Mathematics
Research
Reports

Wentao Fan, Boris Hasselblatt

Strictly convex ergodic billiards

Volume 7 (2026), p. 71-79.

https://doi.org/10.5802/mrr.28

Mathematics Research Reports *is member of the*
Centre Mersenne for Open Scientific Publishing
www.centre-mersenne.org
e-ISSN: 2772-9559

𝕄athematics ℝesearch ℝeports
Volume 7 (2026), 71–79

Strictly convex ergodic billiards

Wentao Fan & Boris Hasselblatt

(Recommended by Svetlana Katok)

Abstract. We produce the first known strictly convex ergodic billiards: the intersection of two circles of which one contains the centers of both.

1. Can a strictly convex planar billiard be ergodic?

Mathematical billiard systems [26] describe a point particle moving freely in a domain with piecewise smooth boundary and reflecting off the boundary the way a photon does in a mirror, with equal angles before and after collision (specular reflection) [27]. These are dynamical systems which Birkhoff introduced as a simplified limit of a range of mechanical systems with two degrees of freedom in which "the formal side, usually so formidable in dynamics, almost completely disappears, and only the interesting qualitative questions need to be considered" [15, p. 288f]. For instance, under contraction of the z-coordinate (say), an ellipsoid will eventually be squashed to an elliptic billiard in the xy-plane. In terms of paradigms for dynamical features and methods of study, billiards can be grouped into convex billiards (elliptic), polygonal billiards (parabolic), and dispersing billiards (hyperbolic) [22]. The latter are the subject of the present article.

1.1. Elliptic and parabolic billiards. Rectangular billiards may naturally come to mind first, and pertinent dynamical questions are whether there are dense orbits (yes on the table, no for the space of unit vectors), periodic orbits (clearly), and how many distinct periodic orbits of up to some given period there might be. More generally, one can consider billiards in polygons, and ask similar questions (and about *mixing*). Until 2026 it was not even known whether every triangular billiard has a periodic orbit [16]. This subject belongs to what is now considered *parabolic dynamics*, and it is related to other subjects, such as interval-exchange maps and Teichmüller theory.

The simplest smooth strictly convex billiard table would be a circle, where it is easy to understand all dynamical possibilities explicitly. For instance, some thought or experimentation reveals that every orbit keeps grazing a smaller concentric circle. This quickly organizes the entire orbit structure (this is a *completely integrable* mechanical system; either the radius of the concentric circle or the collisions angle provide a *constant of motion*/invariant function). Elliptic billiards are less simple, but orbits graze confocal quadrics, so these are also integrable, and the full orbit structure is well understood in qualitative or even computable terms. For smooth but not necessarily algebraic strictly convex boundary curves, some *caustics* remain to play a like role; this is related to KAM theory, which was originally motivated by celestial mechanics (and integrable mechanical systems, of which circular and elliptic billiards are examples). This is elliptic dynamics [22].

Received April 26, 2026.
2020 *Mathematics Subject Classification.* 37C83.
Keywords. Hyperbolicity, billiards, ergodicity.

1.2. Hyperbolic billiards and ergodicity. Finally, billiard tables with *hyperbolic* time-evolution are those with dynamics so complicated that one can hope to establish good statistical properties, beginning with ergodicity, a form of indecomposability (of the natural area measure) [25] and the opposite of complete integrability (there is not even a measurable invariant function). Hyperbolicity, however, relies on dispersing boundary pieces, or on restrictions on focusing boundary arcs (Wojtkowski's criteria for hyperbolicity [29, 28]; for focusing boundary circle arcs, these require that the circle to which such an arc belongs must lie inside the billiard table), and these criteria preclude strict convexity. (They produce a "defocusing" mechanism which traces back to Bunimovich [3], Wojtkowski [29], Markarian [24], Donnay [12], and again Bunimovich [4].)

A review article on "Some new surprises in chaos" [6] presented *asymmetric lemon billiards* (Figure 1) as "a new class of chaotic focusing billiards...that clearly violates the main condition considered to be necessary for chaos in focusing billiards." These billiards are strictly convex, but even that article dared not venture a guess about their ergodicity.

Accordingly, the subject of mathematical billiards has implicitly, but in an ever more pressing way, invited the question: *can a strictly convex planar billiard be ergodic?*

2. A strictly convex ergodic billiard

As we just noted, just over a decade ago, a candidate billiard emerged that combines strict convexity with (nonuniform) hyperbolicity, the (sufficiently asymmetric) *lemon billiard* in Figure 1. We here report that it is, in fact, ergodic [14, 13]. Indeed, it has the strongest of all ergodic properties:

2.1. A strictly convex Bernoulli billiard.

Theorem A. *For almost all* $\zeta \in \left(0, \tan^{-1}(1/3)\right)$*, the lemon billiard in Figure 1 is Bernoulli for large enough* R.

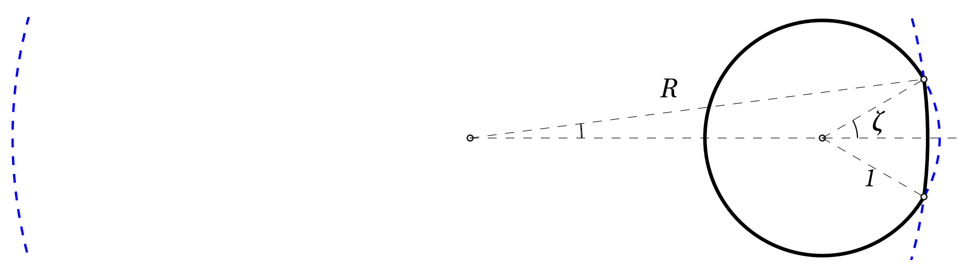

FIGURE 1. Asymmetric lemon billiard

The Bernoulli property means that in a specific probabilistic sense, this billiard system is indistinguishable from a random coin toss. There is a hierarchy of *mixing* properties with ergodicity being the weakest form and the Bernoulli property the strongest. We repeat that this appears to be the first example of a strictly convex ergodic billiard.

The route to ergodicity runs through hyperbolicity, and along the way, our work affirmatively answers a related open question: is there a strictly convex billiard that admits a global section with *uniformly* hyperbolic return map?

It may be of interest that in Theorem A, R needs to exceed 300 000. So if Figure 1 were drawn to scale, then the flatter boundary arc would look straight. When that boundary piece is literally straight, the billiard does, in effect, satisfy the Wojtkowski criteria (by the *unfolding* device shown in Figure 2), but it is not strictly convex. As it happens, this observation does little to help prove ergodicity—the approach is not perturbative, and we hope to make clear why not.

Lemon billiards may have first appeared in [18] in the form of the intersection of two circles of equal radius. Their hyperbolicity was proved in [5]. We here consider *asymmetric lemon billiards*, which consist of the intersection of two disks of which one contains the centers of both and were introduced by Chen–Mohr–Zhang–Zhang, who conjectured that they are ergodic for large R [8, Observation, p. 3].[1] We note the extreme cases of this: $R = 1$ gives the billiard in the unit circle, which we described earlier, and which is neither hyperbolic nor ergodic. "$R = \infty$" gives the aforementioned petal billiard, which has been well understood—it satifies the Wojtkowski criteria.

The proof involves two main stages, each substantial enough for a standalone publication [14, 13].

2.2. Strict convexity. It may be of some interest to see exactly why the usual criteria for hyperbolicity are incompatible with strict convexity.

Proposition 2.1. *A piecewise C^3 planar billiard that satisfies the Wojtkowski criteria for hyperbolicity* [27, p. 144] *is not strictly convex.*

Proof. Consider a strictly convex piecewise C^3 planar billiard. Each boundary arc is focusing (as opposed to dispersing or straight) by strict convexity. One part of the criterion is that at a common end-point of two such arcs the interior angle must exceed π, which is incompatible with convexity. Thus, there is only one arc, i.e., the billiard is C^3. Another part of the criterion is that for any focusing arc (hence in this context for the entire boundary) the radius of curvature as a function of arc length is strictly concave, i.e., has strictly monotone first derivative. This is incompatible with periodicity of the first derivative. □

3. Hyperbolicity

The first major step of the proof of ergodicity is to establish *uniform* hyperbolicity [14].

3.1. The phase space. It is important to clarify what exactly is proved to be uniformly hyperbolic. Our introduction of billiards suggested a flow, and indeed, the billiard flow is the unit-speed (say) geodesic flow on the region in question with respect to the Euclidean metric, plus collisions: in the tangent space at a boundary point, vectors are identified that are related by reflection over the tangent line to the boundary at that point. Instead, we study the billiard *map*, which associates with an inward-pointing (post-collision) vector at a boundary point the vector after the next collision with the boundary.

Thus, the *phase space* (or state space) consists of the product of the boundary (a topological circle) with the interval $[0,\pi]$ of possible angles of inward-pointing vectors with the (positively oriented) tangent vector to the boundary. In fact, wanting a smooth structure, our phase space is the union of two rectangles whose bases are the two circle arcs.

3.2. The section. It is, however, not this billiard map we prove to be uniformly hyperbolic—and indeed, the billiard map is not uniformly hyperbolic. Diameters of the unit arc (the boundary piece which is part of a unit circle) are nonhyperbolic 2-periodic points. We instead choose a subset $\widehat{M}$ for which we prove that the first-return map is uniformly hyperbolic. This first-return map is what it says it is: for a point (i.e., vector) in the section,

[1] Though an "observation" in the publication, it was stated as a conjecture at submission.

apply the billiard map, repeatedly if needed, until an image point is again in the section (for the first time); this is the image under the first-return map.

Here is an example of a section, which is close to the section we actually use: vectors at points of the unit arc pointed at the other boundary arc; these are the exit points from the unit arc. The return map is then determined by tracking however many successive collisions next occur on the flatter arc and then all successive collisions with the unit arc; the last of these is the image under the return map because it again points at the flatter arc. This is a *global* section: All that is missing are periodic points that only collide with the unit arc; these have the shape of regular polygons, and each comes in a 1-parameter family, but the periods (number of sides) are bounded in terms of ζ in Figure 2.

For the section we consider, we establish that (Lebesgue-)almost every point of phase space visits it in past and future (it is a global section), so images of $\widehat{M}$ under the billiard map cover almost all of the phase space (but not, for instance, the aforementioned non-hyperbolic periodic points). Here, the first-return map associates with any point of $\widehat{M}$ the first of its successive images under the billiard map that lies in $\widehat{M}$. "Almost all" is sufficient for our eventual purpose: in the sense of ergodic theory, $\widehat{M}$ is a global section, and ergodicity of the return map implies ergodicity of the natural Liouville volume in the phase space for the full billiard map, and also for the billiard flow.

3.3. Hyperbolicity. By design, lemon billiards do not satisfy the Wojtkowski criteria, so hyperbolicity must be established by arguments closely tied to the specifics of the billiard table. Consequently, our work adapts techniques from the theory of dispersing billiards without being able to invoke ready-made theorems (other than the abundant general toolkit provided by [9]). This work also rests heavily on a pioneering article which

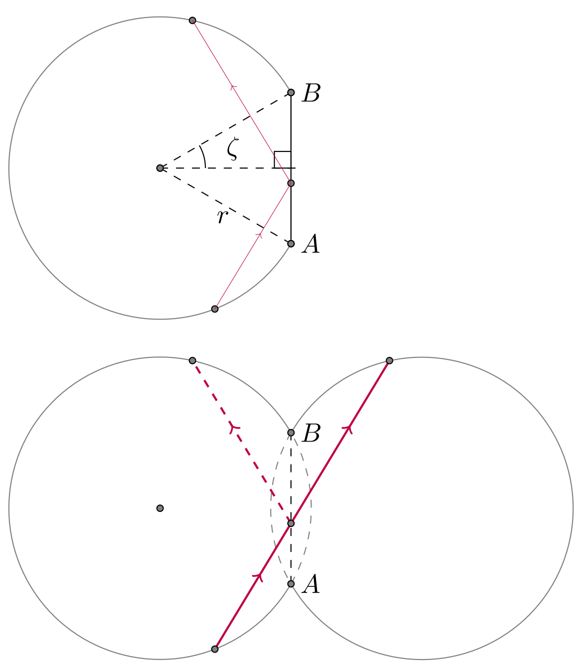

FIGURE 2. The 1-petal billiard and its unfolding to the 2-petal billiard

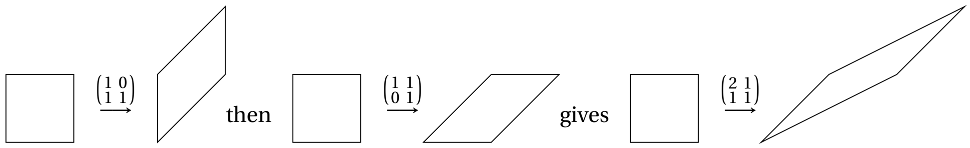

FIGURE 3. Linked twists

chose essentially the section we are using and established that the corresponding return map has an invariant cone family, which implies (nonuniform!) hyperbolicity [21].

This can be expressed in terms of invariant cone families but is defined as the property that the natural measure has nonzero Lyapunov exponents, i.e., pointwise eventual exponential expansion. What remains to do, however, is daunting: to establish *uniform* expansion of vectors in cones of a suitable family under application of (the differential of) the return map to a suitable section. This is *uniform* hyperbolicity: uniform exponential lower bounds for expansion a.e. We note that the tractability of the 1-petal billiard is being used in handling *some* of our arguments but we repeat that this is not (and likely cannot be) the overall strategy.

The definition of hyperbolicity of a map f is that there are two line bundles which are, respectively, contracted and expanded by the differential Df, the stable and unstable subbundles. A foundational fact is that these bundles are then tangent to stable and unstable foliations. An iconic example is the action of $\left(\begin{smallmatrix}2&1\\1&1\end{smallmatrix}\right)$ on the torus $\mathbb{R}^2/\mathbb{Z}^2$, where the subbundles are translates of the eigenspaces, and the foliations are their projection to $\mathbb{R}^2/\mathbb{Z}^2$. The combination of stretching and folding is a fundamental mechanism for complicated long-term behavior which is both deterministic and practically unpredictable.

It is rarely practical to verify the defining condition in examples because that would seem to necessitate finding those subbundles first. Instead, one uses the (necessary and sufficient) Alekseev cone criterion: in the preceding example, the cone defined in each tangent space by the first and third quadrant is (strictly) preserved by the map and under iteration asymptotically shrinks down to the unstable subspace.

As it happens, for planar billiards, the first (and third) quadrant universally provides the natural candidate cone family for establishing hyperbolicity (though for the lemon billiard, its bottom half is itself invariant).

3.4. A mechanism. It may be illuminating to return to the circular billiard, which for a natural coordinatization of the phase space is explicitly given by the twist map $(x, y) \mapsto (x + y, y)$ on $S^1 \times [0,1]$, which rigidly rotates each set $\{y = \text{const.}\}$. This neatly illustrates area-preservation by Cavalieri's Principle (or by computing the determinant) as well as invariance of the first quadrant under the differential. However, that invariance is not strict in our sense—even though it involves proper inclusion, the closed quadrant is not mapped into its interior, which is the requirement for the cone criterion. This is, of course, as it should be, given the predictability of circle billiards.

In the lemon billiard, two circle billiards interact by virtue of the transitions of orbits between the two boundary pieces. Linking two twist maps with orthogonal twist directions can make the first quadrant strictly invariant, as illustrated in Figure 3: each of the two twists moves one cone edge to the interior.

While this makes hyperbolicity believable, it is daunting to establish. Creativity and hard estimates were required to show invariance of cones in [21], and the "upgrade" to uniform expansion is not a small refinement but far more daunting. Indeed, there are parts of the section where there is no expansion, and indeed contraction by as much as a factor of 20 is possible. Accordingly, the "combinatorics" of orbits needs to be understood well enough to establish that any stretch of possible contraction is followed by so much expansion as to compensate for the prior loss.

3.5. Hopf chains. The mechanism that leads from hyperbolicity to ergodicity goes back to Hopf [20] and uses the stable and unstable curves as follows. Ergodicity means that there is no (measurable) invariant function. For the circle billiard, the radius of the concentric circle to which an orbit keeps being tangent is a (smooth) function on the phase space which is constant along orbits (in physics, this is also called a constant of motion). Ergodicity means that there is not even a measurable invariant function, or, rather, that a measurable invariant function is constant.

To see how hyperbolicity implies this, one can run the Hopf argument for *continuous* functions, simplistically. If a continuous function is invariant, then it must be constant on stable curves: one can compare the function values at two points by mapping these forward repeatedly, which by invariance does not change the function values. Since the points are mapped ever closer to each other, (uniform) continuity implies that the function values agree. The same argument running backwards in time shows that the function is constant on unstable curves as well. But then for any two points one can travel first along the stable curve of one of them, then on the unstable of the other to see that the function has the same value at the intermediate point as at either of the two others. (This idea is correct, but there can be much subtlety in handling null sets [10, 7].)

There are also geometric issues, notably with the last step here: it assumes that these curves are transverse to each other and long enough that one can indeed connect two points in this way. This is the reason *uniform* hyperbolicity is the first step.

4. Singularities

The discussion so far elided the fact that the billiard map is not actually well defined: upon hitting a corner, it is not clear how a particle should rebound. Indeed, there are clearly distinct one-sided limits for the rebound from either adjacent boundary arc. These singularities are a deep concern in the study of billiards and in the second major step of the proof of ergodicity [13]. It is worth emphasizing that the choice of section provides a novel way of handling configurations of singularity curves that made them seem intractable before.

4.1. Density. We first describe them a little more. For the lemon billiard, each of the two corners defines a singularity curve in the phase space which projects to the boundary. Specifically, at each boundary point, there is a unique vector pointing to that corner. The two curves thus defined are quite regular, but for the full picture, these need to be pulled back and pushed forward by the billiard map to identify the loci of corner impacts in the more distant future or past. This results in smooth branched curves that eventually fill the phase space densely.

An understanding of this overall picture is needed because singularity curves interfere with the use of stable and unstable curves in the Hopf argument, effectively cutting them—the Hopf-chain argument does not work if one has to cross a singularity to get from one point to another along an unstable curve, say.

Put differently, uniform hyperbolicity "grows" an unstable leaf with each iteration, but that is potentially undone when the leaf hits a corner. Thus, we look for *sufficient* points where there is enough expansion before any close encounter with a singularity.

4.2. Alignment. An alternative description is in [9]: adding ever more distant images of singularity curves to the picture subdivides the phase space ever more finely; indeed, their union is dense. However, if these higher-order singularity curves align well enough, then the addition of ever more of them in effect whittles components of the rest of phase space down to arcs, and these then are (long) unstable leaves.

Accordingly, a core task in our second article [13] is to establish alignment of singularity curves of lemon billiards.

It is somewhat natural that singularity curves would align. Indeed, Figure 4, left, shows these for an infinite-horizon Sinai billiard, where they are far more densely placed than in the present context, but yet align almost as well as a foliation. (This picture was created by Samuel Gomez at Tufts University.) Notably, having positive slope makes

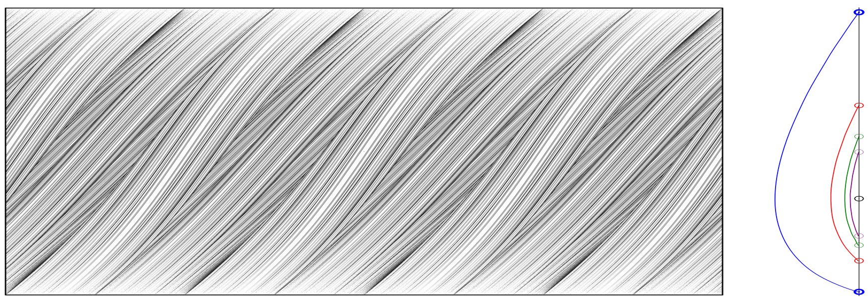

FIGURE 4. Singularity curves, good and bad

sense because it is the first quadrant that is invariant and hence contains the unstable subspace, so unstable curves should also have positive slope. It is this alignment of unstable curves and singularity curves that is needed.

And here we come upon a further reason others shied away from lemon billiards: singularity curves for lemon billiards look quite different in places than they do for other billiards. Figure 4, right, shows, for instance, that their slope is not even always positive. This would seem to stop the entire strategy cold.

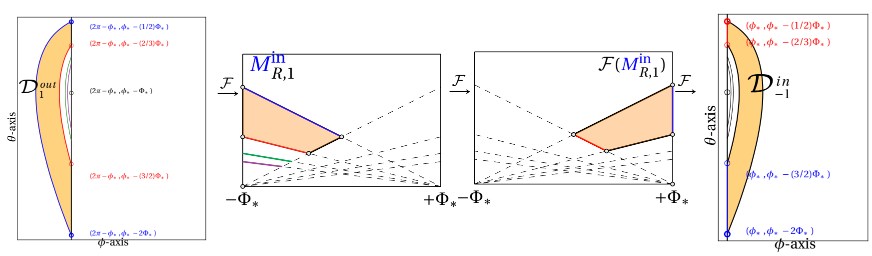

FIGURE 5. Mapping "arch" singularity curves

The remedy is a matter of choosing the right timing. Figure 5 shows that the image or preimage of such worrisome singularity curves has well-controlled slope. Accordingly, we replace the relevant piece of our section by its preimage.

That is to say that this issue is sidestepped by choosing as the global section not the last collisions with the flatter side as previously suggested, but *the set of next to last or sole* collisions with that boundary arc.

5. Ergodicity

The combination of uniform hyperbolicity with regularity and alignment of singularity curves yields ergodicity. On one hand, this step can be implemented following the techniques presented in [9], which remains definitive. We found, however, that the arguments in [11] require the least adaptation, and this adaptation yields *local ergodicity*: the Hopf-chain argument works near any given point to show that an invariant function must be locally constant (and, in fact, a local Bernoulli property). The remaining step is nontrivial but far easier than anything that has come before. It shows that the phase space is not separated into distinct ergodic pieces, i.e., that the section map is ergodic (and then indeed Bernoulli).

It is well to note that these properties translate to the billiard map itself; they are not lost in the passing through the intermediate iterates skipped by the return map. However, uniform hyperbolicity does *not* hold for the billiard map itself—the expansion per time is not uniformly large because there is no bound on how many collisions a trajectory has before returning to the section. This is also the reason it seems impossible to effectively treat this as a perturbation of a petal billiard.

6. Questions

Two questions naturally motivate further work.

Does the section map considered here have exponential decay of correlations? The Bernoulli property implies *mixing*, i.e., that viewing two scalar functions on the phase space as random variables and pulling one of them back with the billiard map makes it asymptotically independent of the other function in the probabilistic sense. This means that correlations tend to zero in the process, and one can ask whether they do so exponentially rapidly. Because of the long return times, this cannot be expected for the billiard map itself but is a reasonable conjecture for the return map.

Can a strictly convex ergodic billiard be C^1? The heuristics above suggested that it is the corner angle that links twist maps into a hyperbolic aggregate, but it is not actually clear that corners are strictly needed—even though the challenges in the absence of corners might be far more daunting yet.

To make the question concrete, one could ask whether Benettin–Strelcyn billiards [2, 1, 19] (consisting of a rectangle whose sides are topped by circle arcs with matching tangents at the corners) admit a global section with uniformly hyperbolic return map. These are historically important, having given rise to [23]; see [17].

References

[1] G. Benettin & J. M. Strelcyn, "Numerical experiments on the free motion of a point mass moving in a plane convex region: Stochastic transition and entropy", *Phys. Rev. A* **17** (1978), p. 773-785.

[2] G. Benettin, L. Galgani & J.-M. Strelcyn, "Kolmogorov entropy and numerical experiments", *Phys. Rev. A* **14** (1976), p. 2338-2345.

[3] L. A. Bunimovich, "The ergodic properties of certain billiards", *Funkcional. Anal. i Priložen.* **8** (1974), no. 3, p. 73-74.

[4] ———, "On absolutely focusing mirrors", in *Ergodic theory and related topics, III (Güstrow, 1990)*, Lecture Notes in Math., vol. 1514, Springer, Berlin, 1992, p. 62-82.

[5] L. Bunimovich, H.-K. Zhang & P. Zhang, "On another edge of defocusing: hyperbolicity of asymmetric lemon billiards", *Comm. Math. Phys.* **341** (2016), no. 3, p. 781-803.

[6] L. A. Bunimovich & L. V. Vela-Arevalo, "Some new surprises in chaos", *Chaos* **25** (2015), no. 9, p. 097614, 11.

[7] K. Burns & A. Wilkinson, "On the ergodicity of partially hyperbolic systems", *Ann. of Math. (2)* **171** (2010), no. 1, p. 451-489.

[8] J. Chen, L. Mohr, H.-K. Zhang & P. Zhang, "Ergodicity of the generalized lemon billiards", *Chaos* **23** (2013), no. 4, p. 043137, 12.

[9] N. Chernov & R. Markarian, *Chaotic billiards*, Mathematical Surveys and Monographs, vol. 127, American Mathematical Society, Providence, RI, 2006, xii+316 pages.

[10] Y. Coudène, *Théorie ergodique et systèmes dynamiques*, Savoirs Actuels (Les Ulis). [Current Scholarship (Les Ulis)], EDP Sciences, Les Ulis; CNRS Éditions, Paris, 2012, vi+198 pages.

[11] G. Del Magno & R. Markarian, "A local ergodic theorem for non-uniformly hyperbolic symplectic maps with singularities", *Ergodic Theory Dynam. Systems* **33** (2013), no. 4, p. 983-1007.

[12] V. J. Donnay, "Using integrability to produce chaos: billiards with positive entropy", *Comm. Math. Phys.* **141** (1991), no. 2, p. 225-257.

[13] W. Fan & B. Hasselblatt, "Strictly convex ergodic billiards", *preprint* (2026).

[14] ———, "Uniform Hyperbolicity of strictly convex billiards", *preprint* (2026).

[15] T. Fisher & B. Hasselblatt, *Hyperbolic flows*, Zurich Lectures in Advanced Mathematics, EMS Publishing House, Berlin, [2019] ©2019, xiv+723 pages.

[16] G. Forni, "Existence of a Periodic Orbit for Billiards in Polygons", 2026, `https://arxiv.org/abs/2606.10102`.

[17] B. Hasselblatt, "Anatole Katok", *Ergodic Theory Dynam. Systems* **42** (2022), no. 2, p. 321-388.

[18] E. J. Heller & S. Tomsovic, "Postmodern Quantum Mechanics", *Physics Today* **46** (1993), no. 7, p. 38-46, `https://arxiv.org/abs/https://pubs.aip.org/physicstoday/article-pdf/46/7/38/11186551/38_1_online.pdf`.

[19] M. Hénon & J. Wisdom, "The Benettin-Strelcyn oval billiard revisited", *Phys. D* **8** (1983), no. 1-2, p. 157-169.

[20] E. Hopf, "Statistik der geodätischen Linien in Mannigfaltigkeiten negativer Krümmung", *Ber. Verh. Sächs. Akad. Wiss. Leipzig Math.-Phys. Kl.* **91** (1939), p. 261-304.

[21] X. Jin & P. Zhang, "Hyperbolicity of asymmetric lemon billiards", *Nonlinearity* **34** (2021), no. 1, p. 92-117.

[22] A. B. Katok, "Billiard table as a playground for a mathematician", in *The collected works of Anatole Katok. Vol. 1*, World Sci. Publishing, Singapore, [2024] ©2024, Reprint of [2166929], p. 1067-1088.

[23] A. Katok, J.-M. Strelcyn, F. Ledrappier & F. Przytycki, *Invariant manifolds, entropy and billiards; smooth maps with singularities*, Lecture Notes in Mathematics, vol. 1222, Springer-Verlag, Berlin, 1986, viii+283 pages.

[24] R. Markarian, "Billiards with Pesin region of measure one", *Comm. Math. Phys.* **118** (1988), no. 1, p. 87-97.

[25] C. E. Silva, "What is … an ergodic transformation?", *Notices Amer. Math. Soc.* **63** (2016), no. 1, p. 26-27.

[26] Y. Sinai, "What is … a billiard?", *Notices Amer. Math. Soc.* **51** (2004), no. 4, p. 412-413.

[27] S. Tabachnikov, *Geometry and billiards*, Student Mathematical Library, vol. 30, American Mathematical Society, Providence, RI; Mathematics Advanced Study Semesters, University Park, PA, 2005, xii+176 pages.

[28] M. Wojtkowski, "Invariant families of cones and Lyapunov exponents", *Ergodic Theory Dynam. Systems* **5** (1985), no. 1, p. 145-161.

[29] ———, "Principles for the design of billiards with nonvanishing Lyapunov exponents", *Comm. Math. Phys.* **105** (1986), no. 3, p. 391-414.

Wentao Fan: Department of Mathematics, Tufts University, Medford, MA 02155
E-mail: `wentao.fan@tufts.edu`

Boris Hasselblatt: Department of Mathematics, Tufts University, Medford, MA 02155
E-mail: `boris.hasselblatt@tufts.edu`