

On spurious detection of linear response and misuse of the fluctuation–dissipation theorem in finite time series

Article

Accepted Version

Creative Commons: Attribution-Noncommercial-No Derivative Works 4.0

Gottwald, G. A., Wormell, J. P. and Wouters, J. ORCID: <https://orcid.org/0000-0001-5418-7657> (2016) On spurious detection of linear response and misuse of the fluctuation–dissipation theorem in finite time series. *Physica D: Nonlinear Phenomena*, 331. pp. 89-101. ISSN 0167-2789 doi: <https://doi.org/10.1016/j.physd.2016.05.010> Available at <https://centaur.reading.ac.uk/81414/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1016/j.physd.2016.05.010>

Publisher: Elsevier

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

On spurious detection of linear response and misuse of the fluctuation-dissipation theorem in finite time series

Georg A. Gottwald^a, J. P. Wormell^a, Jeroen Wouters^{a,b}

^a*School of Mathematics and Statistics, University of Sydney, NSW 2006, Australia*

^b*Universität Hamburg Geowissenschaften, Meteorologisches Institut, Bundesstr. 55, 20146 Hamburg, Germany*

Abstract

Using a sensitive statistical test we determine whether or not one can detect the breakdown of linear response given observations of deterministic dynamical systems. A goodness-of-fit statistics is developed for a linear statistical model of the observations, based on results for central limit theorems for deterministic dynamical systems, and used to detect linear response breakdown. We apply the method to discrete maps which do not obey linear response and show that the successful detection of breakdown depends on the length of the time series, the magnitude of the perturbation and on the choice of the observable.

We find that in order to reliably reject the assumption of linear response for typical observables sufficiently large data sets are needed. Even for simple systems such as the logistic map, one needs of the order of 10^6 observations to reliably detect the breakdown with a confidence level of 95%; if less observations are available one may be falsely led to conclude that linear response theory is valid. The amount of data required is larger the smaller the applied perturbation. For judiciously chosen observables the necessary amount of data can be drastically reduced, but requires detailed *a priori* knowledge about the invariant measure which is typically not available for complex dynamical systems.

Furthermore we explore the use of the fluctuation-dissipation theorem (FDT) in cases with limited data length or coarse-graining of observations. The

Email addresses: georg.gottwald@sydney.edu.au (Georg A. Gottwald), jpwormell@gmail.com (J. P. Wormell), jeroen.wouters@uni-hamburg.de (Jeroen Wouters)

FDT, if applied naively to a system without linear response, is shown to be very sensitive to the details of the sampling method, resulting in erroneous predictions of the response.

Keywords: linear response theory; fluctuation-dissipation theorem; climate science

1. Introduction

An important question in the study of probabilistic properties of dynamical systems is how to determine the response of a system if subjected to a small perturbation. For example, in climate science we would like to know how the global mean temperature changes upon increasing CO₂ levels. This problem was solved in statistical physics in the context of thermostatted Hamiltonian systems, establishing the framework of linear response theory [45, 11, 65, 51]. In essence, linear response theory employs a Taylor expansion of the perturbed invariant measure around the unperturbed equilibrium measure; this then allows to calculate averages of observables in the perturbed system entirely from knowledge of the statistics of the unperturbed system.

The study of linear response involves two issues: proving differentiability of the response and finding an expression for the derivative of the response. To establish linear response, the invariant measure needs to be differentiable with respect to the parameter describing the magnitude of the perturbation. For the existence of an analytical formula for the response in terms of the equilibrium fluctuations of the unperturbed system, which is the statement of the celebrated fluctuation-dissipation theorem (FDT), the invariant measure needs additionally to be differentiable with respect to the phase space variables.

Applying this framework to deterministic dynamical systems, in particular to forced dissipative systems whose dynamics evolves on an attractor of zero Lebesgue measure in the full space, has been a challenge. In a series of papers, Ruelle showed that the response is linear for the class of uniformly hyperbolic Axiom A systems, i.e. the invariant measure is differentiable with respect to the magnitude of the perturbation [59, 60, 61, 62].

Due to the singular nature of the invariant measure of forced dissipative systems the fluctuation-dissipation theorem, however, cannot hold. Heuristically this failure can be understood by realizing that typical perturbations will have a non-zero projection along the stable manifold, generally transverse to the attractor, whereas the invariant measure is supported entirely

on the attractor. Therefore one cannot estimate the response by solely considering correlations of the unperturbed system. A linear response formula can still be expressed, but involves the full linear tangent dynamics and must take into account the evolution of exponentially attenuated perturbations along stable directions rather than just the unperturbed fluctuations along the unstable manifolds as in the FDT.

The hope that linear response theory can be extended to more general chaotic dynamical systems has been dampened by numerical results on the tent map [28] and rigorous analysis by Baladi and co-workers [7, 8, 6, 10, 24]. In particular, it was shown that the logistic map does not obey linear response. This is due to the non-smooth changes of the invariant measure when perturbing from a chaotic parameter value to a periodic one or vice versa. Even worse, even when restricting to the Cantor set of chaotic parameter values the measure is not differentiable in the sense of Whitney. On the other hand, there are numerical simulations suggesting that linear response might exist for some examples of non-uniformly hyperbolic systems [57, 16, 48] including the Lorenz '63 system which involves homoclinic tangencies. Furthermore, the lack of structural stability, which was believed to be an obstruction to linear response theory in the climate system [52], does not preclude the existence of linear response as was rigorously shown in [26]. The current belief in the mathematical community is that a sufficient condition for the existence of linear response is the summability of the correlation function; the summability of the correlation function is, however, shown not to be necessary for general observables [44, 9].

Notwithstanding the lack of rigorous mathematical proofs for its validity for general forced dissipative non-equilibrium systems, linear response theory has been taken up in the climate sciences to predict the response of the climate, as was first proposed by Leith [47]. Linear response theory and the fluctuation-dissipation theorem have since been used with some success by several groups. They have been applied to various toy models related to atmospheric chaos [50, 48, 1, 2, 22, 23], barotropic models [12, 34, 3], quasi-geostrophic models [27], atmospheric models [55, 18, 35, 33, 36, 58, 37] and coupled climate models [46, 43, 29, 56]. These successes have led scientists to believe that high-dimensional complex systems may very well obey linear response. The standard argument is that complex systems involve a multitude of interacting processes on several temporal and spatial scales and behave effectively stochastically with a smooth invariant measure [50]. This point of view seems at least reasonable for observables of the slow dynamics

of complex multi-scale systems which in the limit of infinite time-scale separation are asymptotically stochastic [54, 32, 42]. In the case of stochastic dynamical systems linear response theory can indeed be justified [39, 38]. However, several instances are now known where atmospheric and oceanic dynamics exhibits a rough dependence on parameters [17], and where, even if linear response theory is observed, the fluctuation-dissipation theorem is not valid [21].

On a more fundamental level, however, it is by no means clear that high-dimensional complex systems do obey linear response theory. In this paper we do not attempt to answer this question. Rather, we consider the following practical issue: systems which do not obey linear response theory are observed with *finite* time series. In such cases we seek to show that the breakdown might not be detectable, and the system's observed behavior may appear consistent with linear response theory. Moreover, the choice of the observable is crucial for the detectability of the breakdown of linear response in finite time series. In particular, we will show that global observables are less able to detect the non-smoothness of the invariant measure whereas local observables which hone in on the roughness of the invariant measure will make the non-smoothness apparent for smaller amounts of data. Finally, the perturbation size also impacts on the detectability of breakdown, with smaller perturbations requiring more data for successful breakdown detection.

This work is motivated by the contradiction between the reported success of linear response theory in the climate sciences and rigorous mathematical results proving the non-existence of linear response theory for a large class of dynamical systems.

The paper is organized as follows. In Section 2 we briefly review linear response theory and the fluctuation-dissipation theorem. In Section 3 we propose a goodness-of-fit test to probe for the validity of linear response in time series. In Section 4 we discuss the logistic map, demonstrate the mechanism leading to the breakdown of linear response for this one-dimensional map and show how this breakdown might not be apparent with time series of insufficient length. We show the effect of finite data size as well as how the choice of the observable can either mask or emphasize the non-smoothness of the invariant measure. In Section 6 we show further that an application

of the FDT in situations where linear response does not exist cannot provide any reliable statistical information, not even in an averaged sense. We conclude with a summary in Section 7.

2. Linear response theory

We consider here a family of dynamical systems $f_\varepsilon : M \rightarrow M$ on some space M . We assume that the map f_ε depends smoothly on the parameter ε and that for each ε the dynamical system admits a unique invariant physical measure μ_ε , e.g. absolutely continuous measures or Sinai-Ruelle-Bowen measures (SRB). An ergodic measure is called physical if for a set of initial conditions of nonzero Lebesgue measure the temporal average of a typical observable converges to the spatial average over this measure. Considering an observable $A : M \rightarrow \mathbb{R}$, we are interested in the change of the average of the observable

$$\langle A \rangle_\varepsilon = \int_M A d\mu_\varepsilon$$

upon varying ε . A system is said to have *linear response* if the derivative

$$\langle A \rangle'_{\varepsilon_0} := \frac{\partial}{\partial \varepsilon} \langle A \rangle_\varepsilon |_{\varepsilon_0}$$

exists. It is obvious that a sufficient condition for linear response is that the invariant measure μ_ε is differentiable with respect to ε . If the limit does not exist, we say there is a breakdown of linear response. We assume that the observable captures sufficient dynamic information about the dynamical system; for example, an odd observable on a system symmetric about 0 would be identically zero regardless of whether the system had a linear response or not.

One may further ask whether, if linear response exists, a computable analytical expression for the linear response

$$\langle A \rangle_\varepsilon \approx \langle A \rangle_{\varepsilon_0} + \langle A \rangle'_{\varepsilon_0} \delta\varepsilon \tag{1}$$

can be found for small values of $\delta\varepsilon = \varepsilon - \varepsilon_0$. To write down an expression of the linear response, we introduce a vector field X as $X \circ f_{\varepsilon_0} := \partial_\varepsilon f_\varepsilon |_{\varepsilon=\varepsilon_0}$. Note that the introduction of the vector field X is the standard way of

formulating perturbations in statistical physics as $f_\varepsilon = f_{\varepsilon_0} + \delta\varepsilon X(f_{\varepsilon_0})$. The linear response $\langle A \rangle'_\varepsilon|_{\varepsilon_0}$ can then be formally expressed as

$$\frac{\partial}{\partial\varepsilon}\langle A \rangle_\varepsilon|_{\varepsilon_0} = \sum_{n=0}^{\infty} \langle X(x) \nabla(A \circ f_{\varepsilon_0}^n)(x) \rangle_{\varepsilon_0}, \quad (2)$$

for $x \in M$. Provided the unperturbed invariant measure μ_{ε_0} has a density $\rho_{\varepsilon_0}(x)$ that is differentiable with respect to $x \in M$ and non-vanishing, one can perform partial integration in (2) to rewrite the linear response in terms of an integral of a correlation function. This form of the linear response formula is known as the fluctuation-dissipation theorem [60, 51] and reads as

$$\frac{\partial}{\partial\varepsilon}\langle A \rangle_\varepsilon|_{\varepsilon_0} = - \sum_{n=0}^{\infty} \left\langle \left(\frac{\nabla(\rho_{\varepsilon_0}(x)X(x))}{\rho_{\varepsilon_0}(x)} \right) A \circ f_{\varepsilon_0}^n(x) \right\rangle_{\varepsilon_0}. \quad (3)$$

In the form (3) the response formula is easier to apply to a numerical integration or to experimental data than the original response formula (2), as it can be estimated directly from a long integration. The assumption that the invariant measure is differentiable is, however, a strong limitation, as it fails for dissipative systems with singular measures with support on an attractor as well as for absolutely continuous maps involving singularities such as the logistic map (see Section 4).

Introducing the notation of a divergence operator with respect to a density ρ

$$\operatorname{div}_\rho B(x) = \frac{\operatorname{div}(\rho B)(x)}{\rho(x)},$$

the response formula (3) can be concisely written as

$$\frac{\partial}{\partial\varepsilon}\langle A \rangle_\varepsilon|_{\varepsilon_0} = - \sum_{n=0}^{\infty} C_n(\operatorname{div}_{\rho_{\varepsilon_0}} X, A), \quad (4)$$

with the correlation function C_n between two observables A and B defined as

$$C_n(A, B) = \langle A B \circ f^n \rangle_{\varepsilon_0} - \langle A \rangle_{\varepsilon_0} \langle B \rangle_{\varepsilon_0}.$$

For sufficiently fast decay of correlations one can estimate (3) from a time series $x_{i=1,\dots,N}$ of finite length N via

$$\frac{\partial}{\partial \varepsilon} \langle A \rangle_\varepsilon |_{\varepsilon_0} \approx - \sum_{n=0}^{n_{\max}} \frac{1}{N-n} \sum_{i=1}^{N-n} (\operatorname{div}_{\rho_{\varepsilon_0}} X)(x_i) A(x_{i+n}) \quad (5)$$

with $1 \ll n_{\max} \ll N$. This expression allows for the estimation of the first-order response to a perturbation using a times series of the unperturbed system, provided the unperturbed density ρ_{ε_0} can be estimated from the time series as well. In the climate sciences ρ_{ε_0} is mostly approximated either via a quasi-Gaussian approximation [33, 36] or by kernel smoothing [22].

Before exploring examples where linear response does not exist, we show in Figure 1 an example of linear response for the doubling map $f_\varepsilon(x) = \operatorname{mod}(2x + \varepsilon \sin(4\pi x), 1)$ which for $\varepsilon = 0$ admits the Lebesgue measure as its invariant measure. Here the fluctuation-dissipation formula (5) becomes

$$\frac{\partial}{\partial \varepsilon} \langle A \rangle_\varepsilon |_{\varepsilon_0} = -\pi$$

and accurately reproduces the actual response. We estimate the actual response numerically using a spectral method. In particular, we approximate the transfer operator which propagates densities under the perturbed dynamics f_ε (see, for example, [5]) by projecting onto a finite number of basis function [25, 15, 63]. For the doubling map we choose 100 trigonometric functions. The invariant measure ρ_ε is then approximated by the eigenfunction corresponding to the eigenvalue 1 of the approximated transfer operator. The advantage of spectral methods over using a long but finite time series with subsequent binning is its high accuracy and fast convergence with the number of resolved eigenfunctions [63]. Their applicability, however, is restricted to low-dimensional systems.

3. Testing for linear response in finite time series

In this section we develop a quantitative goodness-of-fit test for the detectability of linear response in time series of finite size which allows to make statements about the significance of an observed linear response. Given a family of chaotic maps f_ε that may or may not obey linear response, we




Figure 1: Response of the doubling map to the perturbation $X(x) = \sin 2\pi x$ for an observable $A(x) = \cos 2\pi x$. The continuous line represents the actual response, the dashed line depicts the result of the fluctuation-dissipation theorem (5). Note that for small values of the perturbation ε the two curves are indistinguishable.

test for linear response at some reference state with parameter $\varepsilon = \varepsilon_0$ by examining the linear dependency of the response

$$\delta A = \langle A \rangle_\varepsilon - \langle A \rangle_{\varepsilon_0} \quad (6)$$

for $M > 2$ different values of the perturbation parameter $\varepsilon_1, \dots, \varepsilon_M$, by sampling N_1, \dots, N_M consecutive values from the equilibrium dynamics of $f_{\varepsilon_1}, \dots, f_{\varepsilon_M}$, respectively. Explicitly, for each $i = 1, \dots, M$, we have time series $x_n^i = f_{\varepsilon_i}(x_{n-1}^i)$ for $n = 1, \dots, N_M$. The initial conditions x_0^i are distributed according to the physical measure associated with f_{ε_i} .

We consider bounded and continuous observables and assume that for each member of the family f_{ε_i} the autocorrelation $C_j(A, A)$ decays sufficiently rapidly, and that the lengths of the time series N_i are large compared to typical decay times of the autocorrelation function; in practice we choose $N_i \gg \tau_{A, \varepsilon_i}$, where τ_{A, ε_i} is the $1/e$ -folding time of A under the dynamics f_{ε_i} . We further set, for simplicity, $N_i = N$ for all i .

For a large class of chaotic dynamical systems, the sample averages of the

observations

$$\bar{A}_i = \frac{1}{N_i} \sum_{n=1}^{N_i} A(x_n^i) \quad (7)$$

obey the central limit theorem and are distributed asymptotically as $\mathcal{N}(\langle A \rangle_{\varepsilon_i}, \sigma_i^2/N_i)$ [53, 20]. The variances σ_i^2 are given by the Green-Kubo formula in terms of lag-correlations of f_{ε_i} as

$$\sigma_i^2 = C_0(A, A) + 2 \sum_{j=1}^{\infty} C_j(A, A). \quad (8)$$

Numerically, the variances are determined as a Monte-Carlo estimate from observations of the observables under the perturbed dynamics using the central limit theorem. According to the central limit theorem

$$\bar{A}_i = \langle A \rangle_{\varepsilon_i} + \frac{\sigma_i}{\sqrt{N}} \xi_i, \quad (9)$$

for $i = 1, \dots, M$ and *iid* noise $\xi_i \sim N(0, I)$. If the dynamical system indeed has linear response and provided the perturbations $\delta\varepsilon_i = \varepsilon_i - \varepsilon_0$ are sufficiently small, the following statistical model holds for \bar{A}_i

$$\bar{A}_i = \alpha_0 + \alpha_1 \delta\varepsilon_i + \frac{\sigma_i}{\sqrt{N}} \xi_i, \quad (10)$$

with $\alpha_0 = \langle A \rangle_{\varepsilon_0}$ and $\alpha_1 = \langle A \rangle'_{\varepsilon_0}$ for some unperturbed reference state with $\varepsilon = \varepsilon_0$. Note that the ξ_i are independent as the samples from each perturbed system are independent.

To determine the parameters α_0 and α_1 of the model (10) from time series we apply a weighted least squares fit to obtain

$$\begin{pmatrix} \hat{\alpha}_0 \\ \hat{\alpha}_1 \end{pmatrix} = (D^T D)^{-1} D^T Y$$

with the design matrix

$$D = \begin{pmatrix} 1/\sigma_1 & \delta\varepsilon_1/\sigma_1 \\ \vdots & \vdots \\ 1/\sigma_M & \delta\varepsilon_M/\sigma_M \end{pmatrix},$$

and the vector of scaled observations

$$Y = \begin{pmatrix} \bar{A}_1/\sigma_1 \\ \vdots \\ \bar{A}_M/\sigma_M \end{pmatrix}.$$

Higher-order responses can naturally be incorporated by adding a quadratic term $\alpha_2 \delta \varepsilon_i^2$ to (10) and employing higher-order regression allowing, in principle, for a larger range of perturbations (in case linear response exists).

To test whether the observations could have been drawn from the linear model (10) with normally distributed errors ξ_i with mean zero and variance 1, we choose a Pearson χ^2 -test to test the goodness-of-fit with statistics

$$\begin{aligned} \chi^2 &= N \sum_{i=1}^M \left(Y_i - \frac{1}{\sigma_i} (\hat{\alpha}_0 + \hat{\alpha}_1 \varepsilon_i) \right)^2 \\ &= N Y^T (I - H) Y, \end{aligned} \tag{11}$$

where the idempotent hat matrix

$$H = D(D^T D)^{-1} D^T$$

maps scaled observations Y to their linear fits, i.e. $HY = D(\hat{\alpha}_0 \ \hat{\alpha}_1)^T$ [14]. If the response of the underlying dynamical system is linear, χ^2 has a χ^2 -distribution with $M - 2$ degrees of freedom and expectation value $\mathbb{E}\chi_{M-2}^2 = M - 2$. We therefore introduce as a measure for the breakdown of linear response the difference between the χ^2 test statistic for the scaled observations $Y_i = \bar{A}_i/\sigma_i$ and the expectation of the test statistic under the null hypothesis of linear response

$$\mathbf{q} = \frac{1}{N} (\chi^2 - \mathbb{E}\chi_{M-2}^2). \tag{12}$$

Defining W as the vector with components $W_i = \langle A \rangle_{\varepsilon_i} / \sigma_i$ we can use the central limit theorem (9), which holds independent of the existence of linear response, to obtain the following expressions for the mean and variance of the breakdown parameter. The mean is calculated as

$$\mathbb{E}\mathbf{q} = \frac{1}{N} (\mathbb{E}\chi^2 - \mathbb{E}\chi_{M-2}^2)$$

$$\begin{aligned}
&= \mathbb{E} \left(\left(W + \frac{1}{\sqrt{N}} \xi \right)^T (I - H) \left(W + \frac{1}{\sqrt{N}} \xi \right) - \frac{1}{N} \mathbb{E} \chi_{M-2}^2 \right) \\
&= \|W - HW\|^2,
\end{aligned} \tag{13}$$

where we used that H is idempotent. Hence \mathbf{q} is a random variable whose expected value measures the difference between the actual response $\langle A \rangle_{\varepsilon_i}$ and an assumed linear response $\alpha_0 + \alpha_1 \varepsilon_i$ as calculated via least square regression. We have $\mathbb{E} \mathbf{q} \geq 0$ with equality only for $W = HW$, i.e. if the actual response is linear. The variance of the breakdown parameter \mathbf{q} is calculated as

$$\begin{aligned}
\mathbb{V} \mathbf{q} &= \mathbb{E} \left(\left(W + \frac{1}{\sqrt{N}} \xi \right)^T (I - H) \left(W + \frac{1}{\sqrt{N}} \xi \right) - \frac{M-2}{N} - \mathbb{E} \mathbf{q} \right)^2 \\
&= \frac{1}{N} \mathbb{E} \left(\xi^T (I - H) (2W + \frac{1}{\sqrt{N}} \xi) - \frac{M-2}{\sqrt{N}} \right)^2.
\end{aligned}$$

This shows that \mathbf{q} is a consistent estimator for the mismatch $\mathbb{E} \mathbf{q} = \|W - HW\|^2$ since $\mathbb{V} \mathbf{q} \rightarrow 0$ for $N \rightarrow \infty$. In the numerical experiments in Section 5 we will consider Monte-Carlo estimates of the mismatch over realizations \mathbf{q}_j differing in their initial condition and set

$$\hat{\mathbf{q}} = \frac{1}{K} \sum_{j=1}^K \mathbf{q}_j. \tag{14}$$

Now, consider a dynamical system which does not obey linear response, i.e. $\mathbb{E} \mathbf{q} \neq 0$. Using Chebyshev's inequality we have that for all $b < N \mathbb{E} \mathbf{q}$,

$$\begin{aligned}
P(N \mathbf{q} < b) &\leq P(|\mathbf{q} - \mathbb{E} \mathbf{q}| > \mathbb{E} \mathbf{q} - b/N) \\
&\leq \frac{\mathbb{V}(\mathbf{q})}{(\mathbb{E} \mathbf{q} - b/N)^2}.
\end{aligned}$$

Since, as we have shown above, $\mathbb{V} \mathbf{q} \rightarrow 0$ as $N \rightarrow \infty$, we conclude that $N \mathbf{q} \rightarrow \infty$ in probability as $N \rightarrow \infty$. Hence, if F is the cumulative distribution function of the χ_{M-2}^2 distribution, the p -value obtained using the χ^2 -test

$$p = 1 - F(\chi^2) = 1 - F(M - 2 + N \mathbf{q}) \tag{15}$$

converges quickly in probability to zero as $N \rightarrow \infty$ [14]. In practical terms this means that the probability of falsely accepting the null hypothesis of

linear response at any significance level can be made arbitrarily small if N is large enough.

For a specified significance level α we can now define

$$\mathfrak{q}_\alpha = \frac{1}{N} (F^{-1}(1 - \alpha) - (M - 2)) . \quad (16)$$

This defines a threshold value for the observed random variable $\hat{\mathfrak{q}}$ such that if $\hat{\mathfrak{q}} > \mathfrak{q}_\alpha$ the null hypothesis of linear response is rejected with significance level α (i.e. with probability $1 - \alpha$); conversely, if $\hat{\mathfrak{q}} < \mathfrak{q}_\alpha$ the null hypothesis of linear response is accepted with significance level α (i.e. with probability $1 - \alpha$).

The detectability of breakdown of linear response is linked to the amount of available data. As $N \rightarrow \infty$, a breakdown will always become detectable at any specified significance level α . Conversely, if the mismatch $\mathbb{E}\mathfrak{q}$ between the true response of the dynamical system and the linear response is too small and there is an insufficient amount of data available, the actual response will be swamped by the sampling noise, and one will not be able to detect the breakdown of linear response with a reasonable significance level.

In Section 5 we will use our goodness-of-fit test to study the detectability of breakdown of linear response in time series of finite length.

4. Breakdown of linear response theory

A standard dynamical system for which linear response fails [6] is the logistic map $f : [0, 1] \rightarrow [0, 1]$ given by

$$f(x) = ax(1 - x) \quad (17)$$

for $a \in [0, 4]$. This family of maps is particularly well-understood [49, 4]: we can decompose the parameter interval according to $[0, 4] = \mathcal{P} \cup \mathcal{C} \cup \mathcal{N}$ where \mathcal{N} has Lebesgue measure zero, and the asymptotic dynamics consists of a periodic attractor for a in the open and dense set \mathcal{P} and of a strongly chaotic attractor for a in the set \mathcal{C} of positive measure. For $a \in \mathcal{C}$ the logistic map admits a unique absolutely continuous invariant measure (a.c.i.m.) [40, 13] and moreover satisfies the Collet-Eckmann condition [19], i.e. there exists $d > 0$, $\lambda > 1$ such that

$$|Df^n(f(\frac{1}{2}))| \geq d\lambda^n \text{ for all } n \geq 1.$$

Then the Lyapunov exponent is positive, the attractor Λ consists of finitely many intervals $\Lambda_1, \dots, \Lambda_q$ permuted cyclically by f , and $f^q|_{\Lambda_i}$ has exponential decay of correlations for Hölder observables for each $i = 1, \dots, q$ [41, 64].

The logistic map is not uniformly expanding and has a critical point at $x = c = 1/2$ with $f'(c) = 0$. The critical point gives rise to a complicated and rough absolutely continuous invariant measure, because f and its iterates f^n compress the phase space around $x = 1/2$ non-uniformly (see Figure 2). We summarize here the analysis given in [62]. Near the critical point c we approximate $y = f(x) \approx c_1 + \frac{1}{2}f''(c)(x - c)^2$ with $c_n = f^n c$ and hence $x - c = \pm\sqrt{2(c_1 - y)/f''(c)} + O(b - y)$. This implies that an initial smooth density $\rho_0(x)$ including the critical point $x = c$ in its support will evolve under the dynamics into a spike with a square-root singularity at $x = c_1$. Propagating the density for a further time step will transport this peak to $x = c_2$ and create a second, new spike at $x = c_1$, and so forth. The expanding action of the logistic map away from the critical point leads to a broadening of the spikes, and thereby consecutive spikes will have smaller amplitudes, preserving the normalization of the initial density $\rho_0(x)$. This is illustrated in Figure 3 and can be formalized to find an explicit formula for the unique a.c.i.m. in terms of its density

$$\rho(x) = \phi(x) + \sum_{n=1}^{\infty} \eta_n(x). \quad (18)$$

Here $\phi(x)$ represents a continuous background density with $\phi(c_1) = \phi(c_2) = 0$. The countably infinite family of spikes η_n are found to be

$$\eta_n(x) \sim \Upsilon_n \frac{1}{\sqrt{x - f^n c}}, \quad (19)$$

with magnitude

$$\Upsilon_n = \rho(c) \left| \frac{1}{2} f''(c) \prod_{i=1}^n f'(c_i) \right|^{-\frac{1}{2}}. \quad (20)$$

For large n the product in (20) is asymptotically α^n , where $1 < \alpha < 2$ denotes the Lyapunov multiplier, hence the magnitude Υ_n of the spikes decays as $\alpha^{-n/2}$. This implies that the widths of the spikes (defined as the distance from the singularity at $x = f^n c$ at which the amplitudes drop to some chosen threshold) scale like α^{-n} .

We will now study the effect of parameter perturbations $a = a_0(1 + \varepsilon)$ onto the logistic map (17) and its a.c.i.m. $\rho(x)dx$. We may ask how fast the spikes move upon increasing ε . Expanding the displacement length $\ell_\varepsilon = |f_\varepsilon^{n+1}(c) - f_0^{n+1}(c)|$ around $\varepsilon = 0$ yields that the speed $v_n = \ell_\varepsilon/\varepsilon$ is proportional to $\partial_\varepsilon f_\varepsilon^{n+1}(c) = \prod_{i=1}^n f'(c_i) + O(\varepsilon)$ and hence is proportional to α^n for large n . Hence the smaller spikes move faster than the larger spikes corresponding to small values of n . This is illustrated in Figure 4 where we overlay the invariant densities corresponding to a small perturbation with $\varepsilon = 6.05 \times 10^{-4}$. The family of perturbed invariant measures can then be formally written in terms of their associated densities as

$$\rho_\varepsilon(x) = \phi_\varepsilon(x) + \sum_{n=1}^{\infty} \eta_n^{(\varepsilon)}(x + \varepsilon\alpha^n),$$

where the spikes are given as in (19)-(20) with f replaced by f_ε , and the magnitude of the perturbed spikes also decays as $\Upsilon_n^{(\varepsilon)} \sim \alpha^{-n/2}$. Differentiation of $\rho_\varepsilon(x)$ with respect to ε produces an exponentially growing term $\alpha^{n/2}$ inside the sum which prevents the differentiability of the a.c.i.m., and hence causes the breakdown of linear response. A different way to see the non-differentiability of the invariant measure is to consider the linear response of an indicator function with support $[0, c_n]$. Without loss of generality we assume that the spike has support to the right of c_n and moves to the left upon perturbation (if this is not the case, take $\varepsilon \rightarrow -\varepsilon$ to change the direction). Upon applying a perturbation ε the spike will enter the support of the observable and the probability mass moving into the interval is proportional to

$$\delta\rho_{\text{spike}} \sim \int_0^{\alpha^n \varepsilon} \frac{\alpha^{-\frac{n}{2}}}{\sqrt{x}} dx \sim \sqrt{\varepsilon},$$

and therefore $\delta\rho_{\text{spike}}/\varepsilon \sim 1/\sqrt{\varepsilon}$. Since spikes are dense on the support of the a.c.i.m., the non-smoothness extends to the whole a.c.i.m. The non-differentiability of the a.c.i.m. is clearly seen in Figure 5 where we show the observable $\langle A \rangle_\varepsilon$ as a function of ε . The results shown are obtained here again employing spectral methods [25, 15, 63].

5. Resolving breakdown of linear response in finite time series

The rigorous theory by Baladi and co-workers [7, 8, 6, 10] shows that certain dynamical systems such as the logistic map do not obey a linear response.




Figure 2: The a.c.i.m. of the logistic map (17) with $a = 3.8$ (black lines), and a cobweb diagram (online orange) relating the spikes of the a.c.i.m. to the first 17 iterates of the forward orbit of the critical point.

In this section we will investigate how the finitude of data may prevent the breakdown to be detectable and how one may falsely be led to believe that linear response was valid.

As seen in Section 4 the non-smoothness of the invariant measure is caused by the rapid displacement of spikes upon perturbation. The smaller and narrower the spike, the faster it moves. This points to an issue of resolution: the faster spikes carry less mass and therefore require a certain amount of data to be reliably resolved; the slower spikes carry more mass but their smaller displacement upon perturbation requires sufficient data to be resolved. This means that a sufficiently large amount of data is needed for the breakdown parameter \mathfrak{q} to accurately estimate the mismatch $\mathbb{E}\mathfrak{q}$, and to determine whether a system obeys linear response or not. This issue of resolving the mismatch $\mathbb{E}\mathfrak{q}$ is an additional finite size issue to the one discussed in Section 3 whereby N needs to be sufficiently large to assure that the observed p -value is properly estimated (cf. (15)).

Throughout the paper we simulate the logistic map (17) with $a = a_0(1 + \varepsilon)$ and $a_0 = 3.8$. In the notation of Section 2 we set $\varepsilon_0 = 0$ from now on. We choose $M = 20$ equidistant values $\varepsilon_i = \frac{2i-M+1}{2}d\varepsilon$ with $i = 0, \dots, M-1$ for some $d\varepsilon > 0$ to determine the breakdown parameter \mathfrak{q} . Note that the breakdown parameter cannot be determined at an exact perturbation size

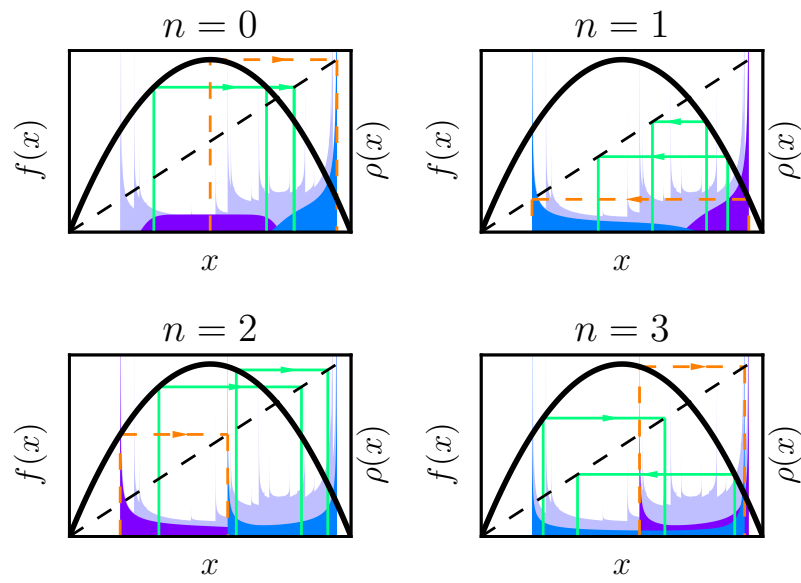


Figure 3: The repeated action of the logistic map (17) with $a = 3.8$ on a continuous density (blue), with a.c.i.m. (light blue). The dashed lines (online orange) show the image of the critical point after n iterations of the logistic map. The continuous lines (online green) show the image of a few points within the support of the initial density at $n = 0$ upon subsequent iteration of the map.




Figure 4: Overlaying two a.c.i.m.s of the logistic map with $a = 3.8$ (maroon) and $a = 3.8 + 6.05 \times 10^{-4}$ (light blue). The displacement of a smaller spike with $n = 17$ is illustrated by the two ovals, and that of a large spike with $n = 3$ with a dashed box.




Figure 5: Observable $\langle A \rangle_\varepsilon$ as a function of the perturbation ε for the logistic map (17) with $a = 3.8(1 + \varepsilon)$ and observable $A(x) = x$. Gaps in the curve correspond to periodic windows.

ε but we determine the validity of a linear approximation over a range of perturbation sizes parametrized by $d\varepsilon$ (for fixed $M = 20$). We restrict the set of perturbations ε_i to include only those which belong to the chaotic Cantor set \mathcal{C} .

The variances σ_i^2 (8) are estimated as a Monte-Carlo estimate from the observed response using the central limit theorem (9) with 200 realizations for each $i = 1, \dots, M$.

5.1. Effect of finite data length N

Figure 6 shows how the breakdown parameter \mathbf{q} behaves with increasing data length N for given perturbation size $d\varepsilon = 10^{-6}$. For each value of N the breakdown parameter is calculated for the above mentioned range of $M = 20$ perturbation sizes ε_i . Shown is the Monte-Carlo estimate $\hat{\mathbf{q}}$ of the expectation value of the breakdown parameter $\mathbb{E}\mathbf{q}$ over $K = 200$ realizations, differing in the initial condition of the logistic map as well as in the threshold value \mathbf{q}_α corresponding to a significance level $\alpha = 0.05$. We show error bars obtained from the ensemble statistics indicating the two-sided 90% prediction interval for \mathbf{q} . We see clearly the saturation of the breakdown parameter with increasing data length N towards the deterministic limit $\mathbb{E}\mathbf{q}$ which eventually leads to detection of the breakdown above a significance level of $p = 0.05$. The breakdown can, however, only be detected reliably with a statistical significance level larger than 0.05 for long time series with $N > 600,000$. The corollary of this is that when analyzing single time series of length $N < 600,000$ at several values of the perturbation size ε the error bars lie below \mathbf{q}_α and the dynamics may be falsely classified as obeying linear response.

For comparison we have included in Figure 6 a plot showing the breakdown parameter as a function of N for the doubling map which does obey linear response with $\mathbb{E}\mathbf{q} \rightarrow 0$ for $\varepsilon \rightarrow 0$. Here the observed breakdown parameter $\hat{\mathbf{q}}$ decreases with N according to the law of large numbers and the sample statistics is consistent with the two-sided 90% prediction interval for the whole range of N . Since the expectation of the breakdown parameter $\mathbb{E}\mathbf{q}$ approaches zero for vanishing perturbation size, the estimator of $\mathbb{E}\mathbf{q}$ is noisy due to sampling errors, and hence may be small and negative¹.

¹In the doubling map we have $\mathbb{E}\mathbf{q} \sim \mathcal{O}(\varepsilon^4)$ since deviations of W_i from a linear fit are $\mathcal{O}(\varepsilon_i^2)$; cf. (13).

A concrete example of how statistical noise may impede the detection of linear response breakdown from observations is shown in Figure 7 for an observable $A(x) = x$. Shown is the observed sample average \bar{A} (7) as a function of the perturbation size. The error bars are calculated from the standard deviation as calculated for the single available time series, which is the situation for scientists analyzing observations. For insufficient data length $N = 10^5$ a linear response is consistent within the available statistical significance levels (top of Figure 7). Only for significantly larger time series with data length $N = 10^6$, does the breakdown become detectable in a statistically significant way (bottom of Figure 7).

5.2. Effect of the perturbation size ε

The critical length of the data N_b above which breakdown of linear response can be detected in a statistically significant way depends on the perturbation size ε . In particular, χ^2 is an increasing function of ε for sufficiently large values of ε , cf. (11). This dependency can be intuitively understood since the response to small perturbations must be distinguished from the variations in the unperturbed system due to the sampling error. This implies that to be able to identify a deviation from linear response at a specified perturbation size ε with a significance level p the perturbation size needs to be sufficiently large. This is illustrated in Figure 8 where we show the Monte-Carlo estimate \hat{q} of the expectation value of the breakdown parameter as a function of the perturbation size which is parametrized by the perturbation interval $d\varepsilon$. For each value of $d\varepsilon$ the perturbed system is sampled at $\varepsilon_i = \frac{2i-M+1}{2}d\varepsilon$ with $i = 0, \dots, M-1$ for fixed data length $N = 10^6$ with $M = 20$. For perturbation sizes $d\varepsilon < 8 \times 10^{-7}$ the observations are consistent with linear response theory and only for $d\varepsilon > 8 \times 10^{-7}$ can the actual breakdown be detected in a statistically significant way. For comparison we have again included in Figure 8 a plot showing the breakdown parameter as a function of $d\varepsilon$ for the doubling map where linear response assures $\mathbb{E}q \rightarrow 0$ for $\varepsilon \rightarrow 0$. Here linear response is consistent with the observations for the whole range of perturbation sizes considered.

Figure 9 illustrates that the smaller the applied perturbation the larger the data length has to be to detect breakdown. Shown is the critical data length N_b above which breakdown can be detected for a given perturbation size. The critical data length N_b was determined to be the value of N such that $\hat{q} = q_\alpha$ for $\alpha = 0.05$. A linear fit suggests $N_b \sim \varepsilon^{-\gamma}$, where γ was estimated in Figure 9 to be 0.91.





Figure 6: Breakdown parameter \hat{q} (solid line) as a function of the data size N , estimated using $d\varepsilon = 10^{-6}$. Top: logistic map (17) with fixed range of perturbations $a = 3.8(1 + \varepsilon)$ and observable $A(x) = x$. Bottom: doubling map with perturbation $X(x) = \sin 2\pi x$ and observable $A(x) = \cos 2\pi x$.

The error bars show the two-sided 90% prediction interval for \hat{q} as estimated from $K = 200$ realizations differing in the initial conditions. The dashed line shows q_α for $\alpha = 0.05$. Note that for the doubling map (bottom) the breakdown parameter \hat{q} assumes values below the plotted range for some values of N .




Figure 7: The observed sample average \bar{A} as a function of the perturbation size ε for an observable $A(x) = x$ and a linear fit (solid line). Top: for data length $N = 10^5$ where $\hat{q} = 5.03 \times 10^{-5}$ and the breakdown is not detectable ($p = 0.148$). Bottom: for data length $N = 10^6$ where $\hat{q} = 5.32 \times 10^{-5}$ and the breakdown is detectable ($p = 1.31 \times 10^{-8}$). Here \bar{A} was obtained from a single simulation.





Figure 8: Breakdown parameter \hat{q} (solid line) as a function of the perturbation size $d\varepsilon$ for fixed $N = 10^6$. Top: logistic map (17) with $a = 3.8(1 + \varepsilon)$ for an observable $A(x) = x$. Bottom: doubling map with perturbation $X(x) = \sin 2\pi x$ for an observable $A(x) = \cos 2\pi x$.

The error bars show the two-sided 90% prediction interval for q as estimated from $K = 200$ realizations differing in the initial conditions. The dashed line shows q_α for $\alpha = 0.05$. Note that for the doubling map (bottom) the breakdown parameter \hat{q} assumes values below the plotted range for some values of $d\varepsilon$.




Figure 9: The critical data length N_b above which the breakdown of linear response is detectable as a function of the perturbation size ε , parametrized by $d\varepsilon$. The slope of the linear fit is -0.91 .

5.3. Effect of the observable

In the previous sections we presented results for a global observable $A(x) = x$ which probes the dynamics over the full support of the invariant measure. The breakdown is caused by the spikes in the a.c.i.m. and their rapid displacement under perturbation. It is therefore natural to expect that observables which locally probe the displacement require less data to see the breakdown of linear response. We now consider localized observables

$$A(x) = \exp\left(-\frac{(x - x_s)^2}{2w^2}\right).$$

This observable allows us to probe the local non-smooth behavior of a spike η_n at location x_n with width w_n . Recall the displacement length $\ell_\varepsilon = v_n \varepsilon$ of the spike upon perturbation with ε where $v_n \sim \alpha^n$ with α being the average expansion rate. Hence, for $x_s \approx x_n$ and $w_n \approx w \leq w^* := \ell_\varepsilon/2$ the spike η_n can be resolved by the observable and the displacement will be detectable when it leaves the effective support of the observable upon perturbation by ε . An example of such a judiciously chosen Gaussian observable is given in Figure 10.

The effect of a localized observable on the ability to detect breakdown of linear response is illustrated in Figure 11. We performed two sets of simula-

tions. In the first we fixed the characteristic scale of the observable $w = w^*$ to equal half the displacement length of the $n = 11$ th spike, $w^* \approx 1.9 \times 10^{-4}$, and varied the centre x_s of the observable. A clear peak of statistically significant values of the breakdown parameter \hat{q} above a significance level with $p = 0.05$ is obtained for x_s close to the location of the 11th spike at $x_s^* := x_{11} = 0.573$. Note that the size of the displacement window within which breakdown is detectable corresponds roughly to the displacement width of the spike (cf. Figure 10). In a second set of simulations we centered the Gaussian observable at the location of the $n = 11$ th spike with $x_s = x_s^*$ and varied the observational scale w . Again, a pronounced peak of the expected value of the breakdown parameter \hat{q} is seen above the significance level for $w = w^*$. The maximum is not obtained exactly at the estimated value of w^* due to approximations made when relating v_n for finite n to its asymptotic value α^n .

Figure 12 shows the breakdown parameter \hat{q} as a function of the available data length N for a given perturbation size $d\varepsilon = 10^{-6}$ for a Gaussian observable where $x_s = x_s^*$ and $w = w^*$ is chosen to focus on the displacement of the 11th spike. It is revealed that a time series with only $N = 30,000$ is needed to reliably detect breakdown of linear response; this should be compared to the required length of $N = 600,000$ when an observable $A(x) = x$ is used (cf. Figure 6). Similarly, Figure 13 shows the breakdown parameter \hat{q} as a function of the perturbation size (here the perturbation interval $d\varepsilon$) for fixed data length $N = 10^6$ for a Gaussian observable with finely tuned $w = w^*$ and $x_s = x_s^*$. Breakdown is reliably detected for perturbation sizes with $d\varepsilon \geq 1.2 \times 10^{-7}$; for an observable $A(x) = x$ with $N = 10^6$ one needs larger perturbation sizes with $d\varepsilon \geq 8 \times 10^{-7}$ to detect breakdown (cf. Figure 8).

The preceding discussion indicates that detailed knowledge of the underlying dynamical system (both the location of a spike x_s^* and its displacement scale w^*) is required for the successful detection of the breakdown of linear response given a time series of finite length N . In particular, these finely tuned observables depend on the magnitude of the perturbation ε . Figure 14 illustrates how a lack of this knowledge may indeed mislead us into deducing the validity of linear response. We show the observed sample average \bar{A} (7) for a Gaussian observable as a function of the perturbation size ε in the case when the characteristic observational scale w and the location x_s are judiciously chosen to probe for a particular spike and in the case when they are not chosen to align with a spike and its displacement length. In the latter case the existence of linear response is consistent with the observations and

a scientist might be misled in believing in a linear relationship between the perturbation and the response. If, however, the location and scale of the observable are tuned to match a particular spike and its least rapid displacement of the perturbation sizes under consideration, the breakdown is clearly detectable. This, of course, as we have seen above, requires the length of the time series to be sufficiently large. The saturation of the response for sufficiently large perturbations in the case of a finely tuned localized observable (cf. bottom plot in Figure 14) is an indication that the length of the time series is insufficient to detect the contribution of the other spikes to the non-smoothness of the invariant measure.

It is pertinent to state that the mere inclusion of a scale $w = w^*$ in the observable to probe the non-smooth dynamic behavior of the spikes of the invariant measure is not sufficient to detect breakdown for smaller values of N (for fixed ε) or for smaller perturbation sizes ε (for fixed N). For example, an observable $A(x) = \cos(2\pi(x - x_s)/w)$ with wave length $w = w^*$ and $x_s = x^*$ finely tuned to capture the displacement of the 11th spike does not exhibit any variation of the expected value of the breakdown parameter \hat{q} as a function of the scale parameter w . Figure 15 reveals that there is no peak in the breakdown parameter near $w = w^*$ for finite $N = 5 \times 10^4$. The failure of the cosine-function to enhance the detectability of breakdown of linear response, despite its wave length matching the characteristic displacement length of a particular spike, is due to the global character of the cosine-function. Although the non-smooth behavior of the fast and narrow spike is resolved by the observable, this is swamped by the dominant contribution of the observable stemming from other parts of the a.c.i.m., in particular from the smooth background and from the larger, slower spikes.

6. The fluctuation-dissipation theorem

The methods presented in Section 5 are based on performing the perturbation experiment by brute force, i.e. by running a numerical experiment for a range of values of ε . However, one of the aspects of linear response theory that has attracted a lot of attention from practitioners is the fact that for many systems, if linear response holds, formulae exist (see Eqs. (2) and (3)) that express the linear response in terms of properties of the unperturbed dynamical system, providing the tantalizing prospect of predicting the linear response without having to perform the kind of brute perturbation experiment used in Section 5.




Figure 10: Gaussian observable probing the displacement of the $n = 11$ th spike upon perturbing the logistic map (17) from $a = 3.8(1 - \varepsilon)$ to $a = 3.8(1 + \varepsilon)$ with $\varepsilon = 2.5 \times 10^{-6}$. The dash-dotted line shows the invariant measure at $a = 3.8(1 - \varepsilon)$ and the dashed line shows the invariant measure at $a = 3.8(1 + \varepsilon)$.

In that vein, the fluctuation-dissipation theorem has been applied to various atmospheric and climate models. It has been mostly applied in the form of the so-called quasi-Gaussian approximation, where the invariant measure is assumed to be Gaussian, resulting in a response in the form of an integrated auto-covariance function [47]. This autocovariance function can be estimated from unperturbed model integrations or from measurements. The assumption of Gaussianity may be reasonable for some large-scale climatic observables, but it is not valid universally, for example for observables related to bi-stable subsystems such as the Kuroshio Extension or the El Niño Southern Oscillation. A more general approach was taken in [22], where the invariant measure was not assumed to be Gaussian, but was obtained by smoothing the observed empirical density with a smoothing kernel. Since we are dealing here with highly non-Gaussian densities, we will investigate this approach rather than imposing Gaussianity.

In this section we consider the situation where one is unaware of the existence or absence of linear response for the system of interest, but only has access to a data set of observations of the unperturbed system. In such a case a practitioner might be led to estimate the right hand side of (3) from data and hope that the obtained quantity gives an indication of the response




Figure 11: Breakdown parameter \hat{q} (solid line) for a Gaussian observable, for the logistic map with $a = 3.8(1 + \varepsilon_i)$ with $d\varepsilon = 10^{-6}$ and $N = 5 \times 10^5$. Top: as a function of the location of the center x_s of the observable with $x_s^* := x_{11}$ the location of the $n = 11$ th spike. Here $w = w^*$ is fixed to equal half the displacement length of the 11th spike. Bottom: as a function of the width ratio w/w^* with $x_s = x_s^*$ fixed. The error bars show the two-sided 90% prediction interval for q as estimated from $K = 200$ realizations differing in the initial conditions. The dashed line shows q_α for $\alpha = 0.05$.




Figure 12: Breakdown parameter \hat{q} (solid line) as a function of the data size N for the logistic map with fixed range of perturbations $a = 3.8(1 + \varepsilon)$ with $d\varepsilon = 10^{-6}$ and Gaussian observable with width $w = w^*$ and location $x_s = x_s^*$ tuned to capture the displacement of the $n = 11$ th spike.

The error bars show the two-sided 90% prediction interval for q as estimated from $K = 200$ realizations differing in the initial conditions. The dashed line shows q_α for $\alpha = 0.05$.




Figure 13: Breakdown parameter \hat{q} (solid line) as a function of the perturbation size $d\varepsilon$ for fixed $N = 10^6$ for the logistic map (17) with $a = 3.8(1 + \varepsilon)$ for a Gaussian observable with width $w = w^*$ and location $x_s = x_s^*$ tuned to capture the displacement of the $n = 11$ th spike.

The error bars show the two-sided 90% prediction interval for q as estimated from $K = 200$ realizations differing in the initial conditions. The dashed line shows q_α for $\alpha = 0.05$.




Figure 14: The observed sample average \bar{A} as a function of the perturbation size ε for an observable $A(x) = \exp((x - x_s)^2/(2w^2))$ with $N = 5 \times 10^4$ where $w = w^*$ is half the displacement length of the $n = 11$ th spike, and a linear fit (solid line). Top: when the location of the observable x_s is not centered at the location of the spike x_s^* . Bottom: when the location of the observable x_s is centered at the location of the spike x_s^* . Here \bar{A} was obtained from a single simulation.




Figure 15: Breakdown parameter \hat{q} (solid line) for the logistic map with $a = 3.8(1+\varepsilon)$ with $d\varepsilon = 10^{-6}$ as a function of the wavelength w of an observable $A(x) = \cos(2\pi(x - x_s^*)/w)$ where x_s^* and w^* denote the location and half the displacement length of the $n = 11$ th spike.

The error bars show the two-sided 90% prediction interval for q as estimated from $K = 200$ realizations differing in the initial conditions. The dashed line shows q_α for $\alpha = 0.05$.

over a certain range of ε . When linear response holds this will be the case, however here we investigate whether such an approximation of the response is possible when the response is non-differentiable.

To this end we perform a perturbation experiment for the logistic map and compare the actual response to the prediction obtained through (3). The actual response is obtained through spectral methods [25, 15, 63], in order to avoid finite sample size effects. We have followed the non-parametric method based on kernel smoothing for estimating the linear response from the FDT as described in [22]. The density of the a.c.i.m. ρ_0 is smoothed by convolution with a Gaussian with smoothing width ω_s . This removes the non-differentiable character of the spikes and allows the derivative to be taken in Eq. (3).

The results of such an estimation of the linear response using the fluctuation-dissipation theorem are shown in Figure 16. The experiment was performed for an observable $A(x) = x$ at two different reference states $a = 3.789$ and $a = 3.805$ and for two different smoothing widths $\omega_s = 0.015$ and $\omega_s = 0.005$. It is evident that the results are sensitive to both these parameters and that the actual response is not well approximated by the slope as constructed




Figure 16: Expectation value $\langle A \rangle$ as a function of the parameter a of the logistic map (17) for an observable $A(x) = x$. Gaps in the curve correspond to periodic windows. At two values of a the predictions of the fluctuation dissipation theorem (3) are shown for two different kernel-smoothing widths ω_s . The continuous lines (online green) are for $\omega_s = 0.005$ and the dashed lines (online blue) are for $\omega_s = 0.015$.

through the FDT.

We also present results showing that kernel-smoothing allows for a reliable, convergent estimation of the fluctuation-dissipation theorem in the case of a topological conjugate of the doubling map in Figure 17. We use the smooth conjugation $h(x) = \frac{1}{2\pi} \cot^{-1}(0.25 + \cot 2\pi x)$, which transforms the doubling map's physical Lebesgue measure into $d\rho(x) = h'(x)dx$. We use a perturbation $X(x) = \sin 2\pi x$, and an observable $A(x) = \cos 2\pi x$. It is seen that using kernel smoothing in the fluctuation-dissipation formula (3) approximates the true linear response well for a variety of kernel widths ω_s . Furthermore the linear response estimated using kernel smoothing converges to the true linear response as $\mathcal{O}(\omega_s^2)$ as $\omega_s \rightarrow 0$. In the appendix we show analytically that the linear response as estimated using kernel smoothing converges for uniformly expanding maps to the true linear response upon decreasing smoothing width and that the error decreases with ω_s^2 .




Figure 17: Top: Expectation value $\langle A \rangle$ as a function of the perturbed conjugation of the doubling map for an observable $A(x) = \cos 2\pi x$, and perturbation $X(x) = \sin 2\pi x$. Around $\varepsilon = 0$, the predictions of the fluctuation dissipation theorem (3) are shown for two different kernel smoothing widths ω_s . The continuous line (online green) is for $\omega_s = 0.05$ and the dashed line (online blue) is for $\omega_s = 0.07$. The dotted straight line shows the true linear response. Bottom: Discrepancy between the linear response estimate from kernel smoothing and the true linear response, as a function of the kernel smoothing width ω_s .

7. Summary and conclusion

We have provided a detailed numerical investigation of the manifestation of breakdown of linear response caused by the non-smooth change of the invariant measure with respect to a small perturbation in a finite time series analysis. We have presented results for the logistic map for which the breakdown of linear response is analytically well understood [7, 8, 6, 10].

The main messages which can be deduced from our results are that in order to detect the breakdown of linear response in time series of finite length, the data length needs to be sufficiently long and furthermore that the detectability of linear response strongly depends on the observable and on the perturbation size. We summarize our key findings:

- The amount of data N required to detect a breakdown of linear response for a given perturbation size can be very large. For the logistic map with a given perturbation size of the order of $\varepsilon = \mathcal{O}(10^{-6})$ one needs at least $N = 600,000$ for a smooth observable $A(x) = x$. Hence, an apparent linear response seen in a given time series might be spuriously caused by an insufficient quantity of data.
- The smaller the perturbation size the longer the data need to be to detect a breakdown in general.
- The global character of an observable may inhibit the detection of breakdown of linear response. For a given finite data length and given perturbation size, suitably localized observables may be needed to probe linear response. This, however, requires either detailed knowledge of the underlying dynamical system or computationally involved scans of the parameters of the observable such as its scale and its location.
- Predicting response using the fluctuation-dissipation theorem is highly sensitive to the applied smoothing needed to assure differentiability of the density, and to the point in parameter space where the response is calculated, negating its predictive value. In the case when the FDT is valid, however, our results suggest that kernel smoothing as applied by climate scientists yields a valid approximation to the true linear response.

These findings can be taken as a word of caution for practitioners interpreting observational or numerical time series. Our results aim to narrow

the gap between the body of rigorous theoretical work and the applied research done, for example, in climate science and in atmosphere and ocean dynamics.

Since we currently have no means of deciding whether the coupled atmosphere-ocean system or the whole climate system satisfies linear response theory or whether it does not, our work does not *per se* question the validity of the many results obtained using linear response theory and FDT.

The *chaotic hypothesis* of Gallavotti-Cohen [31, 30] is often invoked to argue that a high-dimensional chaotic physical system can be treated for all practical purposes as if it were Axiom A. It is, however, pertinent to mention, that the chaotic hypothesis only makes a statement about the existence of time averages computed with a probability distribution capturing the statistics of macroscopic observables and satisfying a large deviation law at *one* parameter value; it does not make any statement about the smoothness of the underlying probability density with respect to changes in this parameter and about whether the invariant measures of the approximating Axiom A systems at nearby parameter values are approximately linearly related, which is what is required for linear response theory. We adhere, however, to the current general belief that large complex systems with multi-scale dynamics behave as stochastic systems and therefore linear response theory is valid for large-scale observables (provided the dynamics is not close to a critical point).

For scientists analyzing time series, we propose the following as a practical guide, which could be drawn from our work. In the case when the time series is obtained by costly numerical simulations, prohibiting the usage of very large time series, or by a limited amount of observational data, scientists could perform an ensemble of (parallel) simulations for several moderate data lengths N or of subsamples. If the number of realizations which produce values of the sample mean of the breakdown parameter \hat{q} exceeding the corresponding threshold value for a specified significance level increases with increasing data length N , then this indicates breakdown of linear response as for example seen in Figure 6. The figure suggests that another indication for a finite value of $\mathbb{E}q$ and breakdown of linear response is the case where \hat{q} either increases or saturates over the available range of N . These two criteria, although far from being decisive, may be used as sufficient conditions for breakdown of linear response.

Appendix

We prove the convergence of the estimated linear response using kernel smoothing to the true linear response in the case of a uniformly expanding one-dimensional map on a compact manifold Λ , under the assumptions that the invariant measure ρ is C^4 , and the smoothing kernel has a zero first moment. We further assume that our observable $A(x)$ is an L^1 function. For simplicity we assume the map has two branches, and each individual branch is at least C^2 . This includes the doubling map as discussed in the main part.

We recall the fluctuation-dissipation theorem (3)

$$\langle A \rangle'_{\varepsilon_0} = - \sum_{n=0}^{\infty} \left\langle \left(\frac{(\rho_{\varepsilon_0} X)'}{\rho_{\varepsilon_0}} \right) (A \circ f_{\varepsilon_0}^n - \langle A \rangle_{\varepsilon_0}) \right\rangle_{\varepsilon_0}. \quad (21)$$

which gives the linear response in terms of correlations.

Analogously, in the case where the density is kernel smoothed, the linear response is written as

$$\langle A \rangle'_{\varepsilon_0, \omega_s} = - \sum_{n=0}^{\infty} \left\langle \left(\frac{(\mathcal{S}\rho_{\varepsilon_0} X)'}{\mathcal{S}\rho_{\varepsilon_0}} \right) (A \circ f_{\varepsilon_0}^n - \langle A \rangle_{\varepsilon_0}) \right\rangle_{\varepsilon_0}, \quad (22)$$

where \mathcal{S} is a convolution by a kernel density ϕ with zero first moment and variance ω_s^2 . In particular, we have for the kernel smoothed density

$$\mathcal{S}\rho(x) = \int \rho(x-y)\phi(y)dy.$$

Dropping the ε_0 subscripts, the difference between the kernel smoothed and the true linear response is given by

$$\begin{aligned} \langle A \rangle'_{\omega_s} - \langle A \rangle' &= - \sum_{n=0}^{\infty} \left\langle \left(\frac{(\mathcal{S}\rho X)'}{\mathcal{S}\rho} - \frac{(\rho X)'}{\rho} \right) (A \circ f^n - \langle A \rangle) \right\rangle \\ &= - \sum_{n=0}^{\infty} \left\langle X \left(\frac{\mathcal{S}\rho'}{\mathcal{S}\rho} - \frac{\rho'}{\rho} \right) (A \circ f^n - \langle A \rangle) \right\rangle, \end{aligned}$$

which can be bounded for C^2 maps with two branches by

$$|\langle A \rangle'_{\omega_s} - \langle A \rangle'| \leq \left\| X \left(\frac{\mathcal{S}\rho'}{\mathcal{S}\rho} - \frac{\rho'}{\rho} \right) \right\|_{\text{Lip}} \frac{C\|A\|_1}{1-\gamma}, \quad (23)$$

for some $C > 0$, $\gamma \in (0, 1)$ independent of A [53]. Here the Lipschitz-norm satisfies

$$\|B\|_{\text{Lip}} = \text{Lip } B + \|B\|_{\infty} \leq \|B'\|_{\infty} + \|B\|_{\infty}.$$

The right-hand-side of (23) can be further bounded by using

$$\begin{aligned} \left\| X \left(\frac{\mathcal{S}\rho'}{\mathcal{S}\rho} - \frac{\rho'}{\rho} \right) \right\|_{\text{Lip}} &\leq (\|X'\|_{\infty} + \|X\|_{\infty}) \left\| \frac{\mathcal{S}\rho'}{\mathcal{S}\rho} - \frac{\rho'}{\rho} \right\|_{\infty} \\ &\quad + \|X\|_{\infty} \left\| \left(\frac{\mathcal{S}\rho'}{\mathcal{S}\rho} - \frac{\rho'}{\rho} \right)' \right\|_{\infty}. \end{aligned} \quad (24)$$

Furthermore, we can bound

$$\left\| \frac{\mathcal{S}\rho'}{\mathcal{S}\rho} - \frac{\rho'}{\rho} \right\|_{\infty} \leq \frac{\|\rho\mathcal{S}\rho' - \rho'\mathcal{S}\rho\|_{\infty}}{\inf \rho \inf \mathcal{S}\rho} \quad (25)$$

and

$$\begin{aligned} \left\| \left(\frac{\mathcal{S}\rho'}{\mathcal{S}\rho} - \frac{\rho'}{\rho} \right)' \right\|_{\infty} &\leq \frac{\|\rho\mathcal{S}\rho' - \rho'\mathcal{S}\rho\|_{\infty} (\|\rho\|_{\infty} \|\mathcal{S}\rho'\|_{\infty} + \|\mathcal{S}\rho\|_{\infty} \|\rho'\|_{\infty})}{(\inf \rho \inf \mathcal{S}\rho)^2} \\ &\quad + \frac{\|\rho\mathcal{S}\rho'' - \rho''\mathcal{S}\rho\|_{\infty}}{\inf \rho \inf \mathcal{S}\rho}. \end{aligned} \quad (26)$$

Note that since we assume the dynamics to be uniformly expanding on a compact manifold the invariant measure is bounded away from zero with $\rho(x) > c$ for some $c > 0$. Since $\phi \geq 0$, we have also $\mathcal{S}\rho(x) > c$, and hence, $\inf \rho > 0$ and $\inf \mathcal{S}\rho > 0$.

Now for any twice-differentiable functions p and q , we have

$$(p \mathcal{S}q - q \mathcal{S}p)(x) = \int (p(x)q(x-y) - q(x)p(x-y)) \phi(y) dy.$$

Taylor expanding p and q in y , we find

$$\begin{aligned} (p \mathcal{S}q - q \mathcal{S}p)(x) &= \int \left(p(x) \left(q(x) - yq'(x) + \frac{y^2}{2}q''(\xi_1(y)) \right) \right. \\ &\quad \left. - q(x) \left(p(x) - yp'(x) + \frac{y^2}{2}p''(\xi_2(y)) \right) \right) \phi(y) dy \\ &= \int (p(x)q''(\xi_1(y)) - q(x)p''(\xi_2(y))) \frac{y^2}{2} \phi(y) dy, \end{aligned}$$

where $\xi_{1,2} \in \Lambda$ and where we have used that $\phi(y)$ has a vanishing first moment. Hence we can bound

$$\|p \mathcal{S}q - q \mathcal{S}p\|_\infty \leq \frac{1}{2} (\|p\|_\infty \|q''\|_\infty + \|p''\|_\infty \|q\|_\infty) \omega_s^2.$$

Applying the last inequality to (25) and (26), we can bound the right-hand-side of (23) and arrive at our final estimate for the difference between the true linear response and the kernel smoothed linear response

$$|\langle A \rangle'_{\varepsilon_0, \omega_s} - \langle A \rangle'_{\varepsilon_0}| < Q \|A\|_1 \omega_s^2,$$

for some $Q > 0$ independent of ω_s and A . Hence the difference between the true linear response and the kernel smoothed linear response scales with the square of the kernel width as observed in Figure 17. We remark that this proof can be readily extended to the case where the kernel depends on x .

Acknowledgement

We would like to thank Viviane Baladi for interesting discussions. J. Wouters' research was funded by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° PEOF-GA-2013-626210.

References

References

- [1] Abramov, R. V. and Majda, A. J. (2007). Blended response algorithms for linear fluctuation-dissipation for complex nonlinear dynamical systems. *Nonlinearity*, **20**(12), 2793.
- [2] Abramov, R. V. and Majda, A. J. (2008). New approximations and tests of linear fluctuation-response for chaotic nonlinear forced-dissipative dynamical systems. *J. Nonlinear Sci.*, **18**(3), 303–341.
- [3] Abramov, R. V. and Majda, A. J. (2009). A new algorithm for low-frequency climate response. *Journal of the Atmospheric Sciences*, **66**(2), 286–309.

- [4] Avila, A. and Moreira, C. G. (2005). Statistical properties of unimodal maps: the quadratic family. *Ann. of Math. (2)*, **161**(2), 831–881.
- [5] Baladi, V. (2000). *Positive transfer operators and decay of correlations*, volume 16 of *Advanced Series in Nonlinear Dynamics*. World Scientific Publishing Co., Inc., River Edge, NJ.
- [6] Baladi, V. (2014). Linear response, or else. In *ICM Seoul 2014, Proceedings, Volume III*, page 52545.
- [7] Baladi, V. and Smania, D. (2008). Linear response formula for piecewise expanding unimodal maps. *Nonlinearity*, **21**(4), 677–711.
- [8] Baladi, V. and Smania, D. (2010). Alternative proofs of linear response for piecewise expanding unimodal maps. *Ergodic Theory and Dynamical Systems*, **30**(01), 1–20.
- [9] Baladi, V. and Todd, M. (2016). Linear response for intermittent maps. *Communications in Mathematical Physics*, pages 1–18.
- [10] Baladi, V., Benedicks, M., and Schnellmann, D. (2015). Whitney-Hölder continuity of the SRB measure for transversal families of smooth unimodal maps. *Invent. Math.*, **201**(3), 773–844.
- [11] Balescu, R. (1975). *Equilibrium and Non-equilibrium Statistical Mechanics*. John Wiley & Sons, New York.
- [12] Bell, T. L. (1980). Climate Sensitivity from Fluctuation Dissipation: Some Simple Model Tests. *Journal of the Atmospheric Sciences*, **37**(8), 1700–1707.
- [13] Benedicks, M. and Carleson, L. (1985). On iterations of $1 - ax^2$ on $(-1, 1)$. *Ann. of Math. (2)*, **122**(1), 1–25.
- [14] Box, G. E. P., Hunter, J. S., and Hunter, W. G. (2005). *Statistics for experimenters : design, innovation, and discovery*. Wiley series in probability and statistics. Wiley-Interscience, Hoboken (N.J.).
- [15] Boyd, J. P. (2001). *Chebyshev and Fourier spectral methods*. Courier Corporation, Mineola, NY.
- [16] Cessac, B. and Sepulchre, J.-A. (2007). Linear response, susceptibility and resonances in chaotic toy models. *Physica D: Nonlinear Phenomena*, **225**(1), 13 – 28.

- [17] Chekroun, M. D., Neelin, J. D., Kondrashov, D., McWilliams, J. C., and Ghil, M. (2014). Rough parameter dependence in climate models and the role of Ruelle-Pollicott resonances. *Proceedings of the National Academy of Sciences*, **111**(5), 1684–90.
- [18] Cionni, I., Visconti, G., and Sassi, F. (2004). Fluctuation dissipation theorem in a general circulation model. *Geophysical Research Letters*, **31**(9), L09206.
- [19] Collet, P. and Eckmann, J.-P. (1983). Positive Liapunov exponents and absolute continuity for maps of the interval. *Ergodic Theory Dynam. Systems*, **3**(1), 13–46.
- [20] Collet, P. and Eckmann, J.-P. (2007). *Concepts and results in chaotic dynamics: a short course*. Springer Science & Business Media, Berlin.
- [21] Cooper, F. and Haynes, P. (2013). Assessment of the fluctuation-dissipation theorem as an estimator of the tropospheric response to forcing. *Quart. J. Roy. Met. Soc.* Submitted.
- [22] Cooper, F. C. and Haynes, P. H. (2011). Climate sensitivity via a non-parametric fluctuation-dissipation theorem. *Journal of the Atmospheric Sciences*, **68**(5), 937–953.
- [23] Cooper, F. C., Esler, J. G., and Haynes, P. H. (2013). Estimation of the local response to a forcing in a high dimensional system using the fluctuation-dissipation theorem. *Nonlin. Processes Geophys.*, **20**(2), 239–248.
- [24] De Lima, A. and Smania, D. (2015). Central limit theorem for the modulus of continuity of averages of observables on transversal families of piecewise expanding unimodal maps. *arXiv:1503.01423 [math.DS]*.
- [25] Ding, J., Du, Q., and Li, T. Y. (1993). High order approximation of the Frobenius-Perron operator. *Appl. Math. Comput.*, **53**(2-3), 151–171.
- [26] Dolgopyat, D. (2004). On differentiability of SRB states for partially hyperbolic systems. *Invent. Math.*, **155**(2), 389–449.
- [27] Dymnikov, V. P. and Gritsoun, A. S. (2001). Climate model attractors: chaos, quasi-regularity and sensitivity to small perturbations of external forcing. *Nonlinear Processes in Geophysics*, **8**(4/5), 201–209.

- [28] Ershov, S. V. (1993). Is a perturbation theory for dynamical chaos possible? *Physics Letters A*, **177**(3), 180 – 185.
- [29] Fuchs, D., Sherwood, S., and Hernandez, D. (2014). An Exploration of Multivariate Fluctuation Dissipation Operators and Their Response to Sea Surface Temperature Perturbations. *Journal of the Atmospheric Sciences*, **72**(1), 472–486.
- [30] Gallavotti, G. and Cohen, E. (1995a). Dynamical ensembles in stationary states. *Journal of Statistical Physics*, **80**(5-6), 931–970.
- [31] Gallavotti, G. and Cohen, E. G. D. (1995b). Dynamical ensembles in nonequilibrium statistical mechanics. *Phys. Rev. Lett.*, **74**, 2694–2697.
- [32] Gottwald, G. A. and Melbourne, I. (2013). Homogenization for deterministic maps and multiplicative noise. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, **469**(2156).
- [33] Gritsun, A. and Branstator, G. (2007). Climate response using a three-dimensional operator based on the fluctuation-dissipation theorem. *Journal of the Atmospheric Sciences*, **64**(7), 2558–2575.
- [34] Gritsun, A. and Dymnikov, V. (1999). Barotropic atmosphere response to small external actions: Theory and numerical experiments. *Izv. Akad. Nauk. Fiz. Atmos. Okeana. Biol.*, **35**, 5651.
- [35] Gritsun, A., Branstator, G., and Dymnikov, V. (2002). Construction of the linear response operator of an atmospheric general circulation model to small external forcing. *Russ. J. Numer. Anal. Math. Modelling*, **17**, 399–416.
- [36] Gritsun, A., Branstator, G., and Majda, A. (2008). Climate response of linear and quadratic functionals using the fluctuation-dissipation theorem. *Journal of the Atmospheric Sciences*, **65**(9), 2824–2829.
- [37] Gritsun, A. S. (2010). Construction of response operators to small external forcings for atmospheric general circulation models with time periodic right-hand sides. *Izvestiya, Atmospheric and Oceanic Physics*, **46**(6), 748–756.
- [38] Hairer, M. and Majda, A. J. (2010). A simple framework to justify linear response theory. *Nonlinearity*, **23**(4), 909.

- [39] Hänggi, P. (1978). Stochastic processes 2: response theory and fluctuation theorems. *Helvetica Physica Acta*, **51**(2), 202–219.
- [40] Jakobson, M. V. (1981). Absolutely continuous invariant measures for one-parameter families of one-dimensional maps. *Comm. Math. Phys.*, **81**(1), 39–88.
- [41] Keller, G. and Nowicki, T. (1992). Spectral theory, zeta functions and the distribution of periodic points for Collet-Eckmann maps. *Comm. Math. Phys.*, **149**(1), 31–69.
- [42] Kelly, D. and Melbourne, I. (2014). Deterministic homogenization for fast-slow systems with chaotic noises. *arXiv:1409.5748 [math.PR]*.
- [43] Kirk-Davidoff, D. B. (2009). On the diagnosis of climate sensitivity using observations of fluctuations. *Atmos. Chem. Phys.*, **9**(3), 813–822.
- [44] Korepanov, A. (2015). Linear response for intermittent maps with summable and nonsummable decay of correlations. *arXiv:1508.06571 [math.DS]*.
- [45] Kubo, R. (1966). The fluctuation-dissipation theorem. *Reports on Progress in Physics*, **29**(1), 255.
- [46] Langen, P. L. and Alexeev, V. A. (2005). Estimating $2 \times CO_2$ warming in an aquaplanet GCM using the fluctuation-dissipation theorem. *Geophysical Research Letters*, **32**(23). L23708.
- [47] Leith, C. E. (1975). Climate response and fluctuation dissipation. *Journal of the Atmospheric Sciences*, **32**(10), 2022–2026.
- [48] Lucarini, V. and Sarno, S. (2011). A statistical mechanical approach for the computation of the climatic response to general forcings. *Nonlinear Processes in Geophysics*, **18**(1), 7–28.
- [49] Lyubich, M. (2002). Almost every real quadratic map is either regular or stochastic. *Ann. of Math. (2)*, **156**(1), 1–78.
- [50] Majda, A. J., Abramov, R., and Gershgorin, B. (2010). High skill in low-frequency climate response through fluctuation dissipation theorems despite structural instability. *Proceedings of the National Academy of Sciences*, **107**(2), 581–586.

- [51] Marconi, U. M. B., Puglisi, A., Rondoni, L., and Vulpiani, A. (2008). Fluctuation–dissipation: Response theory in statistical physics. *Physics Reports*, **461**(4), 111 – 195.
- [52] McWilliams, J. C. (2007). Irreducible imprecision in atmospheric and oceanic simulations. *Proceedings of the National Academy of Sciences*, **104**(21), 8709–8713.
- [53] Melbourne, I. (2015). Fast-slow skew product systems and convergence to stochastic differential equations. Lecture notes, available at http://homepages.lboro.ac.uk/~mawb/Melbourne2_notes.pdf.
- [54] Melbourne, I. and Stuart, A. (2011). A note on diffusion limits of chaotic skew-product flows. *Nonlinearity*, **24**, 1361–1367.
- [55] North, G. R., Bell, R. E., and Hardin, J. W. (1993). Fluctuation dissipation in a general circulation model. *Climate Dynamics*, **8**(6), 259–264.
- [56] Ragone, F., Lucarini, V., and Lunkeit, F. (2015). A new framework for climate sensitivity and prediction: a modelling perspective. *Climate Dynamics*, pages 1–13.
- [57] Reick, C. H. (2002). Linear response of the Lorenz system. *Phys. Rev. E*, **66**, 036103.
- [58] Ring, M. J. and Plumb, R. A. (2008). The response of a simplified GCM to axisymmetric forcings: Applicability of the fluctuation–dissipation theorem. *Journal of the Atmospheric Sciences*, **65**(12), 3880–3898.
- [59] Ruelle, D. (1997). Differentiation of SRB states. *Communications in Mathematical Physics*, **187**(1), 227–241.
- [60] Ruelle, D. (1998). General linear response formula in statistical mechanics, and the fluctuation-dissipation theorem far from equilibrium. *Phys. Lett. A*, **245**(3-4), 220–224.
- [61] Ruelle, D. (2009a). A review of linear response theory for general differentiable dynamical systems. *Nonlinearity*, **22**(4), 855–870.
- [62] Ruelle, D. (2009b). Structure and f-dependence of the a.c.i.m. for a unimodal map f of Misiurewicz type. *Communications in Mathematical Physics*, **287**(3), 1039–1070.
- [63] Trefethen, L. N. (2013). *Approximation theory and approximation practice*. Siam, Philadelphia, PA.

- [64] Young, L.-S. (2002). What are SRB measures, and which dynamical systems have them? *Journal of Statistical Physics*, **108**(5-6), 733–754.
- [65] Zwanzig, R. (2001). *Nonequilibrium Statistical Mechanics*. Oxford University Press, Oxford.