb{R}$,*

$$\int \hat{I}_{n+k}(\hat{x})\, \psi(\hat{x})\, d\hat{\mu}(\hat{x}) = \int \hat{I}_n(\hat{x})\, \psi(\hat{x})\, d\hat{\mu}(\hat{x}) .$$

Proof: Let us write $\psi = \psi_n \circ P \circ \hat{f}^{-n}$, for some $\mathcal{B}$-measurable function ψ_n. Since $\hat{\mu}$ is $\hat{f}$-invariant and $\mu = P_* \hat{\mu}$,

$$\int \hat{I}_n(\hat{x})\, \psi(\hat{x})\, d\hat{\mu}(\hat{x}) = \int I_n(x)\, \psi_n(x)\, d\mu(x) . \tag{18}$$

Analogously, using the relation $\psi = (\psi_n \circ f^k) \circ P \circ \hat{f}^{-(n+k)}$,

$$\int \hat{I}_{n+k}(\hat{x})\, \psi(\hat{x})\, d\hat{\mu}(\hat{x}) = \int I_{n+k}(x)\, \psi_n\big(f^k(x)\big)\, d\mu(x) . \tag{19}$$

By Lemma 3.2, the expression on the right hand side of (18) is equal to

$$\iint I_{n+k}(z)\, d\mu_{k,x}(z)\, \psi_n(x)\, d\mu(x) = \iint I_{n+k}(z)\, \psi_n\big(f^k(z)\big)\, d\mu_{k,x}(z)\, d\mu(x) .$$

According to the relation (13), this last expression is the equal to the right hand side of (19). This proves the claim of the lemma. ∎

Proof of Proposition 3.1: By Lemma 3.3 and the martingale convergence theorem (see Durret [5]), the sequence $\hat{I}_n = \hat{I}_n(g,\cdot)$ converges $\hat{\mu}$-almost everywhere to some measurable function $\mathcal{I}(g,\cdot)$. Notice that $|\hat{I}_n(g,\hat{x})| \le \sup|g|$ for every $n \ge 1$, and so $|\mathcal{I}(g,\hat{x})|$ is also bounded above by $\sup|g|$, for $\hat{\mu}$-almost every $\hat{x} \in \hat{\Sigma}$. Considering a countable dense subset of the space of continuous functions, we find a full $\hat{\mu}$-measure set of points $\hat{x}$ such that

$$\hat{I}_n(g,\hat{x}) = \int g\, d\big(A^n(x^n)_* m_{x^n}\big) \;\to\; \mathcal{I}(g,\hat{x})$$

for every continuous function $g\colon \mathrm{Grass}(\ell,d) \to \mathbb{R}$. Let $\tilde{m}_{\hat{x}}$ be the probability measure on $\mathrm{Grass}(\ell,d)$ defined by

$$\int g\, d\tilde{m}_{\hat{x}} = \mathcal{I}(g,\hat{x}) \qquad \text{for every continuous} \quad g\colon \mathrm{Grass}(\ell,d) \to \mathbb{R}\ .$$

Then the previous relation means that $A^n(x^n)_* m_{x^n}$ converges weakly to $\tilde{m}_{\hat{x}}$. ∎

Corollary 3.4. *For $\hat{\mu}$-almost every $\hat{x} \in \hat{\Sigma}$, the limit of $A^n(x^n)_* m_{x^n}$ coincides with the conditional probability $\hat{m}_{\hat{x}}$ of the measure $\hat{m}$.*

Proof: Taking the limit $k \to \infty$ in Lemma 3.3, and using the dominated convergence theorem, we get that

$$\int \mathcal{I}(g,\hat{x})\,\psi(\hat{x})\, d\hat{\mu}(\hat{x}) = \int \hat{I}_n(g,\hat{x})\,\psi(\hat{x})\, d\hat{\mu}(\hat{x})$$

for every $\mathcal{B}_n$-measurable integrable function ψ. This may be rewritten as

$$\int \psi(\hat{x}) \int g(\xi)\, d\tilde{m}_{\hat{x}}(\xi)\, d\hat{\mu}(\hat{x}) = \int \psi(\hat{x}) \int g\big(A^n(x^n)\,\xi\big)\, dm_{x^n}(\xi)\, d\hat{\mu}(\hat{x})\ .$$

Let $\psi = \mathcal{X}_{[I]}$ be the characteristic function of a generic cylinder $[I]$ in $\mathcal{B}_n$. Changing variables $\hat{x} = \hat{f}^n(\hat{z})$, and using the fact that $\hat{\mu}$ is $\hat{f}$-invariant, we get that the right hand side of the previous equality is equal to

$$\int \mathcal{X}_{[I]}\big(\hat{f}^n(\hat{z})\big) \int g\big(A^n(z)\,\xi\big)\, dm_z(\xi)\, d\hat{\mu}(\hat{z})$$

where $z = P(\hat{z})$. Moreover, since the inner integrand $z \mapsto g(A^n(z)\,\xi)$ is constant on local stable leaves, this may be rewritten as

$$\int \mathcal{X}_{[I]}\big(\hat{f}^n(\hat{z})\big) \int g\big(A^n(\hat{z})\,\xi\big)\, d\hat{m}_{\hat{z}}(\xi)\, d\hat{\mu}(\hat{z}) = \int \mathcal{X}_{[I]}(\hat{x}) \int g(\eta)\, d\hat{m}_{\hat{x}}(\eta)\, d\hat{\mu}(\hat{x})\ .$$

In the last step we changed variables $(\hat{x},\eta) = \hat{F}_A^n(\hat{z},\xi)$ and used the fact that $\hat{m}$ is invariant under $\hat{F}_A$. Summarizing, at this point we have shown that

$$\iint \mathcal{X}_{[I]}(\hat{x})\, g(\xi)\, d\tilde{m}_{\hat{x}}(\xi)\, d\hat{\mu}(\hat{x}) = \iint \mathcal{X}_{[I]}(\hat{x})\, g(\eta)\, d\hat{m}_{\hat{x}}(\eta)\, d\hat{\mu}(\hat{x}) .$$

This relation extends immediately to linear combinations of functions $\mathcal{X}_{[I]} \times g$. Since these linear combinations form a dense subset of all bounded measurable functions on $\hat{\Sigma} \times \mathrm{Grass}(\ell, d)$, this implies that $\tilde{m}_{\hat{x}} = \hat{m}_{\hat{x}}$ for $\hat{\mu}$-almost every $\hat{x}$, as claimed. ∎

4 – Properties of u-states

Let $\hat{m}$ be a probability measure on $\hat{\Sigma} \times \mathrm{Grass}(\ell, d)$ that projects down to $\hat{\mu}$ on $\hat{\Sigma}$, in the sense that $\hat{\pi}_* \hat{m} = \hat{\mu}$. We call $\hat{m}$ a *u-state* if it admits some disintegration $\{\hat{m}_{\hat{x}} \colon \hat{x} \in \hat{\Sigma}\}$ into conditional probabilities along the fibers $\{\hat{x}\} \times \mathrm{Grass}(\ell, d)$ that is invariant under unstable holonomies:

$$\hat{m}_{\hat{y}} = (H_{\hat{x},\hat{y}})_* \, \hat{m}_{\hat{x}} \qquad \text{whenever} \quad y \in W^u_{\mathrm{loc}}(\hat{x}) .$$

We call the u-state *invariant* if, in addition, it is invariant under $\hat{F}_A$. We also call (invariant) u-states the projections $m = (P \times \mathrm{id})_* \, \hat{m}$ down to $\Sigma \times \mathrm{Grass}(\ell, d)$ of the (invariant) u-states $\hat{m}$ on $\hat{\Sigma} \times \mathrm{Grass}(\ell, d)$. Notice that $\pi_* \, m = \mu$, and m is invariant under F_A if $\hat{m}$ is invariant under $\hat{F}_A$.

Here we prove that invariant u-states m do exist. Moreover, every u-state admits some disintegration $\{m_x \colon x \in \Sigma\}$ into conditional probabilities along the fibers $\{x\} \times \mathrm{Grass}(\ell, d)$ varying continuously with the base point x, relative to the weak* topology. The formal statements are in Propositions 4.2 and 4.4. The proofs use the assumption that $\hat{\mu}$ has product structure (recall Section 1.2).

4.1. Existence of invariant u-states

Let $\mathcal{M}$ be the space of probability measures on $\hat{\Sigma} \times \mathrm{Grass}(\ell, d)$ that project down to $\hat{\mu}$ on $\hat{\Sigma}$. The *weak* topology* on $\mathcal{M}$ is the smallest topology such that the map $\eta \mapsto \int \psi \, d\eta$ is continuous, for every bounded continuous function $\psi \colon \hat{\Sigma} \times \mathrm{Grass}(\ell, d) \to \mathbb{R}$. Notice that $\mathcal{M}$ is a compact separable space for this topology. This is easy to see from the following alternative description of the topology. Let $K_n \subset \hat{\Sigma}$, $n \geq 1$, be pairwise disjoint compact sets such that $\hat{\mu}(K_n) > 0$ and $\sum \hat{\mu}(K_n) = 1$. Let $\mathcal{M}_n$ be the space of measures on $K_n \times \mathrm{Grass}(\ell, d)$ that project down to $(\hat\mu\,|K_n)$. The usual weak* topology makes $\mathcal{M}_n$ a compact separable space. Given $\eta \in \mathcal{M}$, let $\eta_n \in \mathcal{M}_n$ be obtained by restriction of η. The correspondence $\eta \mapsto (\eta_n)_n$ identifies $\mathcal{M}$ with $\prod \mathcal{M}_n$ and the product topology on $\prod \mathcal{M}_n$ corresponds to the weak* topology on $\mathcal{M}$ under this identification. Thus, the latter is a compact separable space, as claimed.

Remark 4.1. If η^j converges to η in the weak* topology then

$$\int \psi(\hat x,\xi)\, J(\hat x)\, d\eta^j(\hat x,\xi) \ \to \int \psi(\hat x,\xi)\, J(\hat x)\, d\eta(\hat x,\xi) \tag{20}$$

for any continuous function $\psi\colon \hat\Sigma\times \mathrm{Grass}(\ell,d)\to\mathbb{R}$ and any measurable bounded (or even $\hat\mu$-integrable) function $J\colon\hat\Sigma\to\mathbb{R}$. To prove this it suffices to consider the case when $J=\mathcal{X}_B$ for some measurable set B, because every bounded measurable function is the uniform limit of linear combinations of characteristic functions. Now, using that $\hat\mu$ is a regular measure (see Theorem 6.1 in [18]), we may find continuous functions $J_n\colon\hat\Sigma\to[0,1]$ such that $\hat\mu\big(\{\hat x\in\hat\Sigma\colon J_n(\hat x)\neq J(\hat x)\}\big)$ is arbitrarily small. By the definition of the topology,

$$\int \psi(\hat x,\xi)\, J_n(\hat x)\, d\eta^j(\hat x,\xi) \ \to \int \psi(\hat x,\xi)\, J_n(\hat x)\, d\eta(\hat x,\xi) \qquad \text{as} \quad j\to\infty\ .$$

This implies the convergence in (20), because corresponding terms in these two relations differ by not more than $\sup|\psi|\,\hat\mu\big(\{\hat x\in\hat\Sigma\colon J_n(\hat x)\neq J(\hat x)\}\big)$, which can be made arbitrarily small. $\square$

Remark also, for future use, that in these arguments $\hat\mu$ may be replaced by any other probability in $\hat\Sigma$.

Proposition 4.2. *There exists some invariant u-state $\hat m$ on $\hat\Sigma\times\mathrm{Grass}(\ell,d)$.*

Here is an outline of the proof. The space $\mathcal{U}$ of all u-states is non-empty and forward invariant under the cocycle. Every Cesaro weak* limit of the forward iterates of an element of $\mathcal{U}$ is an invariant u-state. The proposition follows by noting that weak* limits do exist, because $\mathcal{U}$ is compact relative to the weak* topology. The last step demands some caution, because conditional probabilities do not behave well under weak* limits, in general. We fix an arbitrary point $w\in\Sigma$ and observe that, restricted to the cylinder, the space $\mathcal{U}$ may be identified with the space $\mathcal{N}$ of probabilities on $W^s_{\mathrm{loc}}(w)\times\mathrm{Grass}(\ell,d)$ that project down to $\hat\mu_w$. Then it suffices to use that the latter space is weak* compact.

Let us fill the details. Let $\{\hat{\mu}_x : x \in \Sigma\}$ be the disintegration of $\hat{\mu}$ along local stable sets in Section 1.2. Denote by J_x the Radon–Nikodym derivatives of the conditional measure $\hat{\mu}_x$ with respect to $\hat{\mu}_w$, for each $x \in \Sigma$. According to (6), these J_x are uniformly bounded from zero and infinity. We use $\hat{x}$ and $\hat{w}$ to denote generic points in $W^s_{\mathrm{loc}}(x)$ and $W^s_{\mathrm{loc}}(w)$, respectively, with the convention that whenever they appear in the same expression they are related by

$$\hat{w} \in W^s_{\mathrm{loc}}(w) \cap W^u_{\mathrm{loc}}(\hat{x}) \ .$$

Let $\mathcal{N}$ be the space of all probability measures λ on $W^s_{\mathrm{loc}}(w) \times \mathrm{Grass}(\ell, d)$ that project down to $\hat{\mu}_w$ on $W^s_{\mathrm{loc}}(w)$. Recall, from the observation at the beginning of this section, that $\mathcal{N}$ is weak* compact. We denote by $\mathcal{U}$ the space of all u-states, that is, all probability measures η on $\hat{\Sigma} \times \mathrm{Grass}(\ell, d)$ that project down to $\hat{\mu}$ and admit some disintegration $\{\eta_{\hat{x}} : \hat{x} \in \hat{\Sigma}\}$ along the Grassmannian fibers that is invariant under unstable holonomy:

$$\eta_{\hat{x}} = (H_{\hat{w},\hat{x}})_* \, \eta_{\hat{w}} \qquad \text{for all} \quad \hat{x} \in \hat{\Sigma} \ . \tag{21}$$

Lemma 4.3. *$\mathcal{U}$ is homeomorphic to $\mathcal{N}$.*

Proof: Every $\lambda \in \mathcal{N}$ may be lifted to some $\eta \in \mathcal{U}$ in the following natural fashion: choose a disintegration $\{\lambda_{\hat{w}} : \hat{w} \in W^s_{\mathrm{loc}}(w)\}$ of λ and then let η be the measure on $\hat{\Sigma} \times \mathrm{Grass}(\ell, d)$ whose projection coincides with $\hat{\mu}$ and which admits

$$\eta_{\hat{x}} = (H_{\hat{w},\hat{x}})_* \, \lambda_{\hat{w}} \tag{22}$$

as conditional probabilities along the fibers $\{\hat{x}\} \times \mathrm{Grass}(\ell, d)$. This definition does not depend on the choice of the disintegration of λ. Indeed, let $\{\tilde{\lambda}_{\hat{w}} : \hat{w} \in W^s_{\mathrm{loc}}(w)\}$ be any other disintegration. By essential uniqueness, we have

$$\tilde{\lambda}_{\hat{w}} = \lambda_{\hat{w}} \quad \text{ for } \hat{\mu}_w\text{-almost every } \hat{w} \in W^s_{\mathrm{loc}}(w) \ .$$

Since the measures $\hat{\mu}_x$, $x \in \Sigma$, are all equivalent, it follows that $\tilde{\eta}_{\hat{x}} = \eta_{\hat{x}}$ for $\hat{\mu}_x$-almost every $\hat{x} \in W^s_{\mathrm{loc}}(x)$ and every $x \in \Sigma$. So, the lifts constructed from the two disintegrations do coincide. It is clear from the construction that $\eta \in \mathcal{U}$.

Let $\Psi : \mathcal{N} \to \mathcal{U}$, $\Psi(\lambda) = \eta$, be the map defined in this way. We are going to prove that Ψ is a homeomorphism. To prove injectivity, suppose $\Psi(\lambda) = \hat{m} = \Psi(\theta)$. By (22), this means that

$$(H_{\hat{w},\hat{x}})_* \, \lambda_{\hat{w}} \, = \, \hat{m}_{\hat{x}} \, = \, (H_{\hat{w},\hat{x}})_* \, \theta_{\hat{w}}$$

for $\hat{\mu}$-almost every $\hat{x} \in \hat{\Sigma}$. Since the conditional probabilities $\hat{\mu}_x$ are all equivalent, this is the same as $\lambda_{\hat{w}} = \theta_{\hat{w}}$ for $\hat{\mu}_\xi$-almost every $\hat{w} \in W^s_{\text{loc}}(w)$. In other words, $\lambda = \theta$. To prove surjectivity, consider any measure $\eta \in \mathcal{U}$. By definition, η admits some disintegration $\{\eta_x : x \in \Sigma\}$ satisfying (21). Define $\lambda_{\hat{w}} = (H_{\hat{x},\hat{w}})_* \eta_{\hat{x}}$ for any $\hat{x} \in W^u_{\text{loc}}(\hat{w})$, and then let λ be the measure on $W^s_{\text{loc}}(w) \times \text{Grass}(\ell, d)$ that projects down to $\hat{\mu}_w$ and has these $\lambda_{\hat{w}}$ as conditional probabilities along the fibers. Then $\lambda \in \mathcal{N}$ and $\eta = \Psi(\lambda)$.

We are left to check that Ψ is continuous. Let $\psi : \hat{\Sigma} \times \text{Grass}(\ell, d) \to \mathbb{R}$ be any bounded continuous function and let λ^j be any sequence of measures converging to some λ in $\mathcal{N}$. Using Remark 4.1,

$$\int \psi(x, \hat{x}, \xi)\, d\lambda^j_{\hat{x}}(\xi)\, d\hat{\mu}_x(\hat{x}) = \int \psi(x, \hat{x}, \xi)\, J_x(\hat{w})\, d\lambda^j_{\hat{w}}(\xi)\, d\hat{\mu}_w(\hat{w})$$

converges to

$$\int \psi(x, \hat{x}, \xi)\, d\lambda_{\hat{x}}(\xi)\, d\hat{\mu}_x(\hat{x}) = \int \psi(x, \hat{x}, \xi)\, J_x(\hat{w})\, d\lambda_{\hat{w}}(\xi)\, d\hat{\mu}_w(\hat{w})$$

as $j \to \infty$, for every $x \in \Sigma$. Integrating with respect to μ, and using the bounded convergence theorem, we get that

$$\iint \psi(x, \hat{x}, \xi)\, d\lambda^j_{\hat{x}}(\xi)\, d\hat{\mu}_x(\hat{x})\, d\mu(x) \to \iint \psi(x, \hat{x}, \xi)\, d\lambda_{\hat{x}}(\xi)\, d\hat{\mu}_x(\hat{x})\, d\mu(x)$$

as $j \to \infty$. This means that $\Psi(\lambda^j)$ converges to $\Psi(\lambda)$ as $j \to \infty$. ∎

Proof of Proposition 4.2: In view of the previous lemma, $\mathcal{U}$ is non-empty and compact relative to the weak* topology. Moreover, $\mathcal{U}$ is invariant under iteration by $\hat{F}_A$: this follows from the invariance property (b) in Section 1.3 for unstable holonomies, together with the fact that local unstable sets are mapped inside local unstable sets by the inverse of $\hat{f}$. Consider any probability measure $\bar{m} \in \mathcal{U}$. The sequence

$$\hat{m}_n = \frac{1}{n} \sum_{j=0}^{n-1} (\hat{F}^j_A)_* \bar{m}$$

has accumulation points $\hat{m}$ in $\mathcal{U}$. Since $\hat{F}_A$ is a continuous map, the push-forward operator $(\hat{F}_A)_*$ is continuous relative to the weak* topology. It follows that any such accumulation point is $\hat{F}_A$-invariant and, consequently, an invariant u-state. ∎

4.2. Continuity of conditional probabilities

Now we prove that conditional probabilities of u-states along the Grassmannian fibers depend continuously on the base point:

Proposition 4.4. *Any u-state m in $\Sigma \times \mathrm{Grass}(\ell, d)$ admits some disintegration $\{m_x : x \in \Sigma\}$ into conditional probabilities along the Grassmannian fibers varying continuously with $x \in \Sigma$ in the weak* topology.*

This continuous disintegration is necessarily unique, because disintegrations are essentially unique and μ is supported on the whole Σ. For the proof of the proposition we need the following simple observation:

Lemma 4.5. *Let $\{\hat{m}_{\hat{x}} : \hat{x} \in \hat{\Sigma}\}$ be a disintegration along $\{\{\hat{x}\} \times \mathrm{Grass}(\ell, d) : \hat{x} \in \hat{\Sigma}\}$ of some probability measure $\hat{m}$ on $\hat{\Sigma} \times \mathrm{Grass}(\ell, d)$ such that $\hat{\pi}_* \hat{m} = \hat{\mu}$. Then*

$$m_x = \int \hat{m}_{\hat{x}} \, d\hat{\mu}_x(\hat{x})$$

is a disintegration of $m = (P \times \mathrm{id})_* \hat{m}$ along $\{\{x\} \times \mathrm{Grass}(\ell, d) : x \in \Sigma\}$.

Proof: For any $\varphi : \Sigma \times \mathrm{Grass}(\ell, d) \to \mathbb{R}$ and $\hat{\varphi} = \varphi \circ (P \times \mathrm{id})$,

$$\begin{aligned}
\iint \varphi \, dm_x \, d\mu(x) &= \iint \Big(\int \varphi(x, v) \, d\hat{m}_{\hat{x}}(v) \, d\hat{\mu}_x(\hat{x}) \Big) \, d\mu(x) \\
&= \iint \Big(\int \hat{\varphi}(x, v) \, d\hat{m}_{\hat{x}}(v) \Big) \, d\hat{\mu}_x(\hat{x}) \, d\mu(x) \\
&= \int \Big(\int \hat{\varphi}(x, v) \, d\hat{m}_{\hat{x}}(v) \Big) \, d\hat{\mu}(\hat{x}) \ = \int \hat{\varphi} \, d\hat{m} \ = \int \varphi \, dm
\end{aligned}$$

and this proves that $\{m_x : x \in \Sigma\}$ is a disintegration of m. ∎

Remark 4.6. For u-states this gives that, for any measurable set $E \subset \mathrm{Grass}(\ell, d)$,

$$m_x(E) = \int \hat{m}_{\hat{x}}(E) \, d\hat{\mu}_x(\hat{x}) \ = \int \hat{m}_{\hat{y}}\big(H^s_{\hat{x},\hat{y}}(E)\big) \, \frac{d\hat{\mu}_x}{d\hat{\mu}_y}(\hat{y}) \, d\hat{\mu}_y(\hat{y})$$

for any pair of points x and y in the same cylinder. When the cocycle is *locally constant* the stable holonomies $H^s_{\hat{x},\hat{y}} = \mathrm{id}$. In this case it immediately follows that the conditional probabilities m_x and m_y are all equivalent. Moreover, their

Radon–Nikodym derivatives are uniformly bounded, as a consequence of the boundedness condition (1). Starting from this observation, in the appendix of [1] we give a version of Theorem A for locally constant cocycles that does not require the continuity hypothesis (2). □

Proof of Proposition 4.4: Let $\{\hat m_{\hat x} : \hat x \in \hat\Sigma\}$ be a disintegration of $\hat m$ into conditional probabilities that are invariant under unstable holonomies: $\hat m_{\hat y} = (H_{\hat x,\hat y})_* \hat m_{\hat x}$ for every $\hat x$, $\hat y$ in the same local unstable set. Let $\{m_x : x \in \Sigma\}$ be the disintegration of m given by Lemma 4.5. For any continuous $g\colon \mathrm{Grass}(\ell, d) \to \mathbb{R}$ and any points x and y in the same cylinder of Σ, we have

$$\int g(\xi)\, dm_x(\xi) = \iint g(\xi)\, d\hat m_{\hat x}(\xi)\, d\hat\mu_x(\hat x) = \iint g\big(H_{\hat y,\hat x}(\eta)\big)\, d\hat m_{\hat y}(\eta)\, d\hat\mu_x(\hat x)$$

where $\hat y$ denotes the unique point in $W^s_{\mathrm{loc}}(y) \cap W^u_{\mathrm{loc}}(\hat x)$. Fix y and consider the function

$$\Phi(x^s, x^u) = \int g\big(H_{\hat y,\hat x}(\eta)\big)\, d\hat m_{\hat y}(\eta), \qquad \text{where } \hat x = (x^s, x^u)\ .$$

It is clear that Φ is measurable and bounded by the $\sup |g|$. Moreover, it is continuous on x^u for each fixed x^s. To check this it suffices to note that $\hat m_{\hat y}$ does not depend on x^u, while the function g and the holonomies depend continuously on $\hat x$ (recall Definition 1.1). It follows from Lemma 2.5 that

$$x \mapsto \int g(\xi)\, dm_x(\xi) = \int \Phi(x^s, x)\, d\hat\mu_x(x^s)$$

is continuous. This proves the claim of the proposition. ∎

Corollary 4.7. *If m is an invariant u-state and $\{m_x : x \in \Sigma\}$ is the continuous disintegration of m, then*

$$m_x = \sum_{z \in f^{-k}(x)} \frac{1}{Jf^k(z)} A^k(z)_* m_z = \int A^k(z)_* m_z\, d\mu_{k,x}(z)$$

for every $x \in \Sigma$ and every $k \geq 1$.

Proof: The second equality is just the definition of the backward averages, see Section 2.5. As for the first equality, it must hold for every $k \geq 1$ and μ-almost every x, because m is invariant under f. Moreover, all the expressions involved vary continuously with $x \in \Sigma$: this follows from Proposition 4.4, property (7), and our assumption that the cocycle is continuous. Hence, the first equality must hold at every point of $\operatorname{supp} \mu = \Sigma$. ∎

Corollary 4.8. *If $\{\hat{m}_{\hat{x}} : \hat{x} \in \hat{\Sigma}\}$ is a disintegration of an invariant u-state $\hat{m}$ into conditional probabilities invariant under unstable holonomies then*

$$\hat{m}_{\hat{f}^n(\hat{x})} = A^n(x)_* \, \hat{m}_{\hat{x}}$$

for every $n \geq 1$, every $x \in \Sigma$, and $\hat{\mu}_x$-almost every $\hat{x} \in W^s_{\mathrm{loc}}(x)$.

Proof: Since $\hat{m}$ is $\hat{F}_A$-invariant, the equality is true for all $n \geq 1$ and $\hat{\mu}$-almost all $\hat{z} \in \hat{\Sigma}$ or, equivalently, for $\hat{\mu}_z$-almost every $\hat{z} \in W^s_{\mathrm{loc}}(z)$ and μ-almost every $z \in \Sigma$. Consider an arbitrary point $x \in \Sigma$. Since μ is positive on open sets, x may be approximated by points z such that

$$\hat{m}_{\hat{f}^n(\hat{z})} = A^n(z)_* \, \hat{m}_{\hat{z}}$$

for every $n \geq 1$ and $\hat{\mu}_z$-almost every $\hat{z} \in W^s_{\mathrm{loc}}(z)$. Since the conditional probabilities of $\hat{m}$ are invariant under unstable holonomies, it follows that

$$\hat{m}_{\hat{f}^n(\hat{x})} = \Big(H_{\hat{f}^n(z), \hat{f}^n(x)}\Big)_* A^n(z)_* \, \hat{m}_{\hat{z}} \;=\; A^n(x)_* \left(H_{\hat{z},\hat{x}}\right)_* \hat{m}_{\hat{z}} \;=\; A^n(x)_* \, \hat{m}_{\hat{x}}$$

for $\hat{\mu}_z$-almost every $\hat{z} \in W^s_{\mathrm{loc}}(z)$, where $\hat{x}$ is the unique point in $W^s_{\mathrm{loc}}(x) \cap W^u_{\mathrm{loc}}(\hat{z})$. Since the measures $\hat{\mu}_x$ and $\hat{\mu}_z$ are equivalent, this is the same as saying that the last equality holds for $\hat{\mu}_x$-almost every $\hat{x} \in W^s_{\mathrm{loc}}(x)$, as claimed. ∎

5 – Invariant measures of simple cocycles

In this section we prove that invariant u-states of simple cocycles are fairly smooth along the Grassmannian fibers: they give zero weight to every hyperplane section.

Proposition 5.1. *Suppose that $\hat{A}$ is simple. Let m be any invariant u-state in $\Sigma \times \mathrm{Grass}(\ell, d)$ and $\{m_x : x \in \Sigma\}$ be the continuous disintegration of m. Then $m_x(V) = 0$ for every $x \in \Sigma$ and any hyperplane section V of $\mathrm{Grass}(\ell, d)$.*

In Section 5.1 we argue by contradiction to reduce the proof of Proposition 5.1 to Proposition 5.5, a combinatorial result about intersections of hyperplane sections. The latter is proved in Section 5.2. See also Appendix B.

5.1. Smoothness of conditional probabilities

Suppose there is some point of Σ and some hyperplane section of the corresponding Grassmannian fiber which has positive conditional measure. Let $\gamma_0 > 0$ be the supremum of the values of $m_x(V) \geq \gamma$ over all $x \in \Sigma$ and all hyperplane sections V. The supremum is attained at every point:

Lemma 5.2. *For every $x \in \Sigma$ there exists some hyperplane section V of* $\operatorname{Grass}(\ell, d)$ *such that $m_x(V) = \gamma_0$.*

Proof: Fix any cylinder $[J] \subset \Sigma$ and any positive constant $c < \mu([J])/K$, where K is the constant in (15). Let $z \in \Sigma$ and V be a hyperplane section with $m_z(V) > 0$. For each $y \in f^{-k}(z)$, let $V_y = A^k(y)^{-1}(V)$. By Corollary 4.7,

$$m_z(V) = \int m_y(V_y)\, d\mu_{k,z}(y) \leq \mu_{k,z}\big([J]\big) \sup\big\{m_y(V_y) \colon y \in [J]\big\} + \big(1 - \mu_{k,z}([J])\big)\, \gamma_0 .$$

By (15), there exist arbitrarily large values of k such that $\mu_{k,z}([J]) \geq c$. Then

$$m_z(V) \leq c \sup\big\{m_y(V_y) \colon y \in [J]\big\} + (1 - c)\, \gamma_0 .$$

Varying the point $z \in \Sigma$ and the hyperplane section V, we can make the left hand side arbitrarily close to γ_0. It follows that

$$\sup\big\{m_y(V_y) \colon y \in [J]\big\} \geq \gamma_0 .$$

This proves that the supremum over any cylinder $[J]$ coincides with γ_0. Then, given any $x \in \Sigma$ we may find a sequence $x_n \to x$ and hyperplane sections V_n such that $m_{x_n}(V_n) \to \gamma_0$. Moreover, we may assume that V_n converges to some hyperplane section V in the Hausdorff topology. Given any neighborhood U of V, we have $m_{x_n}(U) \geq m_{x_n}(V_n)$ for all large n. By Proposition 4.4, the conditional probabilities m_{x_n} converge weakly to m_x. Assuming U is closed, it follows that

$$m_x(U) \geq \limsup_n m_{x_n}(U) \geq \limsup_n m_{x_n}(V_n) \geq \gamma_0 .$$

Making $U \to V$, we conclude that $m_x(V) \geq \gamma_0$. This proves that the supremum γ_0 is realized at x, as claimed. ∎

Lemma 5.3. *For any $x \in \Sigma$ and any hyperplane section V of* $\operatorname{Grass}(\ell, d)$, *we have $m_x(V) = \gamma_0$ if and only if $m_y\big(A(y)^{-1} V\big) = \gamma_0$ for every $y \in f^{-1}(x)$.*

Proof: This is a direct consequence of Corollary 4.7 and the relation (12): for every $x \in \Sigma$,

$$m_x(V) = \sum_{y \in f^{-1}(x)} \frac{1}{Jf(y)}\, m_y\big(A(y)^{-1}V\big) \quad \text{and} \quad \sum_{y \in f^{-1}(x)} \frac{1}{Jf(y)} = 1 \, .$$

Since γ_0 is the maximum value of the measure of any hyperplane section, we get that $m_x(V) = \gamma_0$ if and only if $m_y(A(y)^{-1}V) = \gamma_0$ for every $y \in f^{-1}(x)$, as stated. ∎

Lemma 5.4. *For any $x \in \Sigma$ and any hyperplane section V of* Grass(ℓ, d), *we have $\hat m_{\hat x}(V) \le \gamma_0$ for $\hat\mu_x$-almost every $\hat x \in W^s_{\mathrm{loc}}(x)$. Hence, $m_x(V) = \gamma_0$ if and only if $\hat m_{\hat x}(V) = \gamma_0$ for $\hat\mu_x$-almost every $\hat x \in W^s_{\mathrm{loc}}(x)$.*

Proof: Suppose there is $y \in \Sigma$, a hyperplane section V, a constant $\gamma_1 > \gamma_0$, and a positive $\hat\mu$-measure subset X of $W^s_{\mathrm{loc}}(y)$ such that $\hat m_{\hat y}(V) \ge \gamma_1$ for every $\hat y \in X$. For each $m < 0$, consider the partition of $W^s_{\mathrm{loc}}(y) \approx \Sigma^s$ determined by the cylinders $[I]^s = [\iota_m, \dots, \iota_{-1}]^s$, with $\iota_j \in \mathbb{N}$. Since these partitions generate the σ-algebra of the local stable set, given any $\varepsilon > 0$ we may find m and I such that

$$\hat\mu_y\big(X \cap [I]^s\big) \ge (1-\varepsilon)\, \hat\mu_y\big([I]^s\big) \, .$$

Observe that $[I]^s \approx [I]^s \times \{y\}$ coincides with $\hat f^n\big(W^s_{\mathrm{loc}}(x)\big)$, where $x = f_I^{-n}(y)$. So, using also Lemma 2.6,

$$\hat\mu_x\big(\hat f^{-n}(X) \cap W^s_{\mathrm{loc}}(x)\big) = (\hat f^n_* \hat\mu_x)\big(X \cap [I]^s\big) = J_\mu f^n(x)\, \hat\mu_y\big(X \cap [I]^s\big) \, .$$

By the previous inequality and Lemma 2.6, this is bounded below by

$$(1-\varepsilon)\, J_\mu f^n(x)\, \hat\mu_y\big([I]^s\big) = (1-\varepsilon)\, (\hat f^n_* \hat\mu_x)\big([I]^s\big) = \hat\mu_x\big(W^s_{\mathrm{loc}}(x)\big) = 1 - \varepsilon \, .$$

In this way we have shown that

$$\hat\mu_y\big(\hat f^{-n}(X) \cap W^s_{\mathrm{loc}}(x)\big) \ge (1-\varepsilon) \, .$$

Fix $\varepsilon > 0$ small enough so that $(1-\varepsilon)\,\gamma_1 > \gamma_0$. Using Corollary 4.8, we find that

$$\hat m_{\hat x}\big(A^n(x)^{-1}V\big) = \hat m_{\hat y}(V) \ge \gamma_1$$

for $\hat\mu_x$-almost every $\hat x \in \hat f^{-n}(X) \cap W^s_{\mathrm{loc}}(x)$. It follows that

$$m_x\big(A^n(x)^{-1}V\big) = \int \hat m_{\hat x}\big(A^n(x)^{-1}V\big)\, d\hat\mu_x(\hat x) \ge (1-\varepsilon)\,\gamma_1 > \gamma_0 \, ,$$

which contradicts the definition of γ_0. This contradiction proves the first part of the lemma. The second one is a direct consequence, using the fact that $m_x(V)$ is the $\hat{\mu}_x$-average of all $\hat{m}_{\hat{x}}(V)$. ∎

Before we proceed, let us introduce some useful terminology. Recall that a hyperplane section V of $\mathrm{Grass}(\ell, d)$ is the image of $\Lambda^\ell_v(\mathbb{C}^d) \cap H$ under the projection π_v, where H is the geometric hyperplane of $\Lambda^\ell(\mathbb{C}^d)$ defined by some non-zero element $\upsilon \in \Lambda^{d-\ell}_v(\mathbb{C}^d)$. Notice that, for any linear isomorphism B of $\mathbb{C}^d$,

$$\Lambda^\ell(B)\,H = \Big\{\omega\colon\, \omega \wedge \Lambda^{d-\ell}(B)(\upsilon) = 0\Big\} \quad \text{and} \quad B(V) = \pi_v\Big(\Lambda^\ell_v(\mathbb{C}^d) \cap \Lambda^\ell(B)\,H\Big).$$

Suppose B is diagonalizable. Then we say V is *invariant* for B if the subspace $\pi_v(\upsilon)$ is a sum of eigenspaces of B. Likewise, we say V *contains no eigenspace* of B if $\pi_v(\upsilon)$ intersects any sum of ℓ eigenspaces of B at the origin only or, equivalently, if H contains no ℓ-vector ω such that $\pi_v(\omega)$ is a sum of eigenspaces of B. A subset J of $\{0, 1, ..., N-1\}$ is called *ε-dense* if $\#J \geq N\varepsilon$.

Proposition 5.5. *For any $\varepsilon > 0$ there exists $N \geq 1$ such that*

$$\bigcap_{j \in J} B^j(V) = \emptyset$$

for every ε-dense set $J \subset \{0, 1, ..., N-1\}$, every linear isomorphism $B\colon \mathbb{C}^d \to \mathbb{C}^d$ whose eigenvalues all have distinct absolute values, and every hyperplane section V of $\mathrm{Grass}(\ell, d)$ *containing no eigenspace of B.*

This proposition will be proved in Section 5.2. Right now let us explain how it can be used to finish the proof of Proposition 5.1.

Fix a periodic point $\hat{p} \in \hat{\Sigma}$ of $\hat{f}$ and a homoclinic point $\hat{z} \in \hat{\Sigma}$ as in Definition 1.2. Let $p = P(\hat{p})$ be the corresponding periodic point of f and let $z = P(\hat{z})$. By Lemma 5.2, we may find a hyperplane section V of $\mathrm{Grass}(\ell, d)$ with $m_p(V) = \gamma_0$. Write $V = \pi_v\big(\Lambda^\ell_v(\mathbb{C}^d) \cap H\big)$, where H is the geometric hyperplane defined by some non-zero $(d{-}\ell)$-vector υ. Let $V^n = A^{-nq}(p)\,V$ and H^n be the geometric hyperplane defined by $A^{-nq}(p)\,\upsilon$. Then, $V^n = \pi_v\big(\Lambda^\ell_v(\mathbb{C}^d) \cap H^n\big)$ for each $n \geq 0$. Since all the eigenvalues of $A^q(p)$ have distinct absolute values, $A^{-nq}(p)\,\upsilon$ converges to some $(d-\ell)$-vector υ_1 such that $\pi_v(\upsilon_1)$ is a sum of eigenspaces of $A^q(p)$. This means that V^n converges to $V_1 = \pi_v\big(\Lambda^\ell_v(\mathbb{C}^d) \cap H_1\big)$, where $H_1 = \big\{\omega\colon\, \omega \wedge \upsilon_1 = 0\big\}$ is the geometric hyperplane section defined by υ_1. On the other hand, using Lemma 5.3 we find that $m_p(V^n) = \gamma_0$ for all $n \geq 0$. By lower semi-continuity of the measure, it follows that $m_p(V_1) = \gamma_0$. Note that V_1 is invariant for $A^q(p)$. This shows that we may suppose, right from the start, that V is invariant for $A^q(p)$.

Now define $W = A^l(z)^{-1} V$. From Lemmas 5.3 and 5.4 we get that $m_z(W) = \gamma_0$ and $\hat{m}_\zeta(W) = \gamma_0$ for $\hat{\mu}_z$-almost every $\zeta \in W^s_{\mathrm{loc}}(z)$. For each $\eta \in W^s_{\mathrm{loc}}(p)$, define $W_\eta = H_{\zeta,\eta}(W)$, where ζ is the unique point in $W^u_{\mathrm{loc}}(\eta) \cap W^s_{\mathrm{loc}}(\hat{z})$. Since $\hat{m}$ is a u-state and the measures $\hat{\mu}_z$ and $\hat{\mu}_p$ are equivalent, we have $\hat{m}_\eta(W_\eta) = \hat{m}_\zeta(W) = \gamma_0$ for $\hat{\mu}_p$-almost every η. For each $j \geq 0$, let

$$W^j_\eta = A^{-jq}(p)\, W_{\hat{f}^{jq}(\eta)} \qquad \big(\text{in particular, } \ W^j_{\hat{p}} = A^{-jq}(q)\,(W_{\hat{p}})\big)\ .$$

Using Corollary 4.8, we get that $\hat{m}_\eta(W^j_\eta) = \hat{m}_{\hat{f}^{jq}(\eta)}(W_{\hat{f}^{jq}(\eta)}) = \gamma_0$ for every $j \geq 0$ and $\hat{\mu}_p$-almost every η. It is clear that every W^j_η is an ℓ-dimensional projective subspace. Moreover, it depends continuously on η, for each fixed j, because unstable holonomies vary continuously with the base points (Definition 1.1). Notice that

$$W_{\hat{p}} = H_{\hat{z},\hat{p}}\, A^l(z)^{-1} V = \psi^{-1}_{p,z} V \qquad \big(\text{recall } H_{\hat{z},\hat{p}} = H^u_{\hat{z},\hat{p}} \text{ and } H^s_{\hat{p},\hat{f}^l(\hat{z})} = \mathrm{id}\big)\,.$$

Thus, the second condition in Definition 1.2 implies that $W_{\hat{p}}$ contains no eigenspace of $A^q(p)$.

Taking $\varepsilon = \gamma_0$, $V = W_{\hat{p}}$, $B = A^q(p)$ in Proposition 5.5 we find $N \geq 1$ such that

$$\bigcap_{j \in J} W^j_{\hat{p}} = \emptyset \qquad \text{for every } \ \gamma_0\text{-dense subset } J \text{ of } \{0, 1, \dots, N-1\}\ .$$

Since the family of sets J is finite, we may use continuity to conclude that

$$\bigcap_{j \in J} W^j_\eta = \emptyset \qquad \text{for every } \ \gamma_0\text{-dense subset } J \text{ of } \{0, 1, \dots, N-1\} \tag{23}$$

and any η in a neighborhood of $\hat{p}$ inside the local stable set. On the other hand, the fact that $\hat{m}_\eta(W^j_\eta) = \gamma_0$ for all $j \geq 0$ implies (use a Fubini argument) that there exists some $\omega \in \mathrm{Grass}(\ell, d)$ such that the set

$$J = \Big\{0 \leq j \leq N-1 :\ \omega \in W^j_\eta\Big\}$$

is γ_0-dense in $\{0, 1, \dots, N-1\}$. This contradicts (23). This contradiction shows that we have reduced the proof of Proposition 5.1 to proving Proposition 5.5.

5.2. Intersections of hyperplane sections

Now we prove Proposition 5.5. We say that $I \subset \mathbb{N}$ is a *k-cube* of sides $c_1, ..., c_k \in \mathbb{N}$ based on $c \in \mathbb{N} \cup \{0\}$ if I is the set of all $x \in \mathbb{N}$ that can be written as $x = c + \sum_i a_i c_i$ with $a_i = 0$ or $a_i = 1$. We shall need the following couple of lemmas on k-cubes.

Lemma 5.6. *Let $H \subset \mathbb{C}^d$ be a codimension 1 subspace, $B\colon \mathbb{C}^d \to \mathbb{C}^d$ be a linear isomorphism, and I be a k-cube for some $1 \le k \le d$. If $H(I) = \bigcap_{i \in I} B^i(H)$ has codimension at most k then there exists a subcube $I' \subset I$ and an integer $l \ge 1$ such that $B^l\big(H(I')\big) = H(I')$.*

Proof: The proof is by induction on k. The case $k=1$ is easy. Indeed, the 1-cube $I = \{c, c+c_1\}$ and so $H(I) = B^c(H) \cap B^{c+c_1}(H)$. Since H has codimension 1, if $H(I)$ has codimension at most 1 then all the subspaces involved must coincide:

$$H(I) = B^c(H) = B^{c+c_1}(H) \ ,$$

and this gives the claim with $l = c_1$ and $I' = \{c\}$. Now assume the statement holds for $k-1$. Let I be a k-cube of sides $c_1, ..., c_k$ based on c. Let I_1 and I_2 be the $(k-1)$-cubes of sides $c_1, ..., c_{k-1}$ based on c and on $c + c_k$, respectively. Then $I = I_1 \cup I_2$. If either $H(I_1)$ or $H(I_2)$ has codimension at most $k-1$, then the conclusion follows from the induction hypothesis. Otherwise, both $H(I^1)$ and $H(I^2)$ have codimension at least k. Since their intersection $H(I)$ has codimension at most k, they must all coincide:

$$H(I) = H(I_1) = H(I_2) = B^{c_k}\big(H(I_1)\big)$$

and the conclusion follows, with $l = c_k$ and $I' = I_1$. ∎

Lemma 5.7. *For every $\varepsilon > 0$ and $k \ge 1$ there exists $\delta > 0$ such that for all sufficiently large $N \ge 1$ the following holds: for every ε-dense subset J of $\{0, 1, ..., N-1\}$ there exist $c_1, ..., c_k \in \mathbb{N}$ and a δ-dense subset J^k of $\{0, 1, ..., N-1\}$ such that for every $c \in J^k$ the set J contains the k-cube with sides $c_1, ..., c_k$ based on c.*

Proof: The proof is by induction on k. Let us start with the case $k = 1$. Let a_j, $j = 1, ..., \#J$ be the elements of J, in increasing order. By assumption, $\#J \ge \varepsilon N$. Then, clearly,

$$\frac{1}{\#J-1} \sum_{i=1}^{\#J-1} a_{i+1} - a_i \ \le \ \frac{N-1}{\#J-1} \ \le \ \frac{2N}{\#J} \ \le \ \frac{2}{\varepsilon}$$

(assume N is large enough so that $\#J \geq \varepsilon N \geq 2$). Then at least half of these differences are less than twice the average: there exists $I' \subset \{1, ..., \#J-1\}$ with $\#I' \geq (\#J-1)/2 \geq \#J/4$ such that $a_{i+1} - a_i \leq 4/\varepsilon$ for all $i \in I'$. Then there must be some $c_1 \geq 1$ and a subset I'' of I' such that

$$a_{i+1} - a_i = c_1 \text{ for all } i \in I'' \quad \text{and} \quad \#I'' \geq \frac{\varepsilon\,\#I'}{4} \geq \frac{\varepsilon\,\#J}{16} \geq \frac{\varepsilon^2}{16} N\,.$$

It follows that $\delta = \varepsilon^2/16$ and $J^1 = \{a_i : i \in I''\}$ satisfy the conclusion of the lemma for $k = 1$.

Now assume the conclusion holds for $k-1$. Then there exists $\delta_{k-1} = \delta(\varepsilon) > 0$, positive integers $c_1, ..., c_{k-1}$, and a δ_{k-1}-dense subset J^{k-1} of $\{0, 1, ..., N-1\}$ such that J contains a $(k-1)$-cube of sides $c_1, ..., c_{k-1}$ based on every $c \in J^{k-1}$. Applying case $k=1$ of the lemma with δ_{k-1} in the place of ε, we find $\delta = \delta(\varepsilon) > 0$, a positive integer c_k, and a δ-dense subset J^k of $\{0,1,...,N-1\}$ such that $\{c, c+c_k\}$ $\in J^{k-1}$ for every $c \in J^k$. Then $c_1, ..., c_k$ and J^k satisfy the conclusion of the lemma. ∎

We now conclude the proof of the proposition. Fix $k = \dim \Lambda^\ell(\mathbb{C}^d) - 1$. Assume N is large enough so that Lemma 5.7 applies. It follows from the lemma that J contains some k-cube I. Let H be the geometric hyperplane corresponding to V. If $\bigcap_{i \in I} B^j(V) \subset \mathrm{Grass}(\ell, d)$ is not empty then $H(I) = \bigcap_{i \in I} B^i(H)$ has positive dimension, that is, its codimension in $\Lambda^\ell(\mathbb{C}^d)$ is at most k. So, Lemma 5.6 implies that there exists a subcube $I' \subset I$ and an integer $l \geq 1$ such that $H(I')$ is invariant under B^l. Thus, $\bigcap_{i \in I'} B^i(V) \subset \mathrm{Grass}(\ell, d)$ is non-empty and invariant under B^l. Since all the eigenvalues of B have different absolute values, for every ℓ-subspace $W \subset \mathbb{C}^d$ we have that $B^j(W)$ converges to a sum of eigenspaces of B as $j \to \infty$. Since $\bigcap_{i \in I'} B^j(V)$ is non-empty, invariant, and closed, we conclude that it contains some sum of eigenspaces of B. In particular, V contains a sum of eigenspaces of B, which contradicts the hypothesis. This contradiction proves Proposition 5.5.

6 – Convergence to a Dirac measure

In this section we prove that, for simple cocycles, the limit of the iterates of any invariant u-state m is a Dirac measure on almost every Grassmannian fiber. Recall that, given any $\hat{x} \in \hat{\Sigma}$, we denote $x^n = P\big(\hat{f}^{-n}(\hat{x})\big)$ for $n \geq 0$.

Proposition 6.1. *If $\hat{A}$ is simple then, for every invariant u-state m and $\hat{\mu}$-almost every $\hat{x} \in \hat{\Sigma}$, the sequence $A^n(x^n)_* m_{x^n}$ converges to a Dirac measure $\delta_{\xi(\hat{x})}$ in the fiber $\{x\} \times \mathrm{Grass}(\ell, d)$ when $n \to \infty$.*

Proof: In view of Proposition 3.1, we only have to show that for $\hat{\mu}$-almost every $\hat{x} \in \hat{\Sigma}$ there exists some subsequence $(n_j)_j$ and a point $\xi(\hat{x}) \in \mathrm{Grass}(\ell, d)$ such that

$$A^{n_j}(x^{n_j})_* m_{x^{n_j}} \to \delta_{\xi(\hat{x})} \quad \text{when } j \to \infty . \tag{24}$$

Let $\hat{p} \in \hat{\Sigma}$ be a periodic point, with period $q \geq 1$, and $\hat{z} \in \hat{\Sigma}$ be a homoclinic point as in Definition 1.2. Denote $p = P(\hat{p})$ and $z = P(\hat{z})$. Let $[I] = [\iota_0, ..., \iota_{q-1}]$ be the cylinder of Σ that contains p. It is no restriction to assume that $z \in [I]$: this may always be achieved replacing $\hat{z}$ by some $\hat{f}^{-qi}(\hat{z})$ which, clearly, does not affect the conditions in Definition 1.2.

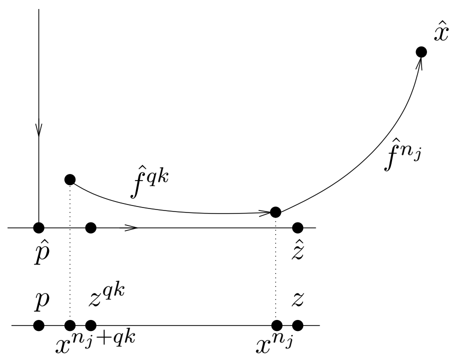

Figure 1 – Proof of Proposition 6.1: case $\xi(\hat{z})$ not in $\ker Q$.

For $\hat{\mu}$-almost every $\hat{x} \in \hat{\Sigma}$ there exists a sequence $(n_j)_j$ such that $\hat{f}^{-n_j}(\hat{x})$ converges to $\hat{z}$. That is because $\hat{\mu}$ is ergodic and positive on open sets. Let $k \geq 1$ be fixed. From Proposition 3.1 we conclude that

$$\begin{aligned} \lim_{j\to\infty} A^{n_j}(x^{n_j})_* m_{x^{n_j}} &= \lim_{j\to\infty} A^{n_j+qk}(x^{n_j+qk})_* m_{x^{n_j+qk}} \\ &= \lim_{j\to\infty} A^{n_j}(x^{n_j})_* A^{qk}(x^{n_j+qk})_* m_{x^{n_j+qk}} . \end{aligned}$$

Note that x^{n_j+qk} converges to z^{qk} when $j \to \infty$. See Figure 1. Then, by Proposition 4.4, the probability $m_{x^{n_j+qk}}$ converges to $m_{z^{qk}}$ when $j \to \infty$. So, since A is continuous,

$$A^{qk}(x^{n_j+qk})_* \, m_{x^{n_j+qk}} \;\to\; A^{qk}(z^{qk})_* \, m_{z^{qk}} \quad \text{when } \; j \to \infty \; .$$

Since the space of quasi-projective maps is compact, up to replacing $(n_j)_j$ by a subsequence we may suppose that $A^{n_j}(x^{n_j})$ converges to some quasi-projective map Q on $\operatorname{Grass}(\ell, d)$. By Lemma 2.3, the kernel of Q is contained in some hyperplane of $\operatorname{Grass}(\ell, d)$. Hence, by Proposition 5.1, the subspace $\ker Q$ has zero measure relative to $A^{qk}(z^{qk})_* \, m_{z^{qk}}$. So, we may apply Lemma 2.4 to conclude that

$$\lim_{j\to\infty} A^{n_j}(x^{n_j})_* \, m_{x^{n_j}} = Q_* \, A^{qk}(z^{qk})_* \, m_{z^{qk}} \; , \tag{25}$$

for any $k \ge 1$ (in particular, the latter expression does not depend on k).

Now, the pinching condition (p) in Definition 1.2 implies that $A^q(p)$ has ℓ eigenvalues that are strictly larger, in norm, than all the other ones. Denote by $\xi(\hat{p}) \in \operatorname{Grass}(\ell, d)$ the sum of the eigenspaces corresponding to those largest eigenvalues, and define $\xi(\hat{z}) = H_{\hat{p},\hat{z}} \cdot \xi(\hat{p})$.

Lemma 6.2. *The sequence $A^{qk}(z^{qk})_* \, m_{z^{qk}}$ converges to $\delta_{\xi(\hat{z})}$ when $k \to \infty$.*

Proof: Using the relations $A^{qk}(p)^{-1} = \hat{A}^{-qk}(\hat{p})$ and $A^{qk}(z^{qk})^{-1} = \hat{A}^{-qk}(\hat{z})$, we find that

$$A^{qk}(z^{qk})_* \, m_{z^{qk}} \; = \; \Big(\hat{A}^{-qk}(\hat{z})^{-1} \cdot \hat{A}^{-qk}(\hat{p})\Big)_* A^{qk}(p)_* \, m_{z^{qk}} \; .$$

By the Definition 1.1 of unstable holonomies, $\hat{A}^{-qk}(\hat{z})^{-1}\hat{A}^{-qk}(\hat{p})$ converges to $H_{\hat{p},\hat{z}}$ when $k \to \infty$. Observe also that $A^{qk}(p)_* \, m_{z^{qk}}$ converges to the Dirac measure at $\xi(\hat{p}) \in \operatorname{Grass}(\ell, d)$ when $k \to \infty$. That is because $m_{z^{qk}}$ converges to m_p, by Proposition 4.4, and m_p gives zero weight to the hyperplane section defined by the sum of the eigenspaces of $A^q(p)$ complementary to $\xi(\hat{p})$, by Proposition 5.1. It follows that $A^{qk}(z^{qk})_* \, m_{z^{qk}}$ converges to $(H_{\hat{p},\hat{z}})_* \, \delta_{\xi(\hat{p})} = \delta_{\xi(\hat{z})}$ when $k \to \infty$, as stated in the lemma. ∎

Suppose, for the time being, that $\xi(\hat{z})$ is in the domain $\operatorname{Grass}(\ell, d) \setminus \ker Q$ of the quasi-projective map Q. From Lemma 2.4 we get that

$$Q_* \, A^{qk}(p)_* \, m_{z^{qk}} \;\to\; Q_* \, \delta_{\xi(\hat{z})} = \delta_{\xi(\hat{x})}$$

when $k \to \infty$, where $\xi(\hat{x}) = Q\big(\xi(\hat{z})\big)$. Combined with the relation (25), this gives that $A^{n_j}(x^{n_j})_* \, m_{x^{n_j}}$ converges to the Dirac measure $\delta_{\xi(\hat{x})}$ when $j \to \infty$. This proves (24) and Proposition 6.1 in this case.

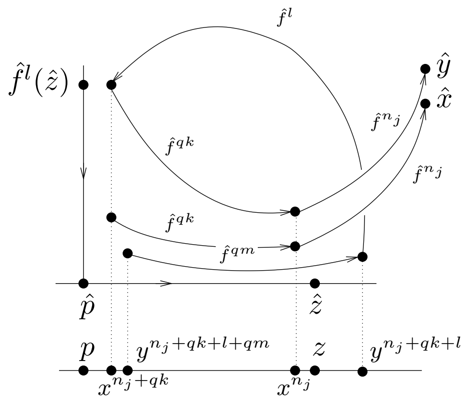

Figure 2 – Proof of Proposition 6.1: avoiding $\ker Q$.

Next, we show that one can always reduce the proof to the previous case. Let $l \geq 1$ be as in Definition 1.2. For each j much larger than k, let $m_j = n_j + qk + l$ and $\hat y = \hat y(j,k)$ be defined by

$$\hat f^{-n_j-qk}(\hat y) \in W^s_{\mathrm{loc}}\big(\hat f^{-n_j-qk}(\hat x)\big) \cap W^u_{\mathrm{loc}}\big(\hat f^l(\hat z)\big)\,. \tag{26}$$

See Figure 2. By construction, y^{m_j+qm} is sent to x^{n_j+qk} by the map f^{l+qm}. Hence, using Proposition 3.1,

$$\begin{aligned}\lim_{j\to\infty} A^{n_j}(x^{n_j})_* \, m_{x^{n_j}} &= \lim_{j\to\infty} A^{n_j+qk}(x^{n_j+qk})_* \, m_{x^{n_j+qk}} \\ &= \lim_{j\to\infty} A^{m_j+qm}(y^{m_j+qm})_* \, m_{y^{m_j+qm}}\ ,\end{aligned} \tag{27}$$

for any fixed k and m. We are going to prove that the limit is indeed a Dirac measure. For this, let $\hat w = \hat w(k)$ be defined by

$$\hat f^l(\hat w) \in W^s_{\mathrm{loc}}\big(\hat f^{-qk}(\hat z)\big) \cap W^u_{\mathrm{loc}}\big(\hat f^l(\hat z)\big)\,. \tag{28}$$

Notice that $\hat w$ is in $W^u_{\mathrm{loc}}(\hat z) = W^u_{\mathrm{loc}}(\hat p)$. Let k and m be fixed, for the time being. As $j \to \infty$, the sequence $\hat f^{-n_j-qk}(\hat x)$ converges to $\hat f^{-qk}(\hat z)$ and so, combining (26) and (28), the sequence $\hat f^{-n_j-qk}(\hat y)$ converges to $\hat f^l(\hat w)$. It follows that y^{m_j} converges to $w = P(\hat w)$, and so

$$A^{m_j}(y^{m_j}) = A^{n_j}(x^{n_j})\, A^{qk+l}(y^{m_j})$$

converges to $\tilde Q = Q \circ A^{qk+l}(w)$ in the space of quasi-projective maps, as $j \to \infty$. Define $\xi(\hat w) = H_{\hat p, \hat w} \cdot \xi(\hat p)$. The key observation is

Lemma 6.3. *Assuming k is large enough, $\xi(\hat{w})$ is not contained in $\ker \tilde{Q}$.*

Proof: From the definitions of $\tilde{Q}$ and $\hat{w}$ we get that

$$\ker \tilde{Q} \,=\, A^{kq+l}(w)^{-1} \cdot \ker Q \,=\, A^{l}(w)^{-1} \cdot A^{qk}(z^{qk})^{-1} \cdot \ker Q \,.$$

By the invariance property of unstable holonomies, we have

$$A^{qk}(z^{qk})^{-1} \,=\, \hat{A}^{-qk}(\hat{z}) \,=\, H_{\hat{p},\hat{f}^{-qk}(\hat{z})} \cdot \hat{A}^{-qk}(\hat{p}) \cdot H_{\hat{z},\hat{p}} \,.$$

So, the previous equality may be rewritten as

$$\ker \tilde{Q} \,=\, \hat{A}^{l}(\hat{w})^{-1} \cdot H_{\hat{p},\hat{f}^{-qk}(\hat{z})} \cdot \hat{A}^{-qk}(\hat{p}) \cdot H_{\hat{z},\hat{p}} \cdot \ker Q \,.$$

Notice that $\hat{f}^{-qk}(\hat{z})$ converges to $\hat{p}$ and so, by (28), the point $\hat{w}$ converges to $\hat{z}$, as $k \to \infty$. By the continuity of the cocycle and the holonomies, it follows that $H_{\hat{p},\hat{f}^{-qk}(\hat{z})}$ converges to the identity and $\hat{A}^{l}(\hat{w})$ converges to $\hat{A}^{l}(\hat{z})$, as k goes to ∞. By Lemma 2.3, the kernel of Q is contained in some hyperplane section of $\mathrm{Grass}(\ell, d)$. Then the same is true for $H_{\hat{z},\hat{p}} \cdot \ker Q$: it is contained in the set of all ℓ-dimensional subspaces that intersect the $(d-\ell)$-dimensional subspace $\pi_v(\upsilon)$ associated to some $(d-\ell)$-vector υ. Since all eigenvalues of $\hat{A}^{q}(\hat{p})$ have distinct absolute values, the backward iterates of $\pi_v(\upsilon)$ under $\hat{A}^{q}(\hat{p})$ converge to some $(d-\ell)$-dimensional sum $\pi_v(\eta)$ of eigenspaces of $\hat{A}^{q}(\hat{p})$. It follows that, as $k \to \infty$, the sequence $\hat{A}^{-qk}(\hat{p}) \cdot H_{\hat{z},\hat{p}} \cdot \ker Q$ converges to some subset V_0 of the hyperplane section V defined by η. Combining these two observations we get that, as $k \to \infty$,

$$\ker \tilde{Q} \,\to\, \hat{A}^{l}(\hat{z})^{-1}(V_0) \subset \hat{A}^{l}(\hat{z})^{-1}(V) \,. \tag{29}$$

It is easy to see that $\xi(\hat{z})$ does not belong to $\hat{A}^{l}(\hat{z})^{-1}(V)$: otherwise,

$$\hat{A}^{l}(\hat{z}) \cdot \xi(\hat{z}) \,=\, \hat{A}^{l}(\hat{z}) \cdot H_{\hat{p},\hat{z}} \cdot \xi(\hat{p}) \,=\, \psi_{p,z} \cdot \xi(\hat{p})$$

would intersect $\pi_v(\eta)$ and, since $\xi(\hat{p})$ and $\pi_v(\eta)$ correspond to sums of eigenspaces with complementary dimensions, that would contradict the twisting condition in Definition 1.2. Using (29) and the fact that $\xi(\hat{w})$ converges to $\xi(\hat{z})$ when $k \to \infty$, we deduce that $\xi(\hat{w})$ is not in $\ker \tilde{Q}$ if k is large enough, as claimed. ∎

We can now finish the proof of Proposition 6.1. The arguments are the same as in the previous case, with n_j and z replaced by $m_j = n_j + qk + l$ and w, respectively, and qm in the role of qk. Indeed, from (26) and (28) we get that $\hat{f}^{-m_j}(\hat{y})$ converges to $\hat{w}$ as $j \to \infty$. Consequently, $A^{qm}(y^{m_j+qm})$ converges to

$A^{qm}(w^{qm})$ and, using also Proposition 4.4, $m_{y^{m_j+qm}}$ converges to $m_{w^{qm}}$ as $j \to \infty$. So, in view of (27),

$$\begin{aligned}\lim_{j\to\infty} A^{n_j}(x^{n_j})_* \, m_{x^{n_j}} &= \lim_{j\to\infty} A^{m_j+qm}(y^{m_j+qm})_* \, m_{y^{m_j+qm}} \\ &= \tilde{Q}_* \, A^{qm}(p)_* \, m_{w^{qm}}\end{aligned}$$

for any $m \geq 1$, which is an analogue of (25). By Proposition 4.4, the measure $m_{w^{qm}}$ converges to m_p as $m \to \infty$. By Proposition 5.1, the measure m_p gives zero weight to the hyperplane section defined by the sum of the eigenspaces complementary to $\xi(\hat{p})$. Therefore, just as in Lemma 6.2, we conclude that $A^{qm}(\hat{w}^{qm})_* \, m_{\hat{w}^{qm}}$ converges to $\delta_{\xi(\hat{w})}$ when $m \to \infty$. Hence, fixing k as in Lemma 6.3 and using Lemma 2.4,

$$\lim_{m\to\infty} \tilde{Q}_* \, A^{qm}(p)_* \, m_{w^{qm}} = \delta_{\xi(\hat{x})} \ ,$$

where $\xi(\hat{x}) = \tilde{Q}_* \, \delta_{\xi(\hat{w})}$. This shows that $\lim_{j\to\infty} A^{n_j}(x^{n_j})_* \, m_{x^{n_j}} = \delta_{\xi(\hat{x})}$. Now the proof of Proposition 6.1 is complete. ∎

In the next proposition we summarize some consequences of the previous results that are needed for the next section:

Proposition 6.4. *Suppose that $\hat{A}$ is simple. Then there exists a measurable section $\xi\colon \hat{\Sigma} \to \operatorname{Grass}(\ell, d)$ such that, on a full $\hat{\mu}$-measure subset of $\hat{\Sigma}$,*

(**1**) *ξ is invariant under the cocycle and under unstable holonomies: $\hat{A}(\hat{x})\,\xi(\hat{x}) = \xi\big(\hat{f}(\hat{x})\big)$ and $\xi(\hat{y}) = H^u_{\hat{x},\hat{y}} \cdot \xi(\hat{x})$ for $\hat{x}$ and $\hat{y}$ in the same local unstable set;*

(**2**) *for any compact set $\Gamma \subset \hat{\Sigma}$, the eccentricity $\mathcal{E}\big(\ell, \hat{A}^n(\hat{f}^{-n}(\hat{x}))\big) \to \infty$, and the image under $\hat{A}^n\big(\hat{f}^{-n}(\hat{x})\big)$ of the ℓ-subspace most expanded by $\hat{A}^n\big(\hat{f}^{-n}(\hat{x})\big)$ converges to $\xi(\hat{x})$, restricted to the subsequence of iterates $\hat{f}^{-n}(\hat{x}) \in \Gamma$.*

Proof: From Corollary 3.4 and Proposition 6.1 we get that the conditional probabilities of the original measure $\hat{m}$ along the Grassmannian fibers coincide with the Dirac measures $\delta_{\xi(\hat{x})}$ almost everywhere. Since $\hat{m}$ is an invariant u-state, it follows that ξ is almost everywhere invariant under the cocycle and under the unstable holonomies, as stated in part 1 of the proposition.

Part 2 follows from Proposition 2.2, with $\mathcal{N} = \{m_{P(\hat{x})} \colon \hat{x} \in \Gamma\}$, $\nu_n = m_{x^n}$, $L_n = A^n(x^n) = \hat{A}^n\big(\hat{f}^{-n}(\hat{x})\big)$, and $\xi = \xi(\hat{x})$. By Proposition 4.4, the family is $\mathcal{N}$ is weak*compact if Γ is compact. It follows that the eccentricity $\mathcal{E}\big(\ell, \hat{A}^n(\hat{f}^{-n}(\hat{x}))\big)$ tends to infinity, and the image under $\hat{A}^n\big(\hat{f}^{-n}(\hat{x})\big)$ of the subspace most expanded by $\hat{A}^n\big(\hat{f}^{-n}(\hat{x})\big)$ converges to $\xi(\hat{x})$, as claimed. ∎

Remark 6.5. In Section 2.6 we replaced the original cocycle $\hat{A}$ by another one conjugate to it,

$$\hat{A}^u(\hat{x}) = H^s_{\hat{f}(\hat{x}),\,\phi^u(\hat{f}(\hat{x}))} \cdot \hat{A}(\hat{x}) \cdot H^s_{\phi^u(\hat{x}),\,\hat{x}} = H^s_{\hat{f}(\phi^u(\hat{x})),\,\phi^u(\hat{f}(\hat{x}))} \cdot \hat{A}\big(\phi^u(\hat{x})\big) ,$$

which is constant on local unstable sets and, consequently, whose stable holonomies are trivial. The statement of Proposition 6.4 is not affected by such substitution. Indeed, if ξ is an invariant section for $\hat{A}^u$ as in the proposition, then

$$H^s_{\phi^u(\hat{x}),\,\hat{x}}\, \xi(\hat{x})$$

is an invariant section for $\hat{A}$, and it is invariant also under the corresponding unstable holonomies. In addition,

$$\hat{A}^n\big(\hat{f}^{-n}(\hat{x})\big) = H^s_{\phi^u(\hat{x}),\,\hat{x}} \cdot (\hat{A}^u)^n\big(\hat{f}^{-n}(\hat{x})\big) \cdot H^s_{\hat{f}^{-n}(\hat{x}),\,\phi^u(\hat{f}^{-n}(\hat{x}))} .$$

Considering only iterates in a compact set, the corresponding conjugating isomorphisms H^s belong to a bounded family. Hence, the claims in part 2 of Proposition 6.4 hold for $\hat{A}$ if and only if they hold for $\hat{A}^u$. □

7 – Proof of the main theorem

We are going to show that $\hat{x} \mapsto \xi(\hat{x}) \in \mathrm{Grass}(\ell, d)$ corresponds to the sum of the Oseledets subspaces of the cocycle associated to the ℓ largest (strictly) Lyapunov exponents. In particular, $\xi(\hat{x})$ is uniquely defined almost everywhere. This will also prove that the invariant u-state is unique if the cocycle is simple.

The first step is to exhibit the sum $\eta(\hat{x})$ of the subspaces associated to the remaining Lyapunov exponents. This is done in Section 7.1, through applying the previous theory to the adjoint cocycle. Then, in Section 7.2 we use the second part of Proposition 6.4 to show that vectors along $\xi(\hat{x})$ are more expanded than those along $\eta(\hat{x})$.

7.1. Adjoint cocycle

Let $\cdot$ be a Hermitian form on $\mathbb{C}^d$, that is, a complex 2-form $(u, v) \mapsto u \cdot v$ which is linear on the first variable and satisfies $u \cdot v = \overline{v \cdot u}$ for every u and v. The *adjoint* of a linear operator $L\colon \mathbb{C}^d \to \mathbb{C}^d$ relative to the Hermitian form is the linear operator $L^*\colon \mathbb{C}^d \to \mathbb{C}^d$ defined by

$$L^*(u) \cdot v = u \cdot L(v) \qquad \text{for every } u \text{ and } v \text{ in } \mathbb{C}^d .$$

The matrix of L^* in any orthonormal basis for the Hermitian form is the conjugate transpose of the matrix of L in that basis: $L^*_{i,j} = \overline{L}_{j,i}$. The eigenvalues of L^* are the conjugates of the eigenvalues of L, and the operator norms of the two operators coincide: $\|L^*\| = \|L\|$.

Let $\hat{B}(\hat{x})\colon \mathbb{C}^d \to \mathbb{C}^d$ be defined by $\hat{B}(\hat{x}) = \hat{A}\big(\hat{f}^{-1}(\hat{x})\big)^*$ or, equivalently,

$$\hat{B}(\hat{x})\, u \cdot v = u \cdot \hat{A}\big(\hat{f}^{-1}(\hat{x})\big)\, v \qquad \text{for every } u \text{ and } v \text{ in } \mathbb{C}^d . \tag{30}$$

Consider the linear cocycle defined over $\hat{f}^{-1}$ by

$$\hat{F}_B\colon \hat{\Sigma}\times\mathbb{C}^d \to \hat{\Sigma}\times\mathbb{C}^d , \quad (\hat{x}, u) \mapsto \big(\hat{f}^{-1}(\hat{x}), \hat{B}(\hat{x})\, u\big) ,$$

as well as the induced Grassmannian cocycle. Notice that

$$\hat{B}^n(\hat{x}) \;=\; \hat{A}\big(\hat{f}^{-n}(\hat{x})\big)^* \cdots \hat{A}\big(\hat{f}^{-2}(\hat{x})\big)^* \hat{A}\big(\hat{f}^{-1}(\hat{x})\big)^* \;=\; \hat{A}^n\big(\hat{f}^{-n}(\hat{x})\big)^* .$$

The choice of the Hermitian form is not important: different choices yield cocycles that are conjugate. For convenience, we fix once and for all such that eigenvectors of $\hat{A}^q(\hat{p})$ form an orthonormal basis.

The integrability condition in the Oseledets theorem holds for $\hat{B}$ if and only if it holds for $\hat{A}$, because $\|\hat{B}(\hat{x})\| = \|\hat{A}\big(\hat{f}^{-1}(\hat{x})\big)\|$ and the measure $\hat{\mu}$ is invariant under $\hat{f}$. It is easy to check that the previous results apply to the cocycle defined by $\hat{B}$. To begin with, our hypotheses on the dynamics (Section 1.1) and on the invariant measure (Section 1.2) are, evidently, symmetric under time reversion. The hypotheses on the cocycle (Section 1.3) are also clearly satisfied: a simple calculation shows that $\hat{B}$ admits stable and unstable holonomies given by

$$H^{u,\hat{B}}_{\hat{x},\hat{y}} = \left(H^{s,\hat{A}}_{\hat{y},\hat{x}}\right)^* \quad \text{and} \quad H^{s,\hat{B}}_{\hat{x},\hat{y}} = \left(H^{u,\hat{A}}_{\hat{y},\hat{x}}\right)^*. \tag{31}$$

Lemma 7.1. *$\hat{B}$ is simple for $\hat{f}^{-1}$ if and only if $\hat{A}$ is simple for $\hat{f}$.*

Proof: Let $\hat{p}$ be a periodic point of $\hat{f}$. For any orthonormal basis of $\mathbb{C}^d$, the matrix of $\hat{B}^q(\hat{p}) = \hat{A}^q(\hat{p})^*$ is the conjugate transpose of the matrix of $\hat{A}^q(\hat{p})$, and the eigenvalues of the former are the complex conjugates of the eigenvalues of the latter. Hence, the pinching condition in Definition 1.2 holds for any of them if and only if it holds for the other. Next, notice that $\hat{z}$ is a homoclinic point for $\hat{f}$ if and only if $\hat{w} = \hat{f}^l(\hat{z})$ is a homoclinic point for the inverse: $\hat{z} \in W^u_{\mathrm{loc}}(\hat{p}, \hat{f})$ and $\hat{f}^l(\hat{z}) \in W^s_{\mathrm{loc}}(\hat{p}, \hat{f})$ if and only if $\hat{f}^{-l}(\hat{w}) \in W^s_{\mathrm{loc}}(\hat{p}, \hat{f}^{-1})$ and $\hat{w} \in W^u_{\mathrm{loc}}(\hat{p}, \hat{f}^{-1})$.

We have chosen the Hermitian form in such a way that eigenvectors of $\hat{A}^q(\hat{p})$ form an orthonormal basis. Then the matrix of $\hat{B}^l(\hat{w}) = \hat{A}^l(\hat{z})^*$ in this basis is the conjugate transpose of the matrix of $\hat{A}^l(\hat{z})$, and so the algebraic minors of the former are the complex conjugates of the algebraic minors of the latter. Thus, the twisting condition in Definition 1.2 holds for $\hat{B}$ if and only if it holds for $\hat{A}$. ∎

This ensures that the previous results do apply to $\hat{B}$. From Proposition 6.4 we obtain that

- **(i)** there exists a section $\xi^* : \hat{\Sigma} \to \mathrm{Grass}(\ell, d)$ which is invariant under the cocycle $\hat{F}_B$ and under the unstable holonomies of $\hat{B}$;
- **(ii)** given any compact $\Gamma \subset \hat{\Sigma}$, restricted to the subsequence of iterates $\hat{f}^n(\hat{x})$ in Γ, the eccentricity $\mathcal{E}\big(\ell, \hat{B}^n(\hat{f}^n(\hat{x}))\big) = \mathcal{E}\big(\ell, \hat{A}^n(\hat{x})\big)$ goes to infinity and the image $\hat{B}^n\big(\hat{f}^n(\hat{x})\big)\,\zeta_n^a\big(\hat{f}^n(\hat{x})\big)$ of the ℓ-subspace $\zeta_n^a\big(\hat{f}^n(\hat{x})\big)$ most expanded by $\hat{B}^n\big(\hat{f}^n(\hat{x})\big)$ tends to $\xi^*(\hat{x})$ as $n \to \infty$.

Let us show that $\xi(\hat{x})$ is outside the hyperplane section orthogonal to $\xi^*(\hat{x})$:

Lemma 7.2. *For $\hat{\mu}$-almost every $\hat{x}$, the subspace $\xi(\hat{x})$ is transverse to the orthogonal complement of $\xi^*(\hat{x})$.*

Proof: Recall, from Section 2.6 and Remark 6.5, that we may take the stable holonomies of $\hat{A}$ to be trivial. Then, by (31), the unstable holonomies of $\hat{B}$ are also trivial. So, the fact that ξ^* is invariant under unstable holonomies just means that it is constant on local unstable sets of $\hat{f}^{-1}$, that is, on local stable sets of $\hat{f}$. Then the same is true about the orthogonal complement of $\xi^*(\hat{x})$. In other words, the hyperplane section of $\mathrm{Grass}(\ell,d)$ orthogonal to $\xi^*(\hat{x})$ depends only on $x = P(\hat{x})$. Denote it as H_x. Using Proposition 5.1 and then Proposition 6.4, we obtain

$$0 \;=\; m_x(H_x) \;=\! \int \delta_{\xi(\hat{x})}(H_x)\, d\hat{\mu}_x(\hat{x}) \;=\; \hat{\mu}_x\Big(\Big\{\hat{x} \in W^s_{\mathrm{loc}}(x) :\; \xi(\hat{x}) \in H_x\Big\}\Big)\,,$$

for μ-almost every x. Consequently, $\hat{\mu}\big(\big\{\hat{x} \in \hat{\Sigma} :\; \xi(\hat{x}) \in H_x\big\}\big) = 0$. This means that, for almost every point, the subspace $\xi(\hat{x})$ intersects the orthogonal complement of $\xi^*(\hat{x})$ at the origin only, which is precisely the claim in the lemma. ∎

Let $\eta(\hat{x}) \in \mathrm{Grass}(d-\ell, d)$ denote the orthogonal complement of $\xi^*(\hat{x})$. Recall that ξ and ξ^* are invariant under the corresponding cocycles:

$$\hat{A}(\hat{x})\,\xi(\hat{x}) = \xi\big(\hat{f}(\hat{x})\big) \quad \text{and} \quad \hat{B}(\hat{x})\,\xi^*(\hat{x}) = \xi^*\big(\hat{f}^{-1}(\hat{x})\big)$$

$\hat{\mu}$-almost everywhere. The latter implies that $\eta(\hat{x})$ is also invariant under $\hat{A}$.

According to Lemma 7.2, we have $\mathbb{C}^d = \xi(\hat{x}) \oplus \eta(\hat{x})$ at almost every point. We want to prove that the Lyapunov exponents of $\hat{A}$ along ξ are strictly bigger than those along η. To this end, let

$$\xi(\hat{x}) = \xi^1(\hat{x}) \oplus \cdots \oplus \xi^u(\hat{x}) \quad \text{and} \quad \eta(\hat{x}) = \eta^s(\hat{x}) \oplus \cdots \oplus \eta^1(\hat{x})$$

be the Oseledets decompositions of $\hat{A}$ restricted to the two invariant subbundles. Take the factors to be numbered in such a way that ξ^u corresponds to the smallest Lyapunov exponent among all ξ^i, and η^s corresponds to the largest Lyapunov exponent among all η^j. Denote $d_u = \dim \xi^u$ and $d_s = \dim \eta^s$, and let λ_u and λ_s be the Lyapunov exponents associated to these two subbundles, respectively.

7.2. Direction of maximum expansion

Given a linear map $L\colon \mathbb{C}^d \to \mathbb{C}^d$ and a subspace V of $\mathbb{C}^d$, we denote by $\det(L, V)$ the determinant of L along V, defined as the quotient of the volumes of the parallelograms determined by $\{Lv_1, ..., Lv_s\}$ and $\{v_1, ..., v_s\}$, respectively, for any basis $v_1, ..., v_s$ of V. Then we define, for each $n \geq 1$,

$$\Delta^n(\hat{x}) = \frac{\det\big(\hat{A}^n(\hat{x}), \xi^u(\hat{x})\big)^{1/d_u}}{\det\big(\hat{A}^n(\hat{x}), W(\hat{x})\big)^{1/(d_u+d_s)}} \quad \text{where} \ \ W(\hat{x}) = \xi^u(\hat{x}) \oplus \eta^s(\hat{x}) \ . \tag{32}$$

According to the theorem of Oseledets [13],

$$\frac{1}{n} \log\det\big(\hat{A}^n(\hat{x}), \xi^u(\hat{x})\big) \to d_u \lambda_u \quad \text{and} \quad \frac{1}{n} \log\det\big(\hat{A}^n(\hat{x}), W(\hat{x})\big) \to d_u \lambda_u + d_s \lambda_s \ .$$

Consequently,

$$\lim_{n\to\infty} \frac{1}{n} \log \Delta^n(\hat{x}) \ = \ \frac{d_s}{d_u + d_s} \, (\lambda_u - \lambda_s) \ . \tag{33}$$

So, to prove that λ_u is strictly larger than λ_s we must show that $\log \Delta^n$ goes linearly to infinity at almost every point. The main step is

Proposition 7.3. *For any compact set $\Gamma \subset \hat{\Sigma}$ and for $\hat{\mu}$-almost every $\hat{x} \in \hat{\Sigma}$,*

$$\lim_{n\to\infty} \Delta^n(\hat{x}) \ = \ +\infty$$

restricted to the subsequence of values of n for which $\hat{f}^n(\hat{x}) \in \Gamma$.

Proof: Let $\xi_n^a(\hat{x}) = \hat{B}^n\big(\hat{f}^n(\hat{x})\big)\,\zeta_n^a(\hat{x})$ be the image of the ℓ-dimensional subspace most expanded by $\hat{B}^n\big(\hat{f}^n(\hat{x})\big) = \hat{A}^n(\hat{x})^*$. Equivalently, $\xi_n^a(\hat{x})$ is the ℓ-dimensional subspace most expanded by $\hat{A}^n(\hat{x})$. Throughout, we consider only the values of n for which $\hat{f}^n(\hat{x}) \in \Gamma$. Then we may use property (ii) in Section 7.1: the eccentricity

$$E_n = \mathcal{E}\big(\ell, \hat{B}^n(\hat{f}^n(\hat{x}))\big) = \mathcal{E}\big(\ell, \hat{A}^n(\hat{x})\big)$$

tends to infinity, and $\xi_n^a(\hat{x})$ tends to $\xi^*(\hat{x})$, as $n \to \infty$. In view of Lemma 7.2, the latter fact implies that the subspace $\xi(\hat{x})$ is transverse to the orthogonal complement of $\xi_n^a(\hat{x})$, with angle uniformly bounded from zero for all large n. Let us consider the orthogonal splitting

$$\mathbb{C}^d = \xi_n^a(\hat{x}) \oplus \xi_n^a(\hat{x})^\perp .$$

Let $\xi_n^u(\hat{x}) \subset \xi_n^a(\hat{x})$ be the image of the subspace $\xi^u(\hat{x}) \subset \xi(\hat{x})$ under the orthogonal projection. We claim that

$$\det\big(\hat{A}^n(\hat{x}) \mid \xi_n^u(\hat{x})\big) \le C_1 \det\big(\hat{A}^n(\hat{x}) \mid \xi^u(\hat{x})\big) , \tag{34}$$

for some constant C_1 independent of n. To see this, observe that any basis α of $\xi_n^u(\hat{x})$ may be lifted to a basis β of $\xi^u(\hat{x})$. This operation increases the volume of the corresponding parallelogram by, at most, some factor C_1 that depends only on a bound for the angle between $\xi(\hat{x})$ and the orthogonal complement of $\xi_n^a(\hat{x})$. Note also that, the $\hat{A}^n(\hat{x})$-images of $\xi_n^a(\hat{x})$ and $\xi_n^a(\hat{x})^\perp$ are orthogonal to each other, because $\xi_n^a(\hat{x})$ is the ℓ-subspace most expanded by $\hat{A}^n(\hat{x})$. Hence, the $\hat{A}^n(\hat{x})$-image of α may be obtained from the $\hat{A}^n(\hat{x})$-image of β by orthogonal projection, an operation that can only decrease the volume of the parallelogram. Combining these observations, we get (34). Next, let $\eta_n^s(\hat{x})$ be the subspace of $\xi_n^a(\hat{x})^\perp$ characterized by

$$W(\hat{x}) = \xi^u(\hat{x}) \oplus \eta^s(\hat{x}) = \xi^u(\hat{x}) \oplus \eta_n^s(\hat{x}) .$$

Equivalently, $\eta_n^s(\hat{x})$ is the projection of $\eta^s(\hat{x})$ to the orthogonal complement of $\xi_n^a(\hat{x})$ along the direction of $\xi(\hat{x})$. Since the angle between $\xi^u(\hat{x})$ and $\eta_n^s(\hat{x})$ is bounded from zero,

$$\det\big(\hat{A}^n(\hat{x}), W(\hat{x})\big) \le C_2 \det\big(\hat{A}^n(\hat{x}), \xi^u(\hat{x})\big) \det\big(\hat{A}^n(\hat{x}), \eta_n^s(\hat{x})\big) \tag{35}$$

where the constant C_2 is independent of n. Furthermore,

$$\det\big(\hat{A}^n(\hat{x}), \eta_n^s(\hat{x})\big) \le \big\|\hat{A}^n(\hat{x}) \mid \eta_n^s(\hat{x})\big\|^{d_s} \le \big\|\hat{A}^n(\hat{x}) \mid \xi_n^a(\hat{x})^\perp\big\|^{d_s} ,$$

$$\det\big(\hat{A}^n(\hat{x}), \xi_n^u(\hat{x})\big) \ge m\big(\hat{A}^n(\hat{x}) \mid \xi_n^u(\hat{x})\big)^{d_u} \ge m\big(\hat{A}^n(\hat{x}) \mid \xi_n^a(\hat{x})\big)^{d_u}$$

because $\eta_n^s(\hat{x}) \subset \xi_n^a(\hat{x})^\perp$ and $\xi_n^u(\hat{x}) \subset \xi_n^a(\hat{x})$. Consequently,

$$E_n = \frac{m\big(\hat{A}^n(\hat{x}) \mid \xi_n^a(\hat{x})\big)}{\big\|\hat{A}^n(\hat{x}) \mid \xi_n^a(\hat{x})^\perp\big\|} \le \frac{\det\big(\hat{A}^n(\hat{x}),\, \xi_n^u(\hat{x})\big)^{1/d_u}}{\det\big(\hat{A}^n(\hat{x}),\, \eta_n^s(\hat{x})\big)^{1/d_s}} . \tag{36}$$

From (34)–(36) we obtain

$$\det\big(\hat{A}^n(\hat{x}), W(\hat{x})\big) \le C\, E_n^{-d_s} \det\big(\hat{A}^n(\hat{x}),\, \xi^u(\hat{x})\big)^{1+d_s/d_u}$$

with $C = C_1^{s/u} C_2$. Consequently,

$$\Delta^n(\hat{x}) = \frac{\det\big(\hat{A}^n(\hat{x}),\, \xi^u(\hat{x})\big)^{1/d_u}}{\det\big(\hat{A}^n(\hat{x}), W(\hat{x})\big)^{1/(d_u+d_s)}} \ge \big(C^{-1} E_n^s\big)^{1/d_u+d_s}$$

and this goes to infinity when $n \to \infty$. The proof of the proposition is complete. ∎

Now we are ready for the proof of Theorem A. Fix any compact set $\Gamma \subset \hat{\Sigma}$ such that $\hat{\mu}(\Gamma) > 0$. By Poincaré recurrence, the first return map

$$g \colon \Gamma \to \Gamma , \qquad g(\hat{x}) = \hat{f}^{r(\hat{x})}(\hat{x})$$

is well defined on a full $\hat{\mu}$-measure subset of $\hat{\Sigma}$. The normalized restriction $\hat{\mu}/\hat{\mu}(\Gamma)$ of the measure $\hat{\mu}$ to Γ is invariant and ergodic for g. Moreover, $\hat{F}_A$ induces a linear cocycle

$$G \colon \Gamma \times \mathbb{C}^d \to \Gamma \times \mathbb{C}^d , \qquad G(\hat{x}, v) = \big(g(\hat{x}), \mathcal{G}(\hat{x})\, v\big)$$

where $\mathcal{G}(\hat{x}) = \hat{A}^{r(\hat{x})}(\hat{x})$. Clearly, this cocycle preserves the subbundles $\xi(\hat{x})$ and $\eta(\hat{x})$, as well as their Oseledets decompositions

$$\xi(\hat{x}) = \xi^1(\hat{x}) \oplus \cdots \oplus \xi^u(\hat{x}) \quad \text{and} \quad \eta(\hat{x}) = \eta^s(\hat{x}) \oplus \cdots \oplus \eta^1(\hat{x}) .$$

It is also clear (see Section A.1) that the Lyapunov exponents of G with respect to $\hat{\mu}/\hat{\mu}(\Gamma)$ are the products of the exponents of $\hat{F}_A$ by the average return time $1/\hat{\mu}(\Gamma)$.

Thus, to show that $\lambda_u > \lambda_s$ it suffices to prove the corresponding statement for G. Define

$$\mathcal{D}^k(\hat{x}) = \frac{\det\big(\mathcal{G}^k(\hat{x}),\, \xi^u(\hat{x})\big)^{1/d_u}}{\det\big(\mathcal{G}^k(\hat{x}), W(\hat{x})\big)^{1/(d_u+d_s)}} \qquad \text{where } W(\hat{x}) = \xi^u(\hat{x}) \oplus \eta^s(\hat{x}) .$$

Notice that, since ξ^u and η^s are both G-invariant,

$$\mathcal{D}^k(\hat{x}) = \mathcal{D}(\hat{x})\, \mathcal{D}\big(g(\hat{x})\big) \cdots \mathcal{D}\big(g^{k-1}(\hat{x})\big)$$

for all $k \geq 1$, where we write $\mathcal{D} = \mathcal{D}^1$. Notice also that $\mathcal{D}^k(\hat{x})$ is a subsequence of the sequence $\Delta^n(\hat{x})$ defined in (32). Since g is a return map to Γ, this subsequence corresponds to values of n for which $\hat{f}^n(\hat{x}) \in \Gamma$. So, Proposition 7.3 may be applied to conclude that

$$\lim_{k\to\infty} \sum_{j=0}^{k-1} \log \mathcal{D}\big(g^j(\hat{x})\big) = \lim_{k\to\infty} \mathcal{D}^k(\hat{x}) = \infty \qquad \text{for } \hat{\mu}\text{-almost every } \hat{x} \in \Gamma\,. \tag{37}$$

We use the following well-known fact (see [10, Corollary 6.10]) to conclude that the growth is even linear:

Lemma 7.4. *Let $T\colon X \to X$ be a measurable transformation preserving a probability measure ν in X, and $\varphi\colon X \to \mathbb{R}$ be a ν-integrable function such that $\lim_{n\to\infty} \sum_{j=0}^{n-1} \big(\varphi \circ T^j\big) = +\infty$ at ν-almost every point. Then $\int \varphi \, d\nu > 0$.* ∎

Applying the lemma to $T = g$ and $\varphi = \log \mathcal{D}$, we find that

$$\lim_{k\to\infty} \frac{1}{k} \log \mathcal{D}^k(\hat{x}) = \lim_{k\to\infty} \frac{1}{k} \sum_{j=0}^{k-1} \log \mathcal{D}\big(g^j(\hat{x})\big) = \int \log \mathcal{D} \, \frac{d\hat{\mu}}{\hat{\mu}(\Gamma)} > 0 \tag{38}$$

at $\hat{\mu}$-almost every point. On the other hand, from (33) and the relation between the Lyapunov spectra of $\hat{F}_A$ and G,

$$\lim_{k\to\infty} \frac{1}{k} \log \mathcal{D}^k(\hat{x}) = \frac{d_s}{d_u + d_s} (\lambda_u - \lambda_s) \frac{1}{\hat{\mu}(\Gamma)}\,. \tag{39}$$

These two relations imply that $\lambda_u > \lambda_s$. In this way, we have shown that there is a definite gap between the first ℓ Lyapunov exponents and the remaining $d-\ell$ ones. Since this applies for every $1 \leq \ell < d$, we conclude that the Lyapunov spectrum is simple. The proof of Theorem A is complete.

Remark 7.5. *A posteriori*, we get from (33), (38), (39) that $\Delta^n(x)$ goes linearly to infinity when $n \to \infty$, that is, we do not need to restrict to $\hat{f}^n(\hat{x}) \in \Gamma$. □

APPENDIX

A – Extensions and applications

In this appendix we check that our methods apply to the Zorich cocycles introduced in [19, 20]. We start with a few simple comments on our hypotheses.

A.1. Inducing

Here we explain how cocycles over more general maps can often be reduced to the case of the full countable shift. We begin by treating the case of subshifts of countable type. In particular, we recover the main results of [4], in a stronger form.

Let $\mathcal{I}$ be a finite or countable set and $T = \big(t(i,j)\big)_{i,j\in\mathcal{I}}$ be a transition matrix, meaning that every entry $t(i,j)$ is either 0 or 1. Define

$$\hat{\Sigma}_T \;=\; \Big\{(\iota_n)_{n\in\mathbb{Z}} \in \mathcal{I}^{\mathbb{Z}}\colon\; t(\iota_n,\iota_{n+1}) = 1 \;\text{ for all } n\in\mathbb{Z}\Big\}$$

and let $\hat{f}_T\colon \hat{\Sigma}_T \to \hat{\Sigma}_T$ be the restriction to $\hat{\Sigma}_T$ of the shift map on $\mathcal{I}^{\mathbb{Z}}$. By definition, the cylinders $[\,\cdot\,]$ of $\hat{\Sigma}_T$ are its intersections with the cylinders of the full space $\mathcal{I}^{\mathbb{Z}}$. One-sided shift spaces $\Sigma_T^u \subset \mathcal{I}^{\{n\geq 0\}}$ and $\Sigma_T^s \subset \mathcal{I}^{\{n<0\}}$, and cylinders $[\,\cdot\,]^u \subset \Sigma_T^u$ and $[\,\cdot\,]^s \subset \Sigma_T^s$ are defined analogously.

Let $\hat{\nu}_T$ be a probability measure on $\hat{\Sigma}_T$ invariant under $\hat{f}_T$ and whose support contains some cylinder $[I] = [\iota_0; \iota_1, \ldots, \iota_{k-1}]$ of $\hat{\Sigma}_T$. By Poincaré recurrence, the subset X of points that return to $[I]$ infinitely many times in forward and backward time has full measure. Let $r(\hat{x}) \geq 1$ be the first return time and

$$\hat{g}(\hat{x}) = \hat{f}^{r(\hat{x})}(\hat{x}), \qquad \text{for } \; \hat{x}\in X\,.$$

This first return map $\hat{g}\colon X \to X$ may be seen as a shift on $\hat{\Sigma} = \mathbb{N}^{\mathbb{Z}}$. Indeed, let $\big\{J(\ell)\colon\, \ell\in\mathbb{N}\big\}$ be an enumeration of the family of cylinders of the form

$$\big[\iota_0; \iota_1, \ldots, \iota_{r-1}, \iota_r, \ldots, \iota_{r+k-1}\big], \qquad \text{with } \;\iota_{r+i} = \iota_i \;\text{ for }\; i = 0, 1, \ldots, k-1 \tag{40}$$

and $r \geq 1$ minimum with this property. Then

$$\mathbb{N}^{\mathbb{Z}} \to X, \qquad (\ell_n)_{n\in\mathbb{Z}} \;\mapsto\; \bigcap_{n\in\mathbb{Z}} \hat{g}^{-n}\big(J(\ell_n)\big)$$

conjugates $\hat{g}$ to the shift map. Let $\hat{\nu}$ be the normalized restriction of $\hat{\nu}_T$ to X. Then $\hat{\nu}$ is a $\hat{g}$-invariant probability measure and, assuming $\hat{\nu}_T$ is ergodic for $\hat{f}_T$, it is $\hat{g}$-ergodic. The measure $\hat{\nu}$ is positive on cylinders, since $[I]$ is contained in the support of $\hat{\nu}_T$. It has product structure if $\hat{\nu}_T$ has. The latter makes sense because every cylinder $[\iota]$ of $\hat{\Sigma}_T$ is homeomorphic to a product of cylinders of Σ_T^u and Σ_T^s.

To each cocycle defined over $\hat{f}$ by some $\hat{A}_T\colon \hat{\Sigma}_T \to \mathrm{GL}(d,\mathbb{C})$ we may associate a cocycle defined over $\hat{g}$ by

$$\hat{B}(\hat{x}) = \hat{A}_T^{r(\hat{x})}(\hat{x}) \ .$$

Notice that $\hat{B}$ is continuous if $\hat{A}_T$ is, since the return time $r(\hat{x})$ is constant on each cylinder as in (40). Also, $\hat{B}$ admits stable and unstable holonomies if $\hat{A}_T$ does: the holonomy maps for the two cocycles coincide on the domain of $\hat{B}$. Furthermore, the Lyapunov exponents of $\hat{B}$ are obtained by multiplying those of $\hat{A}_T$ by the average return time. Indeed, given any non-zero vector v,

$$\lim_{n\to\infty}\frac{1}{n}\log\big\|\hat{B}^n(\hat{x})\,v\big\| = \lim_{n\to\infty}\frac{1}{n}\log\big\|\hat{A}^{S_n r(\hat{x})}(\hat{x})\,v\big\|\,, \qquad S_n r(\hat{x}) = \sum_{j=0}^{n-1} r\big(g^j(\hat{x})\big)\,,$$

and, for $\hat{\nu}$-almost every $\hat{x}$, this is equal to

$$\lim_{n\to\infty}\frac{1}{n}\,S_n\,r(\hat{x})\,\lim_{m\to\infty}\frac{1}{m}\log\big\|\hat{A}^m(\hat{x})\,v\big\| \;=\; \frac{1}{\hat{\nu}_T\big([I]\big)}\,\lim_{m\to\infty}\frac{1}{m}\log\big\|\hat{A}^m(\hat{x})\,v\big\|\,,$$

since $n^{-1}S_n\,r(\hat{x})$ converges $\int r\,d\hat{\nu} = 1/\hat{\nu}([I])$. In particular, the Lyapunov spectrum of either cocycle is simple if and only if the other one is.

Finally, the cocycle $\hat{B}$ is simple for $\hat{g}$ if $\hat{A}_T$ is simple for $\hat{f}_T$. More precisely, suppose $\hat{f}_T$ admits points $\hat{p}$ and $\hat{z}$ satisfying the conditions in Definition 1.2 for the cocycle defined by $\hat{A}_T$ and such that $\hat{p}$ is in the interior of the support of $\hat{\nu}_T$. Let $q \geq 1$ be the minimum period of $\hat{p}$ and $[I] = [\iota_0; \iota_1, \ldots, \iota_{qs-1}]$ be a cylinder that contains $\hat{p}$, with $s \geq 1$. Taking s sufficiently large, we may assume that $[I]$ is contained in the support of $\hat{\nu}_T$. Replacing $\hat{z}$ and $\hat{f}^l(\hat{z})$ by appropriate backward and forward iterates, respectively, we may also assume that they are both in $[I]$. Then $\hat{p}$ is also a periodic point for $\hat{g}$ and $\hat{z}$ is an associated homoclinic point. Since the holonomies of the cocycles defined by $\hat{A}_T$ and $\hat{B}$ coincide, it follows that the pinching and twisting conditions in Definition 1.2 hold also for the cocycle defined by $\hat{B}$.

In this way we have shown that *our simplicity criterion extends directly to cocycles over any subshift of countable type* $f_T\colon \hat{\Sigma}_T \to \hat{\Sigma}_T$. There is also a non-invertible version of this construction, where one starts with a one-sided subshift of countable type $f_T\colon \Sigma_T \to \Sigma_T$ and an invariant probability ν_T on Σ_T, and one constructs a first return map $g(x) = f^{r(x)}(x)$ to some cylinder $[I]$ contained in the support of ν_T. Then g is conjugate to the shift map on $\mathbb{N}^{\{n\geq 0\}}$ and the normalized restriction ν of the measure ν_T to its domain is a g-invariant probability. Moreover, the measure ν is ergodic for g if ν_T is ergodic for f_T. The natural extension $\hat{g}$ of the return map may be realized as the shift map on $\mathbb{N}^{\mathbb{Z}}$. The lift $\hat{\nu}$ of the probability ν is a $\hat{g}$-invariant measure, and it is $\hat{g}$-ergodic if ν is ergodic for g. In Section A.2 we discuss conditions on ν under which the lift has product structure. Given any $A_T\colon \Sigma_T \to \mathrm{GL}(d,\mathbb{C})$, the map $B(x) = A_T^{r(x)}(x)$ defines a cocycle over g. Moreover, B lifts canonically to a cocycle $\hat{B}$ over $\hat{g}$, constant on local stable sets, and having the same Lyapunov exponents. Thus, the Lyapunov spectrum of A_T is simple if and only if the Lyapunov spectrum of $\hat{B}$ (or B) is.

More generally, let $f\colon M \to M$ be a transformation preserving a probability ν_f and assume there exists a return map g to some domain $D \subset \operatorname{supp} \nu_f$ which is a *Markov map*. By this we mean that there exists a finite or countable partition $\{J(\ell)\colon \ell \in \mathbb{N}\}$ of D such that (i) g maps each $J(\ell)$ bijectively to the whole domain D and (ii) for any sequence $(\ell_n)_n$ in $\mathbb{N}^{\{n\geq 0\}}$ the intersection of $g^{-n}(J(\ell_n))$ over all $n \geq 0$ consists of exactly one point. Then g may be seen as the shift map on $\mathbb{N}^{\{n\geq 0\}}$. The normalized restriction ν of ν_f to the domain of g is a g-invariant probability, and it is g-ergodic if ν_f is ergodic for f. As before, to any cocycle over f we may associate a cocycle over g, or its natural extension, such that the Lyapunov spectrum of either is simple if and only if the other one is. This type of construction will be used in Section A.4.

A.2. Bounded oscillation

Let $f\colon \Sigma \to \Sigma$ be the shift map on $\Sigma = \mathbb{N}^{\{n\geq 0\}}$. The *lift* of an f-invariant probability measure μ is the unique $\hat{f}$-invariant measure $\hat{\mu}$ on $\hat{\Sigma} = \mathbb{N}^{\mathbb{Z}}$ such that $P_* \hat{\mu} = \mu$. The *k-oscillation* of a function $\psi\colon \Sigma \to \mathbb{R}$ is defined by

$$\operatorname{osc}_k(\psi) = \sup_I \sup \Big\{ \psi(x) - \psi(y) \colon \ x, y \in [I] \Big\}$$

where the first supremum is over all sequences $I = (\iota_0, \ldots, \iota_k)$ in $\mathbb{N}^k$. We say ψ has *bounded oscillation* if $\sum_{k=1}^{\infty} \operatorname{osc}_k(\psi) < \infty$. This implies $\operatorname{osc}_k(\psi) \to 0$ and so ψ is continuous, in a uniform sense. We are going to prove

Proposition A.1. *If the Jacobian of ν for f has bounded oscillation then the lift $\hat{\mu}$ has product structure.*

Lemma A.2. *Let x and y be in $\Sigma = \Sigma^u$. For each point $\hat{x} \in W^s_{\mathrm{loc}}(x)$, define $\hat{y} \in W^u_{\mathrm{loc}}(\hat{x}) \cap W^s_{\mathrm{loc}}(y)$. Then the limit*

$$J_{x,y}(\hat{x}) = \lim_{n\to\infty} \frac{Jf^n(x^n)}{Jf^n(y^n)}, \qquad \text{where } x^n = P\big(\hat{f}^{-n}(\hat{x})\big) \text{ and } y^n = P\big(\hat{f}^{-n}(\hat{y})\big),$$

exists, uniformly on x, y and $\hat{x}$. Moreover, the function $(x, y, \hat{x}) \mapsto J_{x,y}(\hat{x})$ is continuous and uniformly bounded from zero and infinity.

Proof: The arguments are quite standard. Begin by noting that

$$\log \frac{Jf^n(x^n)}{Jf^n(y^n)} = \sum_{j=1}^{n} \log Jf(x^j) - \log Jf(y^j) \ . \tag{41}$$

Notice that x^j and y^j are in the same cylinder $[\iota_{-j}, \ldots, \iota_{-1}]^u$, for each $j \geq 1$. Hence, the jth term in the sum is bounded in norm by the j-oscillation of $\log Jf$. It follows that the series in (41) converges absolutely and uniformly, and the sum is bounded by $\sum_j \mathrm{osc}_j(\log Jf)$. This implies all the claims in the lemma. ∎

Lemma A.3. *Let $\{\hat{\mu}_x \colon x \in \Sigma\}$ be any disintegration of the lift $\hat{\mu}$ of μ. For a full μ-measure subset of points $x \in \Sigma$, we have*

$$\hat{\mu}_x(\xi_n) = \frac{1}{Jf^n(x^n)}$$

for every cylinder $\xi_n = [\iota_{-n}, \ldots, \iota_{-1}]^s$, $n \geq 1$, and every point $\hat{x} \in \xi_n \times \{x\}$.

Proof: Let F be any measurable subset of Σ. Then $\hat{f}^{-n}(\xi_n \times F) = P^{-1}(F_n)$, where F_n is the subset of $[\iota_{-n}, \ldots, \iota_{-1}]^u$ that is sent bijectively to F by the map f^n. Consequently,

$$\frac{\hat{\mu}(\xi_n \times F)}{\mu(F)} = \frac{\hat{\mu}\big(P^{-1}(F_n)\big)}{\mu(F)} = \frac{\mu(F_n)}{\int_{F_n} Jf^n \, d\mu} \ . \tag{42}$$

On the other hand, for μ-almost any point $x \in \Sigma$ and any cylinder $\xi_n \subset \Sigma^s$,

$$\hat{\mu}_x(\xi_n) = \lim_{F\to x} \frac{\hat{\mu}(\xi_n \times F)}{\mu(F)}$$

where the limit is over a basis of neighborhoods F of x. As $F \to x$, the sets F_n converge to the unique point in $[\iota_{-n}, \ldots, \iota_{-1}]^u$ that is mapped to x by f^n. This point is precisely $x^n = P\big(\hat{f}^{-n}(\hat{x})\big)$, for any choice of $\hat{x} \in \xi_n \times \{x\}$. In view of (42), this gives that

$$\hat{\mu}_x(\xi_n) = \frac{1}{Jf^n(x^n)}$$

for every cylinder ξ_n and any x in some full μ-measure subset. ∎

Lemma A.4. *There exists a disintegration $\big\{\hat{\mu}_x \colon x \in \Sigma\big\}$ of the lift $\hat{\mu}$ such that $\hat{\mu}_y = J_{x,y}\,\hat{\mu}_x$ for every x and y in Σ.*

Proof: Let $\big\{\bar{\mu}_x \colon x \in \Sigma\big\}$ be an arbitrary disintegration. By the previous lemma, there exists a full measure subset S of Σ such that

$$\text{(43)} \qquad \frac{\bar{\mu}_y(\xi_n)}{\bar{\mu}_x(\xi_n)} = \frac{Jf^n(x^n)}{Jf^n(y^n)} \qquad \text{for any} \quad \xi_n = [\iota_{-n}, \ldots, \iota_{-1}]^s \quad \text{and any} \quad x, y \in S \ ,$$

where $x^n = P\big(\hat{f}^{-n}(\hat{x})\big)$ and $y^n = P\big(\hat{f}^{-n}(\hat{y})\big)$ for any $\hat{x} \in \xi_n \times \{x\}$ and $\hat{y} \in \xi_n \times \{y\}$. Define $J_{n,x,y}$ to be the function on $W^s_{\mathrm{loc}}(x)$ which is constant equal to the right hand side of (43) on each $\xi_n \times \{x\}$. Given any cylinder $\eta \subset \Sigma^s$ and any large $n \geq 1$, we may write

$$\bar{\mu}_y(\eta) = \sum_{\xi_n \subset \eta} \bar{\mu}_y(\xi_n) = \sum_{\xi_n \subset \eta} J_{n,x,y}(\hat{x})\,\bar{\mu}_x(\xi_n) = \int_\eta J_{n,x,y}(\hat{x})\, d\hat{\mu}_x(\hat{x}) \ ,$$

where the sum is over all the cylinders ξ_n that form η. Passing to the limit as $n \to \infty$, we obtain from Lemma A.2 that

$$\bar{\mu}_y(\eta) = \int_\eta J_{x,y}\, d\bar{\mu}_x \qquad \text{for any cylinder} \quad \eta \subset \Sigma^s \, .$$

This shows that $\bar{\mu}_y = J_{x,y}\,\bar{\mu}_x$ for every x and y in the full measure set S. Fix any $\bar{x} \in S$ and define $\hat{\mu}_y = J_{\bar{x},y}\,\bar{\mu}_{\bar{x}}$ for every $y \in \Sigma$. Then $\hat{\mu}_y = \bar{\mu}_y$ for every $y \in S$, and so $\{\hat{\mu}_x\}$ is a disintegration of $\hat{\mu}$. Moreover,

$$\hat{\mu}_y = J_{\bar{x},y}\,\bar{\mu}_{\bar{x}} = J_{x,y}\,J_{\bar{x},x}\,\bar{\mu}_{\bar{x}} = J_{x,y}\,\hat{\mu}_x$$

for any $x, y \in \Sigma$, as claimed in the lemma. ∎

Proof of Proposition A.1: Fix an arbitrary point w in Σ and then define

$$r(x^s, x^u) = J_{w,x^u}(x^s, x^u) \qquad \text{for every} \quad \hat{x} = (x^s, x^u) \in \hat{\Sigma} \, .$$

By the previous lemma, $\hat\mu_{x^u} = r(x^s, x^u)\,\hat\mu_w$ for every $x^u \in \Sigma$. The lift $\hat\mu$ projects to $\mu^u = \mu$ on Σ, by definition. The projection μ^s to Σ^s is given by

$$\mu^s \;=\; \hat\mu_w \int_\Sigma r(x^s, x^u)\, d\mu(x^u)\ .$$

It follows that $\hat\mu = \rho(x^s, x^u)\,\mu^s \times \mu^u$, with

$$\rho(x^s, x^u) \;=\; \frac{r(x^s, x^u)}{\int_\Sigma r(x^s, x^u)\, d\mu(x^u)}\ .$$

Since the function $r(x^s, x^u)$ is continuous and uniformly bounded from zero and infinity, so is the density ρ. This implies that $\hat\mu$ has product structure. ∎

A.3. Fiber bunched cocycles

As pointed out in Section 1.3, existence of stable and unstable holonomies is automatic when the cocycle is locally constant. Another, more robust, construction of cocycles with stable and unstable holonomies was given in [3]. Let us recall it briefly here.

Definition A.5. We say that $\hat A : \hat\Sigma \to \mathrm{GL}(d,\mathbb{C})$ is *s-fiber bunched* (or *s-dominated*) for $\hat f : \hat\Sigma \to \hat\Sigma$ if there exist constants $N \ge 1$, $C > 0$, $\nu \in (0,1]$, $\tau \in (0,1)$ and $\theta \in (0,1)$, and a distance d on $\hat\Sigma$, such that

- **(a)** $d\big(\hat f^N(\hat x), \hat f^N(\hat y)\big) \le \theta\, d(\hat x, \hat y)$ if $\hat x$, $\hat y$ are in the same local stable set,
- **(b)** $\big\|\hat A^N(\hat x)^{\pm 1}\big\| \le C$ and $\big\|\hat A^N(\hat x) - \hat A^N(\hat y)\big\| \le C\, d(\hat x, \hat y)^\nu$,
- **(c)** $\big\|\hat A^N(\hat x)\big\|\,\big\|\hat A^N(\hat x)^{-1}\big\|\,\theta^\nu < \tau$,

for every $\hat x, \hat y \in \hat\Sigma$. We say that $\hat A$ is *u-fiber bunched* (or *u-dominated*) for $\hat f$ if $\hat A^{-1}$ is s-fiber bunched for $\hat f^{-1}$. □

Proposition A.6. *If $\hat A$ is s-fiber bunched (respectively, u-fiber bunched) then it admits stable holonomies (respectively, unstable holonomies).*

Proof: Replacing $\hat f$ by $\hat f^N$ in Definition A.5, we may assume $N = 1$. Denote $H_n(\hat x, \hat y) = \hat A^n(\hat y)^{-1} \hat A^n(\hat x)$ for each $n \ge 1$ and $\hat x$ and $\hat y$ in the same local stable set. Then

$$H_{n+1}(\hat x, \hat y) - H_n(\hat x, \hat y) \;=\; \hat A^n(\hat y)^{-1} \hat A\big(\hat f^n(\hat y)\big)^{-1} \Big[\hat A\big(\hat f^n(\hat x)\big) - \hat A\big(\hat f^n(\hat y)\big)\Big]\, \hat A^n(\hat x)\ .$$

By condition (a), we have $d\big(\hat f^n(\hat x), \hat f^n(\hat y)\big) \le \theta^n d(\hat x, \hat y)$. Using condition (b), it follows that

$$\big\|H_{n+1}(\hat x,\hat y) - H_n(\hat x,\hat y)\big\| \;\le\; C^2\, d(\hat x,\hat y)^\nu \prod_{j=0}^{n-1}\Big(\big\|\hat A\big(\hat f^j(\hat y)\big)^{-1}\big\|\,\big\|\hat A\big(\hat f^j(\hat x)\big)\big\|\,\theta^\nu\Big).$$

Fix $\hat\tau \in (\tau, 1)$. By conditions (a) and (b), $\hat A\big(\hat f^j(\hat y)\big)$ is close to $\hat A\big(\hat f^j(\hat x)\big)$ when j is large, uniformly on $\hat x$ and $\hat y$. Combining this with condition (c), we get that there exists $k \ge 1$, independent of $\hat x$ and $\hat y$, such that

$$\big\|\hat A\big(\hat f^j(\hat y)\big)^{-1}\big\|\,\big\|\hat A\big(\hat f^j(\hat x)\big)\big\|\,\theta^\nu < \hat\tau$$

for all $j \ge k$. Thus, the previous inequality implies that

$$\big\|H_{n+1}(\hat x,\hat y) - H_n(\hat x,\hat y)\big\| \;\le\; C^2\, d(\hat x,\hat y)^\nu\, C^{2k}\,\theta^{k\nu}\,\hat\tau^{n-k} \;\le\; \hat C\,\hat\tau^n\, d(\hat x,\hat y)^\nu\,,$$

for some appropriate constant $\hat C > 0$. This implies that H_n is a Cauchy sequence, uniformly on (x, y). Hence, it is uniformly convergent, as claimed. This proves that $\hat A$ admits stable holonomies if $\hat A$ is s-fiber bunched. The dual statement is proved in just the same way. ∎

We say that $\hat A\colon \hat\Sigma \to \mathrm{GL}(d,\mathbb{C})$ is *fiber bunched* if it is simultaneously s-fiber bunched and u-fiber bunched. From Proposition A.6 we immediately get that if $\hat A$ is fiber bunched then it admits stable and unstable holonomies.

Remark A.7. In some cases it is possible to reduce non-fiber bunched cocycles to the fiber bunched case. For instance, let $F = (f, A)$ be a linear cocycle $F = (f, A)$ over a shift map, say, which is not fiber bunched but whose Lyapunov spectrum is narrow, meaning that the difference between all Lyapunov exponents is sufficiently small. Then we may use inducing to construct from F a fiber bunched cocycle. □

A.4. Zorich cocycles

Finally, we are going to explain how the methods in this paper can be applied to Zorich cocycles [19, 20]. We begin by recalling the definition of these cocycles. Motivations and proofs for the results we quote can be found in Kontsevich, Zorich [9], Marmi, Moussa, Yoccoz [11], Rauzy [14], Veech [16, 17], Zorich [19, 20], and references therein. See also [1], where we show that Zorich cocycles are simple, thus proving the Zorich–Kontsevich conjecture that the corresponding Lyapunov spectra are simple.

A.4.1. The Rauzy algorithm

Fix some integer $d \geq 2$. Let $\Pi = \Pi_d$ be the set of all irreducible pairs $\pi = (\pi_0, \pi_1)$ of permutations $\pi_\varepsilon = (\alpha_1^\varepsilon, \alpha_2^\varepsilon, ..., \alpha_d^\varepsilon)$ of the alphabet $\{1, ..., d\}$. By *irreducible* we mean that $\pi_1 \circ \pi_0^{-1}$ preserves no subset $\{1, ..., k\}$ with $k < d$. We shall denote the rightmost symbol α_d^ε simply as $\alpha(\varepsilon)$ for $\varepsilon \in \{0, 1\}$. Let $\Delta = \Delta_d$ be the standard open simplex of dimension $d-1$, that is, the set of all vectors $\lambda = (\lambda_1, \lambda_2, ..., \lambda_d)$ such that $\lambda_j > 0$ for all j and $\sum_{j=1}^d \lambda_j = 1$. We call $g \colon \Delta \to \Delta$ a projective map if there exists a linear isomorphism $G \colon \mathbb{R}^d \to \mathbb{R}^d$ with non-negative coefficients such that

$$g(\lambda) = \frac{G(\lambda)}{\sum_{i=1}^d G(\lambda)_i} = \frac{G(\lambda)}{\sum_{i,j=1}^d G_{i,j}\,\lambda_j}\ . \tag{44}$$

If the coefficients of G are strictly positive then the image of g is relatively compact in Δ. In this case g is a contraction for the *projective metric* defined in Δ by

$$d(\lambda, \lambda') \ = \ \log \max \left\{ \frac{\lambda_i\,\lambda'_j}{\lambda_j\,\lambda'_i} \colon \ i, j = 1, ..., d \right\}.$$

The contraction rate depends only on a lower bound for the coefficients of G or, equivalently, for the Euclidean distance from $g(\Delta)$ to the boundary of Δ.

Let $\mathcal{R} \colon (\pi, \lambda) \mapsto (\pi', \lambda')$ be defined on an open dense subset of $\Pi \times \Delta$, as follows. For each $\pi \in \Pi$ and $\varepsilon \in \{0, 1\}$, let

$$\Delta^\varepsilon(\pi) \ = \ \Big\{ \lambda \in \Delta \colon \ \lambda_{\alpha(\varepsilon)} > \lambda_{\alpha(1-\varepsilon)} \Big\}.$$

We say that (π, λ) has *type* ε if $\lambda \in \Delta^\varepsilon(\pi)$. Then, by definition, $\pi'_\varepsilon = \pi_\varepsilon$ and

$$\pi'_{1-\varepsilon} \ = \ \Big(\alpha_1^{1-\varepsilon}, ..., \alpha_{k-1}^{1-\varepsilon}, \alpha(1-\varepsilon), \alpha_k^{1-\varepsilon}, ..., \alpha_{d-1}^{1-\varepsilon} \Big)$$

where $k \in \{1, ..., d-1\}$ is defined by $\alpha_k^{1-\varepsilon} = \alpha(\varepsilon)$. In other words, $\pi'_{1-\varepsilon}$ is obtained from $\pi_{1-\varepsilon}$ by looking for the position k the last symbol of π_ε occupies in $\pi_{1-\varepsilon}$, leaving all symbols to the left of k unchanged, and rotating the symbols to the right of k one position to the right. Moreover,

$$\lambda'_j = \frac{1}{a}\,\lambda_j \quad \text{for} \quad j \neq \alpha(\varepsilon)\ , \qquad \lambda'_j = \frac{1}{a}\left(\lambda_{\alpha(\varepsilon)} - \lambda_{\alpha(1-\varepsilon)}\right) \quad \text{for} \quad j = \alpha(\varepsilon)$$

where the normalizing factor $a = 1 - \lambda_{\alpha(1-\varepsilon)}$. Notice that $\lambda \mapsto \lambda'$ sends each Δ^ε bijectively onto Δ. Moreover, this map is just the projectivization of the linear isomorphism $R_{\pi,\lambda} \colon \mathbb{R}^d \to \mathbb{R}^d$

$$\Big(\lambda_1, ..., \lambda_{l-1}, \lambda_{\alpha(\varepsilon)}, \lambda_{l+1}, ..., \lambda_d\Big) \ \mapsto \ \Big(\lambda_1, ..., \lambda_{l-1}, \lambda_{\alpha(\varepsilon)} - \lambda_{\alpha(1-\varepsilon)}, \lambda_{l+1}, ..., \lambda_d\Big)\,,$$

in the sense that $\lambda' = (1/a)\, R_\pi(\lambda)$ with $a = \sum_{i=1}^d (R_\pi \lambda)_i$. It is interesting to write this also as $\lambda = a R_{\pi,\lambda}^{-1}(\lambda')$, because the inverse operator

$$\Big(\lambda_1, ..., \lambda_{l-1}, \lambda_{\alpha(\varepsilon)}, \lambda_{l+1}, ..., \lambda_d\Big) \mapsto \Big(\lambda_1, ..., \lambda_{l-1}, \lambda_{\alpha(\varepsilon)} + \lambda_{\alpha(1-\varepsilon)}, \lambda_{l+1}, ..., \lambda_d\Big).$$

has non-negative integer coefficients.

Let us call a Rauzy component of $\Pi\times\Delta$ any smallest set of the form $\Pi_0\times\Delta$ which is invariant under $\mathcal{R}$. From now on we always consider the restriction of the algorithm to some Rauzy component. The map $\mathcal{R}$ admits an absolutely continuous invariant measure ν, that is, an invariant measure such that the restriction to each $\{\pi\}\times\Delta$ is absolutely continuous with respect to Lebesgue measure on the standard simplex. However, ν is usually infinite. This can be overcome by considering the following accelerated algorithm.

A.4.2. The Zorich algorithm

Define $\mathcal{Z}(\pi,\lambda) = (\mathcal{R}^n)(\pi,\lambda)$, where the acceleration time $n = n(\pi,\lambda) \geq 1$ is the largest number of consecutive iterates by the Rauzy algorithm during which the type remains unchanged. In precise terms, $n = n(\pi,\lambda)$ is characterized by (assume $\big(\pi^{(i)}, \lambda^{(i)}\big) = \mathcal{R}^i(\pi,\lambda)$ is defined for all $0 \leq i \leq n$)

$$\big(\pi^{(i)}, \lambda^{(i)}\big) \text{ has type } \varepsilon \text{ for } 0 \leq i < n \quad \text{ and } \quad \big(\pi^{(n)}, \lambda^{(n)}\big) \text{ has type } 1-\varepsilon\ .$$

Since each $\mathcal{R}\colon \{\pi^{(i)}\}\times\Delta^\varepsilon(\pi^{(i)}) \to \{\pi^{(i+1)}\}\times\Delta$ is a projective bijection, the map $\mathcal{R}^n$ sends some sub-simplex $\{\pi\}\times D(\pi,\lambda) \subset \{\pi\}\times\Delta^\varepsilon(\pi)$ containing (π,λ) bijectively onto $\{\pi^{(n)}\}\times\Delta^{1-\varepsilon}(\pi^{(n)})$. Moreover, its inverse is the restriction of a projective map $\{\pi^{(n)}\}\times\Delta \to \{\pi\}\times\Delta$. By definition, $\mathcal{Z} = \mathcal{R}^n$ restricted to $D(\pi,\lambda)$. Let $\mathcal{D}$ be the (countable) family of all these sub-simplices $D(\pi,\lambda)$. The union of its elements has full measure on $\Pi\times\Delta$.

The transformation $\mathcal{Z}$ admits an absolutely continuous invariant probability measure μ on each Rauzy component, and this measure is ergodic. Moreover, the density of μ is a rational function of the form

$$\frac{d\mu}{dm}(\lambda) = \sum_\alpha \frac{1}{\mathcal{P}_\alpha(\lambda)} \qquad \text{on each domain } \{\pi\}\times\Delta \tag{45}$$

where the sum is over some finite set of polynomials $\mathcal{P}_\alpha$ with non-negative coefficients and degree d. In particular, the density is smooth and bounded from zero on every $\{\pi\}\times\Delta$. In general, the density is not bounded from infinity, because the $\mathcal{P}_\alpha$ may have zeros on the boundary of Δ.

A.4.3. Linear cocycles

The *Rauzy cocycle* over $\mathcal{R}$ is defined by

$$F_R\colon \Pi\times\Delta\times\mathbb{R}^d \to \Pi\times\Delta\times\mathbb{R}^d\,, \qquad \big(\pi,\lambda,v\big) \mapsto \big(\mathcal{R}(\pi,\lambda), R_{\pi,\lambda}^{-1*}(v)\big)\,.$$

Notice that this cocycle is constant on each $\Delta^\varepsilon(\pi)$, because $R_{\pi,\lambda}$ depends only on π and the type ε of λ. The *Zorich cocycle* over $\mathcal{Z}$ is defined by

$$F_Z\colon \Pi\times\Delta\times\mathbb{R}^d \to \Pi\times\Delta\times\mathbb{R}^d\,, \qquad F_Z\big(\pi,\lambda,v\big) = F_R^{n(\pi,\lambda)}\big(\pi,\lambda,v\big)\,.$$

Notice that $F_Z\big(\pi,\lambda,v\big) = \big(\mathcal{Z}(\pi,\lambda), Z_{\pi,\lambda}(v)\big)$ where $Z_{\pi,\lambda}$ is constant on each element of $\mathcal{D}$ and its inverse has non-negative integer coefficients. The Zorich cocycle is integrable with respect that the $\mathcal{Z}$-invariant measure μ, meaning that $\log\big\|Z_{\pi,\lambda}^{\pm 1}\big\|$ are integrable functions. Thus, its Lyapunov exponents are well-defined at μ-almost every point. By ergodicity, the exponents are constant μ-almost everywhere.

Consider the linear map $\Omega_\pi\colon \mathbb{R}^d \to \mathbb{R}^d$ defined by

$$\Omega_\pi(\lambda)_i \;=\; \sum_{j\colon\, \pi_1(j)<\pi_1(i)} \lambda_j \;-\; \sum_{j\colon\, \pi_0(j)<\pi_0(i)} \lambda_j\;.$$

This map Ω_π is anti-symmetric (not necessarily an isomorphism), and so

$$\omega_\pi\big(\Omega_\pi(u),\Omega_\pi(v)\big) \;=\; u\cdot\Omega_\pi(v)$$

defines a symplectic form on the range $H_\pi = \Omega_\pi(\mathbb{R}^d)$. In particular, the dimension of H_π is even. The map Ω_π also satisfies

$$\Omega_{\pi'}\cdot R_{\pi,\lambda} = R_{\pi,\lambda}^{-1*}\cdot\Omega_\pi\;. \tag{46}$$

This implies that the Rauzy cocycle leaves invariant the subbundle

$$\mathcal{H}_\Pi \;=\; \Big\{(\pi,\lambda,v)\in\Pi\times\Delta\times\mathbb{R}^d\colon\; v\in H_\pi\Big\}$$

and even preserves the symplectic form ω_π on it. Then the same is true for the Zorich cocycle.

It follows that the Lyapunov spectrum of the Zorich cocycle restricted to the subbundle $\mathcal{H}_\Pi$ has the form

$$\lambda_1\geq\cdots\geq\lambda_g\geq 0\geq -\lambda_g\geq\cdots\geq -\lambda_1 \qquad (\text{where }\; 2g=\dim H_\pi)\;. \tag{47}$$

The other Lyapunov exponents of F_Z, corresponding to directions transverse to $\mathcal{H}_\pi$, vanish identically and are not of interest here. The *Zorich–Kontsevich* *conjecture* states that all the inequalities in (47) are strict or, in other words, *the Lyapunov spectrum of the restricted Zorich cocycle is simple.* We are going to argue that, modulo the simple observations in Sections A.1 and A.2, all the hypotheses of Theorem A are satisfied in the context of Zorich cocycles, and so our methods can be used to prove this conjecture.

A.4.4. Inducing on a compact simplex

Let $\mathcal{D}$ be the family of sub-simplices introduced in the definition of the Zorich algorithm: $\mathcal{Z}$ maps each element of $\mathcal{D}$ bijectively to some $\{\pi'\}\times\Delta^{1-\varepsilon}$, and the inverse is the restriction of a projective map $\{\pi'\}\times\Delta \to \{\pi\}\times\Delta$. Pulling $\mathcal{D}$ back under $\mathcal{Z}$ we obtain, for each $n \geq 1$, a countable family $\mathcal{D}^n$ of sub-simplices each of which is mapped bijectively to some $\{\pi^{(n)}\}\times\Delta^{1-\varepsilon}$ by the iterate $\mathcal{Z}^n$, the inverse being the restriction of a projective map $\{\pi^{(n)}\}\times\Delta \to \{\pi\}\times\Delta$. For μ-almost every (π,λ), there exists some $n \geq 1$ for which this projective map has strictly positive coefficients, and so the image $\{\pi\}\times\Gamma$ is relatively compact in $\{\pi\}\times\Delta$. Let us fix such n, π, λ once and for all, and denote by $\{\pi\}\times D_*$ the corresponding element of $\mathcal{D}^n$. In particular, $D_* \subset \Gamma$ is relatively compact in Δ. It follows that D_* has finite diameter for the projective metric of Δ, and also that the density $d\mu/dm$ is smooth and bounded from zero and infinity on D_*. For notational simplicity, we identify $\{\pi\}\times\Delta \approx \Delta$ and $\{\pi\}\times D_* \approx D_*$ in what follows.

By Poincaré recurrence, there exists a first return map $\mathcal{G}$ of the map $\mathcal{Z}^n$ to the domain D_*. More precisely, using the Markov structure of $\mathcal{Z}^n$, there exists a countable family $\{D_\iota \colon \iota \in \mathbb{N}\} \subset \bigcup_{k\geq 1}\mathcal{D}_{kn}$ of sub-simplices of D_* such that their union has full measure in D_*, each D_ι is mapped bijectively to the whole D_* by $\mathcal{G}$, and the inverse of each $\mathcal{G}\colon D_\iota \to D_*$ is the restriction of a projective map $\Delta \to \Delta$. By construction, the images of these inverse branches are all contained in Γ, and so they all contract the projective metric, with uniform contraction rates. Let $D \subset D_*$ be the (full measure) subset of points that return infinitely many times to D_*. In particular, the map

$$\Phi\colon \mathbb{N}^{\{n\geq 0\}} \to D\,, \qquad (\iota_n)_n \mapsto \bigcap_{n\geq 0} \mathcal{G}^{-n}(D_{\iota_n})$$

is well defined (the intersection consists of exactly one point), and it conjugates $\mathcal{G}\colon D \to D$ to the shift map on $\mathbb{N}^{\{n\geq 0\}}$. Then the natural extension of $\mathcal{G}$ is realized by the shift map on $\mathbb{N}^{\mathbb{Z}}$.

On the one hand, as observed before, the invariant density $d\mu/dm$ is smooth and bounded from zero and infinity on D. It follows that its logarithm is bounded and Lipschitz continuous, for either Euclidean or projective metric, with uniform constants. On the other hand, the inverse branches of $\mathcal{G}$ are all projective maps with range contained in the same relatively compact domain Γ. This implies that the logarithms of their derivatives are also bounded and Lipschitz continuous, for either metric, with uniform constants. Putting these two facts together we get that the logarithm of the Jacobian of $\mathcal{G}$ with respect to the measure μ is uniformly bounded and Lipschitz continuous on each D_ι. Combining this with the previous observation that inverse branches of $\mathcal{G}$ contract the projective metric uniformly, we easily obtain that $\log J\mathcal{G}$ has bounded oscillation in the sense of Section A.2. Consequently, the lift of $\mu | D$ to the natural extension of $\mathcal{G}$ has product structure.

Recall that the Zorich cocycle F_Z is constant on each element of $\mathcal{D}$. It is clear from the construction that points in each D_ι visit exactly the same elements of $\mathcal{D}$ all the way up to their return to D_*. Thus, the linear cocycle F_G induced by F_Z over the return map $\mathcal{G}$ is also locally constant, meaning that it is constant on each D_ι. In particular, the cocycle F_G is continuous for the shift topology, and it admits stable and unstable holonomies.

A.4.5. Pinching and twisting conditions

The only missing ingredient to establish the Zorich–Kontsevich conjecture is to prove that the Zorich cocycles are simple, in the sense of Definition 1.2. This is done in [1]. In fact, the pinching and twisting conditions appear in a slightly different guise in that paper, in terms of the monoid generated by the cocycle.

In this context, a *monoid* is just a subset of $\mathrm{GL}(d, \mathbb{C})$ closed under multiplication and containing the identity. The *associated monoid* $\mathcal{B} = \mathcal{B}(F)$ is the smallest monoid that contains the image of F. We call $\mathcal{B}$ is *simple* if it is both pinching and twisting, where $\mathcal{B}$ is

- *pinching* if it contains elements with arbitrarily large eccentricity $\mathrm{Ecc}(B)$;
- *twisting* if for any $F \in \mathrm{Grass}(\ell, d)$ and any finite family $G_1, ..., G_N$ of elements of $\mathrm{Grass}(\ell, d)$ there exists $B \in \mathcal{B}$ such that $B(F) \cap G_i = \{0\}$ for all $j = 1, ..., N$.

The *eccentricity* of a linear map $B \in \mathrm{GL}(d, \mathbb{C})$ is defined by

$$\mathrm{Ecc}(B) = \min_{1 \le \ell < d} \frac{\sigma_\ell}{\sigma_{\ell+1}}$$

where $\sigma_1^2 \ge \cdots \ge \sigma_d^2$ are the eigenvalues of the self-adjoint operator B^*B, in non-increasing order. Geometrically, the positive square roots $\sigma_1 \ge \cdots \ge \sigma_d$

correspond to the lengths of the semi-axes of the ellipsoid $\{B(v)\colon \|v\|=1\}$. It is evident from the definition that any monoid that contains a pinching submonoid is also pinching, and analogously for twisting.

It is not difficult to see that the two formulations of the definition of simplicity are equivalent, for locally constant real cocycles. Indeed, Lemma A.5 in [1] states that if the associated monoid is simple then there exists some periodic point and some homoclinic point as in Definition 1.2. Conversely, the conditions in Definition 1.2 imply that the associated monoid is simple. Indeed, the first condition implies that $\mathcal{B}$ contains some element B_1 whose eigenvalues all have distinct norms. Then the powers B_1^n have arbitrarily large eccentricity as $n \to \infty$, and so $\mathcal{B}$ is pinching. Moreover, the second condition implies that the monoid contains some element B_2 satisfying $B_2(V) \cap W = \{0\}$ for any pair of subspaces V and W which are sums of eigenspaces of B_1 and have complementary dimensions. Given any F, G_1, ..., G_n as in the definition, we have that $B_1^n(F)$ is close to some sum V of ℓ eigenspaces of B_1, and every $B_1^{-n}(G_i)$ is close to some sum W_i of $d-\ell$ eigenspaces of B_1, as long as n is large enough. It follows that $B_2(B_1^n(F)) \cap B_1^{-n}(G_i) = \{0\}$, that is, $B_1^n B_2 B_1^n(F) \cap G_i = \{0\}$. This proves $\mathcal{B}$ is twisting.

B – Intersections of hyperplane sections

Here we give an alternative proof of Proposition 5.1 under the assumption that the eigenvalues of the cocycle at the fixed point p are real. Observe that this is automatic for real cocycles, since we also assume that the absolute values of the eigenvalues are all distinct. Instead of Proposition 5.5 we use the following result, which has a stronger conclusion.

Proposition B.1. *There exists $N = N(\ell, d)$ such that*

$$B^{-m_1}(V) \cap \cdots \cap B^{-m_N}(V) = \emptyset$$

for any $B\colon \mathbb{C}^d \to \mathbb{C}^d$ whose eigenvalues all have distinct absolute values, any hyperplane section V of $\operatorname{Grass}(\ell, d)$ *containing no eigenspace of B, and any $0 \le m_1 < \cdots < m_N$.*

To deduce Proposition 5.1 from this result, one can use the same arguments as in Section 5, just replacing the paragraph that contains (23) by the following one.

Applying Proposition B.1 with $B = A^q(p)$ and $V = W_{\hat{p}}$ we conclude that the $W_{\hat{p}}^n$ are N-wise disjoint:

$$W_{\hat{p}}^{m_1} \cap \cdots \cap W_{\hat{p}}^{m_N} = \emptyset \quad \text{for all} \ \ 1 \le m_1 < \cdots < m_N \ .$$

Fix $C \geq 1$ such $C\gamma_0 > 1$. By continuity, we have $W_\eta^{m_1} \cap \cdots \cap W_\eta^{m_N} = \emptyset$ for all $1 \leq m_1 < \cdots < m_N \leq CN$ and every η in a small neighborhood of $\hat{p}$ inside the local stable set. Then, for $\hat{\mu}_p$-almost every η in that neighborhood,

$$\hat{m}_\eta \left(\bigcup_{j=1}^{CN} W_\eta^j \right) \geq \frac{1}{N} \sum_{j=1}^{CN} \hat{m}_\eta(W_\eta^j) \ = \ C\gamma_0 \ > \ 1 \ .$$

This is a contradiction, since $\hat{m}_\eta$ is a probability. This contradiction reduces the proof of Proposition 5.1 to proving Proposition B.1.

In the proof of Proposition B.1 we use the following classical fact about Vandermonde type determinants (see Mitchell [12]). Given $N \geq 1$, $x = (x_1, ..., x_N) \in \mathbb{R}^N$, and $m = (m_1, ..., m_N) \in \big(\mathbb{N} \cup \{0\}\big)^N$, define

$$\Delta_m(x) \ = \begin{vmatrix} x_1^{m_1} & \cdots & x_N^{m_1} \\ \cdots & \cdots & \cdots \\ x_1^{m_N} & \cdots & x_N^{m_N} \end{vmatrix} \ .$$

Proposition B.2. *Suppose* $0 \leq m_1 < m_2 < \cdots < m_N$. *Then*

$$\Delta_m(x) \ = \ \mathcal{P}_m(x) \prod_{1 \le i < j \le d} (x_j - x_i)$$

where $\mathcal{P}_m$ *is a positive polynomial, in the sense that all its monomials have positive coefficients. In particular,* $\Delta_m(x)$ *is different from zero whenever the* x_j *are all positive and distinct.* ∎

Notice that the contents of the proposition remains the same if one replaces B by its square. Indeed, it is trivial that the statement for B implies the one for B^2, and the converse is also easy to check: if the B^2-iterates of any hyperplane section V as in the statement are N-wise disjoint then, using this fact both for V and for $B(V)$, the B-iterates of any such hyperplane section V are $2N$-wise disjoint. Thus, we may always assume the eigenvalues of B to be positive.

Let $\{\theta_1, ..., \theta_d\}$ be a basis of eigenvectors of B, in decreasing order of the eigenvalues $b_1 > \cdots > b_d > 0$. Let $V = \pi_v\big(\Lambda_v^\ell(\mathbb{C}^d) \cap H\big)$ be as in the statement, where

H is the geometric hyperplane of $\Lambda^\ell(\mathbb{C}^d)$ defined by some non-zero $(d-\ell)$-vector υ. Let us write

$$\upsilon = \sum_I \upsilon(i_1, ..., i_\ell) \left(\theta_{j_{\ell+1}} \wedge \cdots \wedge \theta_{j_d}\right),$$

where the sum is over all sequences $I = (i_1, ..., i_\ell)$ with $1 \le i_1 < \cdots < i_\ell \le d$, the $\upsilon(I)$ are scalars, and $j_{\ell+1} < \cdots < j_d$ are the elements of $\{1, ..., d\}$ that are *not* in I. The assumption that V contains no eigenspaces of B implies that every $\upsilon(I)$ is non-zero: otherwise, $\upsilon \wedge (\theta_{i_1} \wedge \cdots \wedge \theta_{i_\ell})$ would vanish, that is, $\pi_\upsilon(\upsilon)$ would have a non-trivial intersection with the subspace generated by $\theta_{i_1}, ..., \theta_{i_\ell}$. Likewise, let us write

$$\omega = \sum_I \omega(i_1, ..., i_\ell) \left(\theta_{i_1} \wedge \cdots \wedge \theta_{i_\ell}\right), \tag{48}$$

where the $\omega(I)$ are scalars. Then $B^{-m}(H) = \{\omega\colon \omega \wedge B^{-m}\upsilon = 0\}$, and

$$\omega \wedge B^{-m}\upsilon = \sum_I b_I^{-m} \sigma_I\, \omega(I)\, \upsilon(I),$$

where $b_I = b_{j_{\ell+1}} \cdots b_{j_d} > 0$ and $\sigma_I = \theta_{i_1} \wedge \cdots \wedge \theta_{i_\ell} \wedge \theta_{j_{\ell+1}} \wedge \cdots \wedge \theta_{j_d}$ is either ± 1.

Fix $N = \dim \Lambda^\ell(\mathbb{C}^d)$ and then let $0 \le m_1 < \cdots < m_N$. In view of the previous paragraph, in order to prove that the intersection of all the $B^{-m_u}(H)$ is empty it suffices to show that there does not exist any non-zero $\omega \in \Lambda^\ell_v(\mathbb{C}^d)$ such that

$$\sum_I b_I^{-m_u} \sigma_I\, \omega(I)\, \upsilon(I) = 0 \qquad \text{for all} \quad u = 1, ..., N, \tag{49}$$

that is, such that the vector $\left(\sigma_I\, \omega(I)\, \upsilon(I)\right)_I$ is in the kernel of $X = \left(b_I^{m_u}\right)_{I,u}$. It is useful to consider first the special case when the b_I are all distinct (and positive). Then, by Proposition B.2, the kernel of X is trivial. This means that (49) implies $\sigma_I\, \omega(I)\, \upsilon(I) = 0$ for every I. Since $\sigma_I\, \upsilon(I)$ never vanishes, this means that $\omega(I) = 0$ for every I. This proves Proposition B.1 in this case. Notice that this argument applies to any element ω of $\Lambda^\ell(\mathbb{C}^d)$, not only ℓ-vectors. Hence, it proves that, under this stronger assumption, the relation (49) has no non-zero solution in the whole exterior power $\Lambda^\ell(\mathbb{C}^d)$.

In general, when the products b_I are not all distinct, condition (49) may hold on a subspace of $\Lambda^\ell(\mathbb{C}^d)$ with positive dimension. The main point in the proof of Proposition B.1 is then to show that this subspace intersects the set of ℓ-vectors at the origin only. From Proposition B.2 we do get that the relation (49) implies

$$\sum_{b_J = b_I} \sigma_J\, \omega(J)\, \upsilon(J) = 0 \qquad \text{for any admissible sequence } I \tag{50}$$

(*admissible* means that $1 \le i_1 < \cdots < i_\ell \le d$), where the sum is over all admissible sequences J such that $b_J = b_I$. So, what we really need to prove is

Lemma B.3. *If an ℓ-vector $\omega = \omega_1 \wedge \cdots \wedge \omega_\ell$ is a solution of* (50) *then $\omega(I) = 0$ for every admissible sequence $I = (i_1, ..., i_l)$.*

Proof: Begin by noting that, for an ℓ-vector $\omega = \omega_1 \wedge \cdots \wedge \omega_\ell$, the coefficients $\omega(I)$ in (48) may be expressed in terms of the vectors ω_i, as follows:

$$\omega(I) = \begin{vmatrix} \omega_1^{i_1} & \cdots & \omega_1^{i_\ell} \\ \cdots & \cdots & \cdots \\ \omega_\ell^{i_1} & \cdots & \omega_\ell^{i_\ell} \end{vmatrix},$$

where $\omega_j = (\omega_j^1, ..., \omega_j^d)$. For each $1 \le j \le d$, let $\omega^i = (\omega_1^i, ..., \omega_\ell^i)$ be a column vector. Hence, $\omega(i_1, ..., i_\ell) \ne 0$ if and only if the vectors $\omega^{i_1}, ..., \omega^{i_\ell}$, are linearly independent. More generally, given any $1 \le s \le \ell$ and $j_1, ..., j_s$, we write $\omega(j_1, ..., j_s) \ne 0$ to mean the vectors $\omega^{j_1}, ..., \omega^{j_s}$ are linearly independent.

Consider first $I = (1, ..., \ell)$. Since we assume $b_1 > \cdots > b_d$, we have $b_I > b_J$ for any admissible sequence $J \ne I$. Thus, relation (50) reduces to $\sigma_I \, \omega(I) \, \upsilon(I) = 0$. Since $\sigma_I \, \upsilon(I)$ is non-zero, that gives $\omega(I) = 0$. Now the proof of Lemma B.3 continues by induction: we consider any admissible sequence I, and assume $\omega(J) = 0$ for every admissible sequence J such that $b_J > b_I$. We use the following simple observation:

Lemma B.4. *Suppose $\omega(j_1, ..., j_s, j, j_{s+1}) = 0$ and $\omega(j_1, ..., j_s, j, j_{s+2}) = 0$, but $\omega(j_1, ..., j_s, j) \ne 0$. Then $\omega(j_1, ..., j_s, j_{s+1}, j_{s+2}) = 0$.*

Proof: The assumptions mean that both $\omega^{j_{s+1}}$ and $\omega^{j_{s+2}}$ are linear combinations of $\{\omega^{j_1}, ..., \omega^{j_s}, \omega^j\}$, and so the set $\{\omega^{j_1}, ..., \omega^{j_s}, \omega^{j_{s+1}}, \omega^{j_{s+2}}\}$ is contained in the $(s+1)$-dimensional subspace generated by $\{\omega^{j_1}, ..., \omega^{j_s}, \omega^j\}$. This implies that $\omega(j_1, ..., j_s, j_{s+1}, j_{s+2}) = 0$. ∎

Lemma B.5. *If $\omega(I) \ne 0$ then we have $\omega(j_1, ..., j_s, j) = 0$ for every $0 \le s \le \ell - 1$, every $j \notin \{i_1, ..., i_\ell\}$, and every $\{j_1, ..., j_s\} \subset \{i_1, ..., i_\ell\}$ that contains all $i_t < j$.*

Proof: Consider first the case $\ell - s = 1$. Then $(j_1, ..., j_s)$ misses exactly one element i_t of I, and we have $j < i_t$. Let J be the admissible sequence obtained by ordering $(j_1, ..., j_s, j)$. Notice that $b_J > b_I$, because $b_j > b_{i_t}$. By induction, we get that $\omega(J)=0$, as claimed. Now the proof proceeds by induction on $\ell-s$. Suppose $\ell-s\geq 2$ and let $j_1,...,j_s,j$ be as in the statement. Choose two different elements j_{s+1} and j_{s+2} of $\{i_1,...,i_\ell\}\setminus\{j_1,...,j_s\}$. By induction,

$$\omega\big(j_1,...,j_s,j,j_{s+1}\big)=0 \quad \text{and} \quad \omega\big(j_1,...,j_s,j,j_{s+2}\big)=0\ .$$

Suppose $\omega(j_1,...,j_s,j)\neq 0$. Then, we would be able to use Lemma B.4 to conclude that

$$\omega\big(j_1,...,j_s,j_{s+1},j_{s+2}\big)=0\ .$$

Since the j_i are distinct elements of $\{i_1,...,i_\ell\}$, that would imply $\omega(i_1,...,i_\ell)=0$, which would contradict the hypothesis. This proves that $\omega(j_1,...,j_s,j)=0$, and so the proof of Lemma B.5 is complete. ∎

Remark B.6. Notice that $s=0$ is compatible with the other assumptions only if $i_1>1$. Then the lemma gives that $\omega(j)=0$ or, equivalently, the column vector $\omega^j=0$, for every $1\leq j<i_1$. This means that the ℓ-vector ω really lives inside a lower dimensional space, corresponding to coordinates i_1 through d only. This case could be easily disposed of, just by assuming Lemma B.3 has already been proved for dimensions smaller than d. □

Let $\prec$ be the usual lexicographical order: $(j_1,...,j_r)\prec(i_1,...,i_r)$ if and only if there exists $0\leq s\leq r-1$ such that $j_1=i_1$, ..., $j_s=i_s$, and $j_{s+1}<i_{s+1}$.

Corollary B.7. *If $\omega(I)\neq 0$ then $\omega(J)=0$ for every $J\prec I$.*

Proof: Fix $0\leq s\leq \ell-1$ as in the definition of $J\prec I$, that is, such that $j_1=i_1$, ..., $j_s=i_s$, and $j_{s+1}<i_{s+1}$. By Lemma B.5, we have $\omega(j_1,...,j_s,j_{s+1})=0$. Consequently, $\omega(j_1,...,j_\ell)=0$, as claimed. ∎

Now the inductive step in the proof of Lemma B.3 is an easy consequence. By Corollary B.7, inside the class of all sequences J with $b_J=b_I$ there exists at most one J such that $\omega(J)\neq 0$. Then the relation (50) reduces to $\sigma_J\,\omega(J)\,\upsilon(J)=0$. Since $\sigma_J\,\upsilon(J)$ never vanishes, this gives $\omega(J)=0$. In other words, $\omega(J)=0$ for every J such that $b_J=b_I$. This finishes the proof of Lemma B.3. ∎

The proof of Proposition B.1 is complete. ∎

ACKNOWLEDGEMENTS – We are grateful to A. Arbieto, C. Matheus, and J.-C. Yoccoz, for several useful conversations, and to E. Esteves for explanations on the structure of the Grassmannian manifolds.

REFERENCES

[**1**] Avila, A. and Viana, M. – *Symplicity of Lyapunov spectra: Proof of the Zorich–Kontsevich conjecture*, Preprint. `www.preprint.impa.br` and `www.archiv.org 2005`.

[**2**] Bonatti, C.; Díaz, L.J. and Viana, M. – *Dynamics beyond uniform hyperbolicity*, volume 102 of "Encyclopaedia of Mathematical Sciences", Springer-Verlag, 2005.

[**3**] Bonatti, C.; Gómez-Mont, X. and Viana, M. – Généricité d'exposants de Lyapunov non-nuls pour des produits déterministes de matrices, *Ann. Inst. H. Poincaré Anal. Non Linéaire,* 20 (2003), 579–624.

[**4**] Bonatti, C. and Viana, M. – Lyapunov exponents with multiplicity 1 for deterministic products of matrices, *Ergod. Th. & Dynam. Sys.,* 24 (2004), 1295–1330.

[**5**] Durret, R. – *Probability: Theory and Examples*, 2nd edition, Wadsworth & Brooks Publ., 1996.

[**6**] Furstenberg, H. – *Boundary theory and stochastic processes on homogeneous spaces*, in "Harmonic Analysis in Homogeneous Spaces", volume XXVI of *Proc. Sympos. Pure Math. (Williamstown MA, 1972)*, pp. 193–229, Amer. Math. Soc., 1973.

[**7**] Gol'dsheid, I.Ya. and Margulis, G.A. – Lyapunov indices of a product of random matrices, *Uspekhi Mat. Nauk.,* 44 (1989), 13–60.

[**8**] Guivarc'h, Y. and Raugi, A. – Products of random matrices: convergence theorems, *Contemp. Math.,* 50 (1986), 31–54.

[**9**] Kontsevich, M. – *Lyapunov exponents and Hodge theory*, in "The Mathematical Beauty of Physics (Saclay, 1996)", volume 24 of *Adv. Ser. Math. Phys.*, pp. 318–332, World Sci. Publishing, 1997.

[**10**] Krengel, U. – *Ergodic Theorems*, De Gruyter Publ., 1985.

[**11**] Marmi, S.; Moussa, P. and Yoccoz, J.-C. – The cohomological equation for Roth type interval exchange transformations, *Journal of the Amer. Math. Soc.,* 18 (2005), 823–872.

[**12**] Mitchell, O.H. – Note on determinants of powers, *Amer. J. Math.,* 4 (1881), 341–344.

[**13**] Oseledets, V.I. – A multiplicative ergodic theorem: Lyapunov characteristic numbers for dynamical systems, *Trans. Moscow Math. Soc.,* 19 (1968), 197–231.

[**14**] Rauzy, G. – Échanges d'intervalles et transformations induites, *Acta Arith.,* 34 (1979), 315–328.

[**15**] Rokhlin, V.A. – On the fundamental ideas of measure theory, *A.M.S. Transl.,* 10 (1962), 1–52; Transl. from *Mat. Sbornik,* 25 (1949), 107–150.

[**16**] Veech, W. – Gauss measures for transformations on the space of interval exchange maps, *Ann. of Math.,* 115 (1982), 201–242.

[**17**] Veech, W. – The Teichmüller geodesic flow, *Ann. of Math.,* 124 (1986), 441–530.

[**18**] Walters, P. – *An Introduction to Ergodic Theory*, Springer Verlag, 1982.

[**19**] Zorich, A. – Finite Gauss measure on the space of interval exchange transformations. Lyapunov exponents, *Ann. Inst. Fourier (Grenoble),* 46 (1996), 325–370.

[**20**] Zorich, A. – *How do the leaves of a closed* 1*-form wind around a surface?* In "Pseudoperiodic Topology", volume 197 of *Amer. Math. Soc. Transl. Ser. 2,* pp. 135–178, Amer. Math. Soc., 1999.

Artur Avila,
CNRS UMR 7599,
Laboratoire de Probabilités et Modèles Aléatoires, Université Pierre et Marie Curie,
Boîte Postale 188, 75252 Paris Cedex 05 – FRANCE
E-mail: `artur@ccr.jussieu.fr`

and

Marcelo Viana,
IMPA – Estrada D. Castorina 110, Jardim Botânico,
22460-320 Rio de Janeiro – BRAZIL
E-mail: `viana@impa.br`