Dynamics of periodically kicked oscillators

Kevin K. Lin and Lai-Sang Young

To Steve Smale on the occasion of his 80th birthday

Abstract. We review some recent results surrounding a general mechanism for producing chaotic behavior in periodically kicked oscillators. The key geometric ideas are illustrated via a simple linear shear model.

Introduction

This paper reviews some recent results on a topic with a considerable history: the periodic forcing of limit cycles. Some 80 years ago, van der Pol and van der Mark [26] observed that irregularities developed when certain electrical circuits exhibiting stable oscillations were periodically forced. Their work stimulated a number of analytical studies; see, e.g., [5, 13, 12, 9]. Another classical example of driven oscillators is the FitzHugh–Nagumo neuron model [6]; the response of this and other models of biological rhythms to external perturbations have been extensively studied (see, e.g., [33]). As a topic of mathematical study, the dynamics of forced oscillations is well motivated: Oscillatory behavior are ubiquitous in physical, biological, and engineered systems, and external forcing, whether artificially applied or as a way to model forces not intrinsic to the system, is also commonplace.

In this article, we are not concerned with modeling specific physical phenomena. Instead, we consider a generic dynamical system with a limit cycle and seek to understand its qualitative behavior when the system is periodically disturbed. To limit the scope of the problem, we restrict ourselves to periodic *kicks*, or forcings that are turned on for only short durations, leaving the limit cycle ample time to restore itself during the relaxation period. We are interested in large-time behavior, particularly in questions of stability and chaos.

As we will show, one of the properties of the limit cycle that plays a key role in determining whether the kicked system is stable or chaotic is *shear*, by which we refer to the differential in speed (or angular velocity) for orbits near the limit cycle. A central theme of this article is that under suitable

conditions, the impact of a kick can be substantially magnified by the under-
lying shear in the unforced system, leading to the formation of horseshoes and
strange attractors. This does not always happen, however: some limit cycles
are more vulnerable, and some types of kicks are more effective than others.
These ideas are discussed in [29, 30] and [14], the material which forms the
basis of the present review.

   Even though we seek to provide insight into dynamical mechanisms that
operate under general conditions, we have found the ideas to be most trans-
parent in a very simple linear shear model, to which we will devote a nontriv-
ial part of the paper. Section 1 introduces this example and familiarizes the
reader with the various parameters (including the one which measures shear);
it also reports on results of a numerical study on Lyapunov exponents. Sec-
tions 2 and 3 are organized around explaining these simulation results. Along
the way, we take the opportunity to review a number of mathematical ideas
which clearly go beyond this one example. Some of the rigorous results re-
viewed, notably those on Sinai–Ruelle–Bowen (SRB) measures for a relevant
class of strange attractors [28, 31, 32], are recent developments. With the
main ingredients of the linear shear model and the relevant mathematical
background in hand, we return to a discussion of general limit cycles in the
final section.

1. Increasing shear as a route to chaos

This section introduces the main example we use in this review, and ac-
quaints the reader with the various parameters in the model and how they
impact the dynamics. Of particular interest to us is the effect of increasing
shear. Numerically computed Lyapunov exponents as functions of shear are
presented in Section 1.3. They will serve as a focal point for some discussions
to follow.

1.1. Periodic kicking of a linear shear flow

Our main example is the periodic kicking of a linear shear flow with a hy-
perbolic limit cycle. This two-dimensional (2D) model was studied rigorously
in [29, 30] and numerically in [14, 35]. Though exceedingly simple in appear-
ance, it already exhibits rich and complex dynamical behaviors.

   The model is given by

$$\dot{\theta} = 1 + \sigma y,$$

$$\dot{y} = -\lambda y + A \cdot \sin(2\pi\theta) \cdot \sum_{n=-\infty}^{\infty} \delta(t - n\tau), \tag{1}$$

where $(\theta, y) \in \mathbb{S}^1 \times \mathbb{R}$, $\mathbb{S}^1 \equiv \mathbb{R}/\mathbb{Z}$, and σ, λ, A, and τ are constants with
$\sigma, \lambda, \tau > 0$. We will refer to equation (1) with $A = 0$ as the *unforced equa-
tion*, and the term involving A as the *forcing* or the *kick*. Here δ is the
usual δ-function, that is to say, the kicks occur instantaneously at times
$0, \tau, 2\tau, 3\tau, \dots$.

More precisely, let Φ_t denote the flow corresponding to the unforced equation. It is easy to see that for all $z \in \mathbb{S}^1 \times \mathbb{R}$, $\Phi_t(z)$ tends to the limit cycle $\gamma = \{y = 0\}$ as $t \to \infty$. The precise meaning of equation (1) is as follows. Let $\kappa(\theta, y) = (\theta, y + A\sin(2\pi\theta))$ be the kick map; it represents the action of the forcing term. Then assuming we start with a kick at time 0, the time-τ map of the flow generated by equation (1) is

$$\Psi_\tau = \Phi_\tau \circ \kappa,$$

and the evolution of the system is defined by iterating Ψ_τ. We generally assume that τ is not too small, so that during the relaxation period between kicks, the flow Φ_t of the unforced equation "restores" the system to some degree.

The parameters of interest are

$$\sigma = \text{amount of shear,}$$
$$\lambda = \text{rate of contraction to } \gamma,$$
$$A = \text{amplitude of kicks, and}$$
$$\tau = \text{time interval between kicks.}$$

Our aim in the remainder of this section is to understand—via geometric reasoning and numerical simulations—the meanings of these quantities, and the roles they play in questions of stability and chaos. Our line of reasoning follows [29, 14].

1.2. Geometry of Ψ_τ

A simple way to gain intuition on the geometry of Ψ_τ is to study the Ψ_τ-image of γ, the limit cycle of the unforced system. We will do so by freezing some of the parameters and varying others.

Effects of varying σ, λ, and A. To begin with, let us freeze λ, A, and τ. To fix ideas, let us take λ to be relatively small, so that the rate of contraction is weak, and choose τ large enough that $e^{-\lambda\tau}$ is a nontrivial contraction. This is when the effects of shear are seen most clearly. In Figure 1, $\lambda = A = 0.1$, and $\tau = 10$. Here $\Psi_\tau(0,0) = (0,0)$ because the limit cycle γ has period 1 and τ is an integer multiple of this period; for noninteger τ the picture is shifted horizontally.

Figure 1(b) shows four images of γ under Ψ_τ for increasing shear. The larger σ, the greater the difference in velocity between two points with different y coordinates. This applies in particular to the highest and lowest points in $\kappa(\gamma)$ in Figure 1(a). For σ small enough, order in the θ-direction is preserved; i.e., for $z_1 = (\theta_1, 0)$ and $z_2 = (\theta_1 + \varepsilon, 0)$, $\Psi_\tau(z_1)$ will continue to have a slightly smaller θ-coordinate than $\Psi_\tau(z_2)$. As σ increases, some points in γ may "overtake" others, spoiling this order. As σ gets larger still, the total distances traveled in τ units of time vary even more, and a fold develops. This fold can be made arbitrarily large: we can make it wrap around the cylinder as many times as we wish by taking σ large enough.

(a) The limit cycle γ and its image $\kappa(\gamma)$ after one kick.

$\sigma = 0.05$ $\sigma = 0.25$

$\sigma = 0.5$ $\sigma = 1$

(b) $\Psi_\tau(\gamma)$ for $\tau = 10$.
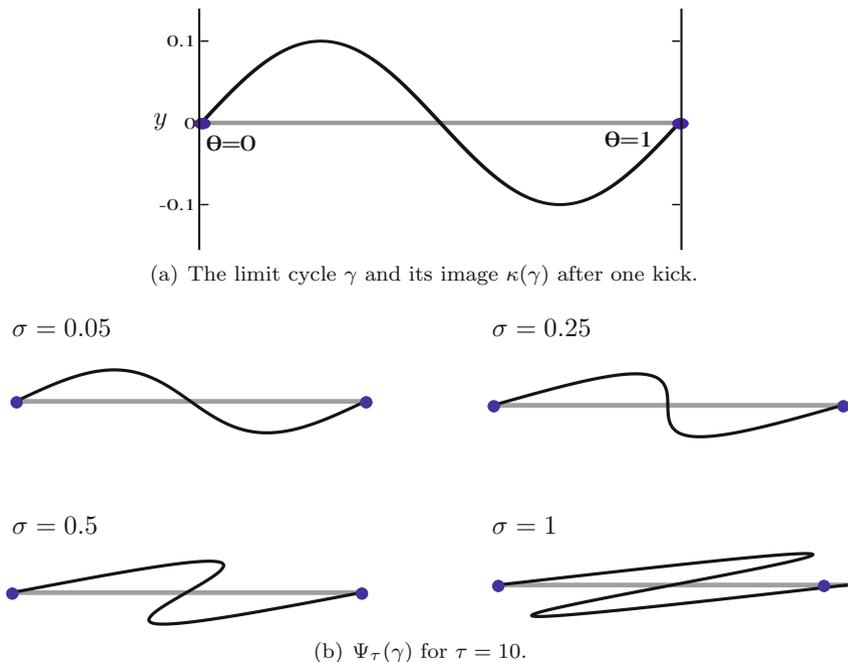
FIGURE 1. Effect of increasing shear. Here, $\lambda = A = 0.1$.

If we had fixed σ instead, and increased λ starting from $\lambda = 0.1$, the resulting sequence of pictures would be qualitatively similar to Figure 1(b) but in reverse order: The smallest λ would correspond to the bottom-right image in Figure 1(b), and the largest λ to the top-left—provided τ is scaled so that $\lambda\tau$ remains constant. This is because for λ small, $\kappa(\gamma)$ returns to γ very slowly, giving the shear a great deal of time to act, while for larger λ, $\kappa(\gamma)$ is brought back to γ more quickly. Thus all else being equal, σ and λ, i.e., shear and damping, have opposite effects.

The consequence of varying A while keeping the other parameters fixed is easy to see: the stronger the kick, the greater the difference in y-coordinate between the highest and lowest points in $\kappa(\gamma)$, and the farther apart their θ-coordinates will be when flowed forward by Φ_τ.

What we learn from the sequence in snapshots in Figure 1 is that A *acts in concert with σ to promote fold creation, while λ works against it.*

Formulas for Ψ_τ. Since the unforced equation is easy to solve, one can in fact write down explicitly the formulas for Ψ_τ. Let $(\theta_\tau, y_\tau) = \Psi_\tau(\theta_0, y_0)$. A simple computation gives

$$\theta_\tau = \theta_0 + \tau + \frac{\sigma}{\lambda} \cdot [y_0 + A\sin(2\pi\theta_0)] \cdot (1 - e^{-\lambda\tau}) \pmod 1,$$

$$y_\tau = e^{-\lambda\tau}[y_0 + A\sin(2\pi\theta_0)].$$

(2)

The reader can easily check that equation (2) is in agreement with the intuition from earlier.

Note the appearance of the ratio $\frac{\sigma}{\lambda}A$, or rather $\frac{\sigma A}{\lambda}(1 - e^{-\lambda\tau})$, in the nonlinear term in the equation for θ_τ: the size of this term is a measure of the tendency for a fold to develop in $\Psi_\tau(\gamma)$. As is well known to be the case, stretch-and-fold is a standard mechanism for producing chaos. One can, therefore, think of the ratio

$$\frac{\sigma}{\lambda}A = \frac{\text{shear}}{\text{contraction}} \cdot \text{kick amplitude}$$

as the key to determining whether the system is chaotic, provided $\lambda\tau$ is large enough that the factor $1 - e^{-\lambda\tau}$ is not far from 1.

Trapping region and attractor. From the above, it is evident that much of the action takes place in a neighborhood around γ. Let $U = \{|y| \leq A(e^{\lambda\tau}-1)^{-1}\}$, so that $\Psi_\tau(U) \subset U$, and define

$$\Gamma = \cap_{n\geq 0}\Psi_\tau^n(U)$$

to be the *attractor* for the system in equation (1). The basin of attraction of Γ is the entire cylinder $\mathbb{S}^1 \times \mathbb{R}$, since every orbit will eventually enter U. This usage of the word "attractor" implies no knowledge of dynamical indecomposability (a condition required by some authors).

1.3. Lyapunov exponents

Another measure of chaos is orbital instability, or the speed at which nearby orbits diverge. In this subsection, we focus on the larger of the two Lyapunov exponents of Ψ_τ, defined to be

$$\Lambda_{\max}(z) = \lim_{n\to\infty} \frac{1}{n} \log \|D\Psi_\tau^n(z)\| = \sup_{v\neq 0} \lim_{n\to\infty} \frac{1}{n} \log \|D\Psi_\tau^n(z) \cdot v\|.$$

Leaving technical considerations for later (see Section 3.1), we compute numerically $\Lambda_{\max}$ for the systems in question, sampling at various points $z \in U$. Notice that $\Lambda_{\max}$ measures the rate of divergence of nearby orbits *per kick*, not per unit time.

Each of the plots in Figure 2 shows $\Lambda_{\max}$ as a function of τ for the values of σ, λ and A specified. In all six plots, we have fixed $\lambda = A = 0.1$, while σ varies from plot to another. The first 4 values of σ used in Figure 2 are the same as those used in Figure 1(b). In each plot, 10 randomly chosen initial conditions are used, the largest and smallest computed values of $\Lambda_{\max}$ are discarded, and the largest and smallest of the remaining 8 values are shown, the smallest as a solid black dot, and the largest, if visibly different than the smallest, as an open square. The plots show $\tau \in [5, 15]$. We give some idea of the rates of contraction and sizes of the trapping regions U for these parameters: at $\tau = 5$, $e^{-\lambda\tau} = e^{-0.5} \approx 0.61$, and $U = \{|y| \lesssim 0.15\}$; at $\tau = 15$, $e^{-\lambda\tau} = e^{-1.5} \approx 0.22$, and $U = \{|y| \lesssim 0.03\}$.

(a) $\sigma = 0.05$ (b) $\sigma = 0.25$

(c) $\sigma = 0.5$ (d) $\sigma = 1$

(e) $\sigma = 2$ (f) $\sigma = 4$
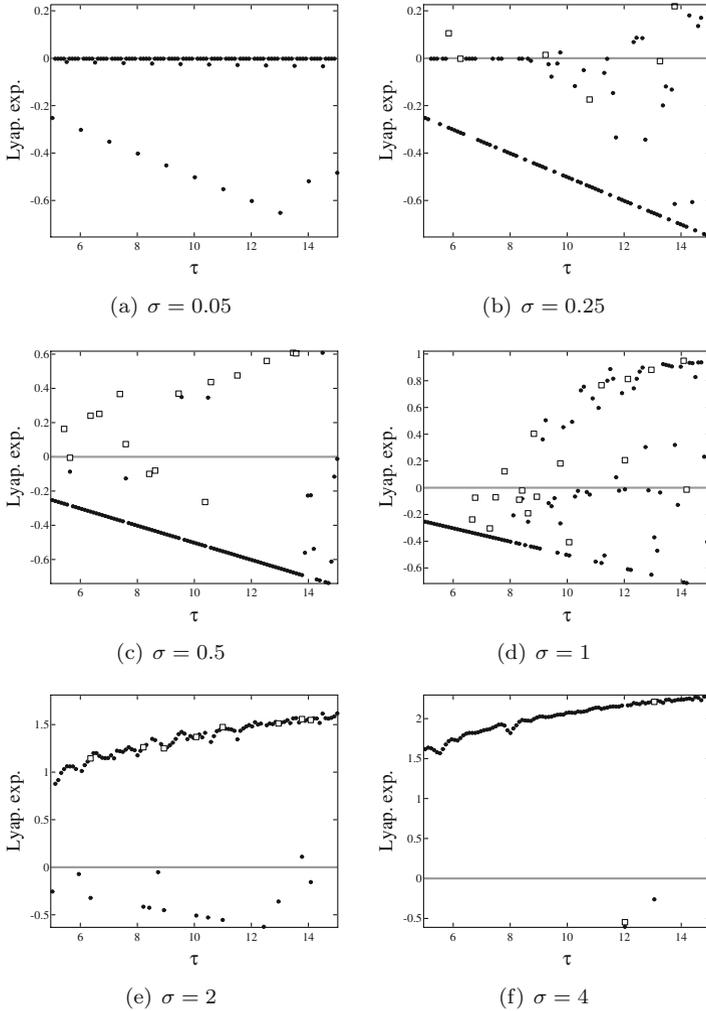
FIGURE 2. Lyapunov exponents $\Lambda_{\max}$ as functions of kick period τ for increasing shear σ. The other parameters are $\lambda = A = 0.1$. For each (σ, τ), we simulate 10 orbits with random initial conditions, iterating each for 4×10^6 steps. We then drop the outliers and plot the remaining estimates, as described in the text.

Observations from simulation results. The discussion in Section 1.2 suggests that as shear is increased with other parameters fixed, the system is likely to get increasingly chaotic. This may lead us to expect $\Lambda_{\max}$ to increase monotonically with σ. As one can see, that may be correct as an overall trend, but the situation is somewhat more complicated.

In the two low-shear regimes, namely $\sigma = 0.05$ and 0.25, with a few exceptions the computed values of $\Lambda_{\max}$ are either zero or negative, with a majority of them at or very near zero for $\sigma = 0.05$ and becoming considerably more negative at $\sigma = 0.25$. That is to say, $\Lambda_{\max}$ *decreases* as σ increases. Notice that $\Lambda_{\max} < 0$ means the trajectory tends to a *sink*, i.e., a stable fixed point or periodic orbit.

Increasing shear, we see in the middle row of Figure 2 that at first sinks dominate the landscape, giving way to more instances of positive Lyapunov exponents, i.e., chaotic behavior, as σ increases. At $\sigma = 1$, the picture is very mixed, with $\Lambda_{\max}$ fluctuating wildly between positive and negative values as τ varies. Notice also the nontrivial number of open squares, telling us that these parameters often support more than one type of dynamical behavior.

In the two higher-shear regimes, $\sigma = 2$ and 4, $\Lambda_{\max}$ becomes more solidly positive, though occasional sinks are still observed. The route has been a messy one, but one could say that the transition to chaos is complete.

As to the dependence on τ, it appears that other things being equal, longer relaxation times between kicks allow the dynamical phenomenon in effect to play out more completely: regardless of the sign of $\Lambda_{\max}$, its magnitude increases with τ in each of the plots.

Finally, it is important to remember that the limit cycles used to produce the results in Figure 2 are *weakly attracting*, making them more vulnerable to the effects of shear. Strongly attracting limit cycles are more robust, and larger kicks and/or shear will be needed to produce chaos.

In Sections 2 *and* 3 *we review some rigorous theory that supports the numerically computed values of* $\Lambda_{\max}$ *shown.* To avoid technical assumptions, we will focus on the model in Section 1.1, leaving generalizations to Section 4. As the reader will see, the mathematical ideas go considerably beyond this one example. On the other hand, *even for this simple model, state-of-the-art understanding is incomplete.* In the next two sections, we will vary σ, λ, A, and τ, and we will show that there are regions in the parameter space for which a clear description of the dynamics is available, and larger regions on which there is partial understanding.

2. Geometric Structures

To analyze a dynamical system, it is often useful to begin by identifying its most prominent structures, those that are a significant part of the landscape. Even when they do not tell the whole story, these structures will serve as points of reference from which to explore the phase space. This section describes structures of this type for the systems defined by equation (1).

2.1. Persistence of limit cycles at very low shear

Proposition 2.1 (see [29]). *Given* $\lambda\tau > 0$, *the following hold for* $\frac{\sigma}{\lambda}$ *and* A *sufficiently small:*

(a) *the attractor Γ is a smooth, closed invariant curve near γ;*
(b) *every $\zeta \in \mathbb{S}^1 \times \mathbb{R}$ lies in the strong stable curve $W_{\Psi_\tau}^{ss}(z)$ for some $z \in \Gamma$.*

Here $W_{\Psi_\tau}^{ss}(z) = \{\zeta \in \mathbb{S}^1 \times \mathbb{R} : \limsup_{n \to \infty} \frac{1}{n} \log d(\Psi_\tau^n(z), \Psi_\tau^n(\zeta)) \le -\lambda'\tau\}$ where λ' is a constant greater than $\frac{1}{2}\lambda$. Proposition 2.1 follows from standard arguments in stable and center manifolds theory; see, e.g., [10]. The idea is simple: From equation (2), one obtains

$$D\Psi_\tau(\theta, y) = \begin{pmatrix} 1 + 2\pi\frac{\sigma}{\lambda} A \cos(2\pi\theta)(1 - e^{-\lambda\tau}) & \frac{\sigma}{\lambda}(1 - e^{-\lambda\tau}) \\ e^{-\lambda\tau} 2\pi A \cos(2\pi\theta) & e^{-\lambda\tau} \end{pmatrix}. \tag{3}$$

Since invariant cones depending on $\frac{\sigma}{\lambda}$ and $\lambda\tau$ clearly exist when $A = 0$, they will persist when $\frac{\sigma}{\lambda}$ and A are small enough.

When Γ is a smooth invariant curve, the dynamics on Γ is given by the theory of circle diffeomorphisms. The situation for a smooth one-parameter family of circle diffeomorphisms $\{f_\omega\}$ can be summarized as follows (see, e.g., [7]). Let $\rho(f_\omega)$ denote the rotation number of f_ω. Then $\omega \mapsto \rho(f_\omega)$ is a *devil's staircase*, the flat parts corresponding to intervals of ω on which the rotation number is rational. Moreover, the set of ω for which $\rho(f_\omega)$ is rational is typically open and dense, while the set of ω for which $\rho(f_\omega)$ is irrational has positive Lebesgue measure. When $\rho(f_\omega) \in \mathbb{Q}$, f_ω typically has a finite number of periodic sinks and sources alternating on the circle; these aside, every orbit converges to a periodic sink. When $\rho(f_\omega) \notin \mathbb{Q}$, f_ω is topologically conjugate to an irrational rotation.

The ideas above capture the spirit of the dynamics when shear is small enough: Suppose $\Lambda_{\max}$ is computed using an initial condition $\zeta \in U$, and $\zeta \in W^{ss}(z)$ for $z \in \Gamma$. Then $\Lambda_{\max}(\zeta) = \Lambda_{\max}(z)$, and from the discussion above, the latter is either strictly negative or zero depending on whether $\rho(\Psi_\tau|_\Gamma)$ is rational or irrational.

Breaking of invariant curves. Now if we fix λ and A and increase the shear σ as is done in Figure 2, the invariant cones—and the invariant curve itself—will break.

Here is how it happens in this model for integer values of τ: For $\tau \in \mathbb{Z}^+$, $(\theta, y) = (\frac{1}{2}, 0)$ is a fixed point of Ψ_τ, and a simple computation shows that as σ increases from 0, the larger eigenvalue of this fixed point decreases from 1 to $e^{-\frac{1}{2}\lambda\tau} = \sqrt{\det(D\Psi_\tau)}$, at which time the eigenvalues turn complex. No invariant curve can exist after that. Geometrically, one can think of the breaking of the invariant curve as being due to too much "rotation" or "twist" at this fixed point.

Taking this observation a step further, one notes from equation (3) that the rotational action of $D\Psi_\tau(\theta, y)$ is strongest at $\theta = \frac{1}{2}$, where $\cos(2\pi\theta) = -1$. This suggests that for fixed σ, λ, and A, invariant curves are the most vulnerable for integer values of τ, where this strongest rotation occurs at a fixed point.

Interpreting Figures 2(a) **and** 2(b). At $\sigma = 0.05$, a majority of the $\Lambda_{\max}$ values computed are at or very slightly below zero. This is consistent with the existence of an invariant curve for those τ. One checks easily that for integer values of $\tau \leq 13$, $\Lambda_{\max} = -\frac{1}{2}\lambda\tau$ and the eigenvalues are complex. These are the first places where the invariant curve is broken as predicted.

In the plot for $\sigma = 0.25$, without pretending to account for all data points, it looks as though many are dropping off the $\Lambda_{\max} = 0$ line to join the $\Lambda_{\max} = -\frac{1}{2}\lambda\tau$ line. The only holdouts for $\Lambda_{\max} = 0$ occur for smaller τ where, as noted earlier, shear has not had enough time to act.

2.2. Increasing shear: Horseshoes and sinks

At first, mostly sinks. Figures 2(c), 2(d) suggest that at $\sigma = 0.5$, a sink with complex conjugate eigenvalues dominates the scene for much of the range of τ considered, and the same is true at $\sigma = 1$ for smaller values of τ.

For $\tau \in \mathbb{Z}^+$, this again is easily checked. The "twist" at $\theta = \frac{1}{2}$ is also eminently visible in the last three pictures of $\Psi_\tau(\gamma)$ in Figure 1(b). With a little bit of work, one can settle these questions rigorously, but an a priori fact that makes plausible the extension of this sink to noninteger values of τ is that fixed point sinks with complex conjugate eigenvalues cannot disappear suddenly as parameters are varied: a bifurcation can occur only when these eigenvalues become real, i.e., $\pm e^{-\frac{1}{2}\lambda\tau}$ (and a fixed point can vanish only when one of its eigenvalues is equal to 1).

Finally, we remark that even though the sinks above clearly exert nontrivial influence on the dynamics, other structures (competing sinks, invariant sets, etc.) may be present. The many open squares in Figure 2(c) suggest that for these parameters, trajectories in different regions of the phase space have distinct futures.

Smale's horseshoes. Horseshoes are likely present starting from σ somewhere between 0.5 and 1 for τ large enough. An example of an easily recognizable horseshoe for $\sigma = 2$ and $\tau = 10$ is shown in Figure 3(a). The larger σ, the easier it is to give examples.

Figure 3(a) illustrates how proofs of horseshoes or uniformly hyperbolic invariant sets are often done: One first "spots" a horseshoe with one's eyes, namely one or more boxes that map across themselves in a characteristic way, and then proves that the set of points that remain in these boxes forever has the required splitting into expanding and contracting directions. When $\frac{\sigma}{\lambda}$ is large, $D\Psi_\tau$ expands strongly in the θ-direction for most values of θ; contraction is guaranteed since $\det(D\Psi_\tau) = e^{-\lambda\tau} < 1$.

Proposition 2.2 (see [29]). *Given $\lambda\tau$, Ψ_τ has a horseshoe if $\frac{\sigma}{\lambda}A$ is sufficiently large.*

The presence of horseshoes is sometimes equated with dynamical complexity or chaos in the literature, and that is entirely justified insofar as one refers to the *existence* of chaotic orbits. One must not confuse the existence
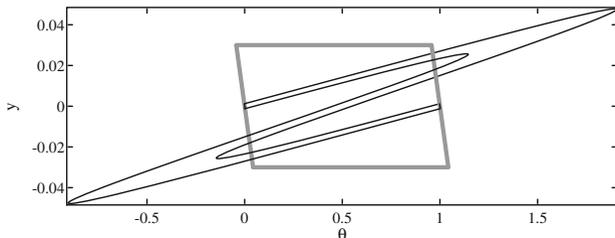
(a) Formation of a horseshoe.
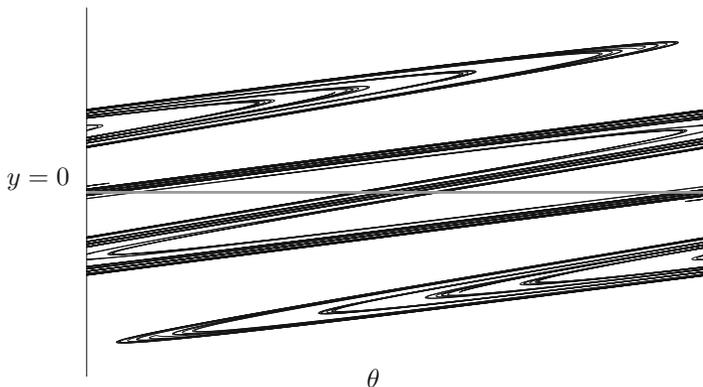
(b) The attractor Γ.

FIGURE 3. The attractor Γ and a horseshoe in it. Panel (a)
illustrates the formation of a horseshoe. Shown are a box R
(thick gray lines) and its image $\Psi_\tau(R)$ (thinner black curves).
The vertical boundaries of R are the stable manifolds of the
fixed point at $(0,0)$; these are mapped into themselves by Ψ_τ.
Panel (b) shows a picture of the attractor Γ. The parameters
are $\lambda = A = 0.1$, $\sigma = 2$, and $\tau = 10$ for both plots.

of these orbits with chaotic behavior starting from "most" or "typical" ini-
tial conditions, however: A system can have a horseshoe (which attracts a
Lebesgue measure zero set), and can have all other points in the phase space
tending to a sink. Or, the horseshoe can be part of a "strange attractor," with
$\Lambda_{\max} > 0$. The presence of a horseshoe alone does not tell us which of these
scenarios will prevail. We will say more about strange attractors *versus* sinks
in Section 3.1. It Suffices to observe here that horseshoes clearly exist for
most of the parameters in Figure 2(d)–2(f), and $\Lambda_{\max}$ is sometimes positive
and sometimes negative.

Sinks from homoclinic tangencies. For larger shear, the attractor can be quite
complicated; see Figure 3(b). Yet in Figures 2(d), 2(e), and even 2(f), sinks
can also occur as noted.

The following is purely theoretical, in the sense that we do not know exactly where the sinks are in these specific systems, but it is a general fact in two dimensions that near homoclinic tangencies of dissipative saddle fixed points ("dissipative" means $|\det(Df)| < 1$), sinks form easily, meaning one can perturb the map and find one near such a tangency; see [17]. Furthermore, tangencies persist once the stable and unstable manifolds of a horseshoe are shown to meet tangentially somewhere. While no results have been proved for this particular model, the "turns" made by unstable manifolds (see Figure 3(b)) suggest the abundance of opportunities for such tangencies.

3. A theory of strange attractors

In this section we focus on the case of *positive* Lyapunov exponents, having discussed negative and zero values of $\Lambda_{\max}$ in Section 2. A combination of geometric and statistical ideas will be used. Since these developments are more recent, we think it may be useful to include more background information: Section 3.1 discusses SRB measures for general chaotic systems. Section 3.2 surveys some recent work on a class of strange attractors called *rank-one attractors*. In certain parameter ranges, the attractors Γ in our kicked oscillator systems are of this type. In Section 3.3, we explain how the general results reviewed in Sections 3.1 and 3.2 are applied to equation (1).

3.1. SRB measures

The setting of this subsection is as follows. Let M be a Riemannian manifold or simply $\mathbb{R}^n$. We consider an open set $U \subset M$ with compact closure, and let f be a C^2 embedding of U into itself with $\overline{f(U)} \subset U$. We will refer to $\Gamma = \cap_{n \geq 0} f^n(U)$ as the *attractor* and U as its *basin of attraction*. Though not a formal assumption, we have in mind here situations that are "chaotic"; in particular, Γ is more complicated than an attracting periodic orbit.

We will adopt the viewpoint that *observable events are represented by positive Lebesgue measure sets,* and we are interested in invariant measures that reflect the properties of Lebesgue measure, which we denote by m. For chaotic systems, the only invariant measures known to have this property are SRB measures. (The terms Lebesgue and Riemannian measures will be used interchangeably in this article.)

Definition 3.1. An f-invariant Borel probability measure μ is called an SRB measure if

(a) $\Lambda_{\max} > 0$ μ-a.e., and
(b) the conditional measures of μ on local unstable manifolds have densities with respect to the Riemannian measures on these manifolds.

Recall from the Multiplicative Ergodic Theorem [18] that Lyapunov exponents, in particular $\Lambda_{\max}$, are defined μ-a.e., so (a) makes sense; in general, these quantities may vary from point to point. The meaning of (b) can be understood as follows. For an invariant measure to reflect the properties of m, it is simplest if it has a density with respect to m, but that is generally not

possible for attractors: All invariant measures in U must live on Γ, and if f is volume decreasing, which is often the case near an attractor, then $m(\Gamma) = 0$. The idea of SRB measures is that if μ cannot have a density, then the next best thing for it is to have a density in unstable directions, the intuition being that the stretching of phase space in these directions leads to a smoothing of distributions.

The main result on SRB measures is summarized in the next proposition, followed by a sketch of its proof. The ideas in the proof will explain how SRB measures, which are themselves singular, are related to Lebesgue measure. Recall that for an ergodic measure μ, $\Lambda_{\max}$ is constant μ-a.e. We will denote this number by Λ_μ.

Proposition 3.1. *Let (f, μ) be an ergodic SRB measure with no zero Lyapunov exponents. Then there is a set $V \subset U$ with $m(V) > 0$ such that the following hold for every $y \in V$:*

(i) $\Lambda_{\max}(y) = \Lambda_\mu$; and
(ii) $\frac{1}{n} \sum_{i=0}^{n-1} \varphi(f^i y) \to \int \varphi d\mu$ for every continuous observable $\varphi : U \to \mathbb{R}$.

The idea of the proof is as follows. Let γ be a piece of local unstable manifold, and let m_γ be the Riemannian measure on γ. By property (b) of Definition 3.1, we may assume m_γ-a.e. $x \in \gamma$ is "typical" with respect to μ. In particular, it has properties (i) and (ii) in the proposition. Let $W^s(x)$ be the stable manifold through x. Properties (i) and (ii) for $y \in W^s(x)$ follow from the corresponding properties for x because $d(f^n x, f^n y) \to 0$ exponentially as $n \to \infty$. It remains to show that the set of points y that are connected to μ-typical points as above has positive m-measure, and that is true by the absolute continuity of the stable foliation [21].

A little bit of history: SRB measures were invented by Sinai, Ruelle and Bowen in the 1970s, when they constructed for every attractor satisfying Smale's Axiom A [25] a special invariant measure with the properties in Definition 3.1 (see [24, 22, 4]).[1] This special invariant measure has a number of other interesting properties; see, e.g., [3, 34] for more information. At about the same time, building on Oseledec's theorem on Lyapunov exponents [18], Pesin [20] and Ruelle [23] extended the uniform theory of hyperbolic systems, also known as Axiom A theory, to an almost-everywhere theory in which positive and negative Lyapunov exponents replace the uniform expansion and contraction in Axiom A. The idea of an SRB measure was brought to this broader setting and studied there by mostly Ledrappier and Young; see, e.g., [11].

The existence problem. While the idea and relevant properties of SRB measures were shown to make sense in this larger setting, existence was not guaranteed. Indeed for an attractor outside of the Axiom A category, no matter

---

[1]Sinai treated first the case of Anosov systems; his results were shortly thereafter extended to Axiom A attractors (which are more general) first by Ruelle and then by Ruelle and Bowen.

how chaotic it appears, there is, to this day, no general theory that will tell us whether or not it has an SRB measure.

Here is where the difficulty lies: By definition, an Axiom A attractor has well-separated expanding and contracting directions that are invariant under the dynamics, so that tangent vectors in expanding directions are guaranteed some amount of growth with every iterate. In general, an attractor that appears chaotic to the eye must expand somewhere; this is how instabilities are created. But since volume is decreased, there must also be directions that are compressed. Without further assumptions, for most points x and tangent vectors v, $\|Df_x^n v\|$ will sometimes grow and sometimes shrink as a function of n. To prove the existence of an SRB measure, one must show that *on balance*, $\|Df_x^n v\|$ grows exponentially for certain coherent families of tangent vectors. The absence of cancellations between expansion and contraction is what sets Axiom A attractors apart from general chaotic attractors.

3.2. Some recent results on rank-one attractors

This subsection reviews some work by Wang and Young [28, 31, 32] on a class of strange attractors. These attractors have a single direction of instability and strong contraction in all complementary directions. Among systems without a priori separation of expanding and contracting directions (or invariant cones), this is the only class to date for which progress has been made on the existence of SRB measures.

The idea is as follows. One embeds the systems of interest in a larger collection, letting b denote an upper bound on their contraction in all but one of the directions (more precisely the second largest singular value of Df_x). One then lets $b \to 0$ in what is called the *singular limit*. If this operation results in a family of well-defined one-dimensional (1D) maps, and *if* some of these 1D maps carry strong enough expansion, then one can try to conclude that for small but positive b, some of the systems have SRB measures. Obviously, this scheme is relevant only for attractors that have a 1D character to begin with. For these attractors, what is exploited here is the fact that 1D objects, namely those in the singular limit, are more tractable than the original n-dimensional maps.

Since it is not illuminating to include all technical details in a review such as this one, we refer the reader to [31, Section 1] for a formal statement, giving only enough information here to convey the flavor of the main result.

Let $M = I \times D_{n-1}$ where I is either a finite interval or the circle $\mathbb{S}^1$ and D_{n-1} is the closed unit disk in $\mathbb{R}^{n-1}$, $n \geq 2$. Points in M are denoted by (x, y) where $x \in I$ and $y = (y^1, \ldots, y^{n-1}) \in D_{n-1}$, and I is sometimes identified with $I \times \{(0, \ldots, 0)\}$. Given $F : M \to I$, we associate two auxiliary maps:

$$F^\sharp : M \to M, \quad \text{where} \quad F^\sharp = (F, 0, \ldots, 0),$$
$$f : I \to I, \quad \text{where} \quad f = F|_{I \times \{(0, \ldots, 0)\}}.$$

We need to explain one more terminology: There is a well-known class of 1D maps called *Misiurewicz maps* [16]. Roughly speaking, a map f is in this

class if it is C^2, piecewise monotonic with nondegenerate critical points, and it satisfies the following conditions: (i) it is expanding away from $C = \{f' = 0\}$, and (ii) the forward orbit of every $\hat{x} \in C$ is trapped in an expanding invariant set (bounded away from C). Maps in this class are known to have positive Lyapunov exponents Lebesgue-a.e.

Theorem 3.1 (see [31]). *Let $F_a : M \to I$ be a 1-parameter family of C^3 maps with the following properties:*

(C1) *there exists a^* such that f_{a^*} is a Misiurewicz map;*
(C2) *$a \mapsto f_a$ satisfies a transversality condition at $a = a^*$ and $\hat{x} \in C(f_{a^*})$;*
(C3) *for every $\hat{x} \in C(f_{a^*})$, there exists j such that $\partial_{y^j} F_{a^*}(\hat{x}, 0) \neq 0$.*

Then there exists $b > 0$ (depending on $\{F_a\}$) such that if $T_a : M \to M$ is a family of C^3 embeddings of M into itself with $\|T_a - F_a^\sharp\|_{C^3} < b$, then there is a positive measure set Δ in a-space such that for all $a \in \Delta$, T_a admits an SRB measure.

That $\|T_a - F_a^\sharp\|_{C^3}$ must be sufficiently small is the *rank-one* condition discussed above, and (C1) is where we require the singular limit maps to have sufficient expansion. Notice that (C1)–(C3) all pertain to behavior at or near $a = a^*$. The set Δ will also be in the vicinity of this parameter. (C2) guarantees that one can bring about changes effectively by tuning the parameter a, and (C3) is a nondegeneracy condition at the critical points.

Remark. The existence of SRB measures is asserted for a positive measure set of parameters, and not for, say, an entire interval of a. This is a reflection of reality rather than a weakness of the result: there are parameters a arbitrarily near Δ for which T_a has sinks. In a situation such as this one where chaotic and nonchaotic regimes coexist in close proximity of one another, it is impossible to say for certain if any given map has an SRB measure. One can conclude, at best, that nearby maps have SRB measures "with positive probability."

Theorem 3.1 was preceded by the corresponding result for the Hénon family

$$T_{a,b} : (x, y) \mapsto (1 - ax^2 + y, bx). \tag{4}$$

The existence of SRB measures for parameters near $a^* = 2$ and $b \ll 1$ was proved in [2] building on results from [1]. This is the first time the existence of SRB measures was proved for genuinely nonuniformly hyperbolic attractors. Even though [1] is exclusively about equation (4), the techniques developed there were instrumental in the proof of Theorem 3.1.

Returning to the setting of Theorem 3.1, let us call $a \in \Delta$ a "good parameter" and $T = T_a$ a "good map." The following two properties of these maps are directly relevant to us. They were proved under the following additional assumption on M:

$$|\det(DT_a)| \sim b^{n-1}. \tag{$\star$}$$

(1) Lebesgue-a.e. $z \in M$ is contained in $W^s(\xi)$ where ξ is typical with respect to an ergodic SRB measure (in general, there may be more than one such measure). It follows that $\Lambda_{\max}(z) > 0$ for Lebesgue-a.e. $z \in M$.

(2) Another condition on f_{a^*} (Lyapunov exponent $> \log 2$ and $f_{a^*}^N$ mapping every interval of monotonicity to all of I for some N) implies the uniqueness of SRB measure. This in turn implies $\Lambda_{\max}$ is constant a.e. in M.

These and a number of other results for "good maps" were proved in [32]. We mention one that is not used here but sheds light on the statistical properties of these attractors: For an SRB measure μ for which (T, μ) is mixing, the system has exponential decay of correlations for Lipschitz observables, i.e., there exists $\tau \in (0, 1)$ such that for all Lipschitz φ, ψ, there exists $C = C(\varphi, \psi)$ such that for all $n \geq 1$,

$$\left| \int (\varphi \circ T^n)\psi \, d\mu - \int \varphi \, d\mu \int \psi \, d\mu \right| \leq C\tau^n.$$

3.3. Application to kicked oscillators

We now return to the model introduced in Section 1.1 and explain how this system can be fitted into the framework of the last subsection. (See [29] for details.) We fix σ, λ, A, and allow τ to vary. Writing $\tau = k + a$ where $k = [\tau]$, the integer part of τ, we let $T_{k,a} = \Psi_\tau$. For each fixed $k \in \mathbb{Z}^+$, we view $\{T_a = T_{k,a}, a \in [0, 1)\}$ as the family of interest, and we discuss if and when the conditions of Theorem 3.1 will hold for this family.

Here, the singular limit maps F_a are well defined. In fact, they are the first components of $\lim_{k \to \infty} T_{k,a}$, i.e.,

$$F_a(\theta, y) = \theta + a + \frac{\sigma}{\lambda} \cdot (y + A\sin(2\pi\theta)),$$

and the restriction of F_a to $\mathbb{S}^1$ is

$$f_a : \mathbb{S}^1 \to \mathbb{S}^1, \quad f_a(\theta) = \theta + a + \frac{\sigma}{\lambda} A\sin(2\pi\theta), \quad a \in [0, 1). \tag{5}$$

Notice immediately that the range of applicability of Theorem 3.1 is limited to $\lambda\tau$ relatively large. This is because $\|T_{k,a} - F_a^\sharp\|_{C^3} = \mathcal{O}(b)$ where $b = e^{-\lambda k}$ is required to be very small. For a given unforced system, where the amount of damping λ is fixed, this means that the kicks must be applied sufficiently far apart in time.

We comment on (C1), which along with the rank-one condition above are the core assumptions for this theorem. For our purposes let us assume that f_a satisfies the Misiurewicz condition if some iterate of f_a sends its two critical points c_1 and c_2 into an unstable periodic orbit or an expanding invariant Cantor set. First, such Cantor sets are readily available for medium size values of $\frac{\sigma}{\lambda}$ such as $\frac{\sigma}{\lambda} \geq 1$, and unstable periodic orbits start to exist for somewhat smaller values of $\frac{\sigma}{\lambda}$. Suppose for some parameter value a that the forward orbit of c_1 is contained in an expanding invariant set K. As we vary a, both the orbit of c_1 and K will move with a. Condition (C2), assuming it holds, implies that for all large enough $k \in \mathbb{Z}^+$, $f_a^k(c_1)$ moves faster, i.e.,

the path traced out by $a \mapsto f_a^k(c_1(a))$ cuts across K as though the latter was stationary. When K is a Cantor set, this guarantees that $f_a^k(c_1) \in K$ for an uncountable number of a's. By symmetry, when that happens to $f_a(c_1)$, the same is automatically true for $f_a(c_2)$. The larger $\frac{\sigma}{\lambda}$, the denser these Cantor sets are in $\mathbb{S}^1$, and the denser the set of parameters a that can be taken to be a^* in (C1).

The checking of (C2), (C3), and ($\star$) are straightforward. Conditions for the uniqueness of SRB measures require that $\frac{\sigma}{\lambda}$ be a little larger.

To summarize, the results in the last subsection imply that for $\frac{\sigma}{\lambda} \geq 1$ (or even smaller), for all large enough k, there exist positive measure sets Δ_k such that for $a \in \Delta_k$ and $\tau = k + a$, Ψ_τ is a "good map" in the sense of the last subsection. In particular, Ψ_τ has an SRB measure. We conclude also that $\Lambda_{\max}(z)$ is well defined and greater than 0 for Lebesgue-a.e. $z \in U$. To ensure that $\Lambda_{\max}(z) = $ constant a.e. (so there are no open squares in Figure 2) one needs to take $\frac{\sigma}{\lambda}$ a little larger. It is in fact not hard to see that $\Delta_k \approx \Delta_{k'}$ for $k \neq k'$ when both are sufficiently large, so that in this parameter range, the set of τ for which the properties above are enjoyed by Ψ_τ is roughly periodic with period 1.

Remarks on analytic results for chaotic systems. Theorem 3.1 is a perturbative result. As is generally the case with perturbative proofs, the sizes of the perturbations (such as b) are hard to control. Consequently, applicability of Theorem 3.1 is limited to regimes with very strong contraction. The results reported in Section 3.2, however, are the only rigorous results available at the present time. Techniques for analyzing maps in parameter ranges such as those in Figure 2 are lacking and currently quite far out of reach.

The situation here is a reflection of the general state of affairs: Due to the cancellations discussed at the end of Section 3.1, rigorous results for the large-time behavior of chaotic dynamical systems tend to be challenging.

When results such as Theorem 3.1 are available, however, they—and the ideas behind them—often shed light on situations that are technically beyond their range of applicability. Our example here is a good illustration of that: Figure 2(d)–2(f) show that as $\frac{\sigma}{\lambda}A$ increases, positive Lyapunov exponents become more abundant among the parameters tested, interspersed with occasional sinks. This is in agreement with the dynamical picture suggested by Theorem 3.1 even though with $\lambda\tau \in [0.5, 1.5]$, the contraction can hardly be considered strong.

4. Generalizations

In Sections 1, 2, and 3.3, we have focused on a concrete model. We now generalize this example in two different ways:

- the unforced equation in equation (1) is replaced by an arbitrary limit cycle;
- the specific kick in equation (1) is replaced by an arbitrary kick.

More precisely, we consider a smooth flow Φ_t on a finite-dimensional Riemannian manifold (which can be $\mathbb{R}^n$), and let γ be a *hyperbolic limit cycle*, i.e., γ is a periodic orbit of period p, and for any $x \in \gamma$, all eigenvalues for $D\Phi_p(x)$ are less than 1 aside from that in the flow direction. The *basin of attraction* of γ is the set $\mathcal{B} := \{x \in M : \Phi_t(x) \to \gamma \text{ as } t \to \infty\}$. We continue to consider forcing in the form of kicks, and assume for simplicity that the kicks are defined by a smooth embedding $\kappa : M \to M$ (as would be the case if, for example, κ represents the result of a forcing defined by $\dot{z} = G(z,t)$ where the vector field $G(z,t)$ is nonzero, or "on," for only a short time). As before, the kicks are applied periodically at times $0, \tau, 2\tau, 3\tau, \dots$, and the time evolution of the kicked system is given by $\Psi_\tau = \Phi_\tau \circ \kappa$.

A new issue that arises in this generality is that one may not be able to isolate the phenomenon, that is to say, the kicking may cause the limit cycle to interact with dynamical structures nearby. Our discussion below is limited to the case where this does *not* happen; i.e., we assume there is an open set U with $\gamma \subset U \subset \mathcal{B}$ such that $\kappa(U) \subset \mathcal{B}$ and $\Phi_\tau(\kappa(U)) \subset U$, and define $\Gamma = \cap_n \Psi_\tau^n(U)$ to be the *attractor* of the kicked system as before.

We further limit the scope of our discussion in the following two ways: (i) Kicks that are too weak will not be considered; such kicks produce invariant curves and sinks for the same reasons given in Section 2.1, and there is no need to discuss them further. (ii) We consider only regimes that exhibit a substantial contraction during the relaxation period, brought about by long enough kick intervals that permit the "shear" to act. As we will see, this is a more tractable situation. Rigorous results can be formulated—and we will indicate what is involved—but will focus primarily on ideas. Precise formulations of results in this generality (see [30]) are unfortunately not as illuminating as the phenomena behind them.

The geometry of folding: Kicks and the strong stable foliation. As we will show, the key to understanding the effects of kicks is the geometric relation between the kick and the strong stable foliation associated with the limit cycle of the unforced system. For $x \in \gamma$, we define the *strong stable manifold* of Φ_t at x, denoted $W^{ss}(x) = W^{ss}_{\Phi_t}(x)$, to be the set $W^{ss}(x) = \{y \in M : d(\Phi_t(y), \Phi_t(x)) \to 0 \text{ as } t \to \infty\}$; the distance between $\Phi_t(x)$ and $\Phi_t(y)$ in fact decreases exponentially; see, e.g., [10]. (This stable manifold is for the flow Φ_t, not to be confused with that for the kicked map Ψ_τ in Proposition 2.1.) Some basic properties of these manifolds are (i) $W^{ss}(x)$ is a codimension-one submanifold transversal to γ and meets γ at exactly one point, namely x; (ii) $\Phi_t(W^{ss}(x)) = W^{ss}(\Phi_t(x))$, and in particular, if the period of γ is p, then $\Phi_p(W^{ss}(x)) = W^{ss}(x)$; and (iii) the collection $\{W^{ss}(x), x \in \gamma\}$ foliates the basin of attraction of γ. An example of a W^{ss}-foliation for a limit cycle is shown in Figure 4.

Figure 4 shows the image of a segment γ_0 of γ under Ψ_τ. For illustration purposes, we assume γ_0 is kicked upward with its end points held fixed and assume $\tau = np$ for some $n \in \mathbb{Z}^+$ where p is the period of the cycle. Since Φ_{np} leaves each W^{ss}-manifold invariant, we may imagine that during relaxation,
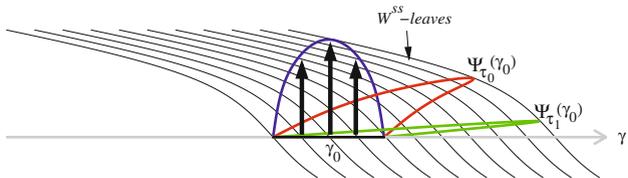
FIGURE 4. Geometry of folding in relation to the W^{ss}-foliation. Images of $\Psi_{\tau_1}(\gamma_0)$ and $\Psi_{\tau_2}(\gamma_0)$ for $\tau_1 < \tau_2$, both multiples of the period of the limit cycle γ, are shown.

the flow "slides" each point of the curve $\kappa(\gamma_0)$ back toward γ along W^{ss}-leaves; the larger n is, i.e., the more times it laps around, the farther down the point slides. In the situation depicted, the folding is quite evident. If τ is not an integer multiple of p, then Φ_τ carries each W^{ss}-manifold to another W^{ss}-manifold. Writing $\tau = np+a$, where $a \in [0, p)$, we can think of the action of Φ_τ as first sliding along the W^{ss}-manifold by an amount corresponding to Φ_{np} and then flowing forward for time a.

The picture in Figure 4 gives insight into what types of kicks are likely to produce chaos. The following observations are intended to be informal but intuitively clear.

(i) Kicks directed along W^{ss}-leaves or in directions roughly parallel to W^{ss}-leaves are not effective in producing chaos, nor are kicks that essentially carry one W^{ss}-leaf to another in an order-preserving fashion. For such kicks, Ψ_τ essentially permutes W^{ss}-leaves, and κ has to overcome the contraction within individual leaves to create chaotic behavior. (This cannot happen in two dimensions.)

(ii) The stretch-and-fold mechanism for producing chaos remains valid: the more $\Psi_\tau(\gamma)$ is folded, the more chaotic the system is likely to be, i.e., the intuition is as in Figure 1. What is different here is that unlike our earlier example, where the propensity for shear-induced chaos is determined entirely by parameters in the unforced equation, namely σ and λ, we see in this more general setting that it matters how the kick is applied. *It is the geometry of the action of κ on the limit cycle γ in relation to the strong stable foliation W^{ss} that determines the stability or chaos of the kicked system.*

(iii) The case of stronger contraction is more tractable mathematically for the following reason: When the contraction in Ψ_τ is weak, as with $\tau = \tau_1$ in Figure 4, one has to deal with the cumulative effects of multiple kicks, which are difficult to treat. When the image $\Psi_\tau(\gamma_0)$ is pressed more strongly against γ, as in the case of $\tau = \tau_2 > \tau_1$, cumulative effects of consecutive kicks are lessened.

We illustrate some of the ideas above in the examples below.

Linear shear-flow examples.

The 2D system in equation (1). The ideas in this section can be seen as

an abstraction of those discussed earlier. To understand that, we compute the W^{ss}-leaves of the unforced system in equation (1) and find them to be straight lines having slopes $-\frac{\lambda}{\sigma}$. We concluded earlier that given A, the larger $\frac{\sigma}{\lambda}$, i.e., the smaller the angle between the W^{ss}-leaves and γ, the more chaotic the system is likely to be. Observation (ii) above corroborates this conclusion: Given that we kick perpendicularly to the limit cycle (as is done in equation (1)), and points in γ are kicked to a given height, the more "horizontal" the W^{ss}-leaves, the farther the points in $\kappa(\gamma)$ will slide when we bring them back to γ. In other words, the $\frac{\sigma}{\lambda}$ part of the ratio from earlier is encoded into the geometry of the W^{ss}-foliation—provided that we kick perpendicularly to the cycle.

Generalization to n dimensions. The n-dimensional analogue of equation (1) with a more general forcing is

$$\dot{\theta} = 1 + \sigma \cdot \mathbf{y},$$

$$\dot{\mathbf{y}} = -\Lambda \mathbf{y} + AH(\theta)\mathbf{v}(\theta) \cdot \sum_{n=-\infty}^{\infty} \delta(t - n\tau), \tag{6}$$

where $\theta \in \mathbb{S}^1$, $\mathbf{y} \in \mathbb{R}^{n-1}$, $\sigma \in \mathbb{R}^{n-1}$ is nonzero, and Λ is an $(n-1) \times (n-1)$ matrix all of whose eigenvalues have strictly positive real parts. For simplicity, we assume the kicks are perpendicular to the limit cycle $\{\mathbf{y} = 0\}$, and to facilitate the discussion, we have separated the following aspects of the kick function: its amplitude is A, variation in θ is $H(\theta)$, and the direction of the kick is $\mathbf{v}(\theta) \in \mathbb{S}^{n-2}$. As a further simplification, let us assume $\mathbf{v}(\theta) \equiv \mathbf{v} \in \mathbb{R}^{n-1}$, i.e., it is a fixed vector.

A computation shows that the W^{ss}-manifolds of the unforced equation are given by

$$W^{ss}(\theta_0, \mathbf{0}) = \{(\theta, \mathbf{y}) : \theta = \theta_0 - \sigma^T \Lambda^{-1} \mathbf{y}\},$$

i.e., they are hyperplanes orthogonal to the covector $(1, \sigma^T \Lambda^{-1})$. As noted in observation (i) above, kick components orthogonal to $(1, \sigma^T \Lambda^{-1})$ are "dissipated" and do not have much effect. If $H \equiv$ constant, then Ψ_τ simply permutes the W^{ss}-planes and again no chaotic behavior will ensue. To produce horseshoes and strange attractors, a sufficient amount of *variation* in θ for $\Psi_\tau(\gamma)$ is needed as noted in observation (ii) above; that variation must come from H. An analysis similar to that in equation (2), Section 1.2, tells us that for large τ, the amount by which the kick is magnified in the θ-direction is $\approx A\, H(\theta)\, \sigma^T \Lambda^{-1} \mathbf{v}$. We remark that the variation in H is far more important than its mean value, which need not be 0.

Finally, given $H \not\equiv$ constant, to maximize the variation of $\Psi_\tau(\gamma)$ in θ for large τ, the discussion above suggests kicking in a direction $\mathbf{v}$ that maximizes $\sigma^T \Lambda^{-1} \mathbf{v}$. Under the conditions above, this direction is unique and is given by $\mathbf{v} = \pm(\mathbf{\Lambda^T})^{-1}\sigma/|(\mathbf{\Lambda^T})^{-1}\sigma|$. Notice that this need not be the direction with the least damping or the direction with maximal shear, but one that optimizes the combined effect of both.

On analytic proofs. When Ψ_τ contracts strongly enough, the system falls

into the *rank-one* category as defined in Section 3.2, independently of the dimension of the phase space. This comes from the fact that Φ_t carries all points back to γ, which is a 1D object.

If the folding (as described in observation (ii) above) is significant enough for the amount of contraction present, then horseshoes can be shown to exist. Proving the existence of horseshoes with one unstable direction is generally not very difficult, and not a great deal of contraction is needed.

To prove the existence of strange attractors or SRB measures, the results in Section 3.2 are as applicable here as in our 2D linear example. It is proved in [30] that the periodic kicking of arbitrary limit cycles fits the general framework of Theorem 3.1, in the sense that as the time between kicks tends to infinity, singular limit maps are well defined. They are given by $f_a : \gamma \to \gamma$, $a \in [0, p)$, where

$$f_a(x) := \lim_{n \to \infty} \Phi_{np+a}(\kappa(x)) \quad \text{for all } x \in \gamma.$$

From our earlier discussion of sliding along W^{ss}-leaves, it is not hard to see that $f_a(x)$ is, in fact, the unique point $y \in \gamma$ such that $\kappa(x) \in W^{ss}(y)$. Whether (C1)–(C3) hold depends on the system in question and hinges mostly on (C1), which usually holds when the variation is large enough. As always, these conditions need to be verified from example to example.

Related results and outlook

We have reviewed a set of results on the periodic kicking of limit cycles. The main message is that the effect of the kick can be magnified by the underlying shear in the unforced system to create an unexpected amount of dynamical complexity. It is an example of a phenomenon known as *shear-induced chaos.*

A similar geometric mechanism is used to prove the existence of strange attractors in (a) certain examples of slow-fast systems [8]; (b) periodic kicking of systems undergoing supercritical Hopf bifurcations (see [30] for details, and [15] for results applicable to evolutionary PDEs); and (c) periodic forcing of near-homoclinic loops [27]. See also [19]. All of these results pertain to strong-contraction regimes; proofs are perturbative and rely on the theory of rank-one attractors reviewed in Section 3.2.

A welcome extension of the results reviewed here is to remove the strong-contraction assumption for strange attractors, but this is likely to be challenging: one has to either develop nonperturbative techniques or go about the problem in a less direct way.

Random forcing is a future direction we believe to be both interesting and promising. Numerical studies of Poisson and white-noise forcing have been carried out [14]. Phenomena similar to those in Section 1 are observed when the kick term in equation (1) is replaced by a term of the form $A \sin(2\pi\theta)\, dB_t$ where B_t is standard Brownian motion, i.e., the setup is a stochastic differential equation. With stochastic forcing, phase space geometry is messier, but $\Lambda_{\max}$ depends continuously on parameters. The absence of wild fluctuations between positive and negative values of $\Lambda_{\max}$ (corresponding

respectively to strange attractors and sinks in the periodic case) gives hope to the idea that the analysis for stochastic forcing may be more tractable than that for periodically forced systems.

References

[1] M. Benedicks and L. Carleson, *The dynamics of the Hénon map.* Ann. of Math. (2) **133** (1991), 73–169.

[2] M. Benedicks and L.-S. Young, *Sinai-Bowen-Ruelle measures for certain Henon maps.* Invent. Math. **112** (1993), 541–576.

[3] R. Bowen, *Equilibrium States and the Ergodic Theory of Anosov Diffeomorphisms.* Lecture Notes in Mathematics 470, Springer-Verlag, New York, 1975.

[4] R. Bowen and D. Ruelle, *The ergodic theory of Axiom A flows.* Invent. Math. **29** (1975), 181–202.

[5] M. L. Cartwright and J. E. Littlewood, *On nonlinear differential equations of the second order.* J. London Math. Soc. **20** (1945), 180–189.

[6] R. FitzHugh, *Impulses and physiological states in theoretical models of nerve membrane.* Biophys. J. **1** (1961), 445–466.

[7] J. Guckenheimer and P. Holmes, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields.* Applied Mathematical Sciences 42, Springer-Verlag, New York, 1983.

[8] J. Guckenheimer, M. Wechselberger and L.-S. Young, *Chaotic attractors of relaxation oscillators.* Nonlinearity **19** (2006), 701–720.

[9] R. Haiduc, *Horseshoes in the forced van der Pol system.* Nonlinearity **22** (2009), 213–237.

[10] M. W. Hirsch, C. C. Pugh and M. Shub, *Invariant Manifolds.* Lecture Notes in Mathematics 583, Springer-Verlag, New York, 1977.

[11] F. Ledrappier and L.-S. Young, *The metric entropy of diffeomorphisms.* Ann. of Math. (2) **122** (1985), 509–574.

[12] M. Levi, *Qualitative analysis of the periodically forced relaxation oscillations.* Mem. Amer. Math. Soc. **32** (1981), 1–147.

[13] N. Levinson, *A second order differential equation with singular solutions.* Ann. of Math. (2) **50** (1949), 127–153.

[14] K. K. Lin and L.-S. Young, *Shear-induced chaos.* Nonlinearity **21** (2008), 899–922.

[15] K. Lu, Q. Wang and L.-S. Young, *Strange attractors for periodically forced parabolic equations.* Preprint.

[16] M. Misiurewciz, *Absolutely continuous measures for certain maps of an interval.* Publ. Math. Inst. Hautes Études Sci. **53** (1981), 17–51.

[17] S. Newhouse, *Lectures on dynamical systems.* In: Dynamical Systems (C.I.M.E. Summer School, Bressanone, 1978), Progr. Math. 8, Birkhäuser Boston, Mass., 1980, 1–114.

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

[19] W. Ott and M. Stenlund, *From limit cycles to strange attractors.* Comm. Math. Phys. **296** (2010), 215–249.

[20] Ya. B. Pesin, *Characteristic Lyapunov exponents and smooth ergodic theory.* Russian Math. Surveys **32** (1977), 55–114.

[21] C. Pugh and M. Shub, *Ergodic attractors.* Trans. Amer. Math. Soc. **312** (1989), 1–54.

[22] D. Ruelle, *A measure associated with Axiom-A attractors.* Amer. J. Math. **98** (1976), 619–654.

[23] D. Ruelle, *Ergodic theory of differentiable dynamical systems.* Publ. Math. Inst. Hautes Études Sci. **50** (1979), 27–58.

[24] Ya. G. Sinai, *Gibbs measure in ergodic theory.* Russian Math. Surveys **27** (1972), 21–69.

[25] S. Smale, *Differentiable dynamical systems.* Bull. Amer. Math. Soc. **73** (1967), 747–817.

[26] B. van der Pol and J. van der Mark, *Frequency demultiplication.* Nature **120** (1927), 363–364.

[27] Q. Wang and W. Ott, *Dissipative homoclinic loops and rank one chaos.* Preprint.

[28] Q. Wang and L.-S. Young, *Strange attractors with one direction of instability.* Comm. Math. Phys. **218** (2001), 1–97.

[29] Q. Wang and L.-S. Young, *From invariant curves to strange attractors.* Comm. Math. Phys. **225** (2002), 275–304.

[30] Q. Wang and L.-S. Young, *Strange attractors in periodically-kicked limit cycles and Hopf bifurcations.* Comm. Math. Phys. **240** (2003), 509–529.

[31] Q. Wang and L.-S. Young, *Toward a theory of rank one attractors.* Ann. of Math. (2) **167** (2008), 349–480.

[32] Q. Wang and L.-S. Young, *Dynamical profile of a class of rank one attractors* Preprint.

[33] A. Winfree, *The Geometry of Biological Time.* 2nd ed., Springer-Verlag, New York, 2000.

[34] L.-S. Young, *What are SRB measures, and which dynamical systems have them?* J. Statist. Phys. **108** (2002), 733–754.

[35] G. Zaslavsky, *The simplest case of a strange attractor.* Phys. Lett. A **69** (1978), 145–147.

Kevin K. Lin
Department of Mathematics and Program in Applied Mathematics
University of Arizona, AZ 85721, USA
e-mail: klin@math.arizona.edu

Lai-Sang Young
Courant Institute of Mathematical Sciences
New York University, NY 10012, USA
e-mail: lsy@cims.nyu.edu