Birkhoff Averages and Rotational Invariant Circles for Area-Preserving Maps

E. Sander and J.D. Meiss [*]

Department of Mathematical Sciences
George Mason University
Fairfax, VA 22030, USA
esander@gmu.edu

Department of Applied Mathematics
University of Colorado
Boulder, CO 80309-0526
James.Meiss@colorado.edu

May 19, 2020

Abstract

Rotational invariant circles of area-preserving maps are an important and well-studied example of KAM tori. John Greene conjectured that the locally most robust rotational circles have rotation numbers that are noble, i.e., have continued fractions with a tail of ones, and that, of these circles, the most robust has golden mean rotation number. The accurate numerical confirmation of these conjectures relies on the map having a time-reversal symmetry, and such high accuracy has not been obtained in more general maps. In this paper, we develop a method based on a weighted Birkhoff average for identifying chaotic orbits, island chains, and rotational invariant circles that does not rely on these symmetries. We use Chirikov's standard map as our test case, and also demonstrate that our methods apply to three other, well-studied cases.

1 Introduction

The dynamics of an integrable Hamiltonian or volume-preserving system is organized by periodic and quasi-periodic motion on invariant tori. When such a system is smoothly perturbed, Kolmogorov-Arnold-Moser (KAM) theory [dlL01] implies that some of these tori persist and some are replaced by isolated periodic orbits, islands, or chaotic regions. On each KAM torus, the dynamics is conjugate to a rigid rotation with some fixed frequency vector. Typically, as the perturbation grows the proportion of chaotic orbits increases and more of the tori are destroyed. Invariant tori can be found numerically by taking limits of periodic orbits [Gre79] and by iterative methods based on the conjugacy to rotation [HdlL06, HdlLS12, HCF$^+$16]. In these methods, one fixes a frequency vector and attempts to find invariant sets on which the dynamics has this frequency. In this paper

---

we explore an alternative technique, based on windowed Birkhoff averages [DSSY16], to distinguish between chaotic, resonant, and quasiperiodic dynamics. Since we do not fix the rotation vector in advance, this method permits us to accurately compute the rotation vector for each initial condition that lies on a regular orbit. As such the method is analogous to Laskar's frequency analysis [LFC92, BBGT96], which uses a windowed Fourier transform to compute rotation numbers. Notably there have been improvements to this technique using collocation [GMS10a] or Richardson extrapolation [LV14].

As an illustrative example, we will primarily study Chirikov's standard map [Chi79], though in the last section we will consider several generalizations. Two-dimensional, area-preserving maps are simplest, nontrivial case of Hamiltonian dynamics (for a review, see [Mei92]). Letting $f : M \to M$, where $M = \mathbb{T} \times \mathbb{R}$, the cylinder, the standard map can be written as $(x_{t+1}, y_{t+1}) = f(x_t, y_t) = f^t(x_0, y_0)$ with

$$
\begin{aligned}
x_{t+1} &= x_t + \Omega(y_{t+1}) \quad \mod 1, \\
y_{t+1} &= y_t + F(x_t).
\end{aligned}
\tag{1}
$$

For Chirikov's case, the "frequency map" and "force" are given by

$$
\Omega(y) = y, \quad F(x) = -\frac{k}{2\pi} \sin(2\pi x),
$$

respectively. When the parameter $k = 0$, the action y is constant, and every orbit lies on a rotational invariant circle with rotation number $\omega = \Omega(y)$. When ω is irrational the orbit is dense on the circle, and the dynamics is conjugate to the quasiperiodic, rigid rotation

$$
\theta \to \theta + \omega \quad \mod 1
\tag{2}
$$

for $\theta \in \mathbb{T}$, under the trivial conjugacy $(x, y) = C(\theta) = (\theta, \omega)$.

As k increases, some of these rotational invariant circles persist, as predicted by KAM theory, but those with rational or "near" rational rotation numbers are destroyed. On each KAM circle, the dynamics is still conjugate to (2), for some irrational ω, under a smooth map $C : \mathbb{T} \to \mathbb{T} \times \mathbb{R}$. As an example, Fig. 1 depicts the dynamics for the Chirikov map for $k = 0.7$. In the top row, we distinguish between non-chaotic and chaotic dynamics, and in the bottom row we distinguish between two types of non-chaotic behavior, namely island chains and rotational invariant circles. The methods for doing this will be discussed in §2-3.

An orbit $\{(x_t, y_t) : t \in \mathbb{Z}\}$ has a rotation number ω if the limit

$$
\omega = \lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} \Omega(y_t)
\tag{3}
$$

exists. Of course, if $k = 0$, ω is simply the value of Ω on the conserved action. If an orbit is periodic, say $(x_n, y_n) = (x_0, y_0) + (m, 0)$ for some integers m, n, then $\omega = \frac{m}{n}$ is rational. Indeed, (1) implies that if we lift x to $\mathbb{R}$, then

$$
x_T - x_0 = \sum_{t=1}^{T} \Omega(y_t)
$$

so for the periodic case $\omega = (x_n - x_0)/n$. Note that ω, as a rotation number, is measured with respect to rotation in x. For an invariant circle within an island chain, the effect of the rotation of the orbit about the island center will average out, and ω will equal m/n, the value for the periodic orbit it encloses. This can be seen in the lower left portion of Fig. 1, where each elliptic island has a single solid color due to having the same value of ω. In particular, the rotational invariant circles are the only non-chaotic orbits with the property that ω is irrational. In §3, we develop a numerical method to determine whether a floating point number is (with high probability) rational or irrational. With this method, we are able to use the rotation number computed with the weighted Birkhoff average to distinguish between rotational and non-rotational invariant circles.

The invariant circles that persist by KAM theory have Diophantine rotation numbers, i.e., there is a $\tau \geq 1$ and a $c > 0$ such that

$$|n\omega - m| > \frac{c}{|n|^\tau}, \quad \forall n \in \mathbb{N}, \, m \in \mathbb{Z}. \tag{4}$$

Such rotation numbers are hard to approximate by rationals (see §3). An invariant circle is said to be locally robust if it has a neighborhood in M in which it is the last invariant circle; i.e., it exists for $0 \leq k \leq k_{cr}(\omega)$ and k_{cr} is a local maximum. It is known from careful numerical studies that invariant circles with "noble" rotation numbers (their continued fractions have an infinite tail of ones) are robust [Gre79, Mac93]. Since these continued fractions are asymptotically periodic, these rotation numbers are quadratic irrationals and satisfy (4) with $\tau = 1$.

John Greene discovered that the last rotational circle of (1) has rotation number given by the golden mean γ,[1] and that it is destroyed at $k = k_{cr}(\gamma) \approx 0.971635$ [Gre79]. It has been proven that the golden circle persists up to $k = .9716$ [FHL17] and that there are no invariant circles for $k > 63/64$ [MP85] and when $k = 0.9718$ [Jun91]. Greene used limits of periodic orbits to find invariant circles, and his method depends on accurate computations of these orbits. Such high accuracy can be obtained because Chirikov's map has a time-reversal symmetry, and every rotational circle is symmetric. Moreover, there are fixed lines of the symmetry (e.g., $x = 0$) that contain symmetric periodic orbits of all rotation numbers [Mac93]. This allows the computation of orbits of high periods, and implies that they alone can be used to approximate the invariant circles. For example, using symmetric orbits up to periods of order 10^5, Haro improved Greene's estimate in his 1998 PhD thesis to obtain $k_{cr} \approx 0.97163540324$ [HCF$^+$16]. Interestingly, conjugacy-based methods can be used to confirm the first 7 digits of k_{cr} using a Fourier series with $524,288$ terms [HCF$^+$16].

The average (3) need not exist for orbits that are neither periodic nor quasiperiodic. For example if an orbit is heteroclinic between two periodic orbits with different rotation numbers, the forward and backward time averages of $\Omega(y)$ will be different. Moreover, when k is large enough, y can be unbounded,[2] and the limit (3) need not even converge. However, if an orbit ergodically covers a bounded region, then Birkhoff's ergodic theorem implies that the time average of Ω does exist.

More generally a finite-time Birkhoff average on a orbit of a map f beginning at a point $z \in M$

---

[1] Or any integer shift of this value by a discrete symmetry of (1).

[2] Often one thinks of y as diffusing in this case, but it can also grow linearly in time due to "accelerator modes" [Chi79]).
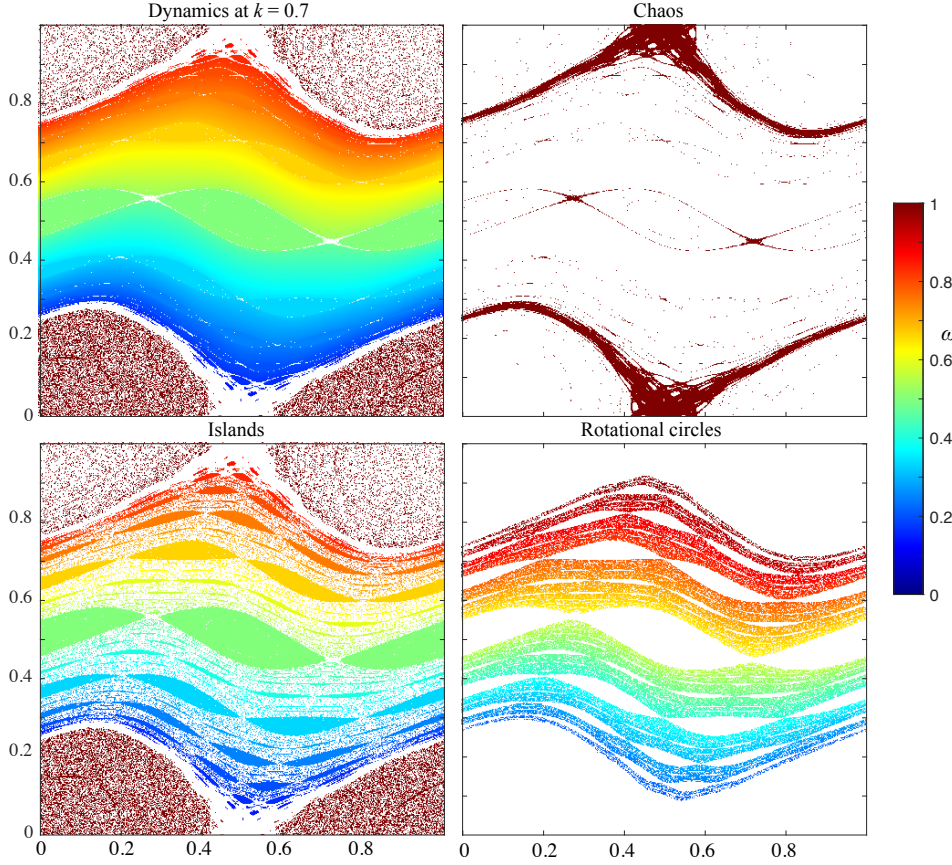
Figure 1: The dynamics of the standard map for $k = 0.7$. Using the weighted Birkhoff method we are able to distinguish chaotic orbits (upper right), islands (lower left), and rotational circles (lower right). The rotation number of each nonchaotic orbit is color-coded (color bar at right). The computations were performed for an evenly spaced grid of 1000^2 points in $[0,1]^2$ with $T = 10^4$, using the distinguishing criteria in (24).

for any function $h : M \to \mathbb{R}$ is given by

$$B_T(h)(z) = \frac{1}{T} \sum_{t=0}^{T-1} h \circ f^t(z). \tag{5}$$

This average need not converge rapidly. Even if the orbit lies on a smooth invariant circle with irrational rotation number, the convergence rate of (5) is $\mathcal{O}(T^{-1})$, due to edge effects at the two ends of the finite orbit segment. By contrast, for the chaotic case, the convergence rate of (5) is observed to be $\mathcal{O}(T^{-1/2})$, in essence as implied by the central limit theorem [LM10].

We can significantly improve the convergence of a Birkhoff sum on a quasiperiodic set by using the method of weighted Birkhoff averages developed in [DSSY16, DDS$^+$16, DSSY17], see §2.1. If the map f, the function h, and z belongs to a C^∞ quasiperiodic set, and the rotation number is Diophantine, this method is superconvergent, meaning that the error decreases faster than any power of T [DSSY17, DY18]. Notice that the weighted Birkhoff method does not speed up the

convergence rate on chaotic sets since these lack smoothness. Therefore weighted Birkhoff averages have two distinct uses: (a) to distinguish chaotic from regular dynamics, and (b) to give a high precision computation of the rotation number.

We comment here on the relationship between the current paper and the previous papers [DSSY16, DDS$^+$16, DSSY17, DY18] on the weighted Birkhoff average mentioned above. Only [DDS$^+$16] and [DSSY17] use (1) as an example; the others focused on other examples and details of the numerical method. In particular, in [DDS$^+$16] the weighted Birkhoff average is first shown to distinguish between regular and chaotic orbits for (1) for one fixed parameter value. Both papers [DDS$^+$16, DSSY17] calculate the rotation number for a single, non-rotational invariant circle inside an island. The papers also present high precision computations of the Fourier series expansion of the conjugacy map for the invariant circle. We do not make use of the conjugacy map in the current paper. None of the previous papers contain a comparison to other methods for distinguishing between chaos and regularity, nor do they contain any calculations for more than one parameter value, and none contain a method distinguishing between rotational orbits and islands.

Other methods for accurate computations of rotation numbers include frequency analysis, which is based on finding the Fourier amplitudes of a quasiperiodic signal using a Hanning window [LFC92]; this has an error that decreases as a power of T. Frequency analysis can be improved by collocation techniques to more accurately find the peaks of the discrete Fourier transform [GMS10a, GMS10b]. Another technique uses a recurrence method based on continued fraction expansions [EV01]. Slater's criterion [Sla50, Sla67, May88, ACP06, ZTRK07] can be used to compute whether an orbit satisfies the same ordering as an irrational rotation; this method can be used to estimate $k_{cr} = 0.9716394$, slightly above Greene's value [AC15]. A conjugacy-based Fourier method for finding Diophantine rotation numbers was applied to circle [SV06] and planar maps [LV09]. This method can also be extended to accurately compute frequency vectors for orbits on higher-dimensional tori in Hamiltonian flows [LV14].

Computations of rotation number, or more generally of frequency vectors have many applications. For example, a finite-time computation of the rotation number has been used to define coherent structures by considering ridges in the finite time sum (3) [SSC$^+$13] . This method also can distinguish between trapped and escaping orbits [SMS$^+$19] by monitoring the gradient of (3) with respect to initial condition, and to determine the break-up of circles in nontwist maps [SMS$^+$18].

Our paper proceeds as follows: We start in §2 with a description of the weighted Birkhoff method in §2.1. In §2.2, we review the two standard methods for distinguishing between regular and chaotic orbits, namely Lyapunov exponents and the 0–1 test of Gottwald-Melbourne. In §2.3 we compare the three methods for distinguishing chaos from regularity for the Chirikov standard map. In §3 we describe how to use the weighted Birkhoff average for non-chaotic orbits to distinguish between rotational circles and island chains. In §4, after removing chaotic orbits and island chains, we are left with the rotational circles. We are able to create the critical function diagram, and describe the number theoretic properties of the rotation numbers for rotational circles, showing that their behavior does not match that of randomly chosen irrational numbers. In §5, we apply our methods to three generalizations of the standard map, namely a symmetric two-harmonic generalized standard map, a standard non-twist map, and an asymmetric two-harmonic map. We conclude in §6 with comments on how these methods can be applied to other maps.

2 Distinguishing chaos and regularity

In this section, we introduce the weighted Birkhoff method, and we compare it to two different methods for distinguishing chaos from regular dynamics, namely Lyapunov exponents and the 0–1 test of Gottwald and Melbourne [GM09].

2.1 The weighted Birkhoff average

We now describe in more detail the method of weighted Birkhoff averages [DSSY16, DDS+16, DSSY17]. Since the source of error in the calculation of a time average for a quasiperiodic set occurs due to the lack of smoothness at the ends of the orbit, we use a windowing method similar to the methods used in signal processing. Let

$$g(t) \equiv \begin{cases} e^{-[t(1-t)]^{-1}} & t \in (0,1) \\ 0 & t \leq 0 \text{ or } t \geq 1 \end{cases} ,$$

be an exponential bump function that converges to zero with infinite smoothness at 0 and 1, i.e., $g^{(k)}(0) = g^{(k)}(1) = 0$ for all $k \in \mathbb{N}$. To estimate the Birkhoff average of a function $h : M \to \mathbb{R}$ efficiently and accurately for a length T segment of an orbit, we modify (5) to compute

$$WB_T(h)(z) = \sum_{t=0}^{T-1} w_{t,T} h \circ f^t(z) , \tag{6}$$

where

$$w_{t,T} = \frac{1}{S} g\left(\tfrac{t}{T}\right), \quad S = \sum_{t=0}^{T-1} g\left(\tfrac{t}{T}\right) . \tag{7}$$

That is, the weights w are chosen to be normalized and evenly spaced values along the curve $g(t)$. For a quasiperiodic orbit, the infinitely smooth convergence of g to the zero function at the edges of the definition interval preserves the smoothness of the original orbit. Indeed it was shown in [DY18] that given a C^∞ map f, a quasiperiodic orbit $\{f^t(z)\}$ with Diophantine rotation number, and a C^∞ function h, it follows that (6) is super-convergent: there are constants c_n, such that for all $n \in \mathbb{N}$

$$\left| WB_T(h)(z) - \lim_{N \to \infty} B_N(h)(z) \right| < c_n T^{-n}. \tag{8}$$

Several papers [GMS10a, LFC92, LV14] include a similar method to compute frequencies with a $\sin^2(\pi s)$ function instead of a bump function, but this function is fourth order smooth rather than infinitely smooth at the two ends, implying that the method converges as $\mathcal{O}(T^{-4})$, see e.g., [DSSY17, Fig. 7]. In addition to converging more rapidly, the weighted Birkhoff average is relatively straightforward to implement. By contrast to the case of regular orbits, when an orbit is chaotic (i.e., has positive Lyapunov exponents), then (6) typically converges much more slowly; in general it converges no more rapidly than the unweighted average of a random signal, i.e., with an error $\mathcal{O}(T^{-1/2})$ [LM10, DSSY17].

A graph of the error in WB_T for $h = \cos(2\pi x)$ as a function of the number of iterates T is shown on the left panel of Fig. 2. Here we have chosen 50 orbits of (1) for the parameter $k = 0.7$

with initial condition $x = 0.45$, and y evenly spaced between 0 and 0.5. For orbits that are independently identified as chaotic (red), WB_T essentially does not decrease with T; however, for orbits that lie on rotational (blue) or island (green) invariant circles, the error for all but three has decreased to machine precision, 10^{-15}, when T reaches 10^4. Further, right panel of Fig. 2 shows the convergence rate as a function of y, this time for 1000 orbits. Note that there is no evidence of superconvergence in Fig. 2: the convergence rate for (8) has $n = 2 - 5.5$. Indeed, superconvergence was only observed in [DSSY17] when extended precision computations were done. Nevertheless, there is a clear distinction between chaotic and regular orbits even for T as small as 10^3. Note that the Richardson extrapolation method used in [LV14] can also compute frequencies to 12-digit accuracy using 2^{16} iterates of a Poincaré map.
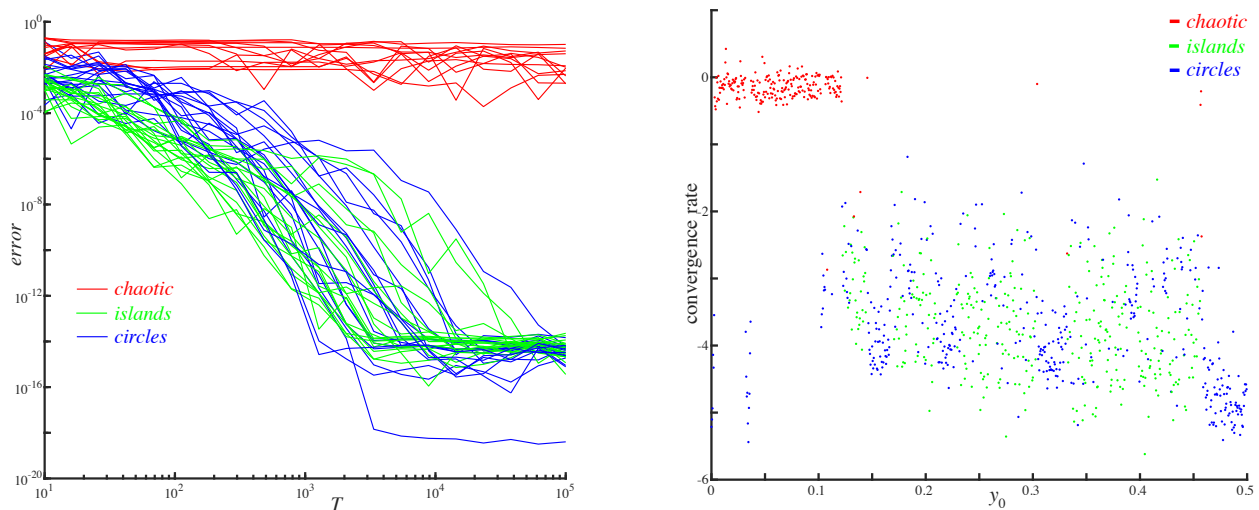

Figure 2: Convergence of the weighted Birkhoff average (6) for orbits of the standard map at $k = 0.7$ for the function $h = \cos(2\pi x)$. On left, the error of the computation is shown as a function of the number of iterates T for 50 initial conditions at $x = 0.45$ with y on a grid in $[0, 0.5]$. On the right, the convergence rate is shown for 1000 initial conditions at the same x and k values, where convergence rate was calculated using the errors before the values flattened out due to floating point errors, measured by where they have dropped below 10^{-13}. In each case, the values are compared with $WB_T(\Omega)$ at $T = 10^5$. Using the distinguishing criteria in (24), the red curves are identified as chaotic, the green curves as islands, and the remaining blue curves are thus the rotational invariant circles.

To distinguish chaotic from regular dynamics, we compute (6) for two segments of an orbit, using iterates $\{1, \ldots, T\}$ and $\{T + 1, \ldots, 2T\}$. In the limit $T \to \infty$, these values should be the same. Therefore we can measure convergence rate by comparing them. In order to distinguish chaotic sets, we compute the number of consistent digits beyond the decimal point in our two approximations of $WB(h)$, which is given by

$$dig_T = -\log_{10} \left| WB_T(h)(z) - WB_T(h)(f^T(z)) \right| . \tag{9}$$

If dig_T is relatively large, then the convergence is fast, meaning the orbit is regular. If dig_T is small, then the convergence is slow, meaning the orbit is chaotic. We will call $2T$, the "total orbit length" as it is the total number of iterates needed to compute (9).

Three examples are shown in Fig. 3(a) for a set of 1000 initial points on a vertical line segment at $x = 0.321$ for three different values of k. For the smallest parameter, $k = 1.0$, a substantial fraction of the orbits are regular, and these have a distribution of dig_{10^4} centered around 14, nearing the maximum possible for a double precision computation. By contrast, when $k = 2.0$ there are only chaotic orbits in the sample, and these have a distribution of dig_{10^4} centered around 2. Note that when $k = 1.0$ there are also orbits with $dig_T \in [6, 13]$, and which seem to represent orbits trapped in islands that are either oscillatory invariant circles or weakly chaotic orbits between a pair of such invariant circles.

In order to determine the cutoff in dig_T between regular and chaotic orbits, we computed a histogram (not shown) of dig_{10^4} for the Chirikov standard map for $500,000$ different starting points: a grid of 500 k-values between 0.1 and 2.5, with 1000 distinct initial conditions for each. This histogram has two large peaks, one at around 2 and the second around 15 (corresponding to the machine epsilon value). As is consistent with the case $k = 1.0$ shown in Fig. 3(a), the lowest probability occurs around $dig_T = 5$. In our calculations of chaos, we wish to err on the side of false positives of chaos, and thus we use a value of 5.5 as our cutoff value to distinguish whether orbits exhibit regular or chaotic behavior.
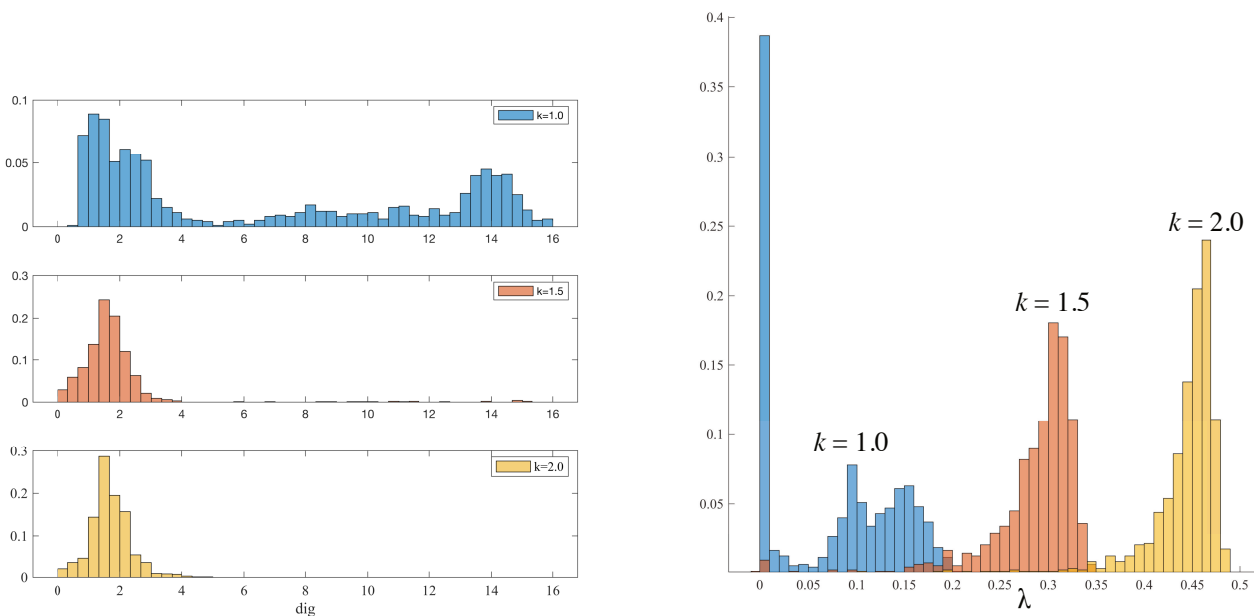

Figure 3: Histograms of (a) Weighted Birkhoff accuracy, dig_T, and (b) finite-time Lyapunov exponent, λ_T, for orbits of the standard map with $k = 1.0$, 1.5 and 2.0. The initial conditions are $(0.321, y)$ with 1000 values of y on a uniform grid in $[0, 1]$. (a) Histograms of dig_T (9), for $h(x, y) = \cos(2\pi x)$ and $T = 10^4$. (b) Histograms of λ_T (10), for $T = 2(10)^4$, and $v = (0, 1)^T$.

Using this cutoff, the putative set of regular orbits with initial conditions along three vertical line segments, $x = 0.0$, 0.321 and 0.5 are shown in Fig. 4 for k ranging from 0.1 to 2.5. When $x = 0.0$, the figure is dominated by the regular region around the fixed point islands surrounding $(0, 0)$ and $(0, 1)$–these points are elliptic up to $k = 4.0$. Other islands can also be seen; for example

for $x = 0.0$ and for $x = 0.5$, we can see the period two orbit $(0, \frac{1}{2}) \mapsto (\frac{1}{2}, \frac{1}{2})$, which is elliptic until $k = 2.0$, where it period doubles. By contrast the line $x = 0.321$ intersects fewer islands, and there appear to be no regular orbits when $k \geq 1.915$ with initial conditions on this line.
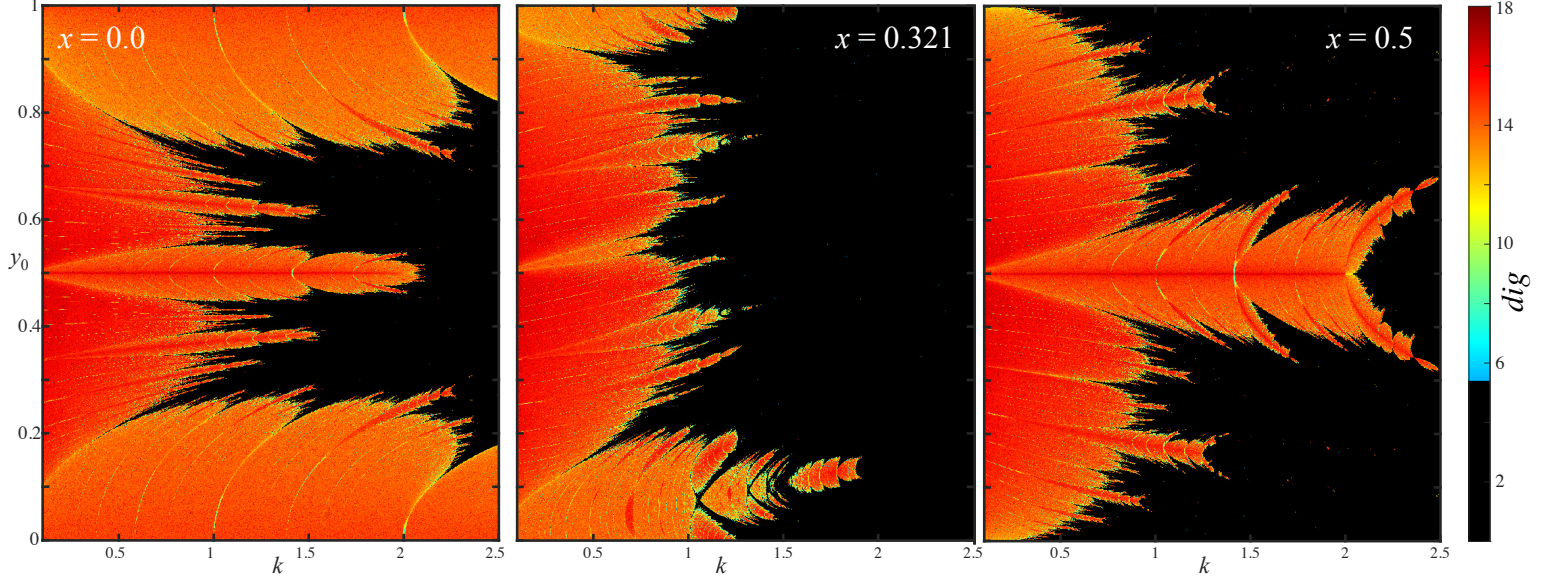

Figure 4: Using the weighted Birkhoff method, this figure shows the number of digits (9) for orbits of the standard map (1) with (a) $x_0 = 0$ and (b) $x_0 = 0.321$ and (c) $x_0 = 0.5$ for $y_0 \in [0, 1]$ and $k \in [0.1, 2.5]$. Here dig_T is computed by computing the average (6) for the function $h(x, y) = \cos(2\pi x)$ for $T = 2(10)^4$ steps. Initial conditions with $dig_T < 5.5$ are colored black. The value of dig_T for the regular orbits is indicated in the color bar.

2.2 Lyapunov Exponent and the 0–1 Test

In this section we recall two other standard tests for chaos: positive Lyapunov exponents and the 0–1 test. The finite-time Lyapunov exponent is defined by

$$\lambda_T(v) = \frac{1}{T} \log \left(|Df^T(x_0, y_0)v| \right), \tag{10}$$

where Df is the Jacobian matrix, and v is a generic deviation vector with $|v| = 1$. Histograms of λ_T for three values of k are shown in Fig. 3(b) using $T = 2(10)^4$, which is the total orbit length used for the weighted Birkhoff method. As noted by [SLV05], when there are regular and chaotic orbits, these histograms are typically bimodal. For example, we observe that when $k = 1.5$ there is a lower peak centered near $\lambda = 0$ with width of order 0.02. This peak is well separated from the broader peak centered near $\lambda = 0.3$. The peaks are less well separated for smaller values of k; for example when $k = 1.0$ about 40% of the orbits have $\lambda_T < 0.01$, and there is a broader peak of presumably chaotic trajectories with $\lambda \in [0.08, 0.2]$. However, these two distributions have some overlap near $\lambda \approx 0.05$. As k grows, the mean value of λ increases and the lower peak of regular orbits disappears.
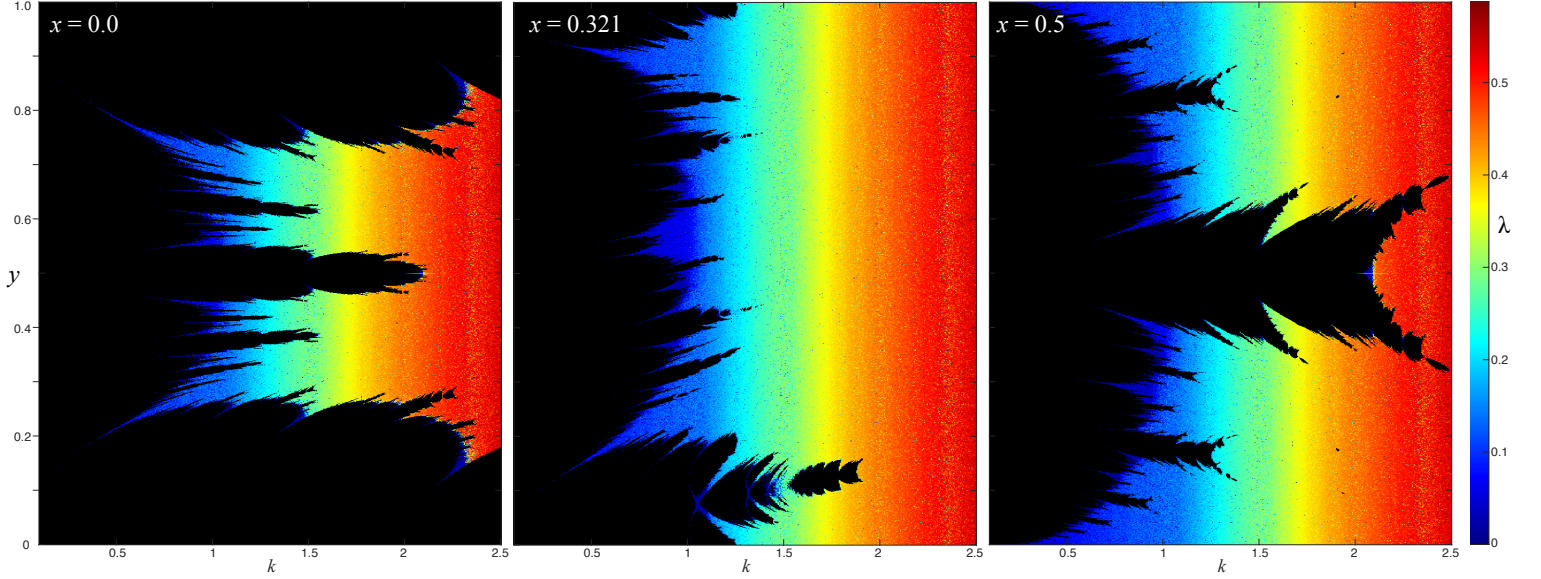
9

Figure 5: Lyapunov exponent λ using $v = (0,1)^T$ for orbits of the standard map (1) with (a) $x_0 = 0$ and (b) $x_0 = 0.321$ and (c) $x_0 = 0.5$ for $y_0 \in [0,1]$ and $k \in [0.1, 2.5]$. To improve convergence, the exponent (10) is computed for $T = 2^{16} = 65,536$. Black corresponds to $\lambda < 0.0022$, the lowest of the 256 colors. The value of λ for the chaotic orbits is indicated in the color bar.

To visualize the dependence of the exponents on k, we chose the same three lines of initial conditions shown in Fig. 4 for the weighted Birkhoff average. The resulting exponent, as a function of y_0 and the parameter k of (1) is shown in Fig. 5. In this figure, orbits with $\lambda < 2.2(10)^{-3}$ are colored black: these correspond to the regular orbits. As k grows, the distribution in the chaotic region is peaked around a growing value that reaches a maximum of $\lambda = 0.588$ when $k = 2.5$. Note that each of the panes of this figure is essentially the negative of the corresponding pane in Fig. 4.

The fraction of chaotic orbits can be estimated by removing orbits with λ_T in the range of the lower peak of Fig. 3(b). Fig. 6 shows the variation of Lyapunov λ_T and its error as a function of T for 30 orbits, where blue depicts regular orbits and red depicts chaotic orbits as determined at the maximum time $T = 2^{16}$. We can see that for the regular orbits, the convergence $\lambda_T \to 0$ is as T^{-1}, which is significantly slower than the convergence of the weighted Birkhoff average. Furthermore, many of the chaotic orbits take on λ_T values very close to zero for large values of T. For example, we calculate that 0.45% of the cases have $\min(\lambda_T) < 0.01$ and $\max(\lambda_T) > 0.02$ for T between 2^{14} and 2^{16}. By careful visual inspection of individual orbits, we find that we must choose our cutoff quite small to avoid the misidentification of chaotic orbits (more detail is given below in §2.3). Based on these considerations, we use the value

$$\lambda_c = 0.0005 \tag{11}$$

as the cutoff between chaos and regularity.

The resulting fraction of "chaotic" orbits as a function of k is shown in Fig. 7. This fraction is strongly dependent on choice of line for the initial conditions. For the lines of symmetry (e.g. $x = 0$
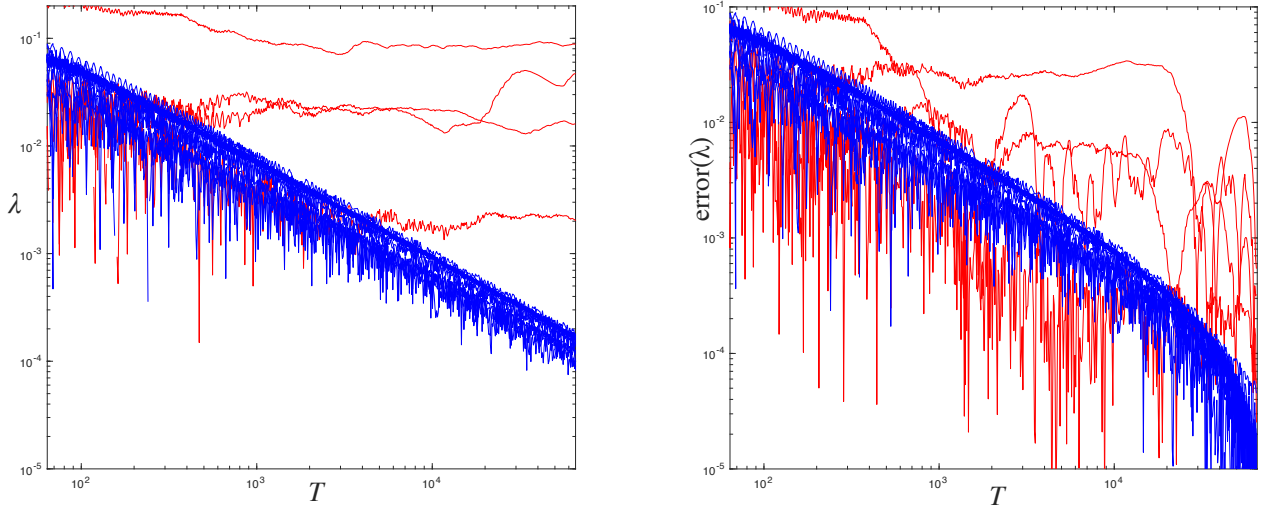
10

Figure 6: (a) Lyapunov exponents (10) computed for $k = 0.55$ for 30 orbits on the line $x_0 = 0.45$ for $T \in [100, 2^{16}]$. If $\lambda_{2^{16}} < \lambda_c = 0.0005$, the orbit is determined using this method to be regular and the curve is colored in blue. The other orbits are determined to be chaotic; they are colored in red. (b) For the same orbits and color scheme, we plot the error at each T value, defined as $|\lambda_T - \lambda_{2^{16}}|$. Note that for "regular orbits" λ_T appears to converge to 0 as T^{-1}.

or 0.5) of the standard map [Mac93], the fraction of orbits trapped in regular islands is larger.

Another test for chaos is the 0–1 test of Gottwald and Melbourne [GM09]. This test involves computing a time series (here we use $\{\sin(2\pi x_t) : t \in [0, T]\}$) from which a supplemental time series (called (p_t, q_t) in [GM09]) is constructed and tested for diffusive behavior. This ultimately gives a parameter, K_{median}, that is ideally either 0, when the orbit is quasiperiodic, or 1 when it is chaotic, and we use the cutoff $K_{median} > 0.5$ for chaos. Implementation of this test requires random samples of a frequency parameter. Using $T = 1000$ and 100 random samples gives an algorithm that is about 200 times slower than computing Lyapunov exponents. The resulting dichotomy between regular and chaotic orbits for this test is shown in Fig. 8 for initial conditions at $x = 0.0$. This figure agrees well with those in Fig. 4(a) and Fig. 5(a), though it appears to identify slightly fewer orbits as chaotic than the Lyapunov test: some orbits designated chaotic by Lyapunov exponent do not have $K_{median} > 0.5$.

2.3 Comparing the Methods

In this section we systematically compare the detection of chaos for the Chirikov standard map using the three different methods: Lyapunov exponents, 0–1 test, and the weighted Birkhoff method. We show that—weighing questions of speed, accuracy, and ease of implementation—the weighted Birkhoff average is better than either of the other two methods.

Figure 9 shows the fraction of orbits identified as chaotic by the three methods for initial conditions on the line $x_0 = 0.0$ with a uniform grid of $y_0 \in [0, 1]$. Note that the Lyapunov and weighted Birkhoff methods are difficult to distinguish on this scale. However, the chaotic fraction from the 0–1 test is uniformly below both that for other two methods. The difference is largest near $k = 1, 1.5$ and 2.3; these values correspond to major bifurcations in which regular islands and
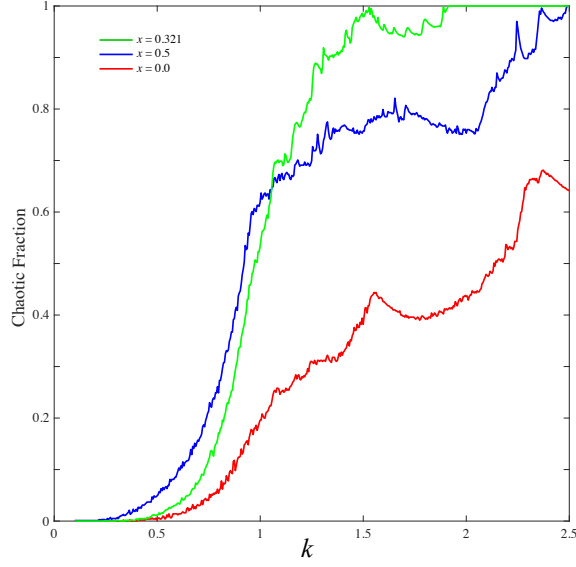
11

Figure 7: Using the Lyapunov method, this figure shows the fraction of orbits that have $\lambda > 0.0005$ for three values of x_0, a uniform grid of 1000 values of y_0, and $T = 2^{16}$. When $k < 4$ there is an island around the elliptic fixed point at $(0, 0)$ that decreases the number of chaotic orbits found for $x_0 = 0$ (red curve), and when $k < 2$ the island around the elliptic period-two orbit through $(0.5, 0.5)$ has a smaller, but similar effect for $x_0 = 0.5$ (blue curve). When $x_0 = 0.321$ (green curve), the regular regions around both of these elliptic orbits are not sampled when $k \geq 1.4$, and when $k \geq 1.914$, at most one of the 1000 y_0 values are deemed to not be chaotic. These variations can be observed in Fig. 4 and Fig. 5.

circles are destroyed. Nevertheless, the mean absolute deviation between the weighted Birkhoff and 0–1 test results is 2.6%.

For a further comparison, we computed the Lyapunov exponent and weighted Birkhoff average using ten different total orbit lengths of $2^7, 2^8, ..., 2^{16} = 65,536$ iterates, for $k \in [0.1, 2.5]$ on an evenly spaced grid of 50 values. We used the same grid of initial conditions for the 0–1 method, but due to its computational burden, we only used 1000 iterates. At each parameter value, we chose initial conditions on the line $x_0 = 0.321$ with $y_0 \in [0, 1]$ on an evenly spaced grid of 500 points: thus there are $25,000$ trials.

Designating one of the methods as the "ground truth," one way to compare another method is by the True Skill Statistic, also known as the Hanssen-Kuiper skill score [Woo78]:

$$TSS = \frac{TP}{TP + FN} - \frac{FP}{FP + TN}.$$

Here TP ("true positive") is the fraction of initial conditions that are classified correctly as chaotic by the test method according to the reference standard, FP is the fraction classified incorrectly as chaotic, TN is the fraction that are correctly as non-chaotic, and FN is the fraction classified incorrectly as non-chaotic. The TSS ranges from -1 for a classification that always disagrees with the reference, to 1 for one that always agrees. An advantage of TSS is that it does not depend upon the number of trials, just on the relative accuracy. However, if we are comparing two predictions,
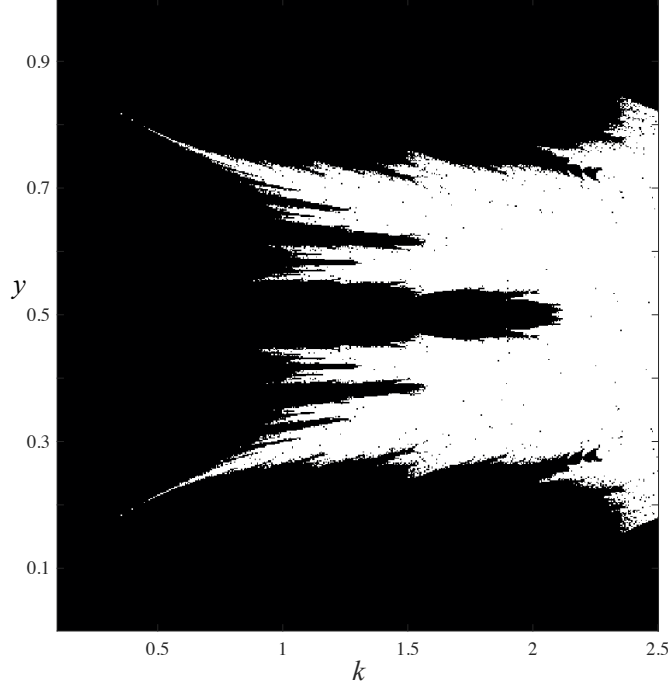
Figure 8: Chaotic region of the standard map with $x_0 = 0$ for $y \in [0, 1]$ and $k \in [0.1, 2.5]$ using a time series $\{\sin(2\pi x_t : 0 < t \leq 1000\}$ and the 0–1 method of [GM09]. An orbit is deemed to be chaotic (colored white) if the 0–1 parameter $K_{median} > 0.5$.

the skill statistic does depend upon which prediction is designated as the "ground truth": changing this designation is equivalent to exchanging $FP \leftrightarrow FN$.

As "ground truth," we declare an orbit to be "chaotic" when $\lambda > \lambda_c = 0.0005$ for an orbit length of 2^{16}. We then find that the weighted Birkhoff method with the same total orbit length, i.e., $2T = 2^{16}$, gives $TSS = 0.997$. Furthermore, whenever $2T \geq 2^{10}$, the weighted Birkhoff method gives $TSS > 0.9$, and even for the small value $2T = 256$, it gives $TSS = 0.80$. These results confirm Fig. 9, which showed the near overlap of the Birkhoff average curve for $2T = 2(10)^4$ with the "ground truth" Lyapunov curve (in that case for $x_0 = 0.0$). In contrast, the Lyapunov method for 2^{15} iterates yields $TSS = 0.77$ as compared to the ground truth, confirming again the convergence of this method is indeed slow.

An alternative comparison measure is to simply count the percentage of correctly classified initial conditions,

$$R = \frac{TP + TN}{TP + FP + TN + FN},$$

the "ratio" of [Woo78]. Using the same definition of ground truth, the weighted Birkhoff method gives $R \geq 0.98$ whenever $2T \geq 2^{12}$. In contrast, the Lyapunov method has $R = 0.8$ for $T = 2^{14}$.

Finally, the 0–1 method for 1000 iterates gave $TSS = 0.91$ and $R = 0.96$ when comparing to the ground truth Lyapunov method. The same values are obtained if we compare the 0–1 method to the weighted Birkhoff method for total orbit length 2^{16}.
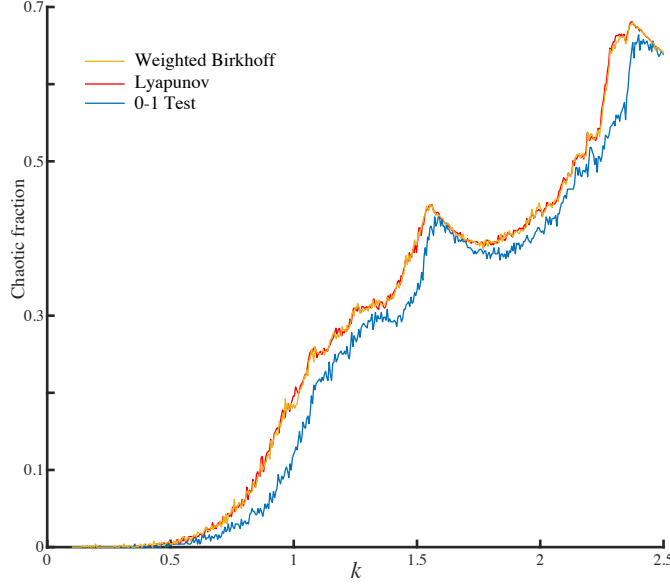
13

Figure 9: (b) The fraction of chaotic orbits as a function of k for initial conditions on the line $x_0 = 0.0$. The red curve shows the fraction of chaotic orbits computed using Lyapunov exponents with $\lambda > 0.0005$, the blue curve shows the fraction with the 0–1 parameter $K_{median} > 0.5$, and the yellow curve shows the fraction of chaotic orbits computed using the weighted Birkhoff average with $dig < 5.5$. For each k value, 1000 initial conditions were used for the Lyapunov and weighted Birkhoff methods. For the Lyapunov method $T = 2^{16}$, and for the weighted Birkhoff method $T = 10^4$ (which involves $2(10)^4$ calculations). For the 0–1 test, 500 initial conditions were used for each k-value and $T = 1000$.

We now comment on the choice of chaos cutoff values dig_T (9) and λ_c (11). The agreement between the weighted Birkhoff and Lyapunov methods changes only slightly if we vary cutoff for dig_T, and the best agreement occurs when $dig_T = 4$. However, R is not very sensitive to variations $3.5 < dig_T < 6$.

In contrast, varying the cutoff value λ_c for the Lyapunov exponent calculation causes significant changes. Indeed, if λ_c is increased, an orbit can be identified as "regular" with a smaller number of iterates. This is explained by the data in Fig. 6: the computed value of λ for regular orbits is seen to decrease as $\frac{1}{T}$, so if λ_c is increased, an orbit can be identified as "regular" sooner. However, choosing a larger cutoff value has the disadvantage of introducing a systematic error in the sense that orbits that are "weakly chaotic" will never be so identified. Moreover, in Fig. 6 one can see that a number of orbits identified as "chaotic" at $T = 2^{16}$ have earlier episodes in which λ_T decreases with time. Such orbits can be trapped in a narrow chaotic layer or can be very close to a boundary of a chaotic layer for a long time. We found by careful examination of the phase space behavior of for over 100 different orbits in which the weighted Birkhoff and Lyapunov methods disagreed, that if we chose $\lambda_c > 0.0005$, the Lyapunov method would systematically misidentify chaotic orbits as compared to the weighted Birkhoff method. The choice $\lambda_c = 0.0005$ minimized these errors when $T = 2^{16}$. Note that because we choose λ_c quite small, this method very rarely misidentifies an orbit as regular when it is chaotic: the number of false positives and false negatives are not equal

14

between the three methods.

For a given total orbit length, the computation time for the Lyapunov exponent was approximately 1.5 times slower than and weighted Birkhoff average. For example computing the Lyapunov exponent for $25,000$ initial conditions and total orbit length 2^{15} took around 10.6 seconds using Matlab on an iMac Pro, whereas the weighted Birkhoff method took around 7 seconds. The 0–1 method was significantly slower; on the same machine with the same initial conditions, and smaller total orbit length 1000, the computation took 45 minutes. In our case, the Jacobian of the map (1) is quite simple, thus we expect that if the derivative were computationally more expensive, then the weighted Birkhoff method would have an even more significant speed advantage over the Lyapunov method.

We note that the comparisons in this section do not include a variety of other efficient methods that use hyperbolic growth to distinguish between chaos and regular behavior including the Fast Lyapunov Indicator (FLI), OFLI, Mean Exponential Growth factor of Nearby Orbits (MEGNO), and the alignment indices SALI and GALI; see [CS00, FLG97, Sko01, SGL16] for discussions of these methods.

The weighted Birkhoff method has another advantage, as we will illustrate in the next section: it gives an accurate calculation of the rotation number ω that we can use to distinguish between rotational invariant circles and island chains.

3   Island chains

The regular orbits of the Chirikov standard map are of two distinct topological types: rotational invariant circles and orbits within the island chains. We are primarily interested in studying the rotational invariant circles, and thus must look for a way to distinguish and remove orbits within island chains.

For a twist map Birkhoff's theorem implies that the rotational invariant circles are graphs, $x \mapsto (x, c(x))$. Generically the dynamics on each such circle is conjugate to an incommensurate rotation, implying that ω in (3) is irrational.

By contrast, around each elliptic period-n orbit there is generically a family of trapped orbits forming a chain of n islands. The regular orbits in these island chains are further partitioned into orbits that are quasiperiodic and those that are periodic relative to the n^{th} power of the map. The latter, if elliptic, can again be the center of chains of islands. This gives rise to the familiar island-around-island structure. Each regular orbit within a period-n island chain that is not itself periodic is generically dense on a family of topological circles: these are *oscillational* invariant circles. Nevertheless, the rotation number ω, (3) will average out the internal dynamics, resulting in a rational value that is the rotation number of the central period-n orbit. Of course if one were to measure the rotation number of an oscillational circle relative to the periodic orbit that it encloses, one would generically find it to be irrational as well.

In §2.1 we developed a highly-accurate method for removing chaotic orbits and for computing ω for regular orbits. In §3.1, we establish a numerical method to remove regular orbits in island chains by determining which of the computed ω values are "rational," and which are "irrational." In §3.2 we use this method to identify orbits within island chains.

3.1 Numerical identification of rational numbers

We are interested in establishing a numerical method to determine whether a numerically computed number determined using floating point arithmetic is representative of a rational or an irrational. At the outset, this is not a well-posed question, since floating point representations of numbers are rational. The question becomes whether a numerical value is—with high probability—the approximation of a rational or an irrational number. In this section, we concentrate on a closely related question, and in the next section we show how the answer can be applied to establish rationality. Our question is: given a number x, and an interval

$$I_\delta(x) \equiv (x - \delta, x + \delta), \tag{12}$$

with some tolerance δ, what is the rational number p/q with the smallest denominator in $I_\delta(x)$?

If, for a small δ, there is a rational $p/q \in I_\delta(x)$ with a sufficiently *small* denominator q, we would expect that x is—to a *good* approximation—given by this rational. Whereas if all such rationals have *large* denominators, we would expect that x is an approximation of an irrational number. Actually, we will argue that if q is *too large*, x is more likely an approximation of a rational number that just missed being in the interval. We will return to the question of what constitutes small, large, and too large for values of q, but first we discuss the question of how to actually find the value p/q in a prescribed interval.

We denote the smallest denominator for a rational in an interval I by

$$q_{min}(I) \equiv \min\{q \in \mathbb{N} : \tfrac{p}{q} \in I, p \in \mathbb{Z}\}. \tag{13}$$

The question of finding q_{min} has been considered previously in [BBdA98, For07, CP16], and a closely related question is considered in [CB09].

Given an interval I in $\mathbb{R}$, one would imagine that there are standard algorithms for finding the rational number p/q in I with $q = q_{min}(I)$. Indeed packages such as Mathematica and Matlab both have commands that appear to do this. However these algorithms use truncations of the continued fraction expansion [HW79], and neither of them work correctly in the sense of finding the smallest denominator [For07]. Recall that the continued fraction expansion for $x \in \mathbb{R}^+$ is

$$x = a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{\cdots}}} \equiv [a_0; a_1, a_2, \ldots], \quad a_i \in \mathbb{N}, \quad a_0 \in \mathbb{N} \cup \{0\}. \tag{14}$$

Truncation of this path after a finite number of terms gives a rational "convergent" of x:

$$\frac{p_k}{q_k} = [a_0; a_1, a_2, \ldots, a_k]. \tag{15}$$

Convergents are *best approximants* in the sense that if

$$\left| \frac{p}{q} - x \right| < \frac{1}{2q^2}, \tag{16}$$

then p/q is a convergent to x [HW79, Theorem 184]. Moreover, at least one of any two successive convergents satisfies (16).

However, the convergents are not necessarily the rationals with the smallest denominators in a given interval. As a simple example, the rational with the smallest denominator within $\delta = 10^{-3}$ of π is $\frac{201}{64}$, i.e., $q_{min}(I_{10^{-3}}(\pi)) = 64$. However, this rational is not a convergent of the continued fraction $\pi = [3; 7, 15, 1, 292, 1, 1, ...]$; indeed, the first convergent in the interval is $\frac{p_2}{q_2} = \frac{333}{106}$.

A correct algorithm (e.g., that proposed by Forisek [For07]), is easiest to explain based on the *Stern-Brocot* or *Farey* tree. Every number in $\mathbb{R}^+$ has a unique representation as a path on this binary tree:

$$x = s_1 s_2 \ldots, \quad s_i \in \{L, R\}. \tag{17}$$

The tree, whose first levels are sketched in Fig. 10, is constructed beginning with the root values $\frac{0}{1}$ and $\frac{1}{0}$. Subsequent levels are obtained by taking the mediants of each neighboring pair:

$$\frac{p_m}{q_m} = \frac{p_l}{q_l} \oplus \frac{p_r}{q_r} \equiv \frac{p_l + p_r}{q_l + q_r}. \tag{18}$$

Level zero of the tree is the mediant of the roots, $\frac{1}{1}$; it is defined to have the null path. If $x < \frac{1}{1}$, then its first symbol is L, and if $x > \frac{1}{1}$, then its first symbol is R. At level ℓ of the tree, 2^ℓ new rationals are added, the mediants of each consecutive pair. The left and right parents are neighboring rationals that have level less than ℓ. Every consecutive pair of rationals at level ℓ are *neighbors* in the sense that

$$p_r q_l - p_l q_r = 1. \tag{19}$$

A consequence is that p_m and q_m are coprime.

For $\ell = 1$, the new mediants $\frac{1}{2} = \frac{0}{1} \oplus \frac{1}{1}$ and $\frac{2}{1} = \frac{1}{1} \oplus \frac{1}{0}$ are added to give the level-two Farey sequence $\frac{0}{1}, \frac{1}{2}, \frac{1}{1}, \frac{2}{1}, \frac{1}{0}$. Then the 2^2 mediants of each neighboring pair are added to give 2^3 level-three intervals, see Fig. 10. Since the level-three rational $\frac{2}{3} > \frac{1}{2}$, to the right of its level-two parent, then $\frac{2}{3} = LR$. Similarly $\frac{3}{2}$ is to the left of its level-two parent $\frac{2}{1} = R$, so $\frac{3}{2} = RL$.
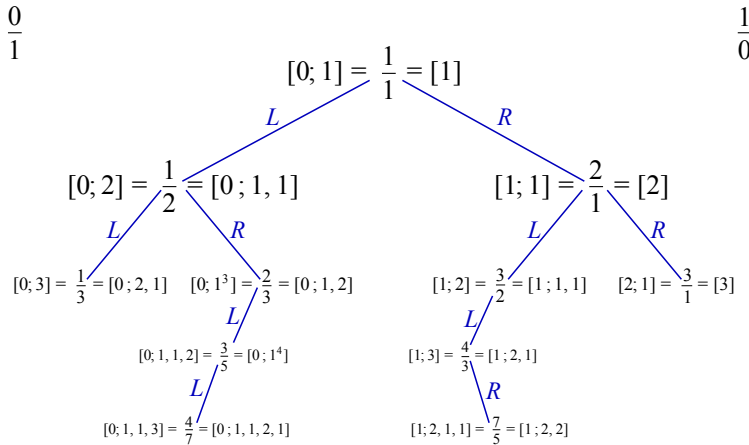

Figure 10: The continued fraction expansion and some entries on the Stern-Brocot or Farey tree. Each rational has two possible finite continued fractions but a unique Farey path. See the appendix for the relationship between the two.

The Farey path (17) for any $x \in \mathbb{R}^+$ is the unique path of left and right transitions that lead to x starting at $\frac{1}{1}$. Every rational has a finite path and every irrational number has an infinite

path [HW79]. Algorithm 1 in the appendix computes the Farey path, up to issues of floating point accuracy and a stopping criterion.

The Farey expansion allows one to find the rational with smallest denominator in any interval:

Lemma 1 (Smallest Rational). *The smallest denominator rational in an interval $I \subset \mathbb{R}^+$ is the first rational on the Stern-Brocot tree that falls in I.*

The proof of this lemma, from [For07], is given in Appendix A. An alternative version of this result using continued fractions can be found in [BBdA98].

An algorithm for finding $q_{min}(I_\delta(x))$, based on Lem. 1 is given in the appendix in Algorithm 2. For example, for $x = 0.12 = \frac{3}{25} = L^8 R^2 = [0; 8, 3]$, the sequence of Farey approximants is

$$\frac{1}{1}, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{5}, \frac{1}{6}, \frac{1}{7}, \frac{1}{8}, \frac{1}{9}, \frac{2}{17}, \frac{3}{25}.$$

Given $\delta = 0.005$ for example, the Farey interval $(\frac{1}{7}, \frac{1}{8}) \supset I_\delta$, and the mediant $\frac{2}{17} \in I_\delta$ since $0.12 - \frac{2}{17} \approx 0.0024 < \delta$. Thus from the algorithm we obtain

$$\text{SmallDenom}(0.12, 0.005) = [2, 17] \Rightarrow q_{min}(I_{0.005}(0.12)) = 17.$$

As noted in [For07], the built-in routines of standard mathematical software do not always compute the smallest rational approximation correctly. For example, the built-in Matlab command "rat" gives $\text{rat}(0.12, 0.005) = [3, 25]$, giving x itself, since the second convergent $\frac{1}{8} = 0.12 + 0.005$, is not in I. The point is that the intermediate convergents of the Farey path can satisfy the approximation criterion before the principal convergent of the continued fraction, and this can happen whenever the Farey path is not alternating $\dots LR \dots$ or equivalently the continued fraction elements are not all 1's.

To determine the "typical" size of a denominator in an interval I, we show in Fig. 11 a histogram of the minimal denominator computed using Algorithm 2 in the appendix for randomly chosen floating point numbers in $(0, 1)$ with a uniform distribution. For this case, when $\delta = 10^{-12}$, the mean minimal denominator appears to be close to $10^6 = \delta^{-1/2}$. The distribution is not log-normal: the data is significantly more concentrated around the mean than a normal distribution with the same standard deviation. Over the range $\delta = [10^{-4}, 10^{-14}]$, the mean log-denominator obeys the relation

$$\langle \log_{10} q_{min} \rangle = -\tfrac{1}{2} \log_{10} \delta - 0.05 \pm 0.001, \tag{20}$$

and in this same range of δ values, the standard deviation is nearly constant,

$$\sigma = 0.2935 \pm 0.0006. \tag{21}$$

Further support for this statement is found in Fig. 12, which shows that for $\delta = 10^{-tol}$, the probability that q_{min} is in the range $10^{tol/2 \pm s}$ does not depend on the choice of *tol*. Indeed, the curves in this graph were obtained from only 10^4 random trials: if more values were randomly chosen, it would be impossible to distinguish between these distribution plots.

The mean of our observations (20) is consistent with the expectation from (16). Indeed, for any δ, then there is a convergent with $|x - p/q| < \delta$, with a denominator that must satisfy $q \geq (2\delta)^{-1/2}$. Since the minimum denominator is no more than this, we expect that $q_{min} \sim (2\delta)^{-1/2}$, and thus

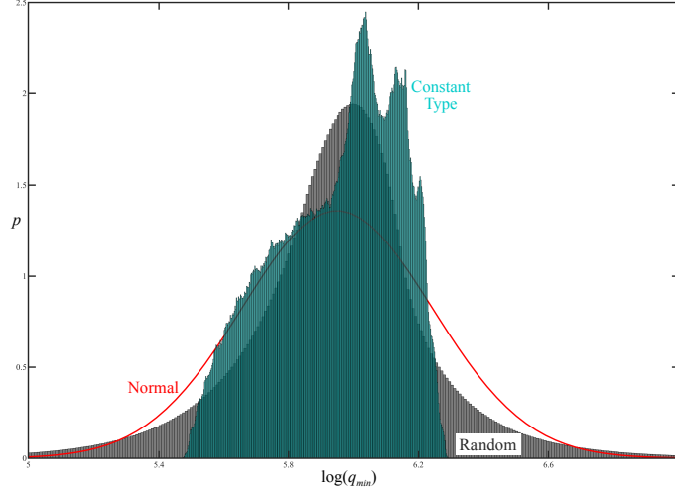$$\log_{10} q_{min} \sim -\tfrac{1}{2} \log_{10} \delta - 0.15,$$

Figure 11: Probability density of $\log_{10}(q_{min})$ computed by appendix Algorithm 2 with $\delta = 10^{-12}$ for 10^8 randomly chosen numbers in $(0, 1)$ (black). This distribution has mean 5.9497, mode 5.9662, standard deviation $\sigma = 0.29333$, and kurtosis 6.3073. The red curve shows the normal distribution with the same mean and standard deviation. Also shown is the histogram for 10^8 randomly chosen numbers of constant type with $A = 10$ (green).
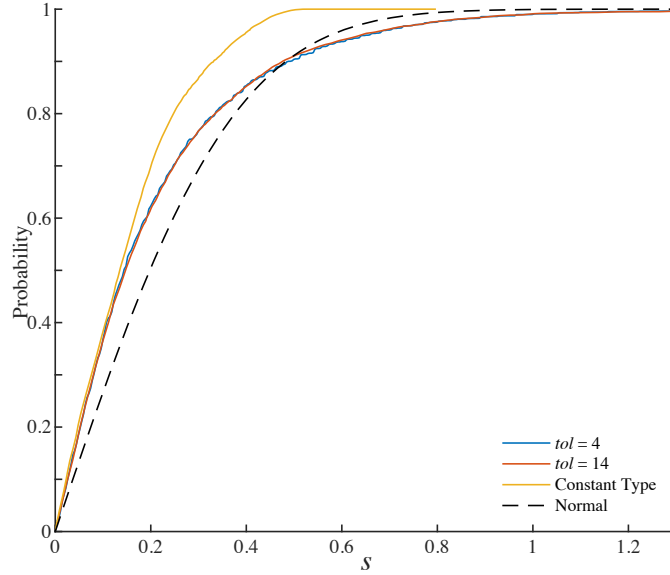

Figure 12: A graph of the probability that $q_{min} \in 10^{tol/2}[10^{-s}, 10^s]$, for $\delta = 10^{-tol}$ for 10^4 randomly chosen $x \in [0, 1]$. The blue curve is for $tol = 4$ and the red for $tol = 14$ (These curves are nearly indistinguishable). The yellow curve shows the probability for numbers of constant type with $A = 10$. The probability for a normal distribution with standard deviation (21) is the dashed curve.

which is not far from the observation (20).

A related result was obtained by [Ste13]: for intervals of the form $J_N = (\frac{i-1}{N}, \frac{i}{N}]$, the mean

19

smallest denominator in grows asymptotically as

$$\langle q_{min}(J_N)\rangle \sim CN^{1/2},$$

with a coefficient $1.35 < C < 2.04$. Since these intervals are of size $\frac{1}{N} = 2\delta$, this gives

$$\log_{10}\langle q_{min}\rangle \sim -\tfrac{1}{2}\log_{10}\delta + K, \quad K \in [-0.020, 0.159].$$

Note that since the minus logarithm is convex, Jensen's inequality implies that (20) is no larger than this result. We are not aware, however, of any results in the literature that imply the validity of (20) or (21).

As a second numerical experiment, we consider numbers of *constant type*; that is numbers that have bounded continued fraction elements: $\sup_k\{a_k\} = A < \infty$. Such numbers can be thought of as "highly irrational" in the sense that they are Diophantine (4), with $\tau = 1$, and $c > \frac{1}{A+2}$. Conversely, if x is Diophantine with constant c then $A < \frac{1}{c}$ [Sha92]. This class of numbers is especially important in the context of area-preserving maps: it was conjectured that invariant circles with constant type rotation numbers are locally robust and that every circle that is isolated from at least one side has constant type [MS92].

For the numerical experiment shown in Fig. 11, we chose rational numbers with continued fractions of length 40, with $a_i \leq 10, i = 1, \dots 40$ chosen as iid random integers. Note that this means that every trial x is rational; however, the denominator of these rationals is at least as large as the case $a_i = 1$, which gives the Fibonacci $F_{40} \approx 1.08(10)^8$. The resulting smallest denominator distribution is the green histogram in Fig. 11. The cumulative distribution of these numbers is also shown in Fig. 12, which shows that the probability that $Prob(|\log_{10}(q_{min}) - tol/2| > 0.728) = 1\%$.

As mentioned previously, rational numbers nearby a given a value of x can result in both extremely small and extremely large values of $q_{min}(I_\delta(x))$. To demonstrate this, Fig. 13 shows a plot of $q_{min}(I_{10^{-5}}(x))$ for evenly spaced x values between 0.095 and 0.105. The dots below the x-axis are centered at each rational with a denominator $q \leq 80$; the size of each dot is inversely proportional to q. Note that in the vicinity of each dot, there is a small region where q_{min} drops to the corresponding small value of q, but additionally, there is a larger interval in which q_{min} becomes much larger than average, with a larger jump near smaller denominators. Dynamically these rationals correspond to orbits that are limiting on the separatrices of islands, and hence are chaotic.

The main takeaway message from the "typical size" experiments in this section is that numbers outside the main peak of the distribution in Fig. 11 correspond to those "close" to rationals. In the next section, we will discard such rotation numbers to filter for candidates for rotational invariant circles.

3.2 Identification of island chains using the weighted Birkhoff average

In this section, we use the weighted Birkhoff method to obtain an accurate computation of the rotation number ω defined for the Chirikov standard map in (3). Namely,

$$\omega(z) = WB(\Omega)(z). \tag{22}$$

Using this, we can distinguish rotational invariant circles from orbits in island chains by determining whether the computed value of ω is an approximation of a rational or irrational number as follows.
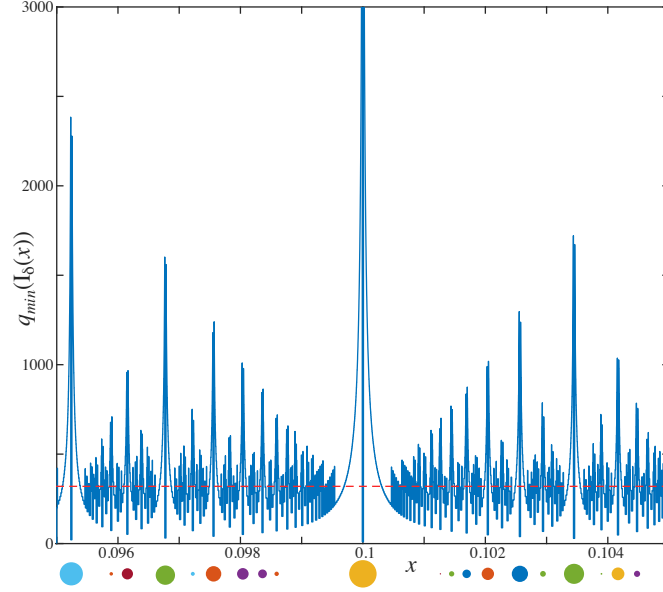
Figure 13: A plot of the smallest denominator, $q_{min}(I_\delta(x))$, in an interval (12) for $\delta = 10^{-5}$ and 10^4 values of $x \in [0.095, 0.105]$. The dots below the x-axis indicate the size of the denominator of each rational number in the interval with a denominator up to 80; larger dots correspond to smaller denominators. There is a spike in denominator size immediately outside the interval around small denominator rationals. The mean log-denominator, (20), is shown by the dashed (red) line.

Fix *tol* and let $\delta = 10^{-tol}$. For a rotation number ω, we find $q_{min}(I_\delta(\omega))$ in (13), the smallest denominator of a rational number within distance δ of ω. In most of our numerics, we have chosen $tol = 8$. To distinguish between rationals and irrationals, for each rotation number ω define the absolute deviation

$$dev_\omega = |\log_{10}(q_{min}(I_\delta(\omega))) - tol/2|. \tag{23}$$

For a fixed cutoff value s, we remove the orbits within island chains as follows. Let z be an initial condition of a regular orbit with associated rotation number ω. If $dev_\omega > s$, then we discard z as a member of an island chain. Note that this is equivalent to saying that q_{min} is outside the range $10^{tol/2 \pm s}$.

It remains to choose a cutoff value s. In our numerics, when we wish to be conservative about identifying rotational circles, we have used the cutoff value $s = 0.3375$, which implies that we have kept slightly above 81% of randomly chosen values, as can be seen in Fig. 12. This corresponds to choosing only irrational numbers that are very badly approximated by rationals with small denominators.

Now that we have established all of our criteria for distinguishing rotation numbers, we summarize the particular values we have used in most of our numerical calculations as two criteria:

$$\begin{array}{lll} \text{Chaos criterion:} & dig_T < 5.5 \text{ for } WB_T(\cos(2\pi x)), \\ \text{Irrationality criterion:} & dev_\omega < 0.3375 \text{ for } tol = 8. \end{array} \tag{24}$$

In each case, for each initial condition (x_0, y_0), we compute an orbit and determine whether the

orbit is chaotic using the above chaos criterion. We also compute the rotation number ω using $WB_T(\Omega)$ and determine whether the orbit is a rotational invariant circle using the irrationality criterion.

4 Rotational invariant circles

Using the strategy of §3.2 for eliminating rationals, we can now remove orbits that are contained in island chains. We show in Fig. 14 the rotation number ω for initial conditions (x_0, y_0) that are identified to lie on rotational invariant circles using distinguishing criteria (24).
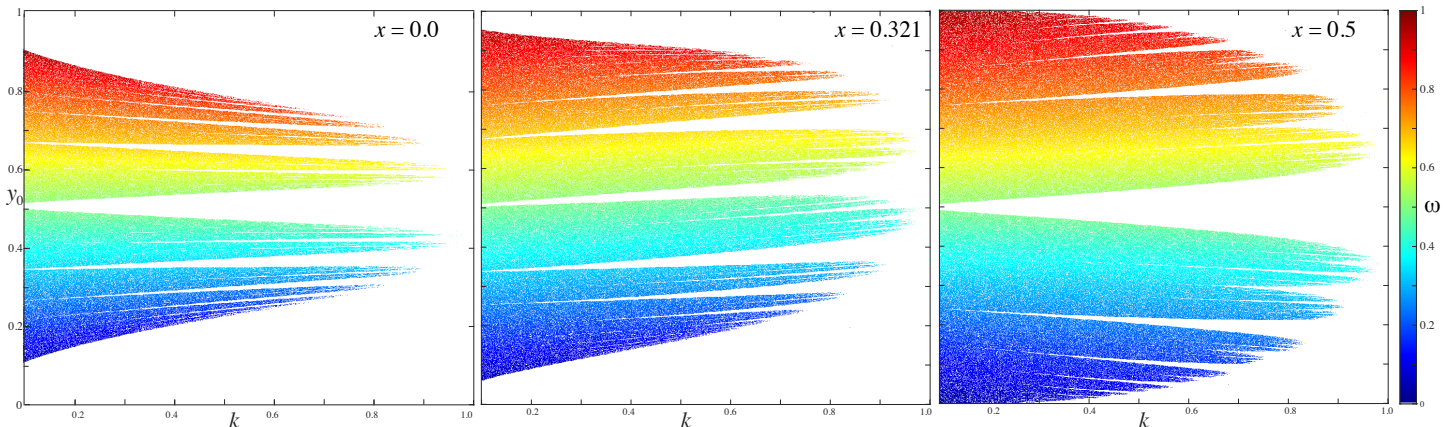

Figure 14: The rotation number computed for rotational invariant circles orbits of the standard map (1) with (a) $x_0 = 0$ and (b) $x_0 = 0.321$ and (c) $x_0 = 0.5$ for $y_0 \in [0, 1]$ and $k \in [0.1, 1.0]$. The computations are done using $T = 2(10)^4$, using the distinguishing criteria in (24).

The panels in Fig. 14 strongly resemble the critical function computed using Greene's method for the standard map [MS91, MS92]. In this method, one typically chooses a set of noble irrational numbers, and finds the threshold of instability for a long periodic orbit that is close to each of these nobles. The periodic orbits used in these computations are those that are symmetric under the reversor for (1); for example, every elliptic, symmetric rotational orbit is observed to have a point on the line $x = 0$. An advantage of our current method is that symmetry is not required.

Recall from §1 that it is believed that there are no rotational invariant circles for the standard map above $k_{cr} = 0.97163540324$ and that the last circle has the golden mean rotation number [Gre79, Mac93, HCF$^+$16]. In Fig. 15, we show how the fraction of initial conditions that are identified as rotational circles in Fig. 14 varies with k. By $k = 0.9685$, 99.9% of the circles are destroyed and the fraction drops to zero at $k = 0.9712$, though there is one orbit misidentified as a circle at $k = 0.9766$. The accuracy of these computations is limited by the fact that the initial conditions are fixed to a grid in y_0.

As a more precise test of the efficacy of the weighted Birkhoff average to determine k_{cr}, we used continuation to find an orbit on the line $(0.321, y_0)$ with the fixed rotation number $\gamma^{-1} = \frac{1}{2}(\sqrt{5} - 1)$ when $T = 2(10)^4$. A computation of dig_T, (9) can then be used to determine if the orbit is not chaotic. For the computation shown in Fig. 16, $dig_T = 12$ at $k = 0.9706$, and it drops to
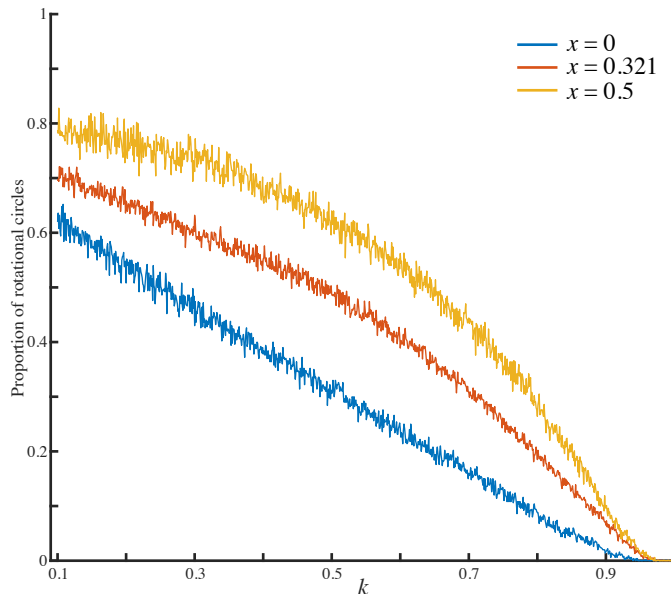
22

Figure 15: Fraction of orbits of (1) from Fig. 14 that are on invariant circles for 1000 initial conditions with $y_0 \in [0, 1]$ on three vertical lines as shown. The largest fraction occurs when $x = 0.5$, as this line tends to avoid many of the larger islands.

6 at $k = 0.9731$, with a precipitous drop just as the curve crosses k_{cr}. As an example, when $k = 9697/9980 \approx 0.971643$, the initial condition with $\omega = \gamma^{-1}$ has $y_0 = 0.676535782378533$. Though this orbit no longer lies on an invariant circle since $k > k_{cr}$, iteration shows that it remains localized to what appears to be a circle for hundreds of millions of iterations.

We now focus on the number theoretic properties of rotation numbers for robust circles. It is thought that the rotation numbers of the more robust invariant circles should have continued fraction elements with more elements $a_i = 1$ [Gre79, MS92]. To test this, we plot the distribution of continued fraction elements, a_n, for the rotation number of invariant circles in Fig. 17. The expected distribution for randomly chosen irrationals is the Gauss-Kuzmin distribution [Sha92], $P(a_i = k) = \log_2(1 + 1/(k(k+2)))$. When k is relatively small, the observed distribution follows the Gauss-Kuzmin distribution closely, at least for $a_n \leq 10$; but for $k = 0.95$, when most circles have been destroyed, the probability of $a_n = 1$ or 2 is larger than would be predicted for random irrational numbers, and the probability that $a_n \geq 8$ is at least four times smaller than the Gauss-Kuzmin value.

5 Generalizations of the Standard Map

The method we have developed to find rotational invariant circles works equally well for other area-preserving maps. As a first example, we consider two-harmonic generalized standard map (1) with the force

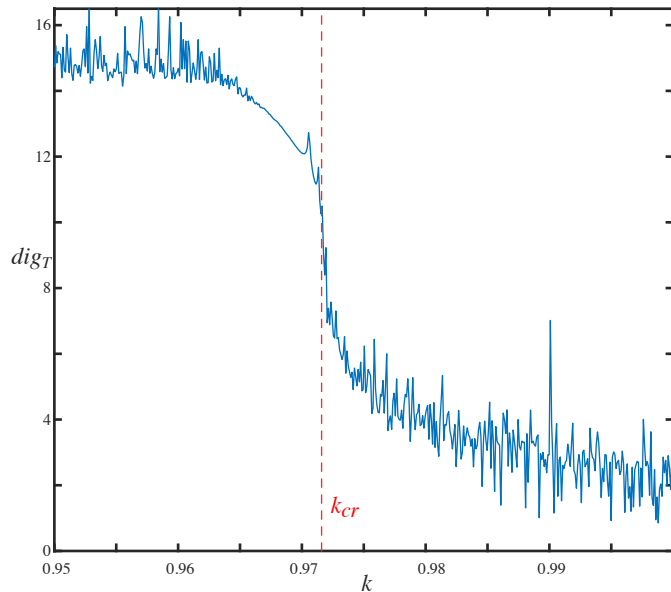$$F(x) = -\frac{k}{2\pi}\left(\sin(\psi)\sin(2\pi x) + \cos(\psi)\sin(4\pi x)\right), \tag{25}$$

23

Figure 16: Computation of dig_T for the golden mean circle for $T = 2(10)^4$ for 500 values of $k \in [0.95, 1.00]$. Continuation is used to find the initial condition $(0.321, y_0)$ that has $\omega = \gamma^{-1}$. The drop of dig_T from 12 to 6 indicates that the circle is destroyed for a parameter value in $(0.9706, 0.9731)$.

that was first studied in [GJSS87] (see [Sim18] for later references). phase portrait of this map, analogous to that shown for the standard map in Fig. 1, is shown in Fig. 18 for the value $\psi = 0.7776$. Note that at these parameters there are invariant circles in four narrow bands. The set of circles as a function of k is shown in Fig. 19. This figure is similar to [FM14, Figure 12(b)], where the critical parameters were computed for a set of 256 noble rotation numbers. In that case the last invariant circle, with $\omega \approx 0.247$, was destroyed at $k \approx 0.613$. The numerical experiment here shows that at least 99.9% of the invariant circles are destroyed when $k > 0.61850$. The last invariant circle in our sample appears to have

$$\omega = 0.239184971708802 = [0; 4, 5, 1, 1, 8, 8, 5, 8, 8, 1, \ldots],$$

with $q_{min}(I_{10^{-8}}) = 13153$.

A similar, well-studied map is the standard nontwist map (see [FM14, SMS$^+$18] for references). This map is of the form (1) with the standard force, but with the frequency map

$$\Omega(y) = y^2 - \delta. \tag{26}$$

Note that $\Omega : \mathbb{R} \to [-\delta, \infty)$, and that the rotation number (3) will also take values in this interval. The phase space of the dynamics for $k = 1.5$ is shown in Fig. 20 for $\delta = 0.3$. At these parameter values, there are large chaotic regions around islands with rotation number 0 (colored green) and a band of rotational circles near the minimum of Ω (colored blue). The most robust circles tend to be the *shearless circles*; they cross the line $y = 0$ where $\Omega'(y) = 0$. The fraction of chaotic orbits and rotational circles as k varies is shown in Fig. 21. For the 1000 initial y values in our experiment,
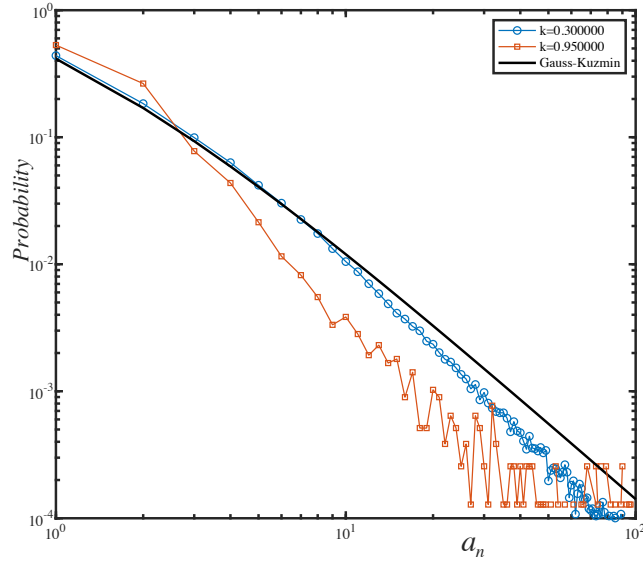
24

Figure 17: Probability distribution for the occurrence of continued fraction elements of the rotation number for rotational invariant circles of the standard map for two values of k. These were computed using $T = 10^4$ iterates, and a grid of $5(10)^4$ initial conditions at $x = 0.321$, using the distinguishing criteria in (24). When $k = 0.3$ (blue), we found $30,176$ invariant circles, and when $k = 0.95$ (red), we found 682. The black curve shows the Gauss-Kuzmin distribution, which is the distribution of elements for a random irrational chosen with uniform probability in $[0, 1]$.

the last detected rotational circle is at $k = 2.7725$ for $(x_0, y_0) = (0.35, -0.2620)$ with the rotation number

$$\omega = -0.255234160728417 = [-1; 1, 2, 1, 11, 5, 4, 7, 19, \ldots],$$

with $q_{min}(I_{10^{-8}}) = 7260$. We verified that there are no rotational circles in the interval $-2 < y < 2$ for $k = 2.7841$ by direct iteration: finding an initial condition whose orbit crossed this region. For example, the initial condition $(x_0, y_0) = (0.0, 2.088767893349248)$ has $y_T < -2$ for $T = 129,072$, though this T value is uncertain due to floating point errors.

Finally, we consider an asymmetric two-harmonic map (1) with the force

$$F(x) = -\frac{k}{2\pi} \left(\sin(\psi) \sin(2\pi x) + \cos(\psi) \cos(4\pi x) \right), \tag{27}$$

studied in [FM14]. This map does not have the usual $x \mapsto -x$ reversor of the standard map (1), and therefore its periodic orbits are not aligned by a symmetry. Phase portraits for $k = 0.2$ and $\psi = 0.7776$ are shown in Fig. 22, and the fraction of circles as a function of k in Fig. 23.

6 Conclusions and future work

The weighted Birkhoff average (6) and the distinguishing criteria (24) have been shown to efficiently categorize orbits as chaotic, trapped in islands, or quasiperiodic on rotational circles. Using only $T = 10^4$ iterations, the rotation number of regular orbits is typically known to machine precision,
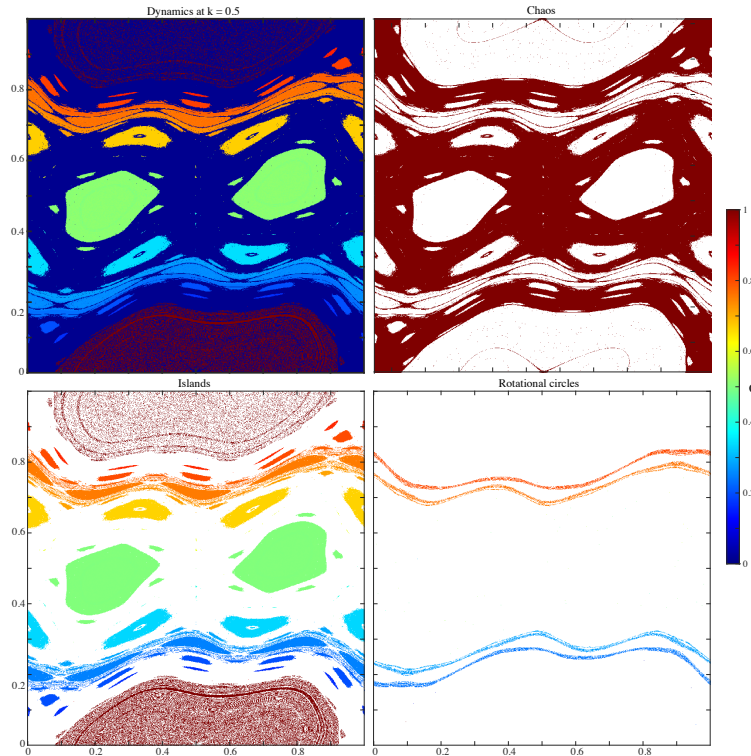
Figure 18: The dynamics of the two-harmonic map, with force (25) for $k = 0.5$ and $\psi = 0.7776$. The weighted Birkhoff method distinguishes chaotic orbits (upper right), islands (lower left), and rotational circles (lower right). The rotation number of each nonchaotic orbit is color-coded (color bar at right). The computations were performed for a grid of 1000^2 initial conditions in $[0, 1] \times [-0.75, 0.75]$ with $T = 10^4$, using the distinguishing criteria in (24).

as shown in Fig. 2. By contrast the weighted Birkhoff average for chaotic orbits converges much more slowly, and this allowed us to identify chaotic trajectories. Note that it requires at least thirty times more iterates to obtain a comparable distinction using Lyapunov exponents. Orbits trapped in islands have rational rotation numbers, and we are able to identify these using the distribution, shown in Fig. 11, of the minimal denominator in an interval of size δ defined by $q_{min}(I_\delta)$ in (13).

The weighted Birkhoff method has the advantages of being extremely simple to compute and that it does not rely on the symmetry used in Greene's residue method. Using a total orbit length of $2(10)^4$, we estimated the break-up parameter for the golden mean invariant circle to 0.3% accuracy, as seen in Fig. 16. While this accuracy does not compete with that of Greene's method nor of conjugacy based methods, it requires much less computation.

Additionally the Birkhoff method does not require fixing the rotation number in advance, nor computing the lengthy Fourier series of conjugacy-based methods. This adds flexibility since, whereas the golden mean is established as the most robust rotation number for the standard map, the rotation number of the most robust invariant circle in a general map is not generally known. For example, the relationship between robustness of invariant circles and noble rotation numbers is less well established for asymmetric maps [FM14], and we have demonstrated in §5 that the
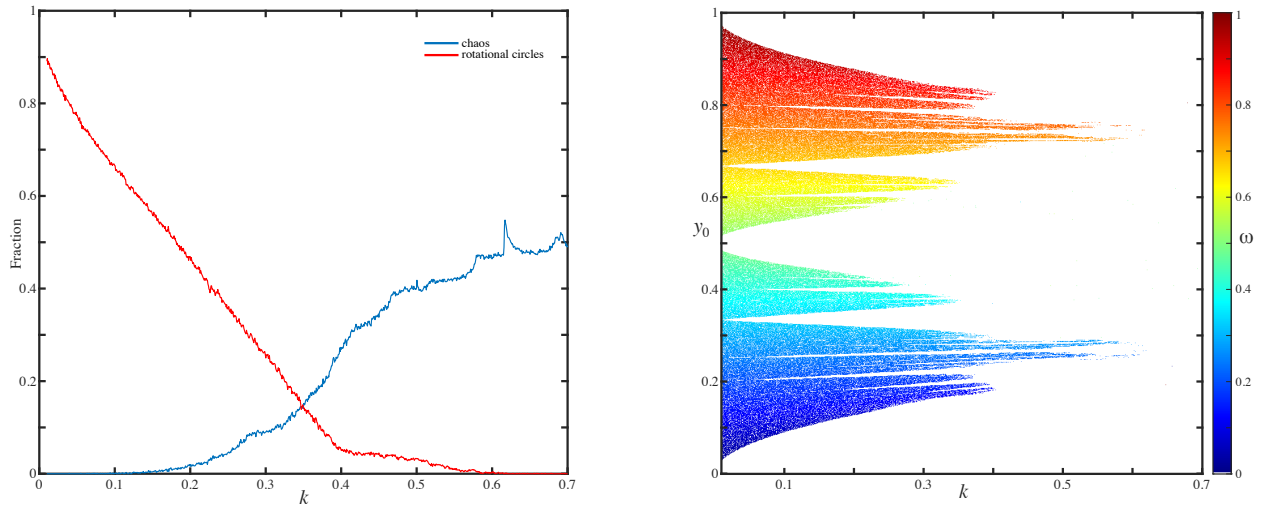
26

Figure 19: (a) Fraction of orbits of the map (1), with two-harmonic force (25) for $\psi = 0.7776$, that are chaotic and are rotational circles for initial conditions along the line $x_0 = 0.35$. (b) The rotation number of the rotational circles as a function of initial y and the parameter k. As in Fig. 18, $T = 10^4$ with distinguishing criteria in (24).

weighted Birkhoff method can compute robustness of invariant circles for such asymmetric maps.

Another potential application of the weighted Birkhoff average is that it can be applied to higher-dimensional maps, with, say, d-dimensional invariant tori. There have been a number of attempts to accurately compute parameters for the break-up of two-dimensional tori by generalizing Greene's criterion [Tom96b, CFL04, FM13]. Though it is known that the residues of a sequence of periodic orbits that limit on a smooth torus do limit to zero [Tom96a, CFL04], this has not led to accurate computations of the parameters at which a torus is destroyed (the best achieved accuracy, about 1%, is for a three-dimensional, volume-preserving map [FM13]). One of the difficulties in any attempt to extend Greene's method is that there is no completely satisfactory continued fraction algorithm for multi-dimensional frequency vectors. To generalize our method will require computing the minimal denominator q_{min} for resonance relations, e.g., finding a minimal $(p, q) \in \mathbb{Z}^{d+1}$ such that $|q\omega - p|$ is small. One possible approach is to use generalized Farey path methods [KO86] that may provide a version of Lem. 1 for this case.

Appendices

A   Farey Paths and the Smallest Denominator

The Farey path (17) for any number x can be computed by the simple method given in Algorithm 1. In a practical calculation, a stopping criterion based on precision must be included. This gives, for
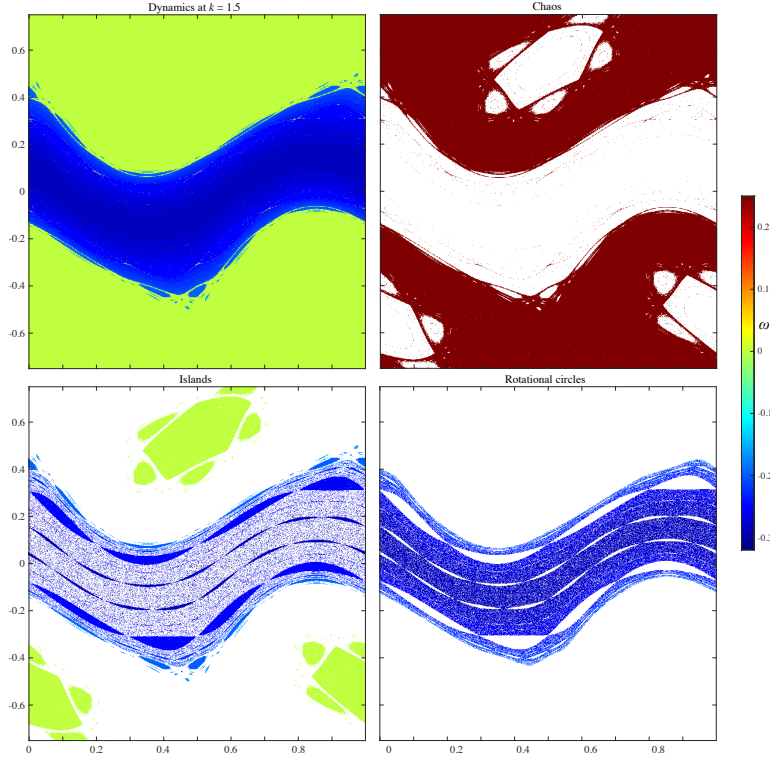
Figure 20: The dynamics of the standard nontwist map, with frequency (26) for $k = 1.5$ and $\delta = 0.3$. The weighted Birkhoff method distinguishes chaotic orbits (upper right), islands (lower left), and rotational circles (lower right). The rotation number of each nonchaotic orbit is color-coded (color bar at right). The computations were performed for a grid of 1000^2 initial conditions in $[0, 1] \times [-0.75, 0.75]$ with $T = 10^4$, using the distinguishing criteria in (24).

example

$$\frac{17}{6} = RRLRRRR = [2; 1, 4, 1],$$
$$\frac{7}{10} = LRRLL = [0; 1, 2, 2, 1],$$
$$e = RRLR^2LRL^4RLR^6LRL^8\ldots = [2; 1, 2, 1, 1, 2, 1, 1, 4, 1, 1, 6, 1, 1, 8, \ldots],$$
$$\pi = R^3L^7R^{15}LR^{292}LRLR^2LR^3L\ldots = [3; 7, 15, 1, 292, 1, 1, 1, 2, 1, 3, 1, \ldots].$$

Note that each element of the continued fraction records the number of repeated Farey symbols. The value of a_0 is nonzero if the Farey path begins with R, otherwise $a_0 = 0$, and a_1 counts the number of leading L's in the path. For the rational case there is an additional last element, which is fixed to be 1.

The Stern-Brocot tree gives a method for finding the rational with the smallest denominator $q_{min}(I)$ (13) in an interval I. Here we prove Lem. 1 to show that q_{min} is the denominator of the first rational on the tree that falls in I:

Proof of Lem. 1. Suppose that for all levels up to ℓ on the Stern-Brocot tree no Farey rational is in I. Since the Farey intervals partition $(0, \infty)$, there must be a Farey interval $J = (\frac{p_l}{q_l}, \frac{p_r}{q_r}) \supset I$
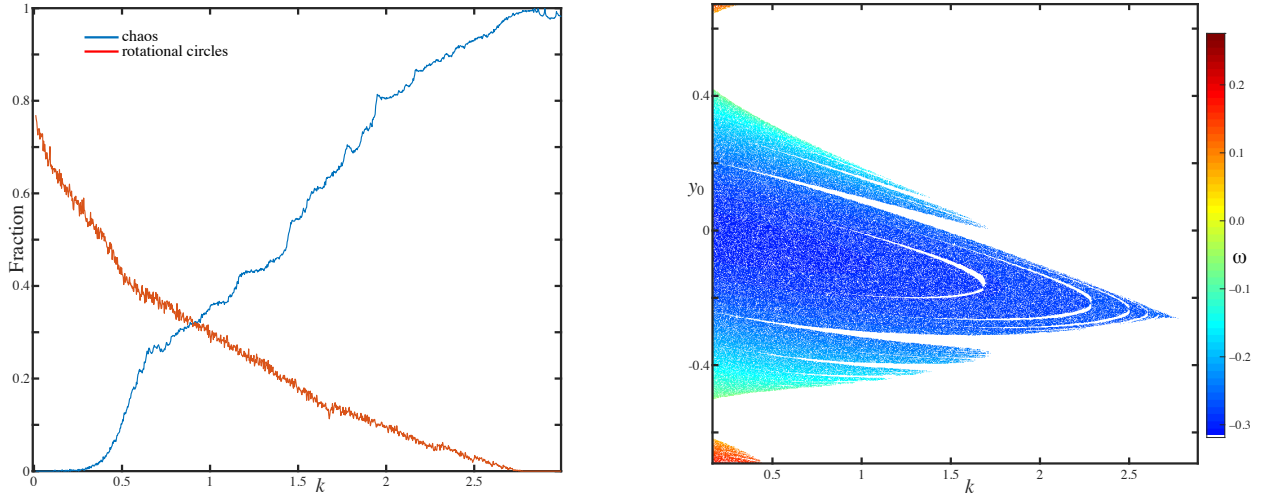
Figure 21: (a) Proportion of orbits of the standard nontwist map with frequency map (26) when $\delta = 0.3$ that are chaotic and are rotational circles. (b) The rotation number of the rotational circles as a function of initial y and the parameter k for orbits with $x_0 = 0.35$. As in Fig. 20, $T = 10^4$ with distinguishing criteria in (24).

for neighbors $\frac{p_l}{q_l}$ and $\frac{p_r}{q_r}$. Note that every number in J and thus every number in I must then be a descendent of these parents. Denote the mediant (18) by p_m/q_m. Without loss of generality, we can assume that $p_m/q_m \in I$. Every rational in the level $\ell + 1$ daughter interval $\left(\frac{p_l}{q_l}, \frac{p_m}{q_m}\right)$ is a descendent of $\frac{p_m}{q_m}$ and since all of these are formed by the mediant operation all of these denominators are larger than q_m. The same is true for the upper interval $\left(\frac{p_m}{q_m}, \frac{p_r}{q_r}\right)$. Since $I \subset \left(\frac{p_l}{q_l}, \frac{p_m}{q_m}\right) \cup \left\{\frac{p_m}{q_m}\right\} \cup \left(\frac{p_m}{q_m}, \frac{p_r}{q_r}\right)$, all remaining rationals in I have denominator greater than q_m. Consequently $q_{min}(I) = q_m$, and, moreover, the rational with minimal denominator is unique. $\qquad \square$

This result is encapsulated in Algorithm 2 to give a computation of the smallest denominator rational in $I_\delta(x)$ (12). For example, this algorithm gives

$$q_{min}(I_{10^{-8}}(\pi)) = 32085, \quad \frac{p_m}{q_m} = [3; 7, 15, 1, 283],$$

$$q_{min}((I_{10^{-10}}(e)) = 154257, \quad \frac{p_m}{q_m} = [2; 1, 2, 1, 1, 2, 1, 1, 4, 1, 1, 6, 1, 1, 8, 1, 1, 8].$$

Neither of these are convergents of the continued fraction expansions. Algorithm 2 ignores issues of finite precision arithmetic, and is not efficient if the Farey path has a long string of repeated symbols. An algorithm that does not have this deficit is given in [CP16].

We can obtain some additional understanding of the smallest denominator for the specific case when the bounds of the interval I are arbitrary rationals [Siv16],

$$I = \left(\frac{p_l}{q_l}, \frac{p_r}{q_r}\right). \tag{28}$$

To find the smallest denominator rational we expand each of the boundary points in their Farey paths:

$$a = \frac{p_l}{q_l} = a_0 a_1 a_2 \ldots a_m, \quad b = \frac{p_r}{q_r} = b_0 b_1 b_2 \ldots b_n,$$

Algorithm 1 Compute the Farey path for $x \in \mathbb{R}^+$ assuming exact arithmetic

---

    procedure FAREYPATH(x)
        $i \leftarrow 1$
        while $x \neq 1$ **do**
            if $x < 1$ **then**
                $s_i = L$
                $x \leftarrow \frac{x}{1-x}$
            else
                $s_i = R$
                $x \leftarrow x - 1$
            end if
            $i \leftarrow i + 1$
        end while
    end procedure

Algorithm 2 Find the smallest rational in the interval $I_\delta(x)$

---

    procedure SMALLDENOM(x, δ)
        $(n, d) = (p_l, q_l) = (0, 1)$
        $(p_r, q_r) = (1, 0)$
        while $|x - \frac{n}{d}| \geq \delta$ **do**
            $(n, d) = (p_l + p_r, q_l + q_r)$ $\triangleright$ Find the mediant
            if $x < n/d$ **then**
                $(p_r, q_r) = (n, d)$ $\triangleright$ $x \in \left(\frac{p_l}{q_l}, \frac{n}{d}\right)$
            else
                $(p_l, q_l) = (n, d)$ $\triangleright$ $x \in \left[\frac{n}{d}, \frac{p_r}{q_r}\right)$
            end if
        end while
        return (n, d) $\triangleright$ The smallest rational is $\frac{n}{d}$
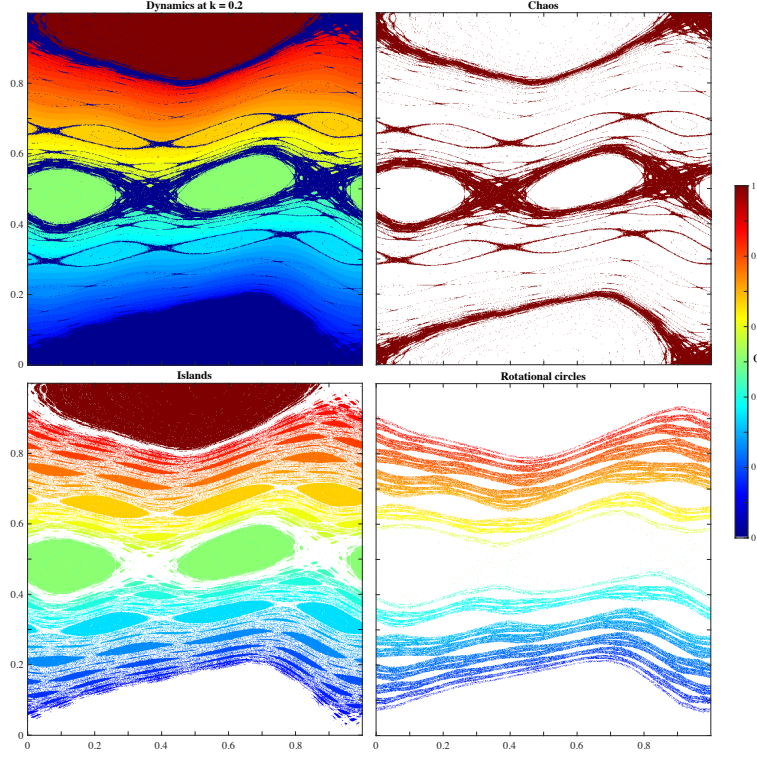   **end procedure**

Figure 22: The dynamics of the asymmetric two harmonic map (27) for $k = 0.2$ and $\psi = 0.7776$. The weighted Birkhoff method distinguishes chaotic orbits (upper right), islands (lower left), and rotational circles (lower right). The rotation number of each nonchaotic orbit is color-coded (color bar at right). The computations were performed for a grid of 1000^2 initial conditions in $[0, 1]^2$ with $T = 10^4$ with distinguishing criteria in (24).

with $a_i, b_i \in \{L, R\}$. Then, as shown by [Siv16, Thm. 1], there are three cases:

1. When the boundary points of (28) are Farey neighbors, the smallest rational in I is the mediant, so $q_{min} = q_l + q_r$.

2. If one Farey path is a subsequence of the other but they are not neighbors, then the smallest rational is a daughter of the shorter path and an ancestor of the longer. For example, if $a = b_1 b_2 \ldots b_n a_{n+1} a_{n+2} \ldots a_m = b a_{n+1} \ldots a_m$, then the smallest rational has the path

$$\frac{p}{q} = b a_{n+1} \ldots a_k,$$

for some $k < m$. This is the appropriate daughter of b and ancestor of a. Note that when $a < b$, then it must be the case that $a_{n+1} = L$. If, for example $a = bLRL\ldots$, then $bL < a < bLR < b$, so then we set $k = n + 2$, and obtain $\frac{p}{q} = bLR$.

3. If neither path is a subsequence of the other, then the smallest rational is the unique rational that is a common ancestor of both on the tree: the longest Farey path for which they agree. For example, if $a_i = b_i$ for $i = 0, \ldots k < \min(m, n)$, and $a_{k+1} \neq b_{k+1}$ then $\frac{p}{q} = a_0 a_1, \ldots, a_k$ is the smallest rational in I.
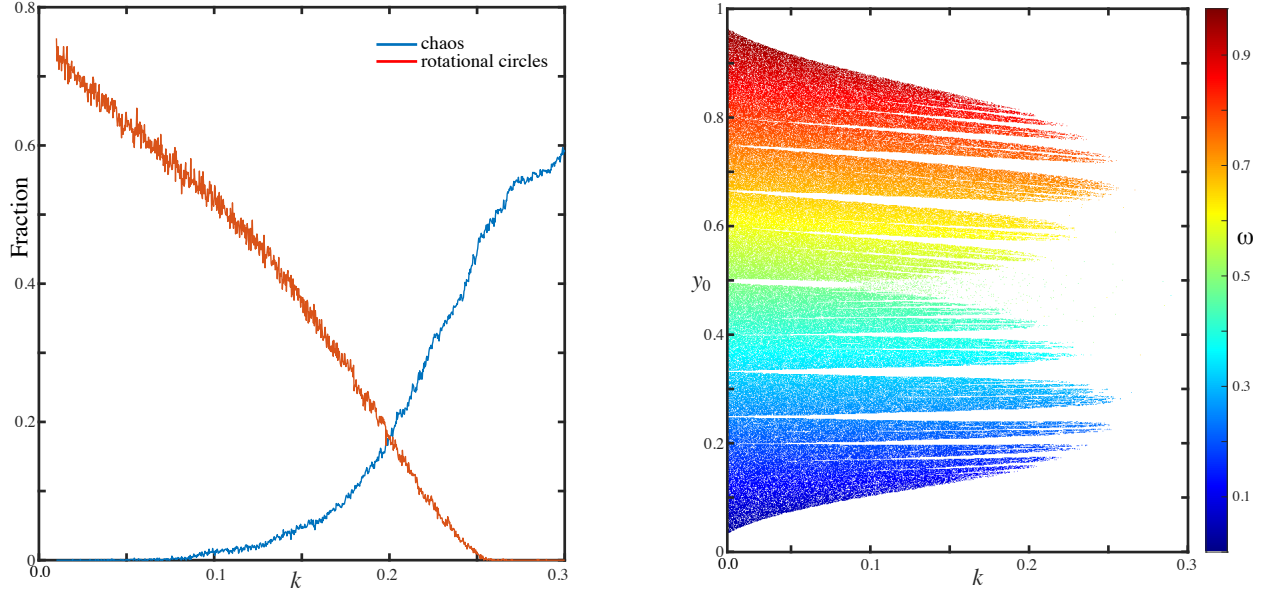
Figure 23: (a) Proportion of orbits of the asymmetric standard map with force (27) and $\psi = 0.7776$ that are chaotic, and rotational circles. (b) The rotation number of the rotational circles as a function of initial y and the parameter k for orbits with $x_0 = 0.35$. As in Fig. 22, $T = 10^4$ with distinguishing criteria in (24).

Finally, for an interval bounded by irrationals we can prove the following lemma.

Lemma 2 (Smallest Rational in an Irrational Interval). *If $I = (a, b)$, $0 < a < b$, $a, b \in \mathbb{R} \setminus \mathbb{Q}$, then $q_{min}(I)$ is the denominator of the common Farey ancestor of a and b, if there is one; otherwise $q_{min}(I) = 1$.*

Proof. Denote the infinite Farey paths of the irrationals by $a = a_1 a_2 \ldots$ and $b = b_1 b_2 \ldots$, where $a_i, b_i \in \{L, R\}$, and let $\ell \in \mathbb{N}$ be chosen so that the common ancestor of a and b is

$$\frac{p_\ell}{q_\ell} = a_1 a_2 \ldots a_\ell = b_1 b_2 \ldots b_\ell, \quad a_{\ell+1} \neq b_{\ell+1}.$$

If ℓ does not exist, then since $a < b$, $a_1 = L$ and $b_1 = R$, which means that $\frac{1}{1} \in I$, so that $q_{min} = 1$.

Now suppose that there is a common ancestor of length $\ell \geq 1$. Then since $a < b$, we must have $a_{\ell+1} = L$ and $b_{\ell+1} = R$ and $a < \frac{p_\ell}{q_\ell} < b$. Denote a "left truncation" of a path as a rational $a_L = a_1 a_2 \ldots a_j < a$ and a "right truncation" as a rational $a_R = a_1 a_2 \ldots a_k > a$, see Fig. 24. For example if $a_{j+1} = R$, and $a_{k+1} = L$ then we know that $a_1 a_2 \ldots a_j < a < a_1 a_2 \ldots a_k$. Note that such truncations always exist for any irrational and any choice of minimal length since the infinite paths with tails $\ldots L^\infty$ and $\ldots R^\infty$ are rationals. Now, by item (3) above [Siv16, Thm. 1], for the interval $I_{outer} = (a_L, b_R)$, the smallest denominator is that of the common Farey ancestor of a_L and b_R: $q_{min}(I_{outer}) = q_\ell$. Thus, whenever these rational truncations are both longer than ℓ, then I_{outer} contains the common Farey ancestor $\frac{p_\ell}{q_\ell}$ and this has the smallest denominator. Note that since $a_L < a$ and $b_R > b$, then $I \subset I_{outer}$. Thus $q_{min}(I)$ is no less than q_ℓ. Moreover since $\frac{p_\ell}{q_\ell} \in I$, then $q_{min}(I)$ is no more than q_ℓ. Thus $q_{min}(I) = q_\ell$. $\qquad \square$
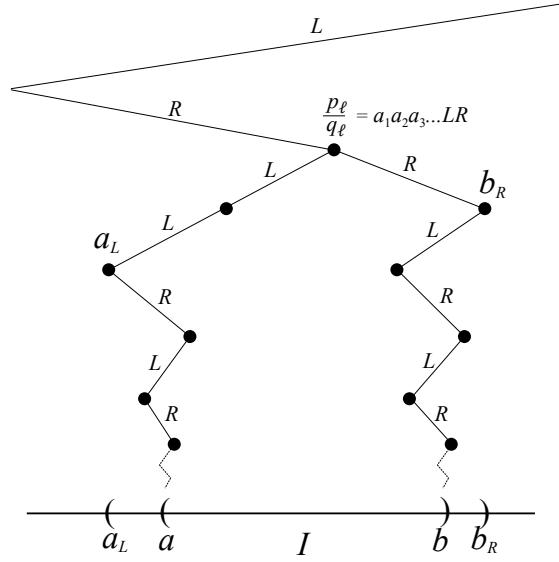
32

Figure 24: An interval $I = (a, b)$ bounded by a pair of irrationals with the outer interval (a_L, b_R), and their common Farey Ancestor $\frac{p_\ell}{q_\ell}$

References

[AC15]   C. V. Abud and I. L. Caldas. On Slater's criterion for the breakup of invariant curves. *Physica D*, 308:34–39, 2015. `https://doi.org/10.1016/j.physd.2015.06.005`.

[ACP06]   E.G. Altmann, G. Cristadoro, and D. Paz. Nontwist non-Hamiltonian systems. *Phys. Rev. E*, 73(5):056201, 2006. `http://link.aps.org/abstract/PRE/v73/e056201`.

[BBdA98] S.J. Beslin, D.J. Baney, and V. de Angelis. Small denominators: No small problem. *Mathematics Magazine*, 71(2):132–138, 1998. `http://www.jstor.org/stable/2691018`.

[BBGT96] R. Bartolini, A. Bazzani, M. Giovanniozzi, and E. Todesco. Tune evaluation in simulations and experiments. *Particle Accelerators*, 52:147–177, 1996. `http://cdsweb.cern.ch/record/292773?ln=en`.

[CB09]   E. Charrier and L. Buzer. Approximating a real number by a rational number with a limited denominator: A geometric approach. *Discrete Applied Mathematics*, 157(16):3473–3484, August 2009. `https://doi.org/10.1016/j.dam.2009.03.005`.

[CFL04]   A. Celletti, C. Falcolini, and U. Locatelli. On the break-down threshold of invariant tori in four dimensional maps. *Regul. Chaotic Dyn.*, 9(3):227–253, 2004. `https://doi.org/10.1070/RD2004v009n03ABEH000278`.

[Chi79]   B.V. Chirikov. A universal instability of many-dimensional oscillator systems. *Phys. Rep.*, 52(5):263–379, 1979. `https://doi.org/10.1016/0370-1573(79)90023-1`.

[CP16]   M. Citterio and R. Pavani. A fast computation of the best k-digit rational approximation to a real number. *Mediterranean Journal of Mathematics*, 13(6):4321–4331, December 2016. `http://link.springer.com/10.1007/s00009-016-0747-z`.

[CS00]   P.M. Cincotta and C. Simó. Simple tools to study global dynamics in non-axisymmetric galactic potentials – I. *Astron. Astrophys. Suppl. Ser.*, 147:205–228, 2000.

[DDS⁺16] S. Das, C.B. Dock, Y. Saiki, M. Salgado-Flores, E. Sander, J. Wu, and J.A. Yorke. Measuring quasiperiodicity. *Euro. Phys. Lett.*, 114:40005, 2016. `https://doi.org/10.1209/0295-5075/114/40005`.

[dlL01]   R. de la Llave. A tutorial on KAM theory. In *Smooth Ergodic Theory and Its Applications (Seattle, Wa, 1999)*, volume 69 of *Proc. Sympos. Pure Math.*, pages 175–292. Amer. Math. Soc., Providence, 2001.

[DSSY16]   S. Das, Y. Saiki, E. Sander, and J.A. Yorke. Quasiperiodicity: Rotation numbers. In C. Skiadas, editor, *The Foundations of Chaos Revisited: From Poincaré to Recent Advancement*, Understanding Complex Systems. Springer, Cham, 2016. `https://doi.org/10.1007/978-3-319-29701-9_7`.

[DSSY17]   S. Das, Y. Saiki, E. Sander, and J.A. Yorke. Quantitative quasiperiodicity. *Nonlinearity*, 30(11):4111, 2017. `https://doi.org/10.1088/1361-6544/aa84c2`.

[DY18]   S. Das and J.A. Yorke. Super convergence of ergodic averages for quasiperiodic orbits. *Nonlinearity*, 31(2):491–501, 2018. `https://doi.org/10.1088/1361-6544/aa99a0`.

[EV01]   K. Efstathiou and N. Voglis. A method for accurate computation of the rotation and twist numbers for invariant tori. *Physica D*, 158:151–163, 2001. `https://doi.org/10.1016/S0167-2789(01)00299-8`.

[FHL17]   J.-L. Figueras, A. Haro, and A. Luque. Rigorous computer-assisted application of kam theory: A modern approach. *Foundations of Computational Mathematics*, 17(5):1123–1193, 2017. `https://doi.org/10.1007/s10208-016-9339-3`.

[FLG97]   Claude Froeschlé, Elena Lega, and Robert Gonczi. Fast Lyapunov indicators. Application to asteroidal motion. *Celestial Mech. Dynam. Astronom.*, 67(1):41–62, 1997. `https://doi.org/10.1023/A:1008276418601`.

[FM13]   A.M. Fox and J.D. Meiss. Greene's residue criterion for the breakup of invariant tori of volume-preserving maps. *Physica D*, 243(1):45–63, 2013. `https://doi.org/10.1016/j.physd.2012.09.005`.

[FM14]   A.M. Fox and J.D. Meiss. Critical invariant circles in asymmetric, multiharmonic generalized standard maps. *Comm. Nonlinear Sci. Numer. Simulat.*, 19(4):1004–1026, 2014. `https://doi.org/10.1016/j.cnsns.2013.07.028`.

[For07]   M. Forisek. Approximating rational numbers by fractions. In P. Crescenzi, G. Prencipe, and G. Pucci, editors, *Fun with Algorithms*, volume LNCS 4475, pages 156–165. Springer-Verlag, Berlin, 2007. `https://doi.org/10.1007/978-3-540-72914-3`.

[GJSS87]   J.M. Greene, H. Johannesson, B. Schaub, and H. Suhl. Scaling anomaly at the critical transistion of an incommensurate structure. *Phys. Rev. A*, 36:5858–5861, 1987. `https://doi.org/10.1103/PhysRevA.36.5858`.

[GM09]   G.A. Gottwald and I. Melbourne. On the implementation of the 0–1 test for chaos. *SIAM J. Appl. Dyn. Sys.*, 8(1):129–145, 2009. `https://doi.org/10.1137/080718851`.

[GMS10a]   G. Gómez, J.M. Mondelo, and C Simó. A collocation method for the numerical Fourier analysis of quasi-periodic functions. I: Numerical tests and examples. *Discrete Contin. Dyn. Syst. Ser. B*, 14:41–74, 2010. `https://doi.org/10.3934/dcdsb.2010.14.41`.

[GMS10b]   G. Gómez, J.M. Mondelo, and C Simó. A collocation method for the numerical Fourier analysis of quasi-periodic functions. II: Analytical error estimates. *Discrete Contin. Dyn. Syst. Ser. B*, 14(1):75–109, 2010. `http://aimsciences.org//article/id/6d6c433a-a1e1-43d2-9953-835620b501de`.

[Gre79]   J.M. Greene. A method for determining a stochastic transition. *J. Math. Phys.*, 20:1183–1201, 1979. `https://doi.org/10.1063/1.524170`.

[HCF⁺16]   A. Haro, M. Canadell, J-L Figueras, A.-L. Josep, and M. Mondelo. *The Parameterization Method for Invariant Manifolds from Rigorous Results to Effective Computations*. Springer International, 2016. `https://doi.org/10.1007/978-3-319-29662-3`.

[HdlL06]   A. Haro and R. de la Llave. A parameterization method for the computation of invariant tori and their whiskers in quasi-periodic maps: Numerical algorithms. *Disc. Cont. Dyn. Sys.*, B6(6):1261–1300, 2006. `https://doi.org/10.3934/dcdsb.2006.6.1261`.

[HdlLS12]   G. Huguet, R. de la Llave, and Y. Sire. Computation of whiskered invariant tori and their associated manifolds: New fast algorithms. *Discrete Contin. Dyn. Syst. Ser. B*, 32(4):1309–1353, 2012. `https://doi.org/10.3934/dcds.2012.32.1309`.

[HW79]   G.H. Hardy and E.M. Wright. *An Introduction to the Theory of Numbers*. Oxford Univ. Press, Oxford, 1979.

[Jun91]   I. Jungreis. A method for proving that monotone twist maps have no invariant circles. *Erg. Th. Dyn. Sys.*, 11:79–84, 1991.

[KO86]   S. Kim and S. Ostlund. Simultaneous rational approximations in the study of dynamical systems. *Phys. Rev. A*, 34:3426–3434, 1986. `https://doi.org/10.1103/PhysRevA.34.3426`.

[LFC92]   J. Laskar, C. Froeschlé, and A. Celletti. The measure of chaos by the numerical analysis of the fundamental frequencies. Application to the standard mapping. *Physica D*, 56:253–269, 1992. `https://doi.org/10.1016/0167-2789(92)90028-L`.

[LM10]   Z. Levnajić and I. Mezić. Ergodic theory and visualization. I. Mesochronic plots for visualization of ergodic partition and invariant sets. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(3):033114, 2010. `https://doi.org/10.1063/1.3458896`.

[LV09]   A. Luque and J. Villanueva. Numerical computation of rotation numbers of quasi-periodic planar curves. *Physica D*, 238(20):2025–2044, 2009. `https://doi.org/10.1016/j.physd.2009.07.014`.

[LV14]   A. Luque and J. Villanueva. Quasi-periodic frequency analysis using averaging-extrapolation methods. *SIAM J. Dyn.Sys.*, 13(1):1–46, 2014. `https://doi.org/10.1137/130920113`.

[Mac93]   R.S. MacKay. *Renormalisation in Area-Preserving Maps*, volume 6 of *Adv. Series in Nonlinear Dynamics*. World Scientific, Singapore, 1993.

[May88]   D.H. Mayer. On the distribution of recurrence times in nonlinear systems. *Lett. Math. Phys.*, 16(2):139–143, 1988. `https://doi.org/10.1007/BF00402021`.

[Mei92]   J.D. Meiss. Symplectic maps, variational principles, and transport. *Rev. Mod. Phys.*, 64(3):795–848, 1992. `https://doi.org/10.1103/RevModPhys.64.795`.

[MP85]   R.S. MacKay and I.C. Percival. Converse KAM: Theory and practice. *Comm. Math. Phys.*, 98:469–512, 1985. `https://doi.org/10.1007/BF01209326`.

[MS91]   S. Marmi and J. Stark. On the standard map critical function. *Nonlinearity*, 5(3):743–761, 1991. `https://doi.org/10.1088/0951-7715/5/3/007`.

[MS92]   R.S. MacKay and J. Stark. Locally most robust circles and boundary circles for area-preserving maps. *Nonlinearity*, 5:867–888, 1992. `http://iopscience.iop.org/0951-7715/5/4/002`.

[SGL16]   C. Skokos, G.A. Gottwald, and J. Laskar. *Chaos Detection and Predictability*, volume 915 of *Lecture Notes in Physics*. Springer, Heidelberg, 2016. `https://doi.org/10.1007/978-3-662-48410-4`.

[Sha92]   J. Shallit. Real numbers with bounded partial quotients: A survey. *L'Enseignement Mathematique*, 38:151–187, 1992. `http://www.unige.ch/math/EnsMath/EM_en/`.

[Sim18]   C Simó. Some questions looking for answers in dynamical systems. *Disc. Cont. Dyn. Sys.*, 38(12):6215–6239, 2018. `https://doi.org/10.3934/dcds.2018267`.

[Siv16]   I. Sivignon. A note on the computation of the fraction of smallest denominator in between two irreducible fractions. *Discrete Applied Mathematics*, 202:197–201, 2016. `http://www.sciencedirect.com/science/article/pii/S0166218X15004448`.

[Sko01]   Ch. Skokos. Alignment indices: a new, simple method for determining the ordered or chaotic nature of orbits. *J. Phys. A*, 34(47):10029–10043, 2001.

[Sla50]   N.B. Slater. The distribution of the integers N for which $\{N\theta\} < \epsilon$. *Proc. Cambridge Philos. Soc.*, 46:525–534, 1950. `https://doi.org/10.1017/S0305004100026086`.

[Sla67]   N.B. Slater. Gaps and steps for the sequence $n\theta \mod 1$. *Proc. Cambridge Philos. Soc.*, 63:1115–1123, 1967. `https://doi.org/10.1017/S0305004100042195`.

[SLV05]   J.D. Szezech, S.R. Lopes, and R.L. Viana. Finite-time Lyapunov spectrum for chaotic orbits of non-integrable Hamiltonian systems. *Phys. Lett. A*, 335(5-6):394–401, 2005. `https://doi.org/10.1016/j.physleta.2004.12.058`.

[SMS+18]   M.S. Santos, M. Mugnaine, J.D. Szezech, A.M. Batista, I.L. Caldas, M.S. Baptista, and R.L. Viana. Recurrence-based analysis of barrier breakup in the standard nontwist map. *Chaos*, 28(8):085717, 2018. `https://doi.org/10.1063/1.5021544`.

[SMS+19]   M.S. Santos, M. Mugnaine, J.D. Szezech, A.M. Batista, I.L. Caldas, and R.L. Viana. Using rotation number to detect sticky orbits in Hamiltonian systems. *Chaos*, 29(4):043125, 2019. `https://doi.org/10.1063/1.5078533`.

[SSC+13]   J.D. Szezech, A.B. Schelin, I.L. Caldas, S R. Lopes, P.J. Morrison, and R.L. Viana. Finite-time rotation number: A fast indicator for chaotic dynamical structures. *Phys. Lett. A*, 377:452–456, 2013. `https://doi.org/10.1016/j.physleta.2012.12.013`.

[Ste13]   C. L. Stewart. On the distribution of small denominators in the farey series of order n. In I.S. Kotsireas and E.V. Zima, editors, *Advances in Combinatorics*, pages 275–286, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. `https://doi.org/10.1007/978-3-642-30979-315`.

[SV06]   T.M. Seara and J. Villanueva. On the numerical computation of Diophantine rotation numbers of analytic circle maps. *Physica D*, 217(2):107–120, 2006. `https://doi.org/10.1016/J.Physd.2006.03.013`.

[Tom96a]   S. Tompaidis. Approximation of invariant surfaces by periodic orbits in high-dimensional maps. some rigorous results. *Experiment. Math.*, 5:197–209, 1996. `https://doi.org/10.1080/10586458.1996.10504588`.

[Tom96b]   S. Tompaidis. Numerical study of invariant sets of a quasiperiodic perturbation of a symplectic map. *Experiment. Math.*, 5:211–230, 1996. `https://doi.org/10.1080/10586458.1996.10504589`.

[Woo78]   F. Woodcock. The evaluation of yes/no forecasts for scientific and administrative purposes. *Monthly Weather Review*, 104(10):1209–1214, 1978. `https://doi.org/10.1175/1520-0493(1976)104<1209:TEOYFF>2.0.CO;2`.

[ZTRK07]   Y. Zou, M. Thiel, M.C. Romano, and J. Kurths. Characterization of stickiness by means of recurrence. *Chaos*, 17:043101, 2007. `http://link.aip.org/link/?CHAOEH/17/043101/1`.