



Equilibration of energy in slow–fast systems

Kushal Shah^a, Dmitry Turaev^{b,c}, Vassili Gelfreich^d, and Vered Rom-Kedar^{e,1}

^aDepartment of Electrical Engineering and Computer Science, Indian Institute of Science Education and Research, Bhopal 462066, India; ^bDepartment of Mathematics, Imperial College, London SW7 2AZ, United Kingdom; ^cLobachevsky University of Nizhny, Novgorod 603950, Russia; ^dMathematics Institute, University of Warwick, Coventry CV4 7AL, United Kingdom; and ^eDepartment of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot 76100, Israel

Edited by Jean-Pierre Eckmann, University of Geneva, Geneva, Switzerland, and accepted by Editorial Board Member Herbert Levine October 29, 2017 (received for review April 17, 2017)

Ergodicity is a fundamental requirement for a dynamical system to reach a state of statistical equilibrium. However, in systems with several characteristic timescales, the ergodicity of the fast subsystem impedes the equilibration of the whole system because of the presence of an adiabatic invariant. In this paper, we show that violation of ergodicity in the fast dynamics can drive the whole system to equilibrium. To show this principle, we investigate the dynamics of springy billiards, which are mechanical systems composed of a small particle bouncing elastically in a bounded domain, where one of the boundary walls has finite mass and is attached to a linear spring. Numerical simulations show that the springy billiard systems approach equilibrium at an exponential rate. However, in the limit of vanishing particle-to-wall mass ratio, the equilibration rates remain strictly positive only when the fast particle dynamics reveal two or more ergodic components for a range of wall positions. For this case, we show that the slow dynamics of the moving wall can be modeled by a random process. Numerical simulations of the corresponding springy billiards and their random models show equilibration with similar positive rates.

Hamiltonian systems | chaos | Fermi acceleration | dynamical billiards | mixed phase space

In classical statistical mechanics, one deals with systems with a large number of degrees of freedom, where the exact knowledge of the state of the system at every moment in time is impossible to obtain or is irrelevant. The state variables are, therefore, declared “microscopic,” and an attempt is made to describe the properties of “macroscopic” variables (certain functions of the microscopic state) in an ensemble of many systems similar to the given one. In other words, statistical mechanics examines averaging over the phase space of a dynamical system, typically a Hamiltonian one. There is no a priori way of choosing the probability distribution over which the averaging is performed: given a dynamical system, its evolution is different for different initial conditions, and the distribution of the initial conditions is not encoded in the system and can be arbitrary. However, as the simplest option, one can use the so-called microcanonical ensemble [i.e., assume that the initial conditions are uniformly distributed in the phase space at a fixed value of the system’s energy (1–4)]. After this choice of the ensemble is made, standard results of statistical mechanics are recovered. In particular, one obtains the equipartition theorem, which describes the distribution of energy among different degrees of freedom (5). For multiparticle systems, this theorem implies that the ensemble averages of the kinetic energy per degree of freedom are equal for all particles, and this is used to derive fundamental thermodynamic properties of the system.

The overwhelming success of statistical mechanics indicates that there must be some deeper rationale beyond the assumption of uniform distribution of initial conditions. Since many physical experiments measure time averages, a classical approach is to invoke the ergodic hypothesis, which implies that the time average of an observable coincides with its phase space average over the microcanonical ensemble at the corresponding energy

level. Thus, the derivation of thermodynamic laws for deterministic multiparticle dynamics relies on establishing ergodicity of the corresponding Hamiltonian system.

The most critical problem is that Hamiltonian systems are usually not ergodic in a range of energy levels, even if the number of degrees of freedom is large. For example, consider the paradigmatic model of the Boltzmann gas of hard spheres. It is, most probably, ergodic (6–8), but replacing the instantaneous collisions of the spheres with mutual repulsion is expected to destroy the ergodicity, even for an arbitrarily steep repulsive potential (9–11). In general, the dynamics in a smooth potential consist of a “chaotic sea” and “stability islands” (12, 13). A stability island contains a set of quasiperiodic motions confined only to a portion of the energy level because of which the ergodicity is broken. If the islands occupy a noticeable portion of the phase space, use of the microcanonical ensemble for averaging is unfounded. In general, this issue remains unresolved.

In this paper, we propose a mechanism for the onset of apparent ergodicity and mixing in slow–fast Hamiltonian systems. It is not based on either the usual assumption of a large number of degrees of freedom or the inherent instability of dispersing geometries (such as the hard spheres models).

We consider an isolated system, in which certain degrees of freedom evolve slower than the rest. One may think of the slow variables as the system’s parameters that evolve in reaction to the dynamics of the fast variables. The system obtained by freezing the values of these parameters is called the fast subsystem. If it is ergodic for every value of the frozen parameters, the evolution of the slow variables over a long time interval is accurately described by the system averaged over the microcanonical

Significance

Do partial energies in slow–fast Hamiltonian systems equilibrate? This is a long-standing problem related to the foundation of statistical mechanics. Altering the traditional ergodic assumption, we propose that nonergodicity in the fast subsystem leads to equilibration of the whole system. To show this principle, we introduce a set of mechanical toy models—the springy billiards—and describe stochastic processes corresponding to their adiabatic behavior. We expect that these models and this principle will play an important role in the quest to establish and study the underlying postulates of statistical mechanics, one of the long-standing scientific grails.

Author contributions: K.S., D.T., V.G., and V.R.-K. designed research; K.S., D.T., V.G., and V.R.-K. performed research; K.S., D.T., V.G., and V.R.-K. contributed new reagents/analytic tools; K.S., D.T., V.G., and V.R.-K. analyzed data; K.S., D.T., V.G., and V.R.-K. conceived work ideas by discussions; K.S., V.G., and V.R.-K. performed simulations; and K.S., D.T., V.G., and V.R.-K. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. J.-P.E. is a guest editor invited by the Editorial Board.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence should be addressed. Email: vered.rom-kedar@weizmann.ac.il.

ensemble in the fast-phase space [a rigorous formulation of this fact is given by the Anosov–Kasuga averaging theorem (14–16)]. This averaged system has an additional conserved quantity, the Gibbs volume entropy of the fast subsystem: the fast subsystem responds adiabatically to the slow change in the parameters, and adiabatic processes are known to keep the entropy constant (5). Therefore, the averaged system is not ergodic.

We see that, on the timescale of the adiabatic approximation, the ergodicity of the fast subsystem prevents the whole slow-fast system from behaving ergodically. A similar phenomenon is observed in the “notorious piston problem,” where the adiabatic compression law implies that the system of two ideal gases at different temperatures contained in a finite cylinder and separated by a heavy piston never comes to equilibrium, which seems to defy the second law of thermodynamics (17–20).

Note that the fast subsystem’s entropy is preserved only approximately, and therefore, although the averaged system is not ergodic, the whole original system may still be. Deriving equations for the energy transfer between degrees of freedom in such systems is a challenging task, with results being sensitive to the structure of the interaction terms and the choice of scaling limit (21–23).

In this paper, we consider a more natural situation, in which the fast subsystem is not always ergodic, and the partition of the fast subspace into ergodic components depends on the slow variables. For example, the fast subsystem can be ergodic for some values of the slow variables and have more than one ergodic component for other values. Then, the adiabatic invariants obtained by averaging over the fast dynamics are destroyed. We show in this setting that ensemble averages of slow observables converge exponentially to the vicinity of the values obtained by averaging over the microcanonical ensemble in the whole-phase space. This suggests, somewhat unexpectedly, that the nonergodic behavior of the fast degrees of freedom can lead to effective equilibration of the whole system!

To elucidate this principle, we consider a special class of slow-fast Hamiltonian systems, springy billiards. These models often allow for a detailed explicit analysis, and their direct numerical simulations do not require integration of stiff differential equations. Recall that a billiard is a dynamical system repre-

senting a point particle, which moves inertially inside a domain and is reflected elastically from the domain boundary. Particle dynamics depend on the billiard domain shape. Thus, the motion may be integrable, or it may be chaotic: weakly chaotic with a zero Lyapunov exponent, chaotic and ergodic with a positive Lyapunov exponent, or chaotic with stability islands (24). The barred rectangle (25), stadium (26), and mushroom (27) billiard domains have been extensively studied as paradigmatic models for the different types of chaotic behavior. We use these models to create slow-fast systems with the controlled properties of the fast chaotic dynamics (Fig. 1). We allow one of the billiard boundaries, hereafter called the bar, to move. We assume that both the particle and the bar are of finite mass and that the bar is attached to a linear spring (refs. 28–31 have related setups). The bar mass M is taken to be much larger than the particle mass m , thus ensuring that the bar typically moves much slower than the particle. The billiard obtained by retaining the bar in a fixed position is called the frozen billiard.

The case of an infinitely heavy bar has been studied under the name of Fermi acceleration (32–43). In this case, the motion of the bar is not affected by the particle, while the particle gains or loses speed at every collision with the moving bar. Two types of behavior have been found. If the frozen billiard is ergodic for all possible positions of the bar (ergodic accelerators; e.g., stadium), then the ensemble-averaged kinetic energy may grow at most quadratically with time (34, 40, 44). If the ergodicity of the frozen billiard is broken for a part of the period of the bar’s motion (a multicomponent accelerator; e.g., barred rectangle and mushroom), then the particle kinetic energy grows exponentially with time for almost every initial condition (40–42, 44, 45). Thus, the nonergodicity of the frozen billiard accelerates the energy transfer from the bar to the particle.

Here, we consider a different case by taking the mass of the bar as finite while keeping the mass ratio m/M small. Then, the question of Fermi acceleration is replaced by the question of equilibration—does the energy transfer between the bar and the particle lead the system to statistical equilibrium? Do ensemble averages converge to averages over the microcanonical ensemble? If so, then by the equipartition theorem, the averaged kinetic energies of the bar and particle should eventually become equal.




Fig. 1. Three springy billiards. A particle with a small mass m bounces elastically within a bounded domain, where one of the walls is an oscillating bar, which moves vertically, has a mass $M = 1$, and is connected to a spring with spring constant k . (A) SBR. In each round, the particle bounces many times either below or above the oscillating bar, exchanging momentum adiabatically with the bar. The frozen billiard corresponds to the vertical motion of the particle. The partition to ergodic components varies as the horizontal position changes. In the left part of the box (that does not contain the bar), the frozen billiard has a single ergodic component. In the right part of the box, the particle can move either above or below the bar, resulting in two ergodic components. Numerical values in all SBR simulations are $\lambda = 1$, $L = h = 2$, $k = 81$, the particle horizontal speed $|u_x| = 18/\sqrt{5}$, and the total energy of the system $E = 1$; therefore, the bar’s maximal amplitude is $y_b^{\max} = \sqrt{2}/9$. (B) Springy stadium. The frozen billiard is ergodic and chaotic for all bar positions (an EFS system). Numerical values in all simulations are $r = 1$, $\tan \theta = 0.17$, and $l = l_0 - y_b$, where $l_0 = 2$ and y_b denotes the vertical position of the bar. The total energy is $E = 1$, and the spring constant is $k = 1$; therefore, $y_b^{\max} = \sqrt{2}$. The slanted walls prevent the appearance of families of horizontal periodic orbits. (C) Springy mushroom. A collar with a throat of radius w is added to the springy stadium. The collar motion is rigidly determined by the bar motion, $w = w(y_b)$ (Eq. 11). For a range of y_b values, we keep $w(y_b) = r$, so that the frozen billiard becomes the stadium, and therefore, ergodic. For another range of y_b values, $w(y_b) < r$, and the fast subsystem is not ergodic—it has one chaotic component consisting of orbits, which visit the stem, and an integrable component consisting of orbits that are trapped in the cap (as in the classical Bunimovich mushroom billiard). Changing the collar’s opening w ensures that particles can enter and leave the integrable component when the bar moves. This is, therefore, a VFS system.

Results from studies of Fermi acceleration suggest that the equilibration depends on the geometry of the frozen billiard (i.e., on the dynamics of the fast subsystem). In general, we divide slow-fast Hamiltonian systems into two groups. The first consists of systems with an ergodic fast subsystem (EFS) for almost all values of the slow variables (EFS systems). An example is given by the springy stadium (Fig. 1*B*). The second is comprised of systems for which the structure of the partition of the fast subspace into ergodic components bifurcates as the slow variables change [examples are the springy barred rectangle (SBR) (Fig. 1*A*) and springy mushroom (Fig. 1*C*)]. Such systems, with variable partition of the fast subspace (VFS), are hereafter called VFS systems.

All of our numerical simulations for nonzero values of m/M show exponential equilibration, suggesting apparent ergodicity and even mixing in both EFS and VFS springy billiards. However, we find that there is a fundamental difference between these two classes: the relaxation time in EFS systems tends to infinity as $m/M \rightarrow 0$, whereas the relaxation time stays bounded in VFS systems.

To analyze this phenomenon, we derive a stochastic model for the slow bar motion. In this model, the force exerted on the bar by the particle is found by averaging over one of the fast subsystem's ergodic components, while the probability of switching between the components is determined by the bar position and velocity. The resulting Markov process of hopping between the ergodic components of the fast subsystem leads to equilibration. Numerics shows that the bar motion is quite accurately represented by this model, and the equilibration rates remain nonzero and approach the Markov process rates as $m/M \rightarrow 0$.

The derivation of the proposed stochastic models is not specific to billiards. Similar Markov processes can be constructed for a general VFS system by a procedure analogous to that used to study the Fermi acceleration in a homogeneous potential (46). Under standard irreducibility and aperiodicity conditions, such processes should converge to a unique stationary measure (cf. refs. 22, 23, 47, and 48), and by uniqueness, it must correspond to the Liouville measure of the whole system. Therefore, we believe that the proposed apparent ergodization mechanism should be universally applicable.

Particle in a Springy Billiard

Consider a particle of mass $m \ll 1$ in a D -dimensional billiard. One of the billiard walls is a bar of mass $M = 1$, which oscillates vertically by virtue of being suspended from a spring of spring constant k . At impact, the bar and the particle undergo an elastic collision, leading to the exchange of momentum and energy between them:

$$v_p^+ = \frac{2v_b^- - (1-m)v_p^-}{1+m}, \quad v_b^+ = \frac{(1-m)v_b^- + 2mv_p^-}{1+m}, \quad [1]$$

where v_b^- , v_p^- and v_b^+ , v_p^+ are the vertical velocities of the bar (b) and the particle (p) just before and just after the collision, respectively. The total energy of the system, $E = E_p + E_b$, is preserved, whereas the particle kinetic energy E_p and the bar energy $E_b = v_b^2 + ky_b^2/2$ (y_b is the bar position) change at impact. Note that, between collisions, the particle moves inertially, while the bar oscillates harmonically. Therefore, the numerical simulations do not require integration to determine the dynamics of the system between the impacts.

The "bar-particle" system is Hamiltonian and has $D + 1$ degrees of freedom. It is slow-fast, since the bar speed $|v_b| \leq \sqrt{2E}$ is of order 1, while the particle speed $|v_p| = \sqrt{2E_p/m}$ is typically large. As the particle moves fast, many collisions with

the bar occur during each short time interval; the averaged motion of the bar is governed by the equation

$$\ddot{y}_b + \frac{dU(y_b)}{dy_b} = F, \quad [2]$$

where $U(y_b) = 1/2ky_b^2$ is the spring potential and F denotes the averaged force exerted on the bar at position y_b by the particle with energy E_p . The averaging is performed over many collisions at a frozen value of y_b . Since the work done by this force corresponds to the change in particle energy, it follows that $F = -dE_p/dy_b$.

If the frozen billiard is ergodic for each value of y_b , the Anosov-Kasuga theorem (20) implies that the phase-space volume at a given energy level is approximately preserved (for most trajectories, for sufficiently small m , and any finite interval of the slow time). Hence,

$$J = E_p^{D/2} V(y_b) \approx \text{const}, \quad [3]$$

where $V(y_b)$ is the volume of the frozen billiard domain. Thus, in the adiabatic approximation,

$$\frac{dE_p}{E_p} = -\frac{2}{D} \frac{dV}{V}. \quad [4]$$

This implies that the force acting on the bar is $F = 2/D E_p/V dV/dy_b$. Since $E_p = E - E_b = E - 1/2\dot{y}_b^2 - U(y_b)$, Eq. 2 becomes

$$\ddot{y}_b + \frac{dU(y_b)}{dy_b} = \frac{2}{D} [E - \frac{1}{2}\dot{y}_b^2 - U(y_b)] \frac{dV(y_b)}{V(y_b)dy_b}. \quad [5]$$

Note that the same formula follows from the ideal gas law.

A straightforward computation shows that solutions of Eq. 5 indeed follow the level sets of J . By noting that the adiabatic law given by Eq. 3 takes the form $(J/V(y_b))^{2/D} = E_p$, we find that the adiabatic bar motion is governed by an effective potential

$$U_{\text{eff}}(y_b) = U(y_b) + \left(\frac{J}{V(y_b)}\right)^{\frac{2}{D}}, \quad [6]$$

where J is determined by the initial condition. A similar effective potential was derived in the context of the piston problem (19, 20). It is easy to see that $0 \leq J \leq J_f = (E - U(y_f))^{2/D} V(y_f)$. Here, $y_b = y_f$ corresponds to the pressure equilibrium, where the spring compensates for the pressure on the bar caused by the collisions with the particle (y_f depends on the total energy E). At the other extreme (i.e., $J = 0$), the particle does not move, and all of the energy is in the oscillating bar.

We see that, when the frozen billiard is ergodic for every position of the bar (an EFS system), the adiabatic approximation predicts that the bar motion will be periodic in time. Consequently, the springy billiard does not equilibrate on the timescale for which the adiabatic invariant is preserved. An example of such a system is a springy stadium (Fig. 1*B*), which is obtained by a modification of the famous Bunimovich stadium. We will discuss the process of equilibration in this billiard later in the paper.

The averaging theory can sometimes be extended to systems with a nonergodic fast subsystem (for example, refs. 41 and 44). Such extension requires detailed knowledge of the partition of the fast subspace into ergodic components. In the general theory of Hamiltonian systems, this partition is not known. However, mathematical billiards provide a rich set of examples, for which the ergodic decomposition can be described in full detail. We use this knowledge to show the accelerated equilibration for two different examples of springy billiards with nonergodic fast subspace (VFS systems). The first is the SBR, and the second is a modification of the Bunimovich mushroom (Fig. 1*A* and *C*).

Springy Barred Rectangle

Consider a particle moving in a rectangle of length L and height $h=2$, which is partially split by a horizontal bar of length λ attached to a spring (Fig. 1A). The particle's horizontal speed $|u_p|$ is preserved, and thus, the horizontal motion is periodic with period $T_p = 2L/|u_p|$. The particle moves with a large vertical speed and moderate horizontal speed, and therefore, the fast subsystem corresponds to the particle's vertical motion, while its horizontal position and the bar position are frozen. The period T_p is divided into two time intervals. In the first, the particle does not hit the bar, and its vertical speed is preserved; the fast subsystem, therefore, has a single ergodic component corresponding to the particle moving in the segment $y_p \in [-1, 1]$. In the second interval of length τT_p , where $\tau = \lambda/L$, the particle enters a chamber above or below the bar, where it gains or loses vertical speed as it hits the moving bar many times. Here, the fast subsystem has two ergodic components, one corresponding to the motion above the bar, $y_p \in [y_b, 1]$, and the other corresponding to the motion below the bar, $y_p \in [-1, y_b]$. During a single passage above or below the bar, the vertical speed $|v_p|$ obeys the adiabatic law given by Eq. 3 with $D = 1$, $E_p = 1/2mv_p^2 = E - 1/2v_b^2 - 1/2ky_b^2$, and $V = V^{up}(y_b) = 1 - y_b$ when the particle is above the bar or $V = V^{down}(y_b) = 1 + y_b$ when the particle is below it. Hence, $J_{up/down} = \sqrt{E - 1/2v_b^2 - 1/2ky_b^2}(1 \mp y_b) = \text{const.}$ Since $E_p = J_{up/down}^2/(1 \mp y_b)^2$, the right-hand side of Eq. 2 becomes $\mp 2J_{up/down}^2/(1 \mp y_b)^3$ for the particle above/below the bar, respectively. Note that, here, E and E_p do not include the particle's horizontal kinetic energy, which is decoupled from the dynamics.

To construct effective equations for the bar's average motion, we approximate the deterministic process by a stochastic one, assuming that the probability of entering the chamber above or below the bar is proportional to the width of the corresponding gap between the bar and the boundary of the rectangle. These probabilities are equal to $(1 \mp y_b)/2$, where y_b is taken at the moment of entrance into the chamber. The same assumptions were successfully used in ref. 45 for the study of Fermi acceleration in a rectangle with an infinitely heavy oscillating bar.

Hence, we suggest that the bar-particle system is well-approximated by the following T_p -periodic probabilistic hybrid model:

$$\ddot{y}_b + ky_b = \begin{cases} F_j(y_b, \dot{y}_b), & \text{for } jT_p \leq t < (j + \tau)T_p, \\ 0, & \text{for } (j + \tau)T_p \leq t < (j + 1)T_p, \end{cases} \quad [7]$$

where $j = \lfloor t/T_p \rfloor$ is the number of the period. For each period, a choice is made between $F_j(y_b) = -2J_{j,up}^2/(1 - y_b)^3$ (with probability β_j) and $F_j(y_b) = 2J_{j,down}^2/(1 + y_b)^3$ (with probability $1 - \beta_j$). Here, $J_{j,up/down} = \sqrt{E - E_b}(1 \mp y_b)$, $E_b = (\dot{y}_b^2 + ky_b^2)/2$, and $\beta_j = (1 - y_b(jT_p))/2$.

The bar dynamics follow the level lines of J_{up} , J_{down} , or $J_0 = 1/2(v_b^2 + ky_b^2)$. At $t = jT_p$, the bar height $y_b(jT_p)$ determines the probability β_j of choosing J_{up} vs. J_{down} , and $(y_b(jT_p), \dot{y}_b(jT_p))$ determines the particular $J_{up/down}$ level set along which the motion will continue.

In Fig. 2A, we plot the projection of direct simulations of the bar-particle system onto the (y_b, \dot{y}_b) plane for a few bar oscillations. Fig. 2A shows that the bar motion indeed follows the level sets of the corresponding J values and that the trajectory switches between level sets (Fig. 2A, Upper Left Inset); while the bar energy changes significantly when the particle is above/below the bar (Fig. 2A, Lower Right Inset), the corresponding J values remain approximately constant for each such interval (Fig. 2A, Lower Left Inset).

To study the equilibration process, we conduct the following numerical experiment. We generate an ensemble of initial conditions and for each initial condition, simulate a single-particle motion in the springy billiard. At each moment in time, we compute the ensemble-averaged difference between the kinetic energy of the bar and the vertical kinetic energy of the particle, $\Delta KE = \langle M/2v_b^2 \rangle - \langle m/2v_p^2 \rangle$. We study the time evolution of $\Delta KE(t)$ for ensembles of 10^4 initial states with a fixed initial bar energy. The initial phase of the bar oscillation is chosen randomly (with the uniform distribution). The position of the particle is also chosen randomly inside the rectangle. Finally, the vertical speed of the particle is determined from the initial particle energy, while the direction of motion (up or down) is chosen randomly. We repeat the experiment for different values of particle mass m and different initial values of the bar energy, keeping the total energy of the system constant, $E = 1$.

In all of our simulations, we find that $\Delta KE_{t \rightarrow +\infty} \rightarrow 0$ exponentially (Fig. 3A). However, for each fixed m , the rate of




Fig. 2. Piecewise adiabatic bar motion for the (A) SBR and (B) springy mushroom. Two trajectories of the springy billiards with particle mass $m = 10^{-8}$, one with $E_p(0) = 0.9$ (at the center) and one with $E_p(0) = 0.1$ (at the periphery), are projected to the bar-phase plane, (y_b, \dot{y}_b) , and plotted (black lines) on top of level lines of the adiabatic invariant. Green indicates the free springy bar motion. Red and blue indicate level lines for the bar motion in the effective potential created when the particle hits the bar from above and below, respectively. Each orbit closely follows the level lines, switching occasionally between them. Lower Left Insets show that the corresponding adiabatic invariants are approximately piecewise constant, and Lower Right Insets show the bar energy oscillations. In A, Upper Left Inset shows a magnification of the transition from a level set of one averaged law of motion to another.

To show that the equilibration observed in the SBR is not a peculiarity of that particular model, in the following, we consider another class of springy billiard systems, which have a different and more general structure of the decomposition of the fast subspace into ergodic components. For this class, we again show that violation of ergodicity in the fast-phase space leads to energy equilibration in the whole slow–fast system.

Springy Mushroom

When there are several fast degrees of freedom, the fast subspace of a slow–fast Hamiltonian system typically includes both regular and chaotic components (12, 32). As the slow variables change, a trajectory can transfer between these components. To study the influence of these transitions on equilibration, we introduce the springy mushroom model.

A mushroom billiard consists of a circular cap and a rectangular stem. It was proposed in ref. 27 to serve as a paradigm for systems with a mixed phase space, as it has the simplest possible mixed-phase space structure: there is a unique chaotic component and one regular (and completely integrable) component—a stability island that corresponds to the particle trapped forever in the cap.

In this paper, we use a variation of the original Bunimovich mushroom (Fig. 1C). Our mushroom billiard is bounded by a semicircular cap, two slanted side walls, and a straight bar at the bottom. This shape is a version of the famous Bunimovich stadium (26); the slanted walls prevent the appearance of horizontal parabolic orbits, which could spoil the statistics (cf. ref. 55). The corresponding billiard has a single chaotic component, which occupies the whole phase space. To create an integrable island, while keeping the billiard area unchanged, we add two horizontal straight segments, which leave a throat of radius w at the cap diameter. In this case, the trajectories that bounce inside the cap close to the circular boundary remain forever in the cap, and their motion is integrable (the absolute value of the angular momentum is preserved for such orbits). The size of this island depends on w , and the island vanishes when $w = r$, where r is the radius of the cap. However, the trajectories that move through the throat between the cap and the stem form a single chaotic component in the phase space (26, 27).

In the springy mushroom, the bar located at the bottom of the stem can move vertically. It has mass $M > 0$ and is attached to a linear spring (Fig. 1C). The bar moves harmonically between collisions with the particle, whereas the phase and amplitude of the oscillations change at every collision.

Finally, we take the radius of the throat change to be a prescribed function of the bar position: $w = w(y_b)$ (e.g., we use Eq. 11). Then, the island size depends on the bar position y_b . At a moment of time when the particle is in the stem, it belongs to the chaotic component of the frozen mushroom. After some collisions with the walls and the bar, the particle typically enters the cap. If the throat shrinks while the particle is in the cap, the particle may be captured in the cap, and it will stay there until the throat opens wide enough to let the particle back into the stem. During the period of capture, the particle follows the integrable dynamics of a semicircular billiard. Since the particle motion is fast compared with the bar, we can interpret these capture and release events as transitions between the chaotic and regular components of the frozen mushroom.

To construct effective equations for the average bar motion, we need to find the effective pressure exerted by the particle on the bar. Since the frozen billiard is not ergodic, the standard adiabatic law given by Eq. 4 is not applicable. Instead, we use the “leaky adiabatic law” of ref. 41:

$$\frac{dE_p}{E_p} = -\frac{dV}{V_c}, \quad [8]$$

where $V(y_b)$ is the phase space volume of the energy level $E_p = 1/2$ in the phase space of the frozen billiard and $V_c(y_b)$ is the phase space volume of the chaotic component for the same energy level. For billiards, V is the area of the billiard domain times 2π (the velocity circle). The effect of collisions between the particle and the bar is averaged over the chaotic component only. To account for this, V is replaced by V_c in the denominator of Eq. 4 (with $D = 2$). The term dV is kept in the numerator, because the work is done only by the bar and is proportional to the change in volume of the entire billiard table.

Expressions for the volumes in Eq. 8 can be found explicitly. Indeed, at the energy level $E_p = 1/2$ for the frozen billiard, we get $V(y_b) = \pi^2 r^2 + 2\pi(2r\ell - \ell^2 \tan \theta)$, where $\ell = \ell_0 - y_b$ is the stem length. The chaotic component is the complement of the integrable one. Thus, $V_c = V - V_{\text{isl}}$, where V_{isl} is the volume of the integrable island. It can be found by using the conservation of the angular momentum, $V_{\text{isl}}(y_b) = 2\pi r^2 \left(\cos^{-1} w/r - w/r \sqrt{1 - (w/r)^2} \right)$, where $w = w(y_b)$ is the width of the cap throat. Ref. 41 has a detailed derivation of these formulas.

When the particle is in the chaotic zone, it acts on the bar with average force $F = -dE_p/dy_b$. However, when the particle is in the cap, it does not collide with the bar. Using Eq. 8, we conclude that the averaged equations for the bar motion are of the form

$$\ddot{y}_b + \frac{dU(y_b)}{dy_b} = \begin{cases} \frac{E_p(y_b, \dot{y}_b)}{V_c(y_b)} \frac{dV(y_b)}{dy_b} & (\text{particle in } V_c), \\ 0 & (\text{particle in } V_{\text{isl}}), \end{cases} \quad [9]$$

where we use energy conservation to express the particle energy $E_p(y_b, \dot{y}_b) = E - 1/2\dot{y}_b^2 - U(y_b)$ as a function of y_b and \dot{y}_b . The first of the equations conserves $J = E_p(y_b, \dot{y}_b)G(y_b)$, where $G(y_b) = \exp\left(\int_0^{y_b} dV(s)/ds ds / V_c(s)\right)$. Thus, the fast collisions of the particle with the bar create, on average, an effective potential: $U_{\text{eff}}(y_b; J_0) = U(y_b) + J_0/G(y_b)$, where J_0 is the value of J at the moment the particle enters the chaotic zone. While the particle is in the chaotic component, the bar motion occurs along the level line $J = J_0$, which is also a level line of the Hamiltonian defined by the effective potential $U_{\text{eff}}(y_b; J_0)$. Similarly, when the particle is in the integrable component, the bar motion occurs along the level lines of the Hamiltonian defined by the free spring potential $U(y_b)$.

We conclude that, similar to the SBR case, the bar’s averaged motion switches between level sets of different functions. Numerical simulations of the bar–particle dynamics confirmed this property for small mass ratios: Fig. 2B shows that trajectories of the springy mushroom follow a level set of J (Fig. 2B, red circles) and then, by following a level set of the free motion (Fig. 2B, green circles), switch to follow another level set of J . Fig. 2B, *Lower Left Inset* shows that J is approximately constant during the time intervals in which the particle remains in the chaotic zone, and Fig. 2B, *Lower Right Inset* shows the corresponding oscillations of E_b .

To complete the derivation of the stochastic model for the motion of the bar, we describe the transition rules between the two possible states of the particle. A transition from the chaotic component to the integrable one can be modeled by a random process as follows. Suppose that the particle is in the chaotic component at time $t = t_c$. Then, the motion is ergodic, and the probability of transfer during a time interval $(t_c, t_c + dt)$ coincides with the ratio of the transferred phase space volume to the chaotic zone volume. The transferred volume is positive only when the throat is shrinking (i.e., when $\dot{w} < 0$ at $t = t_c$). In this case, the probability is given by

$$P_{\text{cha} \rightarrow \text{isl}}(t_c) = \frac{V'_{\text{isl}}(y_b) \dot{y}_b}{V_c(y_b)} dt = d \ln \frac{G(y_b)}{V_c(y_b)}. \quad [10]$$

If $w > 0$, then $P_{\text{cha} \rightarrow \text{isl}}(t_c) = 0$.

The transition from the integrable component back to the chaotic one is described by the following deterministic rule (41). Suppose that the particle is captured in the integrable component at time t_c . While the particle is captured, the bar moves harmonically. Therefore, the bar position and the throat size change periodically with time. Let $t_r > t_c$ be the first instance of time, such that $w(y_b(t_r)) = w(y_b(t_c))$. Then, $w(y_b(t)) > w(y_b(t_r))$ for $t_c < t < t_r$. Since typically, the particle trajectory in the integrable zone densely fills a strip between the semicircle of radius $w(y_b(t_c))$ and the cap boundary, the particle is released from the cap at a time interval $(t_r, t_r + dt)$. In the adiabatic limit, the indeterminacy in the exit moment, dt , vanishes.

These arguments provide the transition rules for the stochastic model of Eq. 9: when the throat shrinks, the particle can move from the chaotic component to the integrable one at time t_c with probability $P_{\text{cha} \rightarrow \text{isl}}(t_c)$ given by Eq. 10. If captured, the particle is released at time t_r when the throat restores its size [i.e., when $w(y_b(t_r)) = w(y_b(t_c))$ for the first time].

Now we are ready to describe the mechanism for the ergodization in the springy mushroom. Choosing a nonmonotonic function $w(y_b)$ allows for the possibility that the bar positions at the moments of capture and release are different: $y_b(t_c) \neq y_b(t_r)$. Then, the monotonicity of the function $G(y_b)$ implies that the particle returns to the chaotic zone with a different value of the adiabatic invariant $J(y_b(t_r), \dot{y}_b(t_r)) \neq J(y_b(t_c), \dot{y}_b(t_c))$. Consequently, this capture–release event moves the particle to a new level line of the adiabatic invariant. The destruction of the adiabatic invariant enhances the energy transfer between the fast and slow degrees of freedom, and therefore, repeated capture–release events lead to faster equilibration.

For example, in our numerical experiments, we use the non-monotonic function

$$w(y_b) = \begin{cases} \min \{1, 0.7 + 0.6(y_b - y_0)^2\} & \text{for } y_b < y_0, \\ \min \{1, 0.7 + 6.0(y_b - y_0)^2\} & \text{for } y_b \geq y_0, \end{cases} \quad [11]$$

where $y_0 = -0.2436$. This choice of w eliminates some evident obstacles for ergodicity in the springy mushroom and its random model. In particular, we have $0 \leq w(y) \leq 1$ for all y and $w(y) = 1$ for $y \geq 0$. The latter property ensures that the integrable component in the fast subspace disappears at $y_b \geq 0$, and therefore, no initial condition can correspond to a particle trapped in the cap forever. In addition, we require the function w to have a unique minimum at a point y_0 , such that $y_f \leq y_0 < 0$ (where y_f is the pressure equilibrium for the energy level that we are studying). One can check that this inequality ensures that, on every level line of J , there exist capture moments, such that necessarily $y_b(t_c) \neq y_b(t_r)$, the crucial condition for the destruction of the adiabatic invariant. Finally, taking $w(y_0) \geq 0.7$ eliminates vertical parabolic orbits in the chaotic component. We expect that the conclusions presented below will apply to other functions fulfilling these requirements.

To test that the proposed mechanism correctly describes the relaxation to the energy equipartition, we conducted a series of numerical experiments with both the springy mushroom and the corresponding stochastic model. Fig. 3B shows that the kinetic energies of the springy mushroom equilibrate. Defining $\Delta KE(t) = \langle M/2v_b^2 \rangle - 1/2 \langle E_p \rangle$, we find that, in all our simulations, $\Delta KE_{t \rightarrow +\infty} \rightarrow 0$. Here, the ensembles with fixed bar energy are chosen with a random initial phase of the bar motion, a random position of the particle in the stem (we do not start in the cap), and random direction of the particles velocity. Calculating the equilibration rates using the procedure outlined for the SBR model, we again find, as with the SBR model, the sensitivity to

the initial energy of the bar and the choice of fitting interval. Nonetheless, for every fixed setting, we find that the rates are well-defined and depend linearly on the square root of the mass ratio, \sqrt{m} . Extrapolating these rates to $m = 0$, we obtain strictly positive limits 1.75×10^{-4} for the ensemble with $E_b(0) = 0.9$ and 3.08×10^{-4} for the ensemble with $E_b(0) = 0.1$.

Simulating the stochastic model for the corresponding ensembles of initial conditions (fixed bar energy, random bar phase, motion starting in the chaotic phase) (more details are in *Numerical Algorithms*), we find that the random model equilibrates in a similar fashion, with positive equilibration rates close to those obtained by the springy billiard simulations when extrapolated to $m = 0$ (Fig. 4B). Indeed, the rates for 10 runs of 10^4 particles each are $(1.42 \pm 0.03) \times 10^{-4}$ and $(2.71 \pm 0.11) \times 10^{-4}$ for $E_b(0) = 0.9$ and $E_b(0) = 0.1$, respectively.

We conclude that the stochastic models, with no adaptable parameters, provide a reasonable approximation of the billiard dynamics observed in our numerical experiments. These results support our claim that slow–fast VFS systems are well-approximated by processes with random switching between several different equations for the slow variables.

Springy Stadium

As a control, we examine the equilibration process in a springy billiard, where the frozen billiard is chaotic for all positions of the bar (an EFS system). Specifically, we consider a springy half-stadium (Fig. 1B), which can be obtained from the springy mushroom by removing the collar (i.e., by using $w \equiv 1$ instead of Eq. 11) and keeping the values of all other parameters. Ergodicity of the frozen billiard is proven in ref. 26.

We show that, in such springy billiards, the presence of the adiabatic invariant slows down the relaxation to energy equipartition and that the equilibration rates approach zero as the mass ratio is decreased.

Indeed, direct simulations of the springy stadium show that the adiabatic invariant J defined by Eq. 3 remains approximately constant and that the bar energy behaves approximately periodically for long time intervals. However, these simulations also show that, for finite m , small fluctuations ultimately destroy the adiabatic conservation law. In particular, defining $\Delta KE(t)$ as for the springy mushroom and repeating exactly the same direct numerical simulations for the springy stadium case, we observe again that, for all initial ensembles, $\Delta KE(t) \rightarrow 0$ as t grows (Fig. 3B). For small values of m , we observe smaller equilibration rates (but not much smaller) than in the springy mushrooms with comparable parameters (Fig. 4B). Using the best fit line to extrapolate the computed rates to $m = 0$, we see that the limit rate vanishes within the accuracy of our computations. Repeating the computations with increased precision, larger ensemble size, and smaller m , we obtain even smaller limit values for the equilibration rates. We conclude that the limit equilibration rates of the springy stadium are at least two orders of magnitude smaller than the corresponding limit rates for the springy mushroom.

This is our main finding, as it reveals the profound difference between EFS and VFS systems. In agreement with adiabatic theory, the limit equilibration rates vanish in EFS systems. The positive limit rates achieved for VFS systems support our claim of efficient mixing induced by the hopping between ergodic components of the fast subsystem.

Discussion

Springy billiards show an important principle in slow–fast Hamiltonian systems: ergodicity of the fast subsystem impedes ergodization and equilibration of the whole slow–fast system, whereas its violation can lead to equilibration and equipartition of energies. For EFS systems, the averaged lower-dimensional system is known to have an adiabatic invariant, and it is thus

not ergodic and does not equilibrate. In contrast, for VFS systems, where the ergodicity of the fast subsystem is broken for a range of values of the slow variables, we showed that the adiabatic behavior is well-approximated by a random process. This process is defined by random switching among several systems that govern the evolution of the slow variables. Importantly, the magnitude and probabilities of the jumps in the stochastic model remain strictly positive in the adiabatic limit. Under mild conditions, such random processes can lead to equilibration of energies with positive rates. We constructed both EFS- and VFS-type springy billiard systems and derived the random models for the VFS springy billiards. We computed the equilibration rates for both the springy billiards and the random models. Extrapolating the rates found at decreasing m values, we showed that the equilibration rate for the considered EFS system (springy stadium) vanishes in the limit $m = 0$, whereas for the VFS systems (SBR and springy mushroom), it remains positive and matches the equilibration rate of the corresponding stochastic model (Fig. 4).

Our equilibration mechanism is also expected to appear in smooth slow–fast systems. The chaotization mechanism that we propose for VFS systems is reminiscent of the phenomena of adiabatic chaos—chaotization of smooth slow–fast systems in which the fast dynamics is integrable, yet the structure of the fast-phase space changes as the slow variables are changed (56–58). We suggest that this mechanism is universal and not restricted to fast subsystems that are integrable. In fact, if some of the ergodic components of the fast system are chaotic and additionally, if there exists a range of slow variables for which the chaotic ergodic component occupies a large portion of the fast subsystem phase space, the equilibration process may be particularly fast.

The proposed framework naturally opens several avenues of research. (i) We saw that the EFS system can equilibrate, despite the presence of an adiabatic invariant. Can a comprehensive theory be built for the behavior beyond the adiabatic timescale? (ii) In the stochastic model for a VFS system, we should typically have convergence to a stationary measure. What if the irreducibility or aperiodicity conditions are violated? For example, the fast subsystem can have a stability island, which persists for all values of slow variables. This could create an invariant domain in the phase space, which has to be excluded from the microcanonical ensemble. (iii) In both EFS and VFS systems, we observe ensemble-dependent rates of convergence to equilibrium. How typical is this phenomenon, and how does it affect the equilibration? (iv) A slow–fast system ceases to be slow–fast near stationary points of the fast subsystem (the fast system freezes for some time). How can one incorporate the freezing incidents into the theory?

Finally, we propose a broader viewpoint on this work. Here, our slow and fast systems were just two mechanical components, and the timescale separation stemmed from the mass ratio. In the broader statistical mechanics context, the fast system governs the motion of many particles (and is, thus, high-dimensional), and the slow macroscopic variables are defined as certain averages over the fast microscopic system. When the structure of the fast system changes, for example, from a gas to a liquid state, a phase transition occurs. Usually, for macroscopic values near the phase transition, the microscopic phase space structure is complex, with long-lasting structures in which the two states coexist. We may say that, in this range of macroscopic variables, the fast subsystem has chaotic (say, gas phase) and ordered (say, liquid state) components. Therefore, we conjecture that phase transitions play a central role in the equilibration process between microscopic and macroscopic variables (e.g., consider the analogue of the notorious piston problem in a multi-phase gas).

Numerical Algorithms

Numerical simulations of smooth slow–fast systems are challenging; by considering springy billiards, we avoid the need for numerical integration of stiff equations. Nevertheless, the small mass ratio, $m/M \in [10^{-7}, 10^{-4}]$, leads to a large number of billiard reflections per unit time, and therefore, we have to deal with long-time simulations of chaotic dynamics. Here, we describe the algorithms used for simulations of the springy billiard dynamics and the associated stochastic models. We also include a list of various consistency checks used to verify the accuracy of the computations.

Direct Simulations of Springy Billiards. We use an explicit formula for the harmonic motion of the bar and the inertial motion of the particle between collisions. The time and place of the collisions are evaluated with the maximal precision allowed by the computer (we use at least double precision). Most of the collisions are with the static walls, and these are computed explicitly. To compute a collision of the particle with the moving bar, we find explicit expressions for the monotonicity intervals of the distance between the particle and the bar as a function of time. Then, we pick out the first monotonicity interval, such that the function changes sign at its ends. This interval contains a single zero of the distance function, which is found using a combination of the bisection and Newton method.

In every series of numerical experiments, we investigate energy equilibration on a fixed energy surface (typically $E = 1$) in a time interval $[0, T_1]$. We construct an ensemble of initial conditions that share the same initial value of the bar energy E_b (the details of the ensembles are in the relevant parts of the paper). For each initial condition, we compute the trajectory of the springy billiard for $0 \leq t \leq T_1$, evaluating the energy components (kinetic and potential energies) at equally spaced moments in time (typically $t_k = kh$ with $h = 10^{-3}$). Then, the ensemble-averaged values of the energy components are computed for every t_k and stored to be used for the computation of the equilibration rate. The rate is defined as the best fitted slope to $\log |\Delta KE|$ vs. t over a time interval $[0, T_0]$, so that $\log |\Delta KE(0)/\Delta KE(T_0)| \approx 1$.

Theoretically, for an infinitely large ensemble, T_0 can be precisely defined by setting it to be the first time where $\log |\Delta KE(0)/\Delta KE(T_0)| = 1$. For finite ensembles, this definition leads to noisy T_0 for different realizations with a small systematic underestimate of T_0 . In the case of the SBR, for all m , we find that $T_0 \approx 10.5$, and its variations do not affect the consistency of the results, as different realizations produce similar rates (Fig. 44). In the case of the springy mushroom and stadium, a higher accuracy is needed to study the behavior of the m -dependent rates in the limit $m \rightarrow 0$. Therefore, to help reduce uncertainty in the definition of the rates, we choose an ensemble-independent T_0 . To set T_0 as a function of m only, we first find the m -dependent rates using ensemble-dependent fitting time intervals. Then, we define a linear function $R(\sqrt{m})$ by the best linear fit for these rates as a function of m , and let $T_0(m) = 1/R(\sqrt{m})$. Finally, we recompute the rates using $T_0 = T_0(m)$. This procedure leads to only a small change in the obtained rate values (the difference is about 2% of their first estimate). Nevertheless, the extrapolation to $m = 0$ is improved; as an example, for the stadium, the extrapolated rates get closer to zero, which is the theoretical rate.

The computation of a trajectory segment for $t \in [0, T_1]$ ($T_1 = 21$ for the SBR; $T_1(m) \in [200, 3,000]$ for the mushroom and the stadium) includes 10^3 – 10^6 collisions. To maintain a high level of accuracy, the following checks of the code are carried out. (i) The total energy of the system for the full duration of the simulations is conserved up to at least nine significant digits. (ii) When the code is run for a static billiard (i.e., the position of the bar is fixed), the particle energy is preserved with at least nine

significant digits. (iii) For selected parameter values, the code is run several times to estimate the size of the fluctuations caused by the finite size of the ensemble (e.g., the reported SDs for the equilibration rates). (iv) The experiments are repeated with different ensemble sizes.

Stochastic Model Simulations. In the stochastic models, the motion of the bar is described by several explicitly defined Hamiltonian systems with one degree of freedom (Eqs. 7 and 9) and a probabilistic law, which determines transitions between these systems. The stochastic models correspond to the limit of $m = 0$ and are independent of the mass ratio.

Stochastic simulations for the SBR. In this case, the bar motion is described by Eq. 7, which is integrated with the help of the standard Runge–Kutta method of order 4 and a time step size of $h \approx 10^{-4}$. When $t = 0 \pmod{T}$, a random choice of F_j is made with the prescribed probabilities, and at $t = \tau \pmod{T}$, the right-hand side of Eq. 5 is reset to zero. The initial conditions for the bar motion are set to reflect the properties of the ensembles used in the simulations of the SBR: the initial energy of the bar takes the same value, while its phase on the (y_b, \dot{y}_b) plane is random. To mimic the uniform distribution of the particle position in the box, a random phase is chosen from which one starts the integration of Eq. 7. If the phase is in $(0, \tau T_p/T)$, a random choice of the up or down system is made with probabilities β_0 and $1 - \beta_0$, respectively.

Stochastic simulations for the springy mushroom. In contrast to the SBR model, the switching time between the two forces in Eq. 9 is stochastic and governed by the probabilistic law of Eq. 10. We wrote two different numerical codes to carry out these simulations. The rates obtained by these two methods coincide within the accuracy determined by the size of the ensembles.

Algorithm 1—Direct integration and switching: We integrate Eq. 9 in the interval $[t, t + dt]$ and find $y_b(t + dt)$ and $\dot{y}_b(t + dt)$. Then, we switch to the harmonic motion of the bar with probability $P(t)dt$ given by Eq. 10, where y_b is replaced by $y_b(t + dt)$ and $\dot{y}_b dt$ is replaced by $dy_b = y_b(t + dt) - y_b(t)$. After the particle is captured, the position and velocity of the bar at the time of release are calculated explicitly using the equations of harmonic motion.

The initial ensemble for the bar position is the same as in the direct simulations, and since in the direct simulations, the particle always starts from the stem, the chaotic simulation always starts with the nontrapped particle. To test the accuracy of our code, we checked that the achieved rates converge as the integration step is decreased (we typically used $dt = T_1/2,000$ in the stochastic simulations). We also checked that repeated runs produce statistically similar results.

Algorithm 2—Switching without integration: We begin by preparing tables, which provide, for each value of the adiabatic invariant J , the period of the bar's oscillation in the effective potential of Eq. 9, the probability of the particle being captured in the cap during one period, and when the particle is captured, the cumulative distribution function (CDF) for the bar position y_b at the moment of capture. After that, for any initial condition (y_b, \dot{y}_b) (assuming that the particle is in the chaotic component), we generate the number of completed rounds before capture and the value of y_b at the moment of capture as random variables, with the CDFs obtained by interpolation of the precomputed tables. After the capture point is generated, the time of release and the new position and velocity of the bar are computed from the harmonic motion law. Then, the process is repeated until the time exceeds T_1 .

The outcome of this stochastic algorithm is a list of capture and release points (time, bar position, and bar velocity). From this list, we find the bar energy at fixed time intervals: it is constant in the capture phase and found by using the adiabatic invariance in the chaotic phase. This algorithm enables fast simulation without integration. Agreement with the results of the simulations run using *Algorithm 1* provides additional verification of those results.

ACKNOWLEDGMENTS. Part of this work was done while K.S. was at the Indian Institute of Technology Delhi. K.S. was supported by Science and Engineering Research Board, Government of India File SR/FTP/PS-108/2012. K.S., D.T., and V.G. acknowledge the financial support and hospitality of the Weizmann Institute of Science, where part of this work was done. D.T. was supported by Russian Science Foundation Grant 14-41-00044 for this research and Engineering and Physical Sciences Research Council (EPSRC) Grant EP/P026001/1 and thanks the Royal Society. The research of V.G. was supported by EPSRC Grant EP/J003948/1. V.R.-K. acknowledges support from Israel Science Foundation Grant 1208/16.

- Landau LD, Lifshitz EM (1980) *Statistical Physics* (Pergamon, Oxford), Vol 5.
- Huang K (1987) *Statistical Mechanics* (Wiley, New York).
- Ruelle D (1999) *Statistical Mechanics: Rigorous Results* (World Scientific, Singapore).
- Dorfman JR (2001) *An Introduction to Chaos in Nonequilibrium Statistical Mechanics* (Cambridge Univ Press, Cambridge, UK).
- Hilbert S, Hänggi P, Dunkel J (2014) Thermodynamic laws in isolated systems. *Phys Rev E Stat Nonlin Soft Matter Phys* 90:062116.
- Sinai YG (1970) Dynamical systems with elastic reflections. *Russ Math Surv* 25: 137–189.
- Sinai YG, Chernov N (1987) Ergodic properties of certain systems of two-dimensional discs and three-dimensional balls. *Russ Math Surv* 42:181–207.
- Simanyi N, Szasz D (1999) Hard ball systems are completely hyperbolic. *Ann Math* 149:35–96.
- Kaplan A, Friedman N, Andersen M, Davidson N (2001) Observation of islands of stability in soft wall atom-optics billiards. *Phys Rev Lett* 87:274101.
- Rapoport A, Rom-Kedar V, Turaev D (2008) Stability in high dimensional steep repelling potentials. *Comm Math Phys* 279:497–534.
- Rom-Kedar V, Turaev D (2012) Billiards: A singular perturbation limit of smooth Hamiltonian flows. *Chaos* 22:026102.
- Zaslavsky GM (1985) *Chaos in Dynamic Systems* (Harwood Academic, New York).
- Lichtenberg AJ, Leiberman M (1992) *Regular and Chaotic Dynamics* (Springer, New York).
- Anosov DV (1960) Averaging in systems of ordinary differential equations with rapidly oscillating solutions. *Izv Akad Nauk SSSR Ser Mat* 24:721–742.
- Kasuga T (1961) On the adiabatic theorem for the hamiltonian system of differential equations in the classical mechanics, I. *Proc Jpn Acad* 37:366–371.
- Lochak P, Meunier C (1988) *Multiphase Averaging for Classical Systems* (Springer, New York).
- Lieb EH (1999) Some problems in statistical mechanics that I would like to see solved. *Physica A* 263:491–499.
- Lebowitz JL, Piasecki J, Sinai YG (2000) Scaling dynamics of a massive piston in an ideal gas. *Hard Ball Systems and the Lorentz Gas*, ed Szasz D (Springer, Berlin), pp 217–227.
- Neishtadt AI, Sinai YG (2004) Adiabatic piston as a dynamical system. *J Stat Phys* 116:815–820.
- Wright P (2007) The periodic oscillation of an adiabatic piston in two or three dimensions. *Commun Math Phys* 275:553–580.
- Dolgopyat D, Liverani C (2011) Energy transfer in a fast-slow Hamiltonian system. *Commun Math Phys* 308:201–225.
- Grigo A, Khanin K, Szasz D (2012) Mixing rates of particle systems with energy exchange. *Nonlinearity* 25:2349–2376.
- Balint P, Gilbert T, Nándori P, Domokos S, Toth IP (2017) On the limiting Markov process of energy exchanges in a rarely interacting ball-piston gas. *J Stat Phys* 166: 903–925.
- Chernov N, Markarian R (2006) *Chaotic Billiards* (American Math Society, Providence, RI).
- Hannay JH, McCraw RJ (1990) Barrier billiards—a simple pseudo-integrable system. *J Phys A Math Gen* 23:887–900.
- Bunimovich LA (1979) On the ergodic properties of nowhere dispersing billiards. *Commun Math Phys* 65:295–312.
- Bunimovich LA (2001) Mushrooms and other billiards with divided phase space. *Chaos* 11:802–808.
- Schmidt G, Chen Q (1994) Apparent transition to ergodicity in the autonomous Fermi system. *Physica D* 71:18–22.
- Lamba H (1995) Chaotic, regular and unbounded behaviour in the elastic impact oscillator. *Physica D* 82:117–135.
- Bievre SD, Parris PE, Silviu A (2005) Chaotic dynamics of a free particle interacting linearly with a harmonic oscillator. *Physica D* 208:96–114.
- Luo ACJ, Guo Y (2009) Motion switching and chaos of a particle in a generalized Fermi-acceleration oscillator. *Math Probl Eng* 2009:1–40.
- Lichtenberg AJ, Lieberman MA, Cohen RH (1980) Fermi acceleration revisited. *Physica D* 1:291–305.

33. Jarzynski C (1993) Energy diffusion in a chaotic adiabatic billiard gas. *Phys Rev E* 48:4340–4350.
34. Loskutov A, Ryabov AB, Akinshin LG (2000) Properties of some chaotic billiards with time-dependent boundaries. *J Phys A* 33:7973–7986.
35. de Carvalho RE, Souza FC, Leonel ED (2006) Fermi acceleration on the annular billiard. *Phys Rev E* 73:066229.
36. Lenz F, Diakonov FK, Schmelcher P (2008) Tunable Fermi acceleration in the driven elliptical billiard. *Phys Rev Lett* 100:014103.
37. Dolgopyat D (2008) Fermi acceleration. *Contemp Math* 469:149–166.
38. Shah K (2011) Energy growth rate in smoothly oscillating billiards. *Phys Rev E* 83:046215.
39. Itin AP, Neishtadt AI (2012) Fermi acceleration in time-dependent rectangular billiards due to multiple passages through resonances. *Chaos* 22:026119.
40. Gelfreich V, Rom-Kedar V, Turaev D (2012) Fermi acceleration and adiabatic invariants for non-autonomous billiards. *Chaos* 22:033116.
41. Gelfreich V, Rom-Kedar V, Turaev D (2014) Oscillating mushrooms: Adiabatic theory for a non-ergodic system. *J Phys A Math Theor* 47:395101.
42. Batistic B (2014) Exponential Fermi acceleration in general time-dependent billiards. *Phys Rev E* 90:032909.
43. Turaev D (2016) Exponential energy growth due to slow parameter oscillations in quantum mechanical systems. *Phys Rev E* 93:050203.
44. Gelfreich V, Rom-Kedar V, Shah K, Turaev D (2011) Robust exponential acceleration in time-dependent billiards. *Phys Rev Lett* 106:074101.
45. Shah K, Turaev D, Rom-Kedar V (2010) Exponential energy growth in a Fermi accelerator. *Phys Rev E* 81:056205.
46. Pereira T, Turaev D (2015) Exponential energy growth in adiabatically changing Hamiltonian systems. *Phys Rev E* 91:010901.
47. Khanin K, Yarmola T (2013) Ergodic properties of random billiards driven by thermostats. *Commun Math Phys* 320:121–147.
48. Mokhtar-Kharroubi M, Rudnicki R (2017) On asymptotic stability and sweeping of collisionless kinetic equations. *Acta Appl Math* 147:19–38.
49. Eckmann J, Young L (2011) Rattling and freezing in a 1D transport model. *Nonlinearity* 24:207–226.
50. Konishi T, Yanagita T (2010) Slow relaxation to equipartition in spring-chain systems. *J Stat Mech* 2010:P09001.
51. Burton R, Easton RW (1980) Ergodicity of linked twist maps. *Global Theory of Dynamical Systems*, eds Nitecki Z, Robinson C (Springer, Berlin), pp 35–49.
52. Sivaramakrishnan A (1989) Linked twist maps in Hamiltonian systems. *Celestial Mech Dyn Astron* 46:35–48.
53. Berger P, Carrasco P (2014) Non-uniformly hyperbolic diffeomorphisms derived from the standard map. *Comm Math Phys* 329:239–262.
54. Blumenthal A, Xue J, Young LS (2017) Lyapunov exponents for random perturbations of some area-preserving maps including the standard map. *Ann Math* 185:285–310.
55. Primack H, Smilansky U (1994) Quantal consequences of perturbations which destroy structurally unstable orbits in chaotic billiards. *J Phys A Math Gen* 27:4439–4448.
56. Tennyson JL, Cary JR, Escande DF (1986) Change of the adiabatic invariant due to separatrix crossing. *Phys Rev Lett* 56:2117–2120.
57. Vainshtein DL, Vasiliev AA, Neishtadt AI (1998) Adiabatic chaos in a two-dimensional mapping. *Chaos* 6:514–528.
58. Neishtadt AN, Chaikovskii DK, Chernikov AA (1991) Adiabatic chaos and particle diffusion. *Sov Phys JETP* 72:423–430.