Continuity of Lyapunov exponents

Marcelo Viana
(joint work with C. Bocker and with A. Avila)

Instituto de Matemática Pura e Aplicada

Lyapunov exponents

Consider $A_1, \ldots, A_N \in \mathrm{GL}(d)$ and $p_1, \ldots, p_N > 0$ with $\sum_j p_j = 1$.

Let $(B_n)_n$ be identical independent random variables in $\mathrm{GL}(d)$ with probability distribution $\mu = \sum_j p_j \delta_{A_j}$. The Lyapunov exponents

$$\lambda_+(\mu) = \lim_n \frac{1}{n} \log \|B_n \cdots B_1\|$$

$$\lambda_-(\mu) = \lim_n -\frac{1}{n} \log \|(B_n \cdots B_1)^{-1}\|$$

exist almost surely.

Lyapunov exponents

Consider $A_1, \ldots, A_N \in \mathrm{GL}(d)$ and $p_1, \ldots, p_N > 0$ with $\sum_j p_j = 1$.

Let $(B_n)_n$ be identical independent random variables in $\mathrm{GL}(d)$ with probability distribution $\mu = \sum_j p_j \delta_{A_j}$. The Lyapunov exponents

$$\lambda_+(\mu) = \lim_n \frac{1}{n} \log \|B_n \cdots B_1\|$$

$$\lambda_-(\mu) = \lim_n -\frac{1}{n} \log \|(B_n \cdots B_1)^{-1}\|$$

exist almost surely.

Theorem (Carlos Bocker, MV)

For $d = 2$, the functions $(A_{i,j}, p_j)_{i,j} \mapsto \lambda_\pm$ are continuous.

Measure space - Stationary measures

What about $d > 2$? (ongoing project, with Artur Avila).

Measure space - Stationary measures

What about $d > 2$? (ongoing project, with Artur Avila).

$\mathcal{G}(d) = \{$compactly supported probability measures on $\mathrm{GL}(d)\}$
with the following topology:
ν is close to μ if ν is weak*-close to μ and $\operatorname{supp}\nu \subset B_\epsilon(\operatorname{supp}\mu)$.

Measure space - Stationary measures

What about $d > 2$? (ongoing project, with Artur Avila).

$\mathcal{G}(d) = \{$compactly supported probability measures on $\mathrm{GL}(d)\}$
with the following topology:
ν is close to μ if ν is weak*-close to μ and $\operatorname{supp}\nu \subset B_\epsilon(\operatorname{supp}\mu)$.

"Theorem"

For any $d \geq 2$, the functions $\mu \mapsto \lambda_\pm(\mu)$ are continuous on $\mathcal{G}(d)$.

"Counterexamples"

Example

For $A_1 = \begin{pmatrix} 2^{-1} & 0 \\ 0 & 2 \end{pmatrix} \qquad A_2 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$

we have $\lambda_+ = 0$ if $p_2 > 0$ but $\lambda_+ = \log 2$ if $p_2 = 0$.

"Counterexamples"

Example

For $A_1 = \begin{pmatrix} 2^{-1} & 0 \\ 0 & 2 \end{pmatrix}$ $\quad A_2 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$

we have $\lambda_+ = 0$ if $p_2 > 0$ but $\lambda_+ = \log 2$ if $p_2 = 0$.

Mañé-Bochi-V:

For any ergodic measure preserving transformation $f : M \to M$, the continuity points for Lyapunov exponents in the space of $\mathrm{GL}(d)$-cocycles over f are very special cocycles: the Oseledets splitting is dominated.

Representation of exponents

Given $\mu \in \mathcal{G}(d)$:

$L < \mathbb{R}^d$ is μ-invariant if $g(L) = L$ for every $g \in \operatorname{supp} \mu$.

A probability η in $\mathbb{PR}^d$ is μ-stationary if $\int g_* \eta \, d\mu(g) = \eta$.

Representation of exponents

Given $\mu \in \mathcal{G}(d)$:

$L < \mathbb{R}^d$ is μ-invariant if $g(L) = L$ for every $g \in \operatorname{supp} \mu$.

A probability η in $\mathbb{PR}^d$ is μ-stationary if $\int g_* \eta \, d\mu(g) = \eta$.

Equivalently, $\Psi(x, g_1, g_2, \ldots, g_n, \ldots) = (g_1(x), g_2, \ldots, g_n, \ldots)$ preserves $\eta \times \mu^{\mathbb{N}}$. We call η ergodic if $\eta \times \mu^{\mathbb{N}}$ is ergodic for Ψ.

Representation of exponents

Given $\mu \in \mathcal{G}(d)$:

$L < \mathbb{R}^d$ is μ-invariant if $g(L) = L$ for every $g \in \operatorname{supp} \mu$.

A probability η in $\mathbb{P}\mathbb{R}^d$ is μ-stationary if $\int g_* \eta \, d\mu(g) = \eta$.

Equivalently, $\Psi(x, g_1, g_2, \ldots, g_n, \ldots) = (g_1(x), g_2, \ldots, g_n, \ldots)$ preserves $\eta \times \mu^{\mathbb{N}}$. We call η ergodic if $\eta \times \mu^{\mathbb{N}}$ is ergodic for Ψ.

Theorem (Furstenberg-Kifer)

There exist $r \geq 0$ and numbers $\beta_0 > \beta_1 > \cdots > \beta_r$ and μ-invariant subspaces $L_0 > L_1 > \cdots > L_r > L_{r+1}$, with $\beta_0 = \lambda_+$ and $L_0 = \mathbb{R}^d$ and $L_{r+1} = 0$, such that for every $v \in L_i \setminus L_{i+1}$ and $0 \leq i \leq r$,

$$\lim_n \frac{1}{n} \log \|(g_n \cdots g_1)(v)\| = \beta_i \quad \mu^{\mathbb{N}}\text{-almost surely.}$$

Representation of exponents

Define $\phi : \mathrm{GL}(d) \times \mathbb{PR}^d \to \mathbb{R}$, $\phi(g,v) = \log\left(\|g(v)\|/\|v\|\right)$. Then:

The β_i are the possible values of $\int \phi \, d\mu d\eta$ when η varies in the set of all μ-stationary ergodic measures.

L_i = largest subspace such that $\eta(L_i) = 0$ for every μ-stationary ergodic measure η with $\int \phi \, d\mu d\eta > \beta_i$.

$\int \phi \, d\mu d\eta > \beta_i$ for every μ-stationary measure η, ergodic or not, such that $\eta(L_i) = 0$.

Examples

Example

For $A_1 = \begin{pmatrix} 3^{-1} & 0 \\ 0 & 3 \end{pmatrix}$ $A_2 = \begin{pmatrix} 2 & 1 \\ 0 & 2^{-1} \end{pmatrix}$ $p_1 = p_2 = 1/2$
we have $r = 1$ and $L_1 = X -$ axis and $\beta_1 = \log 2/3$

Example

For $A_1 = \begin{pmatrix} 3 & 0 \\ 0 & 3^{-1} \end{pmatrix}$ $A_2 = \begin{pmatrix} 2^{-1} & 1 \\ 0 & 2 \end{pmatrix}$ $p_1 = p_2 = 1/2$
we have $r = 0$.

So: $r > 0$ means that there exists some μ-invariant subspace (reducibility) which, in addition, is "mostly contracting".

A partial result

Theorem (Furstenberg-Kifer)

If $r = 0$ then μ is a continuity point for λ_+.

Proof: Given $\mu_n \to \mu$, take μ_n-stationary ergodic measures η_n such that $\lambda_+(\mu_n) = \int \phi \, d\mu_n d\eta_n$.
Suppose that $\eta_n \to \eta$. Then η is μ-stationary and $\int \phi \, d\mu_n d\eta_n$ converges to $\int \phi \, d\mu \, d\eta$.
The hypothesis $r = 0$ implies that $\int \phi \, d\mu d\eta = \lambda_+(\mu)$. □

A partial result

Theorem (Furstenberg-Kifer)

If $r = 0$ then μ is a continuity point for λ_+.

Proof: Given $\mu_n \to \mu$, take μ_n-stationary ergodic measures η_n such that $\lambda_+(\mu_n) = \int \phi \, d\mu_n d\eta_n$.
Suppose that $\eta_n \to \eta$. Then η is μ-stationary and $\int \phi \, d\mu_n d\eta_n$ converges to $\int \phi \, d\mu \, d\eta$.
The hypothesis $r = 0$ implies that $\int \phi \, d\mu d\eta = \lambda_+(\mu)$. □

If there exists at most one μ-invariant subspace then $r = 0$ either for the cocycle or for its inverse, and the conclusion follows just the same.

What about the general case?

Probabilistic repellers

Example

For $A_1 = \begin{pmatrix} 3^{-1} & 0 \\ 0 & 3 \end{pmatrix} \qquad A_2 = \begin{pmatrix} 2 & 0 \\ 0 & 2^{-1} \end{pmatrix} \qquad p_1 = p_2 = 1/2$

there are two μ-stationary ergodic measures in $\mathbb{PR}^2$, namely, the Dirac masses at the X-axis and the Y-axis. They correspond to different β_j.

Probabilistic repellers

Example

For $A_1 = \begin{pmatrix} 3^{-1} & 0 \\ 0 & 3 \end{pmatrix} \qquad A_2 = \begin{pmatrix} 2 & 0 \\ 0 & 2^{-1} \end{pmatrix} \qquad p_1 = p_2 = 1/2$

there are two μ-stationary ergodic measures in $\mathbb{PR}^2$, namely, the Dirac masses at the X-axis and the Y-axis. They correspond to different β_j.

The invariant subspace $L_1 = X$-axis is a probabilistic repeller. The ideology of the proof is that such probabilistic repellers should be *unstable* under most perturbations of the probability distribution μ.

Instability of probabilistic repellers

"Theorem"

Suppose $r > 0$. For every $\epsilon > 0$ there is $\delta > 0$ and a neighborhood $V \subset \mathcal{G}(d)$ of μ such that for every $\nu \in V$ and every ν-stationary ergodic measure η, either $\eta(B_\delta(L_1)) < \epsilon$ or $\eta(B_\delta(L_1)) = 1$.

Idea: in the last case, η is not a candidate for realizing $\lambda_+(\nu)$.

Instability of probabilistic repellers

"Theorem"

Suppose $r > 0$. For every $\epsilon > 0$ there is $\delta > 0$ and a neighborhood $V \subset \mathcal{G}(d)$ of μ such that for every $\nu \in V$ and every ν-stationary ergodic measure η, either $\eta(B_\delta(L_1)) < \epsilon$ or $\eta(B_\delta(L_1)) = 1$.

Idea: in the last case, η is not a candidate for realizing $\lambda_+(\nu)$.

The original Bocker-V approach in $d = 2$ is based on a careful discretization of the phase space $\mathbb{PR}^2$.

With Artur Avila, we have been trying with a more direct analysis of the random walk in continuum space, based on certain *energy estimates*.

A nonlinear setting

Let M be a compact Riemannian manifold (examples: $\mathbb{PR}^d$, Grassmannian manifolds) and $\mathcal{M}$ be the space of probability measures on M.

Let $G < \mathrm{Diff}^1(M)$ (e.g. $G = \mathrm{GL}(d)$) and $\mathcal{G}$ be the space of compactly supported probability measures on G.

A nonlinear setting

Let M be a compact Riemannian manifold (examples: $\mathbb{PR}^d$, Grassmannian manifolds) and $\mathcal{M}$ be the space of probability measures on M.

Let $G < \mathrm{Diff}^1(M)$ (e.g. $G = \mathrm{GL}(d)$) and $\mathcal{G}$ be the space of compactly supported probability measures on G.

A point $v \in M$ is μ-invariant if $g(v) = v$ for every $g \in \operatorname{supp}\mu$.

Then, $\mu^{\mathbb{N}}$-almost surely, $L(\mu, \dot{v}) = \lim_n \frac{1}{n} \log \|D(g_n \cdots g_1)(v)\dot{v}\|$ exists for every non-zero $\dot{v} \in T_v M$.

We call v μ-expanding if $L(\mu, \dot{v}) > 0$ for every $\dot{v} \neq 0$.

Instability of μ-expanding points

Theorem (Artur Avila, MV)

Suppose that v is μ-expanding and $(\mu_n)_n$ converges to μ in $\mathcal{G}$. For each n, let $\eta_n \in \mathcal{M}$ be a μ_n-stationary measure having no atoms in a fixed neighborhood of v, and assume that $(\eta_n)_n$ converges to some $\eta \in \mathcal{M}$. Then $\eta(\{v\}) = 0$.

This proves continuity of λ_+ for all $d \geq 2$ when $\dim L_1 = 1$.

Energy

Given $\beta > 0$, the β-energy of a measure ξ on $M \times M$ is

$$E_\beta(\xi) = \int d(x,y)^{-\beta}\, d\xi(x,y).$$

The map $\xi \mapsto E_\beta(\xi)$ is lower semicontinuous.

Energy

Given $\beta > 0$, the β-energy of a measure ξ on $M \times M$ is

$$E_\beta(\xi) = \int d(x,y)^{-\beta}\, d\xi(x,y).$$

The map $\xi \mapsto E_\beta(\xi)$ is lower semicontinuous.

Let η_1, η_2 be measures on M with $\eta_1(M) = \eta_2(M)$:

A coupling of (η_1, η_2) is a measure ξ on $M \times M$ that maps to η_j on the jth coordinate, for $j = 1, 2$.

Given $\beta > 0$, define $e_\beta(\eta_1, \eta_2) =$ infimum of β-energy $E_\beta(\xi)$ over all couplings ξ. The infimum is attained.

Optimal self-couplings

Given a measure η on M, define $e_\beta(\eta) = e_\beta(\eta, \eta)$. The infimum is attained at some symmetric self-coupling, that is, one invariant under $(u, v) \mapsto (v, u)$. We call this a β-optimal self-coupling.

Optimal self-couplings

Given a measure η on M, define $e_\beta(\eta) = e_\beta(\eta, \eta)$. The infimum is attained at some symmetric self-coupling, that is, one invariant under $(u, v) \mapsto (v, u)$. We call this a β-optimal self-coupling.

The energy $e_\beta(\eta)$ is finite *iff* η has no fat atoms:

$$\eta(\{x\}) < \tfrac{1}{2}\eta(M) \text{ for every } x \quad \Rightarrow \quad e_\beta(\eta) < \infty$$

$$e_\beta(\eta) < \infty \quad \Rightarrow \quad \eta(\{x\}) \leq \tfrac{1}{2}\eta(M) \text{ for every } x$$

Optimal self-couplings

Lemma

If v is μ-expanding then there exists a neighborhood V of v, a weak* neighborhood $\mathcal{V}$ of μ and a constant $c > 0$ such that

$$\int d(g(x), g(y))^{-\beta}\, d\nu(g) < (1 - c\beta) d(x, y)^{-\beta}$$

for every $x \neq y$ in V, every $\nu \in \mathcal{V}$ and every small $\beta > 0$.

Optimal self-couplings

Lemma

If v is μ-expanding then there exists a neighborhood V of v, a weak* neighborhood $\mathcal{V}$ of μ and a constant $c > 0$ such that

$$\int d(g(x), g(y))^{-\beta} \, d\nu(g) < (1 - c\beta) d(x, y)^{-\beta}$$

for every $x \neq y$ in V, every $\nu \in \mathcal{V}$ and every small $\beta > 0$.

Suppose $\eta(\{v\}) > 0$. Fix $U \subset V$ such that $\eta(\{v\}) > 0.9\,\eta(U)$.

Notice: $e_\beta(\eta \mid U) = \infty$.

Energy estimates

For each n, let ξ_n be a β-optimal self coupling of $\eta_n \mid U$ and let $\tilde{\xi}_n$ be its push-forward:

$$\tilde{\xi}_n(A \times B) = \int \xi_n(g^{-1}(A) \times g^{-1}(B))\, d\mu_n(g).$$

Energy estimates

For each n, let ξ_n be a β-optimal self coupling of $\eta_n \mid U$ and let $\tilde{\xi}_n$ be its push-forward:

$$\tilde{\xi}_n(A \times B) = \int \xi_n(g^{-1}(A) \times g^{-1}(B))\, d\mu_n(g).$$

Lemma $\Rightarrow$ $E_\beta(\tilde{\xi}_n) < (1 - c\beta) E_\beta(\xi_n) = (1 - c\beta) e_\beta(\eta_n \mid U)$.

Moreover, $e_\beta(\eta_n \mid U) \leq C + E_\beta(\tilde{\xi}_n)$.

Energy estimates

For each n, let ξ_n be a β-optimal self coupling of $\eta_n \mid U$ and let $\tilde{\xi}_n$ be its push-forward:

$$\tilde{\xi}_n(A \times B) = \int \xi_n(g^{-1}(A) \times g^{-1}(B))\, d\mu_n(g).$$

Lemma $\Rightarrow$ $E_\beta(\tilde{\xi}_n) < (1 - c\beta)E_\beta(\xi_n) = (1 - c\beta)e_\beta(\eta_n \mid U)$.

Moreover, $e_\beta(\eta_n \mid U) \leq C + E_\beta(\tilde{\xi}_n)$.

Combining these inequalities: $e_\beta(\eta_n \mid U) \leq C/(c\beta)$ for all n.

Then $e_\beta(\eta \mid U) \leq C/(c\beta)$. Contradiction.