On typicality and adaptation in driven dynamical systems

by

Pavel Chvykov

B.S., University of Michigan (2011)
M.S., Perimeter Institute (2012)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2019

Author . . . .
Department of Physics
May 26, 2019

Certified by. . . . .
Jeremy England
Cabot Career Development Associate Professor of Physics
Thesis Supervisor

Accepted by . . . .
Nergis Mavalvala
Curtis and Kathleen Marble Professor of Astrophysics
Associate Department Head of Physics

On typicality and adaptation in driven dynamical systems

by

Pavel Chvykov

Submitted to the Department of Physics
on May 26, 2019, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

In this work, I consider the possibility of using typicality-type arguments for under-standing intractably complex damped-driven dynamical systems. By approximating such dynamics with appropriately constrained random process, I illustrate quantitative predictive power for some aspects of the motion. In particular, I argue that local dynamical stability, or exit rate, of a state is typically sufficient to predict steady-state probability in such systems – circumventing the classic no-go theorems via our disorder approximation. I then focus on one consequence of this result: that the most likely long-time configurations should also be the dynamically stable ones. In a strongly-driven system, however, such stability may be hard to achieve, and therefore has interesting implications about the corresponding configurations: they must be well-adapted to the details of the driving forces, their dynamical robustness may be viewed in the context of self-healing, and depending on the drive, they can require substantial collective fine-tuning among the system's degrees of freedom. I confirm the emergence of such adapted states in several example systems, both in simulation and in experiment, and verify a quantitative agreement with the predicted scaling between their steady-state probability and local stability. I then explore several arguments and test-cases suggesting further generality of this framework. While it is not yet clear what the precise limits of applicability are for this approach, our results suggest that the intuition it builds can help with prediction and design in a broad class of complex dynamics.

Thesis Supervisor: Jeremy England
Title: Cabot Career Development Associate Professor of Physics

Acknowledgements

This thesis is as much a culmination of my own work over the last 6 years as it is of the mentorship, discussions, advice, encouragement, guidance and friendship of all the people who have supported me along the way. First I want to thank the MIT physics department, its administrators, professors, students and broader community for providing the fertile environment that allowed this work to germinate and develop thus far. The freedom granted by my first-year fellowship, followed by the guidance in finding an advisor, and the flexibility to take a year academic leave later, as well as the invaluable seminars, discussions and courses offered, were all integral to the way I have developed as a scientist.

Next, I want to sincerely thank my adviser, Prof. Jeremy England, whom I am very glad to have reached out to five years prior. His uniquely inspired scientific inquiries, rooted in thorough philosophical foundations, re-ignited my excitement about research at a time when I was feeling that all the interesting questions have largely been answered. Working with him has reassured me that the romantic vision of science I've imagined is not a fantasy, and there is yet much foundational work for us, requiring deep insights no less than it does relevant tools, just as these are now beginning to emerge. Moreover, his fostering of the personal research styles of his students, helped me to develop my own approach in thinking about these questions, that I feel sets me apart in valuable ways. I'm also very grateful that he fully funded me throughout this time – a luxury few theorists get to have – while still allowing much intellectual freedom and encouraging exploration.

It is also hard for me to overestimate the impact that my friends and colleagues have had on me over this time. In our group, and through my research, I've met some of the most interesting people I know. The regular discussions with Jeremy Owen about physics and philosophy, and his feedback on this draft as well as on other work, have been invaluable throughout. I'm grateful for the discussions with Hridesh Kedia about society and academia, and his suggestions and encouragement in my work. The insightful questions from Weishun Zhong, and our mutual exploration of glassy

Contents

Chapter 1

Introduction

Encouraged by the successes of equilibrium statistical mechanics in predicting behaviors of complicated interacting many-body systems, there has been broad effort throughout physics community to extend similar tools to far-from-equilibrium phenomena. Indeed, the organization of a crystal lattice may, at first, seem like a highly engineered, fine-tuned configuration that should be sensitive to all the details of its formation process and component interactions. With the study of equilibrium phase transitions, however, and the advent of RG universality, we learned to view crystals instead as the maximally-typical configurations of the constituent atoms, once these are viewed in a combined ensemble with their thermal bath at the appropriate temperature. This helped establish such examples of equilibrium organization as robust and independent of most microscopic details, allowing for precise control and design of such structures. Inspired by the power and elegance of this framework, many physicists hoped that by generalizing it to externally driven systems, we may similarly explain organization in out-of-equilibrium scenarios – and ultimately, perhaps even the organization of life [20, 47, 36, 25].

Relaxing the equilibrium assumption, however, proved difficult as this left us with no robust structures that we could leverage for prediction. In a sense, making general claims about nonequilibrium systems is tantamount to making general claims about all systems, and thus the task requires some restriction of scope by clarifying our assumptions. One approach to this came from perturbative extensions of

equilibrium results into weakly-driven regimes. In this regime, perturbation theory often allows for tractable calculations of many quantities, producing such ideas as thermodynamic geometry [71], or time-scale dependent effective temperature [16]. However, within the radius of convergence of such perturbative expansions, we don't expect any qualitatively new nonequilibrium phenomena. Thus, much work focused on another approach: assuming some subdivision of the whole system into conservative components with clearly defined energy flows between them. This produced a number of far-from-equilibrium generalizations of some powerful equilibrium results, such as generalized fluctuation theorems [14, 33] and extremization principles [6, 30], as well as general bounds set by the energy currents [27]. This way, much of the effort in non-equilibrium thermodynamics has been devoted to finding remnants of the predictive power of energy conservation in out-of-equilibrium contexts. However, with time-translation symmetry broken, the concept of energy itself may no longer be relevant in general interacting far-from-equilibrium systems, and we may need to look for new structures to organize our analysis around.

Besides energy conservation, there is another pillar that equilibrium statistical mechanics rests on: ergodicity. As integrable systems rely on having so many conserved quantities that dynamics become fully constrained, statistical mechanics assumes that everything besides its one (or few) conserved quantity is totally random with a uniform prior[34]. This profound insight – saying that complicated interacting molecular dynamics can be simply approximated as randomly sampling all allowed configurations – then gives rise to more precise statements such as molecular chaos, or Landau theory, which underlie all of statistical mechanics[35]. But more than that, since this idea is not tied down to energy conservation, it is easily generalized to a variety of contexts, simply by letting all relevant degrees of freedom be in one of two extremes: conserved or maximally random. The crux of the difficulty, and richness of the resulting theories, then lies in choosing the conserved quantities, as the relevant symmetries can turn out to be highly non-obvious and system-specific. Nonetheless, the appeal of making predictions about intractable systems by finding just a few simple constrains, has been gaining more appreciation in the last decade in a variety of

contexts, giving rise to maximum-entropy modeling[45, 46, 18], typicality results for microbial communities[17, 67], network science[48], etc.

In the study of dynamical systems, a set of studies in the 80s looked at randomly-constructed systems, starting from Sinai diffusion [61] and with a pedagogic overview in [4]. This helped develop an understanding of the conditions giving rise to anomalous diffusion processes, which have since been studied in a variety of contexts. This directions has, however, been largely divorced from the recently growing interest in self-organization phenomena in driven systems. Most prominently, this is considered in the study of active matter: many-body systems driven on the individual-particle level, which often leads to collectively organized patterns or behaviors. While many such examples have been found numerically[11, 44], and some even solved analytically [28, 62], the calculations have largely been system-specific and detailed, with a unifying theory still lacking. In fact, this is not surprising – as driven systems allow for nearly any possible behavior, there may not be any reason why a unified framework should exist. Nonetheless, many find a certain universality to the flavor of self-organization found in these studies, and so statistical [54, 63] and field-theoretic generalizations have been entertained [15, 43].

A few studies considered stochastically-induced organization, where ordering is caused by system settling into more constrained [9], or effectively decoupled [13] configurations. One work even proposed noise-reduction as a mechanism for emergence of cooperative behaviors in nature [51]. In this thesis, I develop these directions by studying self-organization from the perspective of random, or "typical," complex dynamical systems. In that light, I propose a general dynamical ordering mechanism, of which the above stochastic-organization examples, and several more presented here, may be seen as specific instances, with possible further application to the active matter examples mentioned above. The challenge with any such proposal is that its generality is accompanied by certain vagueness. In an attempt to make it more precise, I begin by stating my central conjecture in three (not mutually independent) claims:

1. High-dimensional complex dynamical systems (e.g., many-body interacting sys-

tems) can often be approximated by a multiplicative diffusion process in their full configuration space:

$$\dot{x}_i = \sqrt{2\,T_{eff}(\vec{x})} \cdot \xi_i \qquad \text{[Itô]} \qquad (1.1)$$

where $\xi_i(t)$ – univariate Gaussian white noise, and the effective temperature $T_{eff}(\vec{x})$ depends only on local information about the dynamics near $\vec{x}$.

2. If viewed as a conservative system in the presence of an external drive, then configurations $\vec{x}$ with lower $T_{eff}(\vec{x})$ can be said to be better adapted to that drive, within a reasonable definition of "adaptation."

3. In steady-state, the system tends to end up in configurations of low T_{eff}. In particular, $p_{s.s.}(\vec{x}) \sim 1/T_{eff}(\vec{x})$.

Throughout the thesis, these claims are supported by different general considerations and perspectives, analytically shown to hold in restricted classes of dynamical systems, and verified, both numerically and experimentally, on specific examples and toy models. Along the way, we consider the two key questions that would make these claims into a concrete physical theory: a) what are the minimal assumptions a dynamical system must satisfy to make eq.1.1 a good approximation? and b) how do we practically evaluate T_{eff} in all such system? While this work does not provide a complete satisfactory answer to either of these, it builds substantial intuition for what these answers might be.

Let us briefly elaborate on the three claims above. Claim 1 can be viewed as a dynamical-systems counterpart of the molecular chaos assumption of equilibrium thermodynamics: just as molecular chaos is assumed to destroy correlations between two recently-collided particles before they collide again, so here we assume that a trajectory loses memory of where it was faster than it can return to that same configuration. Explicitly, the idea is that if the dynamics are sufficiently high-dimensional and disordered, with no symmetries or fine-tunings, then we may approximate them as just random exploration of the system's configuration space. The only feature of

the original dynamics that we keep in this approximation is the amplitude of the local random motion, or the exit rate, at a state $\vec{x}$. Crucially, $T_{eff}(\vec{x})$ is a local property of $\vec{x}$, and can be practically determined by studying that state and its short-term dynamics.

The 2nd claim may be viewed as simply defining precisely what we mean by adaptation. To see that this definition agrees with conventional wisdom, note first that the notion of organization of a driven system is much less apparent than that of a crystal lattice. Indeed, while it seems obvious to us that trees are highly organized and fine-tuned arrangements of matter, how could an alien notice that? It seems that this fine-tuning may only become apparent when the tree is considered in relation to its environment: its leaves are matched to the color of its sun, its height allows the leaves to be closer to the sun than other objects typically found nearby, its trunk is appropriately robust to withstand typical winds it encounters, and its flowers and fruit are attractive to the useful animals in its habitat. If this tree was studied out of context, perhaps there would be no way to objectively distinguish it from a pile of coal. If we accept this, then any metric of non-equilibrium organization or adaptation must refer to the environment, or driving forces, that caused it. More precisely, here we frame adaptation as a reduction of "surprise" in response to that environment, which has recently been gaining traction as an organizing principle for learning[21, 60] I.e., for us, a well-adapted system will respond to its environment in a predictable and organized fashion, without much exploration of new configurational states, or significant uncorrelated jumps. This is then just the inverse of the above notion of local random motion, captured in $T_{eff}(\vec{x})$ – albeit it may be difficult to precisely define in general, as it must not depend on any periodic or predictable motion in response to the driving forces (hence we sometimes refer to this measure of random motion as "rattling"). Alternatively, we can intuitively think of $T_{eff}(\vec{x})$ as exit rate, or "death" rate, of a state: such a departure, the return from which may only happen by chance.

Note that T_{eff} will only provide a reasonable measure of adaptation if decoupling from the environment altogether is hard or impossible within the system. Barring

that, to achieve a low T_{eff}, the system must "attenuate" fluctuations of the driving environment, doing which will require it to share some features with that drive. We will try to clarify these points by looking at specific example systems (e.g., Appendix 1.A).

The 3rd claim, which follows directly from the approximation in claim 1, is then seen as the counterpart of the Boltzmann distribution in equilibrium. While at first it seems almost trivial – just saying that the system tends to spend more time in configurations with lower exit-rates – its power comes up as it allows estimating the steady-state probability of a state based only on its local properties. This is not generally possible outside of equilibrium, as illustrated by Landauer's classic "blowtorch theorem" [38]. In the spirit of typicality, however, we argue that the blowtorch effect is very unlikely in large complicated systems, and will have at most a weak effect on the overall trend provided by the local exit rates.

Finally, to make claims about nonequilibrium adaptation in real contexts, we must also comment on the time-scales and nature of relaxation processes to the steady-state. Just as in equilibrium we can have a glassy energy landscape that won't immediately tell us about the practically achievable system configurations, as Boltzmann distribution isn't ever reached in finite time, similarly here a claim that the steady-state peaks in well-adapted fine-tuned configurations doesn't mean you will practically find them. Indeed, no one is surprised by the stability of living matter once it emerges – instead the challenge of abiogenesis lies in getting the right time-scales of the transient. While unlike glassy landscapes, trapping barriers are not a problem in our framework, the time-scales of exploration of a large configuration space to achieve strong fine-tuning can be unreasonably long. Thus, in practice, we need not only to have low T_{eff} at well-adapted configurations, but also T_{eff} gradients to guide us there. This can be achieved by a hierarchy of partial adaptation steps, and thus requires a modular structure of configuration space so that subsystems can adapt independently. This is akin to nucleated crystal formation, where tremendous fine-tuning that could never be achieved all at once, is made possible by local organization building out sequentially. Even then, the fact that our dynamics is guided by inhomogeneous noise, rather than

16

a directed force, means that transients aren't as predictable, and fine-tuning is hard-won, in practice requiring exponential T_{eff} changes – see app.1.B. Such dramatic T_{eff} variation, nonetheless, can be typical when talking about adaptation to driving forces, as we will see in some of our worked-out examples.

In Appendix 1.A of this chapter, we show a very simple toy example system that helps to illustrate these claims in a more concrete context. Appendix 1.B then discusses some typical transient behaviors of eq.1.1 and comments on how non-ergodic fine-tuning may be possible only via temperature gradients.

Chapter 2 goes through a concrete analytical derivation of the above claims for systems with strong time-scale separation. In that context, the diffusive approximation is referring to the dynamics of the slow degrees of freedom, which are shown to perceive the influence of the driven fast dynamics as an effective inhomogeneous bath. This comes up as coarse-graining in time renders most details of the fast motion irrelevant, and by central limit theorem, their dominant effect is seen as a series of independent Gaussian random kicks. In chapter 3, we illustrate this effect on an explicit toy model, and verify our analytical predictions against numerical simulations. We show how scale-separation allows solving the fast and the slow dynamics separately, and then accounting for their interdependence in a simple feedback loop: as the fast variables set up a diffusive environment for the slow ones, so the slow determine the response properties of the fast motion to the drive. This way as the slow variables evolve to lower T_{eff}, the fast must either decouple from the drive, or adapt so as to couple in a deterministic or periodic manner (this work was published in [12]).

In chapter 4, we then further discuss and motivate our diffusive approximation of eq.1.1 in the more general context introduced in claim 1 above. In particular, we focus on the question of how to practically define and measure "rattling," or $T_{eff}(\vec{x})$, for a complex dynamical system that we don't know much about. We present Bayesian-inference and maximum-entropy constructions motivating and clarifying these questions. As these methods give rise to an effective temperature tensor, we then discuss how to construct a scalar quantity from it that could be optimally predictive of the steady-state probability, per claim 3.

Chapter 5 then goes through a number of explicit constructions of random dynamical systems, where we, either numerically or analytically, check our proposals, with the particular focus on claim 3: $p_{s.s.}(\vec{x}) \sim 1/T_{eff}(\vec{x})$. We see that this trend approximately holds in all considered scenarios, as long as there is a wide diversity of T_{eff} amplitudes present in the system.

Finally, chapter 6 checks this trend in a set of real swarm-robotic experiments, done in collaboration with Dan Goldman's and Kurt Wiesenfeld's groups at GA Tech. We studied an ensemble of simple 3-link robots, each programmed to move in a predetermined fashion, thus pushing each other around inside a confining ring. In this regime, the system can be seen as a shape-shifting active matter ensemble, and turns out to self-organize into globally orchestrated regular dynamical patterns. We studied this phenomenon in a range of conditions (explored both in experiments and numerical simulations of the system), and were able to robustly recover the correlation of claim 3 above, showing the self-organization to be a low rattling phenomenon. Moreover, we could use our framework to then quantitatively predict how we can gradually destroy this self-organization by reducing friction in the system, which we then checked in simulations and experiments. [Contribution: For this projects, the robots and tracking were designed and built in Prof. Goldman's lab over the last few years. The specific experimental setup presented here, and in particular the discovery of self-organization in this system, was proposed and tested by myself together with two of his students, Will Savoie and Akash Vardhan. The simulation and data analysis used here I then developed independently.]

Appendix

1.A   Dynamic self-organization in kicked rotor

In this appendix, we look at the simplest dynamical system that illustrates some of the claims presented above, helping to illuminate the sort of structures we have in mind. This system also comes up as a building block for the toy model presented in chapter 3. The kicked rotor (also known as the Chirikov standard map) is a well-studied prototypical example of a mixed chaos dynamical system – a system whose phase space has both integrable and chaotic regions for some sets of parameters. The dynamics are those of a pendulum, but where gravity, rather than being constant, is zero most of the time, and gets turned on in periodic bursts:

$$\dot{\theta} = v$$

$$\dot{v} = -K \sin\theta \, \delta(t - n)$$

where all variables are made dimensionless by setting pendulum arm length, pendulum mass and kicking period to one (θ is 2π-periodic), leaving only one dimensionless parameter K. As the kicking is instantaneous and inter-kick dynamics are integrable, the system is often written as a discrete map, making it very amenable to numerical simulations as no discretization errors are produced. Moreover, in dynamical systems literature, v is often also made periodic in 2π, but as we are viewing this as a physical system rather than just an abstract map, we will keep $v \in \mathbb{R}$.

The key feature of this map is that for $1 \lesssim K \lesssim 8$, both integrable and chaotic regions are present in its phase space. For K below this window, the system well
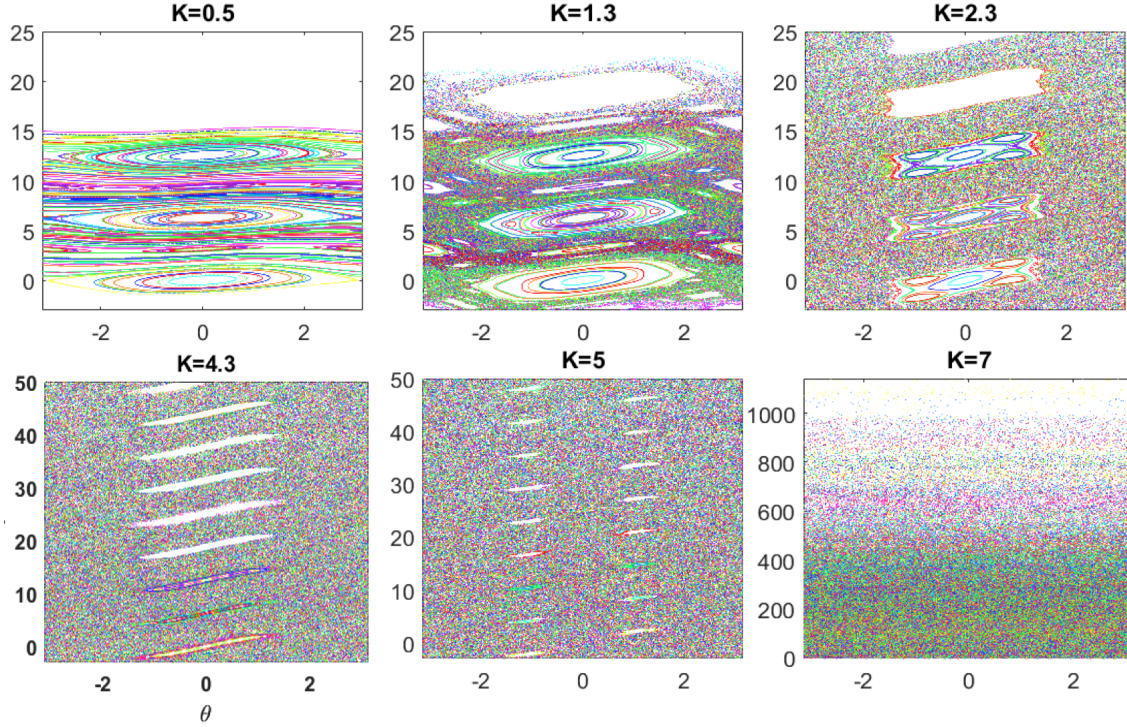
Figure 1-1: Kicked Rotor phase space for different kicking strengths. In each case, initial conditions are chosen randomly in $0 < v < 15$, and the subsequent evolution for each is plotted in its own color. While all regular trajectories are confined to their respective regions, the chaotic ones explore the entire accessible phase space, thus spreading to unbounded velocities over time

approximates a pendulum, making the dynamics regular for all initial conditions, while for larger K, last islands of regularity disappear in the sea of chaos (or rather become too small for the dynamics to typically find them) – see fig. 1-1.

Looking back at our claims in the main text, we can clearly identify the regular islands in configuration space as regions of low T_{eff}, as all the motion there is perfectly periodic, while all configurations in the chaotic sea are then naturally high-rattling. The configurations on the boundaries between these have an interesting role: these are near-critical, where correlations decay as power-laws rather than exponentially, and thus it isn't clear whether we should classify such dynamics as "random noise" or as "organized motion." In fact, this precisely illustrates the scenario where our claim of diffusive approximation will break-down altogether, and so for such situations we need not define T_{eff}. The key observation, however, is that such critical states are

measure-zero in this configuration space, and so while they may play an important role in understanding further details of the behavior, they can be ignored at a first pass.

So why then do we not see trajectories ending up in the low-rattling regular islands? Well thus far the dynamics are highly constrained by being time-reversible and deterministic, and so do not fit our "typicality" criterion. To consider a more realistic physical system, we proceed to weakly couple the above Hamiltonian dynamics to a heat bath by adding a damping and corresponding noise terms:

$$\dot{\theta} = v$$
$$\dot{v} = -K \sin\theta \, \delta(t - n) - b\,v + \sqrt{2\,T\,b}\,\xi \tag{1.2}$$

the noise amplitude is fixed to $\sqrt{2\,T\,b}$ by Einstein relation (with T being the bath temperature), and weak coupling to the bath means we take $b \ll 1$ (so that the qualitative features of the kicked rotor system still remain). As first observed by Feudel et. al. in 1996 ([23, 22, 37] and [53] for recent review), such small damping turns each regular region of the Hamiltonian system into a dynamical attractor – see fig. 1-2 a. This is exciting as an arbitrarily small amount of damping results in completely integrable long-time behavior for an otherwise mostly chaotic system. Furthermore, the addition of noise shows that each of these attractors acts effectively just like an energy well with a depth $\propto \delta v^2$, where δv is the particular attractor width in velocity space (here of $O\,[1]$ for the main attractors). This way, the system can hop between different attractors (see [37]) as long as the bath temperature $T \gtrsim \delta v^2$ (note: the only b-dependence here comes from that of the attractor shape and size). Thus, any particular regular region will only become a sink at long times if $T \ll \delta v^2$ for that attractor, (nearly) *independent* of the magnitude of the bath coupling (given by b).

Thus, we have a picture fully consistent with the claims in the main text. If background thermal bath temperature T is hight, then the ratio of effective temperatures $(T_{regular} + T)/(T_{chaos} + T)$ is not sufficient to create a substantial variation in $p_{s.s.}$.
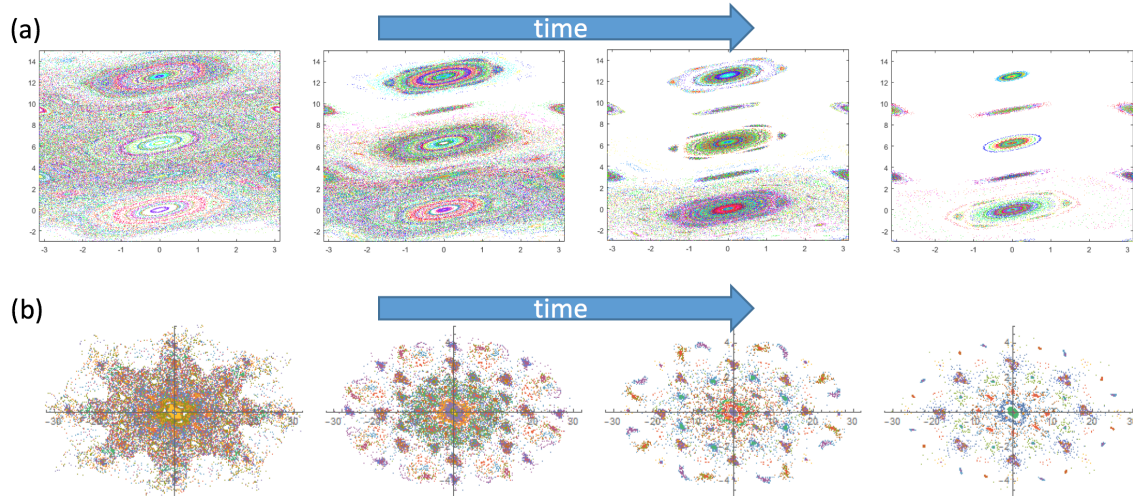
Figure 1-2: a) Time-evolution of phase space for Kicked Rotor ($K = 1.3$) weakly coupled to a thermal bath ($b = 10^{-3}$, $T = 10^{-2}$). Damping now confines velocities in the chaotic region within a Gaussian envelope, while regular regions become attractive sinks. b) Similar phenomenon observed for the web-map ("kicked wind-up" – spring-loaded kicked rotor), under minimal bath-coupling.

The correlation $p_{s.s.} \sim 1/T_{eff}$ isn't very interesting to plot here as we essentially just have two possible effective temperature values in this system.

This example suggests another conjecture, related to the one discussed above, but more precise and restricted: for any Hamiltonian system exhibiting mixed chaos, a weak coupling to a thermal bath at a sufficiently low temperature, causes all phase-space density to collapse into the regular regions over time. Since coupling to a bath is typical for macroscopic systems, this could give some intuition for why chaotic dynamics are rarely observed in our daily experience. To support the generality of this phenomenon, in fig.1-2 b we show this same effect in a slightly richer dynamical system – the web-map, which is essentially a spring-loaded kicked rotor:

$$\dot{\theta} = v$$

$$\dot{v} = -k\,\theta - K\,\sin\theta\,\delta(t-n) - b\,v + \sqrt{2\,T\,b}\,\xi$$

In [22], the authors observe similar dynamic self-organization in three different mixed-chaotic systems, and make some arguments for generality, but to our knowledge, no

clear statements of such conjecture, let alone rigorous results, are available. One might hope to arrive at such general conclusions using tools of dynamical systems theory – e.g., considering the dissipative version of KAM theorem, etc. However, the fact that it is the bath temperature (rather than the direct model coefficients b or $\sqrt{2\,b\,T}$) that serves as a crossover parameter between qualitatively different dynamical regimes, indicates that a thermodynamic approach might be more appropriate.

For now, the main purpose of this example was to illustrate our general proposal in a more concrete context, showing diffusive-like behavior of the dynamics with an inhomogeneous T_{eff} landscape, and the resulting self-organization into regular dynamical patterns. However, being very simple and low-dimensional, it doesn't have much richness of possible behaviors, and so no diversity of T_{eff} amplitudes – thus only giving us a qualitative illustration.

1.B   Drift in multiplicative noise

In this appendix we discuss some transient properties of stochastic processes with multiplicative noise, such as the one in eq. 1.1. The key result that we focus on is that a "typical" behavior in a finite number of process realizations is to drift towards the less-noisy regions (in absence of other forces). Intriguingly, this can vastly differ from the *mean* behavior expected from solving the stochastic differential equation (SDE) due to the typical heavy tails. To our knowledge (surprisingly), this effect was only first discussed in the finance literature in 2011 in [50] and then in [52], where ergodicity breaking properties of such processes were highlighted.

Here we focus on a practical question: given a multiplicative noise Ito process, what is the right intuition about its typical dynamics. To begin, we consider the simple dynamics $\dot{x} = x \cdot \xi$ (where ξ – univariate white noise) starting from initial distribution $x(0) = \delta(x - 1)$. This can be solved exactly to give: $\ln x(t) = -t/2 + \mathcal{N}(0, t)$, where $\mathcal{N}(0, t)$ is the normal distribution with variance t. The resulting log-normal distribution has the mean: $x(t) = 1$ and median: $x(t) = \exp[-t/2]$. However, as the distribution is heavy-tailed, the mean is sensitive to rare large fluctuations,

and become increasingly so as time goes on. In particular, for an ensemble of N realizations, the log-normal distribution is sufficiently sampled, and ensemble mean reproduces the above exact value, only for $t \lesssim O[\log N]$. At later times, the large-x tail of the distribution is insufficiently sampled, and so ensemble mean ultimately decreases towards 0. All the while the median is (as it is in general) insensitive to outliers, and thus reliably reproduces the expected exponential decay - see figure 1-3a.

The easiest way to make this precise is to consider the convergence to the Central Limit Theorem of the sum of N i.i.d. log-normal random variables. The leading correction to Gaussian distribution if of $O\left[\frac{C_3}{\sqrt{N}C_2^3}\right]$, where C_n denotes the n-th cumulant of the individual distributions. As the cumulants of the above log-normal distribution $\sim e^t$, an exact calculation shows that this correction only becomes small when $N \gg e^{3t}$ (for large t), which quickly fails for any practical ensemble of realizations. Moreover the sign of the corrections can be used to say that the deviation from the mean will be towards 0. Note finally that in-so-far as this result constitutes non-commutativity of the limits $N \to \infty$ and $t \to \infty$, this process is non-ergodic.

We thus conclude that for all practical purposes, the mean of the distribution $\langle x(t) \rangle = 1$ is utterly uninformative about the typical behavior. On the other hand, since for any monotonic function $f(x)$ (here being $\log(x)$), $\text{median}[p(f(x))] = f(\text{median}[p(x)])$, the fluctuations of the median of our finite ensemble are no different than for a normal Gaussian. Note also that for Stratonovich multiplicative noise process $\dot{x} = x \circ \xi$, the situation is of the same flavor, except now the mean is $e^{t/2}$ and median is constant at 1.

Besides being more useful for thinking about typical ensemble behavior, the median is also simpler to predict. This is because by virtue of being insensitive to outliers of the distribution, its dynamic are "more local." I.e., if in the above Itô process $\dot{x} = x \cdot \xi$, starting at $x(0) = 1$ we add a far-away hard wall at $x = 10$, the mean, though never coming close to that, is strongly affected, while the median is not - see fig. 1-3c. Note that if we have a more complicated multiplicative noise function – with multiple wells and gradients – then the full probability distribution really becomes necessary to comment on typical trajectories. Nonetheless, the me-
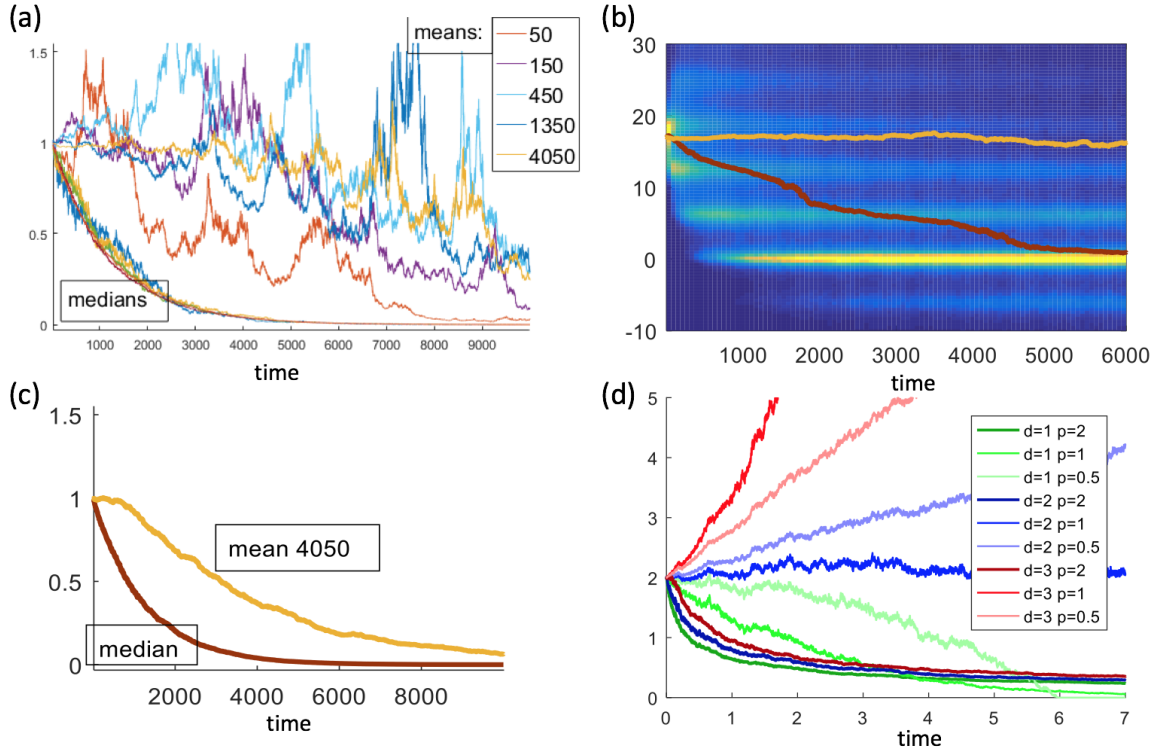
Figure 1-3: a) ensemble mean and median for multiple realizations of $\dot{x} = x \cdot \xi$ process. Ensemble sizes, shown in the legend, strongly affect the mean, but not the median. b) Evolution of probability over time for $\dot{x} = (|x|/2 - 3\cos x + 5) \cdot \xi$, along with the mean and median curves. Though the distribution is complicated, median still usefully captures the bulk probability behavior. c) the effect of introducing a hard wall at $x = 10$ on the mean and median curves for process from (a). d) Medians for multiplicative noise process in higher dimension, as illustrated by $\dot{x}_i = |x_i|^p \cdot \xi_i$ process, with $i = \{1, ..., d\}$ for various values of d and p. To control stability of the simulation, a wall was introduced at $|x_i| = 50$, which should not have much effect according to results shown in (c).

dian behavior can often still be useful as a rough approximation of the bulk of the probability density (and typically much better than the mean), and its dynamics still remain largely local – see fig. 1-3b. Note that for the process in the figure, the noise is nowhere zero, and so the effect is not explained simply by adsorption.

Finally, as we are ultimately interested in problems where effective dynamics 1.1 take place in a high-dimensional space, it is worthwhile to look at multi-dimensional simple multiplicative noise processes. The key point we want to illustrate here is how the effect of median drift in multiplicative noise competes against entropic effects countering fine-tuning. For this, we consider the d-dimensional process: $\dot{x}_i = |x_i|^p \cdot \xi_i$, with $i = \{1, ..., d\}$, and p controlling the noise amplitude scaling. Using rotational symmetry, we can then explicitly write the SDE for radial distance from the origin (applying the Itô chain rule): $\dot{r} = \frac{d-1}{2} r^{2p-1} + r^p \cdot \xi = \frac{d-p-1}{2} r^{2p-1} + r^p \circ \xi$. Though this is difficult to solve explicitly in general, we can immediately note the trend set up by the drift term in the Stratonovich version of the equation: $d - p - 1 < 0$ will typically result in a drift to the origin. To check this explicitly, we simulate this process for different values of d and p and plot the medians for each ensemble of realizations in fig. 1-3d. The key thing to note from this plot is that for a sufficiently high value of p, the fine-tuning required to find the origin is typically possible. Note that this might be somewhat surprising for $d = 3$, where the usual random walk is never expected to hit the origin, and the only difference here is that we have modulated the step size in a radius-dependent way. Perhaps a useful way to get intuition about this is as a random walk in a Riemannian manifold with a metric given by the inverse noise variance – this way low-noise regions become more "attractive" simply because they have a larger proper volume (while this analogy is only qualitative, it can help to build intuition).

This way, guided by $d - p < 1$ as a rule of thumb, we can say that T_{eff} scaling should generally be commensurate to the degree of the fine-tuning we want to get in order for the process to typically converge. In particular, this means that if T_{eff} drops exponentially towards the fine-tuned configuration, then we can expect convergence no matter how strong the desired fine-tuning is.

Chapter 2

Systems with time-scale separation

A broad range of many-body nonequilibrium systems have in common that different degrees of freedom within them undergo motion on two, well-separated time-scales, and that the faster degrees of freedom are the only ones directly subject to external driving. Such separation can occur if a faster set of active particles acts as a bath for a heavier, more slowly relaxing set of larger, extended degrees of freedom, such as in the example of a polymer immersed in a mixture of self-propelling particles [49]. Alternatively, in many systems one can usefully identify coarse-grained variables describing global features of the many-body dynamics, which may relax more slowly than the coordinates of individual particles. Such order parameters might then be thought of as a set of slowly-varying constraints on the driven fast dynamics, as for example in [57].

In all such cases, it is possible in principle for the particular configuration of a set of slow variables to have a significant influence on the specific nonequilibrium steady-state reached by the fast variables. Thus, in general, a feedback loop can arise in which the slow variables first establish the features of the fast steady-state, and then the statistics of this steady-state in turn determine the stochastic dynamics of the resulting local motion in slow variable space. The goal of this paper is to characterize the dynamical attractors of slow variable evolution in terms of the particular, special properties of the fast steady-states to which they give rise.

Nonequilibrium systems with time-scale separation have been extensively studied

over the last several decades. The most common context where they have come up is in formalizing the concept of a "thermal bath" – explicitly modelling the fast bath degrees of freedom as a Hamiltonian system, and studying their effects on the slow variables. In this way, one can in some cases recover the effective friction tensor [2], and the corresponding noise term, related by fluctuation-dissipation theorem [32]. There is also extensive literature studying the conditions and effects of deviations from this basic result, which are generally termed "anomalous diffusion" – see e.g., [39]. Within this context, the "slow" degrees of freedom lack their own dynamics, and are considered only as probes of the fast bath. More recent studies have considered the minimal dissipation required from an external agent to slowly move such probes. A geometric interpretation of this bound was presented in [71], and extended to nonequilibrium baths in [72], as well as to reversible external protocols in [40]. Systems where slow variables have their own dynamics under a conservative coupling to the fast bath have received relatively little attention, excepting notable recent work for a simple harmonic oscillator probe in [19], and a more general exploration in [41, 1], where some formal results relating dissipation and forces on the slow variables were derived.

Most of this previous work has relied on the projection operator technique to adiabatically eliminate fast variables and obtain the reduced Fokker-Planck equations for the slow variables, as in Ch6.4 of [26], or see [3] for a recent review. The straightforward implication of this approach is that at long times, probability density in slow variable space is expected to accumulate in locations where inward mean drift is strong, and where local diffusion is low. Here, we first derive this effect for a general class of Langevin systems using a response-field path integral framework that makes clear the relationship between the reduced Fokker-Planck parameters and the absorption and thermalization of drive energy in the fast steady-state. Some related path-integral system reduction techniques have been studied before (e.g., [24, 5]), but in substantially different contexts. Appendix 2.A helps illustrate our technique by going through two explicit calculations related to the Itô-Stratonovich dilemma, re-deriving some well-known results in a new way. Having established a means of explicitly calculating the parameters of the multiplicative noise stochastic process

governing the slow variables, we then proceed to analyse the implications for what we term "least rattling feedback," in which slow variables dynamically finely-tune themselves to bring about fast variable steady-states that attenuate force fluctuations so as to lower the slow variable effective temperature.

The tendency of slow variables in driven systems to move thermophoretically towards regions of lower effective temperature has been noticed in the past, most commonly in situations where the slow variables find a way to reduce the influx of energy from the drive (as in [42, 13]. As we shall see here, however, a striking alternative can arises if the fast variables are capable of exhibiting regular, integrable dynamics; in such a case, least rattling stability can co-exist with strongly coupling to and absorbing work from the external drive.

In the rest of this chapter, we will present the derivation of our main analytical result, which establishes a relationship between force fluctuations in fast driven variables and the resulting effective temperature experienced by the slow variables in a driven system. In chapter 3, we will carry out a numerical analysis of the kicked rotor on a cart – a time-scale separated, damped, driven dynamical system that is ideally suited for demonstrating the predictive power of the "least rattling" framework. Not only will this analysis draw clear connections to methods of equilibrium statistical physics and show how they generalize in such a nonequilibrium scenario, but it will also underline how "least rattling" helps to explain the non-trivial relationship between dissipation rate and local kinetic stability in driven systems.

2.1   Setup

We will model "slow" variables x_a and "fast" variables y_i that evolve according to a coupled system of Langevin equations. Our approach will be to integrate out the fast degrees of freedom and develop an effective theory for the dynamics of the slow variables that is controlled by a small number ϵ which quantifies the time-scale separation between fast and slow. As we carry out this integration, we will show that the effects on x_a from the fast steady-state of the y_i variables at leading order in ϵ

are an average force and, more subtley, a random force and renormalized drag that are calculated from the two-point correlation function of the forces acting between x_a and y_i. These latter effects are identified as an emergent, position-dependent effective temperature experienced by the slow coordinates.

While the method we present here is not restricted to this context, it is easiest to illustrate on systems whose dynamics can be given by first order equations, as below. In particular, it works the same way for other types of fast dynamics – such as inertial, or discrete – as long as there is a fast relaxation to a steady-state.

$$\eta \, \dot{x}_a = F_a(x_a, y_i, t) + \sqrt{2 \, T \, \eta} \, \xi_a$$
$$\mu \, \dot{y}_i = f_i(x_a, y_i, t) + \sqrt{2 \, T \, \mu} \, \xi_i. \tag{2.1}$$

Here the noise ξ is usual Gaussian white noise: $\langle \xi_\alpha(t) \rangle = 0$ and $\langle \xi_\alpha(t) \xi_\beta(s) \rangle = \delta_{\alpha,\beta} \delta(t - s)$. Taking the limit $\mu/\eta \equiv \epsilon \ll 1$ amounts to explicitly separating x_a as slow modes, and y_i as fast ones (a, b, c index the slow configuration space, and i, j, k – the fast one). The natural physical interpretation of this system as overdamped dynamics in a thermal bath of temperature T, with two different damping coefficients μ and η, the noise amplitudes given by Einstein's relation, and with the forces F_a and f_i will be implied from now on for concreteness, but is not at all necessary. With a slight adjustment the system could as well represent underdamped dynamics, such as in the kicked rotor model system we characterize below.

2.2 Results

The detailed derivation of the effective slow dynamics is relegated to Appendix 2.B. Here we mention only the key steps in the derivation. First, we rescale time $t \to \mu \, t$, making the slow dynamics obey $\dot{x}_a = \epsilon \, F_a + \sqrt{2 \, T \, \epsilon} \, \xi_a$, while the relaxation time of fast variables becomes of $O[1]$. Second, we express probability of slow trajectories in terms of the Martin-Siggia-Rose path integral (also termed the response-field formalism) [8],

and third, we do a cumulant expansion controlled by ϵ:

$$P[x(t)] = \frac{1}{Z_x} \int \mathcal{D}\tilde{x} \left\langle \exp\left[-\int dt \left\{ i\tilde{x}_a \left(\dot{x}_a - \epsilon F_a(x, y, t) \right) + \epsilon T \tilde{x}_a^2 \right\} \right] \right\rangle_{y|x(t)}$$

$$= \frac{1}{Z_x} \int \mathcal{D}\tilde{x} \exp\left[-\int dt \left\{ i\tilde{x}_a\dot{x}_a + \epsilon T \tilde{x}_a^2 - i\epsilon\tilde{x}_a \langle F_a \rangle_y + \frac{\epsilon^2}{2}\tilde{x}_a\tilde{x}_b \langle F_a, F_b \rangle_y + O\left[\epsilon^3\right] \right\} \right].$$

$$(2.2)$$

where Z_x is the normalization, and $\tilde{x}(t)$ is the auxiliary "response" field. In the last line, we see that the $O\left[\epsilon^2\right]$ term in the expansion, like temperature T, comes in $\propto \tilde{x}^2$, and thus gives a correction to the noise on the slow dynamics – this is the effect that we will focus on throughout the rest of this paper. Doing this more carefully (as shown in the Appendix 2.B) the resulting slow dynamics, which is our main analytical result, are

$$\gamma_{ab} \cdot \dot{x}_b = \epsilon \langle F_a \rangle_{y|fix\ x} + \sqrt{2\,\epsilon}D_{ab} \cdot \xi_b \qquad (2.3)$$

$$\gamma_{ab}(x) = \delta_{a,b} + \epsilon \int dt'\ (t - t') \left\langle i\tilde{y}_i\,\partial_b f_i\,\big|_{t'}, F_a\,\big|_t \right\rangle_{y|fix\ x}$$

$$D_{ab}(x) = T\,\delta_{a,b} + \frac{\epsilon}{2} \int dt'\ \left\langle F_a\,\big|_{t'}, F_b\,\big|_t \right\rangle_{y|fix\ x}.$$

where the matrix square root is defined by $B \equiv \sqrt{D} \iff B.B^T = D$. Dots denote Itô products, which will be typical here (see sec.2.B.2). Note that only the connected components of the expectations appear in the expressions for $\gamma(x)$ and $D(x)$ (denoted by commas), and thus are insensitive to any deterministic motion of the fast variables. Further note that there is also an $O\left[\epsilon\right]$ correction of the damping coefficient, which, for a fully conservative (undriven) system, matches the noise correction to preserve Einstein's relation, as it must (see sec.2.B.4). For non-conservative forces, however, this will not be the case, and the ratio of the effective noise to damping amplitudes can be used to define an effective temperature tensor $T_{eff}(x_a) \equiv \gamma^{-1}.D.\left(\gamma^{-1}\right)^T$, which will generally depend on the slow coordinates – i.e., the noise on slow variables becomes multiplicative.

31

2.3   Least Rattling

The significance of the above formal result is that to extract the effective slow dynamics we need not know everything about the fast modes, but only the mean and variance of the force fluctuations F_a in the y_i (fast) steady-state at fixed x_a (slow d.o.f.). All other details of the fast dynamics become irrelevant by the same mechanism as for the central limit theorem. The slow dynamics thus follow the simple equation 2.3, which can often be solved analytically. Its qualitative behavior is guided by a competition between the mean drift along the average force $\langle F(x) \rangle$ and a median drift down the effective temperature gradients $T_{eff}(x)$. While the former effect is larger by a factor $1/\epsilon$, it is a vector quantity, and as such, may be suppressed by averaging in case of high-dimensional disordered fast dynamics. This is in contrast to T_{eff}, which comes in as a positive-definite tensor, making it robust to averaging-out. Without rigorously exploring this trade-off for now, in this chapter we simply choose focus on the effect of $T_{eff}(x)$, which guides the slow variables towards regions in their configuration space that yield more orderly, less chaotic, or less "rattling" fast dynamics (see sec.2.B.5). We suggest that this effect might result in the self-organization lately studied in many non-equilibrium systems [55, 58].

We now expand on a few of the points mentioned above. First, how general is this method? Its scope is basically inherited from the regime of applicability of the Central Limit Theorem (CLT): our requirement of strong time-scale separation amounts to the condition that fast fluctuations decorrelate faster than dynamical time-scale of slow variables. This way their effect on the slow coordinates adds up as i.i.d. random variables, satisfying the conditions of CLT. Thus any fast fluctuations must either decorrelate quickly (e.g., due to thermal noise or chaos) – thus contributing to the Gaussian noise amplitude, or not decorrelate at all (as with integrable behavior) – contributing to the mean force $\langle F \rangle$. This requirement could notably be broken if some fast fluctuations decay slower than exponentially – a scenario that leads to effective colored noise and anomalous diffusion, but retains much of the general intuition from eq.2.3.

This framework is particularly useful in cases where fast dynamics can be in several qualitatively different dynamical phases, controlled by the slow variables. E.g., if a fast variable undergoes a transition from chaotic to integrable behavior as a function of some slow coordinate, then we will typically expect its effect to transition from a noise contribution to an average force contribution respectively – as we will see in the toy system of chapter 3. Making this precise and describing the relevant universality classes of these transitions based on their symmetry structure can be done within the broader framework of renormalization group flow. This could allow extracting the effective slow dynamics, much like it allows finding large scale physics for quantum or statistical fields [35].

Finally, we mentioned above that while the average force $\langle F \rangle$ causes the *mean* of the x_a-ensemble (slow variables) to drift, the multiplicative Itô noise given by the effective temperature bath $T_{eff}(x)$, affects only a drift of the *median* of that same ensemble. This latter effect is realized by virtue of the $p(x_a)$ probability distribution growing increasingly heavy-tailed with time (e.g., log-normal distributions are typical), and so while the mean remains fixed, the median will drift towards the low-noise regions. This means that any finite ensemble of trajectories will also settle in the low-noise region, and the mean will never be realized experimentally. Some aspects of this ergodicity-breaking phenomenon were discussed in [52], and a similar problem considered in [59], with the relevant details further developed here in Appendix 1.B. The key for us is that the least-rattling effect is inherently non-ergodic, and is observed only by monitoring the system over time.

Appendix

2.A   Path-integral technique for marginalizing fast variables

In this appendix, we use two examples to illustrate a general path-integral technique for solving dynamical systems exhibiting motion on two strongly separated time-scales. In particular, we show how the scale-separation can be used to write down a diagramatic perturbation series governing the effective slow dynamics. Our first example solves Brownian motion in the presence of both temperature and viscosity gradients, showing how depending on these, we can get either effective Itô or Stratonovich multiplicative noise for the slow variables. The second example recovers the classic result that correlated noise leads to effective Stratonovich white noise when viewed on sufficiently slow time-scales.

2.A.1   Brownian motion in inhomogeneous medium

Despite its simple statement, different aspects of this problem have caused confusion in the stochastic thermodynamics community for a remarkably long time: from the first studies in the 60s (e.g. [70]), and until the latest publications in [10](2012) and [69](2013). The crux of the confusions is in the distinction between the Ito and Stratonovich interpretations of the multiplicative white noise that arises upon taking the overdamped limit of the motion. In the following, we show how in a field-theoretic description of the problem, all its subtleties are automatically accounted for by the mathematical framework itself, and as an example, we reproduce the correct Langevin

equation for $x(t)$, derived by other means in [69].

The well-posed unambiguous starting point for the problem is stated by including velocity degrees of freedom in the dynamics:

$$\begin{cases} \dot{x} & = v \\ m\,\dot{v} & = -\gamma(x)\,v + F(x) + \chi(x)\,\xi(t) \end{cases}$$

where $\langle \xi(t) \rangle = 0$ and $\langle \xi(t)\xi(s) \rangle = \delta(t-s)$ gives the univariate white noise, and $\chi(x) = \sqrt{2\,\gamma(x)\,T(x)}$ defines the temperature $T(x)$ via the Einstein relation. While the following derivation holds in arbitrary dimension, here we restrict to the 1D case for simplicity of notation. In the above system, due to a finite particle inertia, its position $x(t)$ does not respond instantaneously to noise kicks $\xi(t)$, and thus $\chi(x)$ remains constant over the "duration of the kick" – making the Itô and Stratonovich descriptions equivalent. As long as we are using the mathematical idealization of white noise, this fact remains true for any finite mass m, but fails when $m \equiv 0$ resulting in the ambiguous equation of form $\dot{x} = \chi(x)\,\xi(t)$. The challenge in deriving the Langevin equation for $x(t)$ thus lies in correctly integrating out the fast velocity degrees of freedom $v(t)$ in the limit of small but finite mass m. Most of the literature on the problem either works with Fokker-Planck equations (thus missing trajectory-specific information), or uses a projection operator technique (see [26]), which gets intractable even for fairly simple problems. In some cases, as in [69], some additional physical constraints are used to resolve the ambiguity.

Path integral and scaling

Here, we will begin with writing down the Martin-Siggia-Rose (MSR) path integral (also known as response-function formalism) for the system, constructing it in the Ito

interpretation:

$$Z = \int \mathcal{D}\tilde{x}\mathcal{D}x \int \mathcal{D}\tilde{v}\mathcal{D}v \; e^{-S[\tilde{x},x,\tilde{v},v]}$$

$$\text{with} \quad S = \int dt \left[i\,\tilde{x}\,(\dot{x} - v) + i\,\tilde{v}\,(m\,\dot{v} + \gamma(x)\,v - F(x)) + \frac{1}{2}\,(\chi(x)\,\tilde{v})^2 \right] \qquad (2.4)$$

As this is Gaussian in the v-fields, we could integrate them out exactly without any assumptions. The resulting x-dynamics would then have colored noise and non-Markovian forces (i.e., of form $\dot{x} = \int ds\; g(t-s)\,F(x(s)) + \xi_{colored}(t)$). Taking the $m \to 0$ limit would then give us the desired Markovian equation, but would be no simpler than (and follow along the same lines as) the procedure described below. Instead, it is more instructive, and illustrates all the key aspects relevant for more general scale-separation problems, to treat the interactions between the two fields perturbatively in the UV.

As such, we first consider Taylor-expanding all functions of x in the action to generate all kinds of allowed vertices. This is sufficient to evaluate any observables about the fast time-scale, as we expect x to stay close to its initial value x_0 in that regime. To get a first feel for the behavior of the resulting theory, consider the canonical scaling dimensions of all the generated couplings. Time $[t] = -1$ by definition. For the length scaling $[x]$ we have a choice of fixed points: in the UV, the motion is ballistic and $[x] = -1$ is appropriate, while in the IR, it's diffusive and so $[x] = -1/2$. As we are interested in the effective IR behavior, we want to see the relevance of operators with respect to the diffusive fixed point. Thus: $[x] = -1/2 \to [\tilde{x}] = 1/2 \to [v] = 1/2$. Moreover, since the response field $\tilde{v}(t)$ takes on the values of the noise realizations $\xi(t)$, the physically sensible scaling for it is $[\tilde{v}] = [\xi] = 1/2$. Together, these choices then fix $[m] = -1$, $[\gamma] = 0$, $[F] = 1/2$, $[\chi] = 0$, and consequentially $[\partial_x{}^n \gamma] = n/2$, $[\partial_x{}^n F] = (1+n)/2$, $[\partial_x{}^n(\chi^2)] = n/2$.

We see that all these interactions are relevant, and those with higher powers of ∂_x progressively more so, thus precluding any perturbative calculations towards the IR. This is physically intuitive as on the slow time-scale, the variations of $\gamma(x), F(x), \chi(x)$ strongly affect the dynamics, which thus depend on the particular form of these

37

functions. Instead, we are interested in integrating out the fast degrees of freedom, and thus require that on the relevant scale τ in the UV, the interaction vertices are perturbatively small, and decay with n. The only UV scale in the problem is the correlation length of $v(t)$ and so we let $\tau = m/\gamma(x_0)$ at some reference point x_0 (that we also use for all Taylor expansions). This then gives us a precise definition for our assumption of time-scale separation in the problem: rescaling the theory into the UV by a factor τ, we require $\tau^{n/2} \partial_x{}^n \gamma \gg \tau^{(n+1)/2} \partial_x{}^{n+1} \gamma \quad \forall n \geq 0$ – and same for $F(x), \chi(x)$. This can be easily achieved if γ, F, χ vary smoothly on a scale of $x \sim O[1]$, and $m \to 0$, as we intuitively expect.

Two potential points of confusion should be addressed. For one, clearly as we want $\tau = m/\gamma(x_0) \ll 1$, we must ensure that this holds for all points x_0 – else if γ ever gets too small, the particle's mean free path can be appreciable even for small m. The other point is that while we could have chosen any arbitrary τ as our UV scale, the choice of correlation length is convenient as it makes all propagators of $O[1]$, thus allowing for easy organization of the perturbation series, and so there is nothing fundamental about it.

Perturbative expansion in the UV

This way, the theory takes on a tractable form in the UV, akin to what happens with asymptotic freedom in QCD. Reading off the Feynman rules from eq.2.4, we get:

While typically at this point we might need to write down the RG evolution for the entire perturbatively coupled system, in this case the loop integrals can be evaluated exactly, and thus we can completely integrate out the fast variables in this tractable UV regime. Note that this feature will, in some form, have to hold generally in scale-separated systems (see below) – else the different scales would not really separate. This way, we need only to consider the renormalization of the slow-variable couplings.

At leading order $O\left[\tau^0\right]$, we have only one diagram generating a noise term for x – diagram (a) below. At the next order $O\left[\tau^{1/2}\right]$, there are two diagrams introducing source-terms, and another giving a velocity-dependent noise – b,c,d:
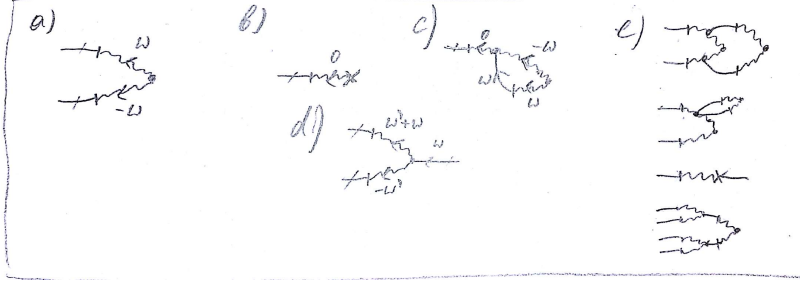
At $O\left[\tau^1\right]$, we have more diagrams, some of which are shown in (e). While we will not include these here, note that one of these is a renormalization of the propagator, which, upon Taylor-expansion in ω, will give corrections $\propto \dot{x}$, $\ddot{x}$, etc. The resulting mass-correction from $\ddot{x}$ term was the focus of the [19] paper (though using other techniques). Evaluating the diagrams (a-d):

$$(a) = -\frac{\chi^2/\gamma^2}{2\left|i\omega + 1\right|^2}$$

$$(b) = \frac{i\sqrt{\tau}\,F/\gamma}{i\,0 + 1}$$

$$(c) = \frac{i^7}{\gamma^3}\frac{\chi^2\sqrt{\tau}\,\partial_x\gamma}{i\,0 + 1}\int_{-\infty}^{\infty}\frac{d\omega}{2\pi}\frac{1}{\left|i\omega + 1\right|^2\,i\omega} = -\frac{i\sqrt{\tau}\,\chi^2\,\partial_x\gamma}{2\,\gamma^3}$$

$$(d) = \frac{-\sqrt{\tau}\,\partial_x(\chi^2)/\gamma^2}{2(i(\omega + \omega') + 1)(-i\omega' + 1)}$$

$$\approx -\frac{\sqrt{\tau}\,\partial_x(\chi^2)/\gamma^2}{2\left|i\omega + 1\right|^2}\,(1 - i\omega)$$

Note that all ω-dependence in the generated noise vertices makes the noise colored. Since we are only interested in effective slow dynamics, we can assume that relevant frequencies ω are below some cutoff Λ, which at our present UV scale is of $O\left[1\right]$. This way we get the approximation in the last step of (d) above, and similarly

$$(a) = -\frac{\chi^2/\gamma^2}{2(\omega^2 + 1)} \approx -\frac{\chi^2}{2\,\gamma^2}\left(1 - \omega^2 + O\left[\left(\frac{\omega}{\Lambda}\right)^4\right]\right)$$

39

which gives the desired white-noise term, and a correction $\sim \dot{\tilde{x}}^2$. Scaling back to our (IR) scale of interest according to the canonical dimension of the couplings, all the $\sqrt{\tau}$ factors disappear, while the UV cutoff $\Lambda \to \Lambda/\tau$, which then gives the effective action for the slow variables:

$$S_{eff} = \int dt \, \left[i\, \tilde{x} \left(\dot{x} - \frac{F}{\gamma} + \frac{\chi^2\, \partial_x \gamma}{2\,\gamma^3} \right) + \right.$$
$$\left. \frac{\chi^2}{2\,\gamma^2} \left(\tilde{x}^2 \left(1 - 2\, \partial_x(\ln \chi)\, \dot{x}\, \tau\right) - \dot{\tilde{x}}^2\, \tau^2 \right) \right]$$

Taking $\tau \to 0$ limit then gets rid of all the noise-color corrections. Note that this limit is separate from the one taken above to truncate the UV perturbative expansion. Whereas that assumption ensured that the parameters in the problem don't vary much over the mean free path of the Brownian particle, the limit here requires that the mean free time is much shorter than observation times. I.e., even with constant parameters, we would have a colored-noise effect if we sampled the dynamics on a time-scale comparable to the v relaxation times. Thus, while the diagram (d) does not contribute in the $\tau \to 0$ (long-time) limit, it provides the first correction when that limit is relaxed – i.e., when we sample the dynamics on faster time-scales by observing large $\dot{x}$.

Finally, we can write down the Langevin equation corresponding to the effective action above, giving:

$$\dot{x} = \frac{F}{\gamma} - \frac{\chi^2\, \partial_x \gamma}{2\,\gamma^3} + \frac{\chi}{\gamma} \cdot \xi$$

in the Ito interpretation – as our entire path-integral construction has been carried through in Ito. This reproduces the result in [69]. Note here that if $\chi(x) = \chi$ is constant, then this can be written simpler in the Stratonovich notation: $\dot{x} = \frac{F}{\gamma} + \frac{\chi}{\gamma} \circ \xi$ – to be contrasted to the case where $\gamma(x) = \gamma$ is constant: $\dot{x} = \frac{F}{\gamma} + \frac{\chi}{\gamma} \cdot \xi$

2.A.2 Delta-correlated limit of colored noise process

Having seen how an Ito equation can emerge for the effective slow dynamics, we might wonder under what circumstances a Stratonovich equation would be appropriate, and where would its derivation diverse from the one presented above. To illustrate this, we take the prototypical case that reduces to Stratonovich in the slow time-scale: that of colored noise whose correlation approaches delta-function: $\dot{x} = a(x) + b(x)\,\zeta(t)$ where $\langle\zeta(t)\rangle = 0$ and $\langle\zeta(0),\zeta(t)\rangle = \frac{1}{\epsilon}g(\frac{t}{\epsilon})$ is the colored noise with kernel $g(t)$ such that $\int_{-\infty}^{\infty} dt\,g(t) = g(\omega = 0) = 1$ and $g(t) \to 0$ for $t \gg 1$. This way the limit $\epsilon \to 0$ makes the noise $\zeta(t)$ white.

On the other hand, we can view $x(t)$ and $\zeta(t)$ as coupled slow and fast fields respectively, with the action in the MSR path integral: $S[x(t), \tilde{x}(t), \zeta(t)] = \int dt\,\tilde{x}\,(\dot{x} - a(x) - b(x)\,\zeta)$ and canonical RG scalings: $[t] = -1$, $[x] = -1/2$ (diffusive), $[\zeta] = 1/2$ (from the above correlation function) $\Rightarrow [\tilde{x}] = 1/2$, $[a] = 1/2$, $[b] = 0$, $[\partial_x{}^n b] = 1/2^n$. We assume that $\partial_x{}^n b \sim O[1]$, and $\zeta(t)$ have very short-range correlations, which means that we are currently in the IR sector of our theory. We can rescale our theory to the UV by factor ϵ using the canonical scalings to get:

$S = \int dt\,\tilde{x}\,[\dot{x} - (\sqrt{\epsilon}\,a + \epsilon\,x\,\partial_x a + ...) - (b + \sqrt{\epsilon}\,x\,\partial_x b + ...)\,\zeta]$. Perturbatively in ϵ, this gives the following diagrams (solid lines – x-propagators, wavy lines – ζ propagators):

Note that here we exploited our knowledge that $\zeta(t)$ is Gaussian and independent of x, thus allowing us to capture any arbitrary kernel $g(t)$ with just a single propagator line (which accounts for all self-interactions). At $O[1]$ we get a propagator and noise term, while at $O[\sqrt{\epsilon}]$ – two drift forces:

$$= \int \frac{d\omega}{2\pi}\left[\tilde{x}\,x\,i\omega + \tilde{x}^2\,g(\omega) - \sqrt{\epsilon}\,\tilde{x}\,a - \sqrt{\epsilon}\,\tilde{x}\,b\,\partial_x b \int \frac{d\omega'}{2\pi}\,\frac{g(\omega')}{i\omega' + \epsilon\,\partial_x a} + O[\epsilon]\right]$$

arranged in the same order as in the diagrems. In the last integral, we used the $\epsilon\,x\,\partial_x a$ generated mass term to regularize the propagator. Since that integral is given by the

residue at its only pole at $\omega' \sim O\,[\epsilon]$, it evaluates to $g(\omega' \sim 0) \sim 1$ as above. Scaling back to the IR now that the "fast field" ζ has been integrated out, we read off the Ito Langevin equation: $\dot{x} = a(x) + \frac{1}{2} b\, \partial_x b + b(x) \cdot \xi(t)$, which we then recognize can be simpler expressed in Stratonovich convention: $\dot{x} = a(x) + b(x) \circ \xi(t)$.

2.B Derivation of effective thermal bath

Starting with the explicitly time-scale separated dynamics given by the system 2.1, the first step is to explicitly bring our small parameter $\epsilon \equiv \mu/\eta \ll 1$ into the equations by rescaling time $t \to \mu\, t$, which gives:

$$\dot{x}_a = \epsilon\, F_a(x_a, y_i, t) + \sqrt{2\,\epsilon\, T}\, \xi_a$$
$$\dot{y}_i = f_i(x_a, y_i, t) + \sqrt{2\, T}\, \xi_i \qquad (2.5)$$

This will allow us to do a systematic expansion in ϵ below. For later convenience, let us explicitly introduce the time-scales of slow $\tau_S \sim O\,[1/\epsilon]$ and fast $\tau_F \sim O\,[1]$ relaxation. Next we want to integrate out the fast degrees of freedom y_i, which we can explicitly do by writing down the Martin-Siggia-Rose (MSR) path integral expression for the probability of a particular slow trajectory $x_a(t)$, as given in the first line of eq. 2.2. For clarity of notation, we have represented the full path-integration over the fast dynamics as the average:

$$\langle \mathcal{O} \rangle_{y|x(t)} \equiv \frac{1}{Z_y[x(t)]} \int \mathcal{D}y\, \mathcal{D}\tilde{y}\, \mathcal{O} \exp\left[-\int dt\, \left\{ i\tilde{y}_i\, (\dot{y}_i - f_i(x(t), y, t)) + T\, \tilde{y}_i^2 \right\} \right] \quad (2.6)$$

Note that so far, this is defined for a specific fixed slow trajectory $x(t)$. With this notation set, we now observe that the only y-dependence that the average can act on in eq. 2.2 is that in $F(x, y, t)$. Thus, all the other terms can be taken out of the average, while the remaining small $\epsilon\, F_a$ exponent can be treated with a cumulant

expansion:

$$
\left\langle \exp\left[i\,\epsilon \int dt\ \tilde{x}_a\, F_a(x,y,t) \right] \right\rangle_{y|x(t)} \approx
$$
$$
\exp\left[i\,\epsilon \int dt\ \tilde{x}_a(t)\, \langle F_a(x,y,t)\rangle_{y|x(t)} - \epsilon^2 \int dt \int dt'\ \tilde{x}_a(t)\, \tilde{x}_b(t')\, \langle F_a(x,y,t), F_b(x,y,t')\rangle_{y|x(t)} \right]
$$

$$(2.7)$$

to $O\left[\epsilon^3\right]$. As this is the only part of the path integral in eq 2.2 that carries the coupling to fast dynamics, it will source all the interesting emergent effects (i.e., coupling renormalizations) for the slow modes, and we thus focus on this for most of this Appendix.

2.B.1 Averages over fast dynamics

Before getting into the physical implications of the different terms in the expansion, let's discuss how to go about calculating the averages $\langle \mathcal{O}\rangle_{y|x(t)}$. Indeed, as defined in eq. 2.6, these averages are just short-hand for path integrals over the full fast dynamics in the presence of arbitrarily time-varying slow variables $x_a(t)$, and hence at this point, merely formal, but not very informative, quantities. On the other hand, intuitively we know that at the lowest order in ϵ, these averages should reduce to averages over the y-steady states under fixed x: $p_{ss}(y|x)$. To derive this result as well as the first correction in ϵ, we need to once again develop a systematic expansion. Besides $\mathcal{O}$ (see below), the only dependence on the trajectory $x(t)$ in eq 2.6 comes in through the force f_i (and similarly in the partition function), which by time-scale separation assumption, we know will vary only slightly on the fast time-scale τ_F: $f_i(x(t),y,t) = f_i(x(t_0),y,t) + (t-t_0)\,\dot{x}_a(t_0)\,\partial_a f_i(x(t_0),y,t) + O\left[\epsilon^2\right]$. Plugging this expansion into eq. 2.6 and Taylor expanding both numerator and denominator

43

(normalization Z_y) in ϵ, we get:

$$\langle \mathcal{O} \rangle_{y|x(t)} = \langle \mathcal{O} \rangle_{y|x(t_0)} + \left\langle \mathcal{O}, \int dt\, i\tilde{y}_i\,(t - t_0)\,\dot{x}_a(t_0)\,\partial_a f_i(x(t_0), y, t) \right\rangle_{y|x(t_0)} + O\left[\epsilon^2\right]$$

$$= \langle \mathcal{O} \rangle_{y|x(t_0)} + \dot{x}_a(t_0) \int dt\, (t - t_0) \left\langle \mathcal{O},\, i\,\tilde{y}_i\,\partial_a f_i\,\big|_t \right\rangle_{y|x(t_0)} + O\left[\epsilon^2\right] \qquad (2.8)$$

The averages here $\langle \mathcal{O} \rangle_{y|x(t_0)}$ are at a fixed x, and thus are precisely the averages over y steady-states $p_{ss}(y|x(t_0))$. Note also that at this order, the possible x-dependence of $\mathcal{O}$ is accounted for at the slow time-scale and does not give any additional contributions here. At the next order in ϵ, however, the variations of $\mathcal{O}$ and f_i on the relaxation time-scale τ_F begin to interact, giving new contributions. Generally, Feynman diagrams are the only practical way to go to higher orders as the number of correction terms potentially becomes large. We do not employ diagrams here because they are not practical for the general context we are working with – but they do become very useful in specific examples.

Applying the result in eq. 2.8 to our cumulant expansion 2.7, we get

$$\langle F_a(x, y, t) \rangle_{y|x(t)} = \left\langle F_a\,\big|_t \right\rangle_{y|fix\,x(t)} - \dot{x}_b(t) \int dt'\, (t - t') \left\langle i\,\tilde{y}_i\,\partial_b f_i\,\big|_{t'},\, F_a\,\big|_t \right\rangle_{y|fix\,x(t)} + O\left[\epsilon^2\right]$$

$$(2.9)$$

Remembering the form of eq. 2.2, we recognize the correction term here as a correction (or renormalization) of the damping coefficient of the original slow dynamics. Crucially, this correction comes in at the same order in ϵ as the $O\left[\epsilon^2\right]$ term in eq. 2.7, and must thus be kept in our expansion. By the same token, the equivalent correction of the $\langle F_a(t), F_b(t') \rangle_{y|fix\,x(t)}$ term only comes in at a higher order and is thus ignored at this stage. Nonetheless, it is interesting to note that including higher order corrections would also introduce an inertia-like term into our slow dynamics (mass renormalization, $\propto \ddot{x}$), as discussed in [19]), as well as higher derivative at progressively higher orders.

2.B.2 Noise correction

The key thing to note from the above results is that at $O\left[\epsilon^2\right]$, the only thing we need to compute the slow dynamics are the one- and two-point functions of the various variables in the y-steady-states. Thus we need neither the full form of the steady-state distribution of the fast variables, nor the deviations from this steady-state under dynamic $x(t)$. This result should be thought of as (and really is a form of) the central limit theorem.

Now that we have a sense of what the different terms in the cumulant expansion 2.7 mean mathematically, we can turn to their physical implications. We already mentioned the correction of the x-damping coefficient that we get by resolving the $\langle F_a \rangle_{y|x(t)}$ term in terms of y-steady-states. The only other contribution at this same order is the second term in the expansion 2.7. This will contribute an additional noise term to the resulting slow dynamics, as it will enter the path integral along with T, correcting the $\tilde{x}^2$ operator. However, this noise term would only be white if $\langle F_a(t), F_b(t') \rangle_y \sim \delta(t - t')$, which in general need not be the case, hence making the noise correction colored. Intuitively, we see that because of time-scale separation, y fluctuations will decorrelate much faster ($\delta t \sim O\left[1\right]$) than the slow time-scale we are sampling by observing x ($\tau_S \sim O\left[1/\epsilon\right]$). This makes the short-range correlations of the noise correction unimportant for the slow evolution, allowing us to approximate it by white noise.

More formally, this situation is precisely identical to having a UV cutoff in a field theory given by, e.g., finite lattice spacing. Similarly, taking the white-noise approximation here corresponds to sending such a cutoff to infinity, which is justified as long as all our observables are confined to energy-scales (or here time-scales) far lower than said cutoff. Explicitly, the approximation we are making (which formally comes from the assumption of RG universality in the fast dynamics):

$$\int dt \int dt' \, \tilde{x}_a(t) \, \tilde{x}_b(t') \, \langle F_a(x, y, t), F_b(x, y, t') \rangle_{y|x(t)} \approx$$
$$\int dt \, \tilde{x}_a(t) \, \tilde{x}_b(t) \int dt' \, \langle F_a(x, y, t), F_b(x, y, t') \rangle_{y|fix \, \hat{x}} \qquad (2.10)$$

One reason why we must be careful in taking this white-noise limit is that the precise limiting procedure will determine whether the correct interpretation of the resulting multiplicative white noise is Itô or Stratonovich, resulting in observable consequences on the slow time-scale. This question is related to the choice of $\hat{x}$ where to evaluate the y-steady-state in the RHS of eq. 2.10, as well as, independently, where to evaluate the explicit dependence $F_a(x)$. To avoid very messy notation, we assume away the latter point by restricting the form of $F_a(x, y, t) = \tilde{F}_a(x) + F_a(y, t)$, which then makes the above expression 2.10 depend on x only via the steady-state $p_{ss}(y|\hat{x})$ ($\tilde{F}$ drops out altogether as it only contributes to the disconnected cumulant): define $\delta T(\hat{x}) \equiv \frac{\epsilon}{2} \int dt' \, \langle F_a(y, t), F_b(y, t') \rangle_{y|fix \, \hat{x}}$. Again, this restriction is not necessary and is taken here for convenience.

Thus, we see that if we discretely change $\hat{x}$, two separate time-scales (both fast, $\sim O[1]$) control the relaxation of $\delta T(\hat{x})$: τ_F, on which $p_{ss}(y|\hat{x})$ globally relaxes to its new form (i.e., relaxation time of one-point functions), and τ_{F2}, on which the two-point function $\langle F_a(y, t), F_b(y, t') \rangle_{y|fix \, \hat{x}}$ decays. If $\tau_F \ll \tau_{F2}$, then we have the usual result that the white-noise limit of multiplicative colored noise should be interpreted as Stratonovich (see Ch.6.5 in [26]). On the other hand, for $\tau_F \gg \tau_{F2}$ we see that $p_{ss}(y|\hat{x})$ remains essentially fixed while the noise correlations decay, and so the noise amplitude must be evaluated according to the value of x at the beginning of the τ_{F2} interval, i.e., in non-anticipating Itô convention. Note that while both limits are possible, the latter is typical, especially for many-body systems, since the relaxation of $p_{ss}(y|\hat{x})$ proceed via relaxations of two-point functions throughout the system. Finally, note that the same Itô / Stratonovich ambiguity occurs in the expression for the damping correction 2.9 and is resolved in exactly the same way as here.

2.B.3 Compiling results

Finally, we are in a position to put everything together. We use our final expressions for damping 2.9 ($\delta\gamma_{ab}(x) \equiv \int dt' \, (t - t') \left\langle i \, \tilde{y}_i \, \partial_b f_i \big|_{t'} \, , \, F_a \big|_t \right\rangle_{y|fix \, x(t)}$), and noise correction 2.10 ($\delta T(\hat{x}) \equiv \frac{\epsilon}{2} \int dt' \, \langle F_a(y, t), F_b(y, t') \rangle_{y|fix \, \hat{x}}$) in the cumulant expansion 2.7, and plug that into the full expression for the probability distribution over slow

paths 2.2 to get:

$$P[x(t)] \approx \frac{1}{Z_x} \int \mathcal{D}\tilde{x} \exp\left[-\int dt \left\{ i\,\tilde{x}_a \dot{x}_a + \epsilon\,T\,\tilde{x}_a^2 - i\,\epsilon\,\tilde{x}_a \left(\langle F_a \rangle_y - \delta\gamma_{ab} \star \dot{x}_b \right) + \tilde{x}_a\,\delta T_{ab}\,\tilde{x}_b \right\} \right]$$

$$= \frac{1}{Z_x} \int \mathcal{D}\tilde{x} \exp\left[-\int dt \left\{ i\,\tilde{x}_a \left(\gamma_{ab}(x,t) \star \dot{x}_b - \epsilon\,\langle F_a \rangle_y \right) + \epsilon\,\tilde{x}_a \star D_{ab}(x,t) \star \tilde{x}_b \right\} \right].$$

where
$$\gamma_{ab}(x,t) \equiv \delta_{a,b} + \epsilon \int dt'\,(t-t') \left\langle i\,\tilde{y}_i\,\partial_b f_i \big|_{t'}, F_a \big|_t \right\rangle_{y|fix\,x}$$

$$D_{ab}(x,t) \equiv T\,\delta_{a,b} + \frac{\epsilon}{2} \int dt' \left\langle F_a \big|_{t'}, F_b \big|_t \right\rangle_{y|fix\,x}$$

to $O\left[\epsilon^3\right]$. The resulting path integral can then be used to extract the corresponding Langevin equation for the slow dynamics:

$$\gamma_{ab}(x) \star \dot{x}_b = \epsilon\,\langle F_a \rangle_{y|fix\,x} + \sqrt{2\,\epsilon D(x)}_{ab} \star \xi_b \qquad (It\hat{o})$$

$$\star \equiv \begin{cases} \cdot & [It\hat{o}] & \tau_{F2} \ll \tau_F \\ \circ & [Strat] & \tau_{F2} \gg \tau_F \end{cases}$$

where ξ_a is the usual white noise: $\langle \xi_a(t), \xi_b(t') \rangle = \delta_{ab}\,\delta(t-t')$, and the square root of the matrix D_{ab} is defined by the condition $\sqrt{D} \cdot \sqrt{D}^T = D$. Finally the notation $\star$ is used to denote the Itô or Stratonovich dot according to the conditions described in the last section: τ_F and τ_{F2} are the decay time-scales for the one- and two-point functions of the fast dynamics respectively. This is then the main analytical result of our work, shown in eq. 2.3 for the more common Itô case.

2.B.4   Equilibrium: sanity check

Now that we have the effective slow dynamics for general stochastic systems with time-scale separation, we want to check that in the equilibrium case, we recover the expected fluctuation-dissipation relation: $D_{ab}(x) = T\,\gamma_{ab}(x)$. Equilibrium in our original system will corresponds to lack of any driving forces: thus all the forces must come from gradients of a single potential landscape $U(x_a, y_i)$: $F_a = -\partial_a U$ and $f_i = -\partial_i U$. Focusing on the expression for γ_{ab} above we note that in this case

$$\partial_b f_i = -\partial_b \partial_i U = \partial_i F_b$$

$$\gamma_{ab}(x,t) = \delta_{a,b} + \epsilon \int dt' \, (t - t') \left\langle i\, \tilde{y}_i \, \partial_i F_b \big|_{t'}, F_a \big|_t \right\rangle_{y | fix \; x}$$

We then note that the response field for the force F_b is given by $\tilde{F}_b = \tilde{y}_i \, \partial_i F_b$ when x is fixed, as it is here. Finally in MSR we know that $\left\langle i\, \tilde{F}_b, F_a \right\rangle$ gives the linear response function for F, and so by fluctuation-dissipation theorem $\left\langle i\, \tilde{F}_b \big|_{t'}, F_a \big|_t \right\rangle = \partial_{t'} \left\langle F_b \big|_{t'}, F_a \big|_t \right\rangle / T$ for $t' < t$ (zero otherwise). Using this in the above expression and integrating by parts:

$$\gamma_{ab}(x,t) = \delta_{a,b} + \frac{\epsilon}{2T} \int dt' \left\langle F_b \big|_{t'}, F_a \big|_t \right\rangle_{y | fix \; x}$$

$$= D_{ab}/T$$

(The factor of two dividing the integral comes from the fact that while the correlator is time-symmetric, the response function is causal.) We thus recover the desired result.

2.B.5   Fast dynamics and $T_{eff}(x)$

The last question we must address is why does the effective temperature $T_{eff}(x)$ found above, in general correlate with how chaotic the fast variables are? In the case where our fast dynamics undergo a phase transition, we clearly see that under integrable dynamics (zero Lyapunov exponents), the connected correlator $\langle F_a(y,t) F_b(y,t') \rangle - \langle F_a(y,t) \rangle \langle F_b(y,t') \rangle = 0$ vanishes (or is proportional to the small thermal bath temperature). By the same token, in the chaotic phase (Lyapunov exponents comparable to inverse characteristic time), the averages $\langle F_a(y,t) \rangle$ are insensitive to the amplitude of the chaotic fluctuations (by symmetry), and thus we get a high T_{eff}. This is the case in the toy model we studied – as illustrated in fig.3-5.

The issue is more subtle, however, when we are not explicitly considering a phase transition in the fast behavior. For example, consider a system that can have chaotic behavior, as well as regular self-oscillations, but with slow random phase-drift. This way, the steady-state probability in both cases is distributed throughout the accessible

configuration space, with the only distinguishing feature being the correlation decay time τ_{F2} – scaling inversely with the Lyapunov exponents λ_{Lyap}. It turns on that in this case also, $T_{eff} \propto 1/\tau_{F2} \propto \lambda_{Lyap}$ is higher for more chaotic systems – as long as $\tau_{F2} > \tau_{char}$ – characteristic return time of fast dynamics.

We can motivate this claim by first realizing that if the fast steady-state is confined to a finite configuration-space region, then it must have some cyclicity with a finite characteristic return-time τ_{char}. This means that the correlator $\langle F_a(y,t), F_b(y,t') \rangle$, besides exponentially decaying, will also fluctuate (not necessarily periodically), with persistence time $\leq \tau_{char}$ (depending on the details of fast-slow coupling). Thus, we can Fourier transform the force correlator as: $\langle F_a(y,t), F_b(y,t') \rangle \sim e^{-t/\tau_{F2}} \int_{\omega_0}^{\infty} d\omega\, f(\omega) \cos(t\,\omega)$, where the infra-red cutoff $\omega_0 = 2\pi/\tau_{char}$ is given by the fact that fast dynamics have no time-scales longer than τ_{char}. Integrating this correlator to recover the effective temperature, we see that $T_{eff}(x) \propto \int_{\omega_0}^{\infty} d\omega\, f(\omega)\,\omega^2/\tau_{F2}$ as long as $\tau_{F2} > \tau_{char}$, as stated above. Of course, all this assumes that the amplitude of the force fluctuations stays roughly the same as their correlation time changes – but the systems we are interested in are those that exhibit a qualitative change in their Lyapunov exponents, thus making this the dominant effect.

Chapter 3

Toy model: kicked rotor on a cart

To illustrate the results of chapter 2, we consider a toy model that is designed to be a simplest possible example capturing all the qualitative features we might expect of more general scale-separated driven systems of interest. Specifically, we take a kicked rotor on a cart setup shown in fig.3-1a. The fast kicked rotor (Chirikov standard map) dynamics here are chosen as the simplest system that can realize both the chaotic and integrable behaviors under different parameter regimes. Essentially, the system is a rigid pendulum that experiences no external forces except for periodic kicks of a uniform force field (as though gravity gets turned on in brief bursts), and is given by the first two lines in eq.3.1. We modelled the system to be immersed in a thermal bath by adding a small damping and noise (see third line in eq.3.1), whose effects have been studied in [23, 37], and presented in further detail in Appendix 1.A. The point relevant for the following analysis is that when the driving force amplitude (henceforth called "kicking strength") is large, the rotor dynamics are fully chaotic, but if the kicking strength drops below a critical value ($K \lesssim 5$), periodic orbits appear in the configuration space, and are made globally attractive in the presence of damping, thus quickly making the dynamics integrable (we refer to this phenomenon below as "dynamical regularization"). Thus, by controlling the effective drive strength, it is possible to switch between chaotic and regular regimes of fast dynamics.

We then fasten the pivot of the fast kicked rotor on a slow cart that can slide back and forth in a highly viscous medium, perpendicular to the direction of the kick

accelerations (i.e., along the symmetry axis of the rotor dynamics – see fig. 3-1a). The cart is pulled by the tension in the rod, which depends on the fast dynamics, while the global cart position x can feed back on the fast dynamics by having a kicking field that varies along the cart's track $K(x)$. This way we have slow variables conservatively coupled to driven fast dynamics, and a feedback loop controlled through the arbitrary form of $K(x)$ – providing a flexible testing ground. Overall, we argue that, while vastly simplified, this model captures essential physical features of many multi-particle nonequilibrium systems of potential interest.

3.1   Model Setup

The toy model explored here is presented in fig. 3-1a: the kicked rotor is attached to a massless cart moving on a highly-viscous track, which ensures that cart's velocity is much smaller than the rotor's. The exact equations of motion for the system can be derived from a force-balance, and in their dimensionless form become:

$$c\,\dot{x} = -\partial_x U(x) + \sqrt{2\,T\,c}\,\xi_x + \underbrace{\sin\theta\,\left(v^2 - \ddot{x}\,\sin\theta\right)}_{\equiv F_x}$$

$$\dot{\theta} = v$$

$$\dot{v} = -K(x)\,\sin\theta\,\delta(t-n) - b\,v + \sqrt{2\,T\,b}\,\xi_v - \ddot{x}\,\cos\theta \qquad (3.1)$$

where all lengths are measured in units of rotor arm length, time in units of kicking period, and the angle θ is 2π-periodic. Note that for practical reasons (see Appendix 3.A), we also assumed that the cart is momentarily pinned down during each kick, so as to remove the term $\frac{1}{2}K(x)\sin 2\theta\,\delta(t-n)$ that should otherwise be included in F_x due to the direct coupling of the kicks to the cart. For now, we can motivate this by saying that the interesting problem is where the driving force affects the slow dynamics only by means of the fast ones, and not directly, while this chosen implementation can simply be viewed as an additional component of the drive protocol. Additionally, to provide more modelling freedom, we can include an arbitrary potential $U(x)$ acting
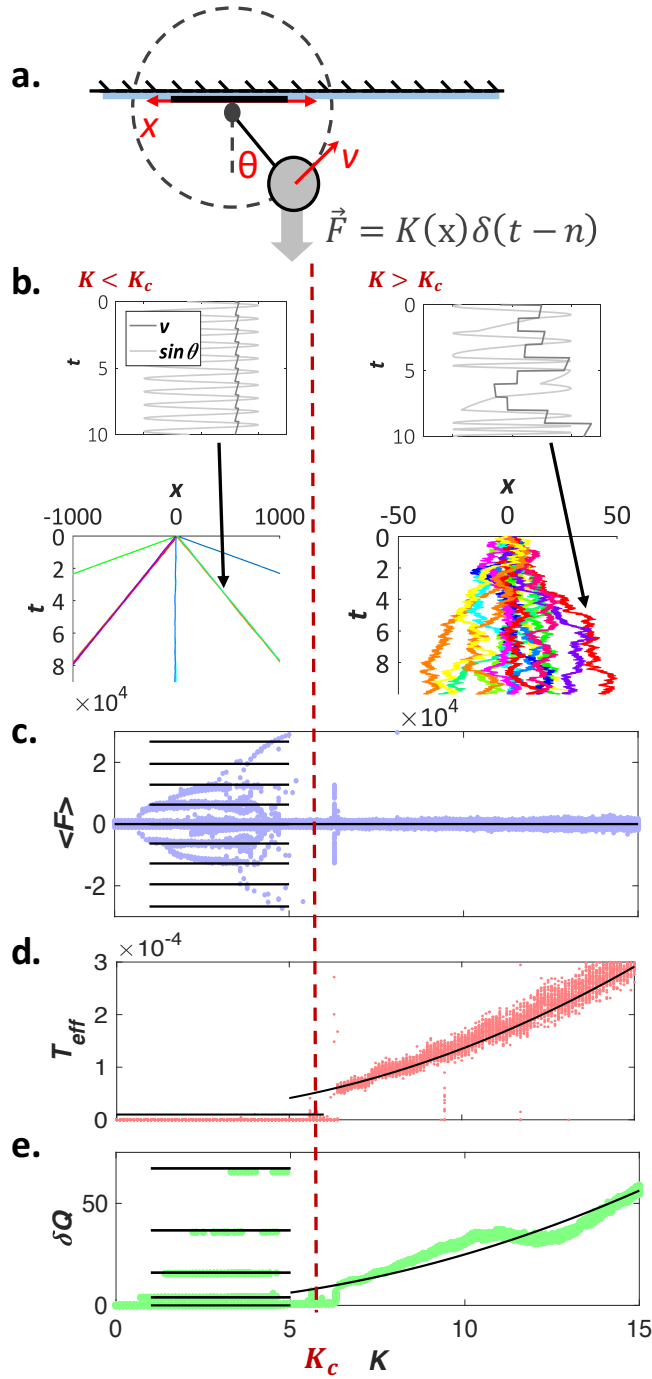
Figure 3-1: a. schematic of the kicked-rotor-on-a-cart toy model. b. numerical realizations of typical cart trajectories over time for driving force below ($K = 3$: regular / ordered regime) and above ($K = 8$: chaotic regime) the ordering transition (at $K_c \sim 5$), along with samples of the corresponding fast (θ, v) dynamics. c,d,e. average force, fluctuations, and dissipation rates in the cart dynamics, measured from trajectories as in panel (b), for the various values of K, along with analytical predictions (in black) from eq.3.3

directly on the cart to produce a conservative force. Time-scale separation in this model implies that back reaction from cart dynamics on the rotor is small – i.e., here that $\ddot{x} \ll K$ (by differentiating the last line, we see that indeed $\ddot{x} \sim O\left[v^3/c\right] \ll$ 1 for $c \gg 1$). Thus the leading-order feedback from the slow variables onto fast dynamics comes from x-dependence of K, which we have full control over, making for a convenient toy-model. We also independently assume that $b \ll 1$ so that fast dynamics are close to the ideal kicked rotor and retain its features.

3.2 Analytical Evaluation

For large K (above the dynamical regularization threshold, i.e. $K \gtrsim 5$), the steady-state of the fast dynamics is fully chaotic, and thus thermal – i.e., we assume thermalization of the entirety of drive energy among the fast fluctuations, as happens in [66] for example. This way, the steady-state distribution is Boltzmann, which is here uniform over θ and Gaussian over v, whose variance we can call T_R (rotor temperature). The symmetry of this state over θ and v gives $\langle F_x \rangle_{s.s.} = 0$, making the fluctuations dominant. The only remaining parameter we need to find is then T_R, which is fully constrained by energy balance as follows. In general, to keep an ergodic system at an effective temperature that is higher than that of its bath requires dissipation [31]:

$$\delta Q = \int dt\, v \circ \left(b\, v - \sqrt{2\, T_0\, b}\, \xi_v\right) = b\left(\langle v^2 \rangle - T_0\right) \delta t$$

$$\mathcal{P} \equiv \frac{\partial Q}{\partial t} = b\left(T_{eff} - T_0\right). \tag{3.2}$$

(for 1D systems with mass=1). Moreover, we can find the power exerted by the kicking force to be $\mathcal{P} = K^2/4$ in the chaotic regime, which in the steady-state must balance the dissipated power. This allows us to extract the effective rotor temperature: $T_R \sim T_0 + \frac{K^2}{4b} + O\left[\frac{1}{c}\right]$ (see sec.3.A.1 for details).

This, however, only gives us information about the fast behavior, while the x-noise correction that we want will also depend on the nature of the rotor-cart coupling. This way, we need to evaluate the x-force correlations and $\delta T_x = \frac{1}{2c} \int dt\, \langle F_x(t), F_x(s) \rangle$,

where as above, $F_x = v^2 \sin\theta - \ddot{x}\sin^2\theta = $ (centripetal F_c) $-$ (inertia F_i) is the force on the cart. The calculation is relatively straightforward and detailed in sec.3.A.2, where we also show that γ damping-coefficient correction is 0 by symmetry of the (θ, v) distribution. We find that, while the inertial term can be ignored at leading order, the correlations of the centripetal force give us $\delta T_x = \frac{1}{2c}\int dt\, \langle F_c, F_c\rangle = K^2/16c$. Note here that this multiplicative noise correction should be interpreted in the Itô sense, as the $\langle F_c, F_c\rangle$ correlations decay on a time-scale faster than kick-period (see Appendix,2.B.2).

For $K \lesssim 5$, on the other hand, the rotor moves periodically in one of the regular attractors. This means that the cart experiences no additional stochasticity other than that from the thermal bath, giving a low $T_{eff} = T_0$, but as some of these attractors spontaneously break the left-right symmetry of the problem, we get $\langle F_x\rangle_{s.s.} \neq 0$. As the motion in most of these attractors is very simple $- n \in \mathbb{Z}$ full revolutions of the rotor per kick $-$ we can estimate this force explicitly: $\langle F_x\rangle_{ss} = \int_0^1 dt\, v^2(t)\sin\theta(t) \sim b\,v_n + v_n^3/2c$, where $v_n \equiv 2\pi n$, and $\theta(t)$ and $v(t)$ were estimated by integrating the equations of motion 3.1 at leading order (see sec.3.A.3).

Compiling the resulting predictions for the cart motion, we get:

$$c\,\dot{x} = -\,\partial_x U(x) + \langle F\rangle + \sqrt{2\,c\,T_{eff}}\cdot\xi \tag{3.3}$$

$$\langle F\rangle = \begin{cases} v_n\,(b + v_n^2/2c) & K \lesssim 5 \\ 0 & K \gtrsim 5 \end{cases}$$

$$T_{eff} = \begin{cases} T_0 \\ T_0 + K^2/16c \end{cases}$$

with $v_n \equiv 2\pi n$ and n some random integer, typically smaller than $O\left[\sqrt{T_R}\right]$ (since the rotor first explores its phase-space thermally before finding one of the regular attractors). Another quantity we can easily estimate for the two phases is the energy

55

dissipation rate:

$$\dot{Q} = \begin{cases} v_n^2 \left(b + v_n^2/2c\right) & K \lesssim 5 \\ K^2/4 & K \gtrsim 5 \end{cases} \tag{3.4}$$

Numerical simulations confirm these predictions in fig.3-1 c, d, and e respectively.

3.3 Numerical Tests

To verify the above analytical results, we can run numerical simulations of the full system dynamics in eq. 3.1. To begin, we check the cart dynamics for different values of (x-independent) K (and $U(x) = 0$). Fig. 3-1b shows typical cart trajectories for K in the regular and chaotic regimes. More systematically, plotting the apparent average drift $\langle F \rangle$ and fluctuations T_{eff} for multiple realizations at each K, we get plots in c and d of fig. 3-1 respectively. We thus see quantitative agreement between the prominent features of these plots and the results of eq. 3.3 – shown here as black lines. Finally fig. 3-1e shows the heat dissipation rate in the different possible steady-states, showing that while lowering T_{eff} corresponds to decreased dissipation within the chaotic phase, this rule is violated if we enter a regular dynamic attractor.

Note that as the original problem is stated exactly, and our method allows for full analytical treatment of the slow variables, there are no fitting parameters in any of the curves we are comparing against throughout the numerical study. We use $c = 5 \times 10^4$, $b = 0.1$ for all simulations, and to emphasize the effects from fluctuations of the fast variables, we take the actual thermal bath to be at a vanishingly low temperature $T_0 \sim 0$, unless otherwise stated.

While fig. 3-1 shows agreement of one- and two-point functions of cart position with our analytical prediction, we have yet to check that the fast dynamics can really be approximated by an effective thermal bath. One convincing way to do this is to introduce a non-trivial potential landscape $U(x)$ acting on the cart's position x, and check the resulting steady-state distribution $p(x)$ against Boltzmann statistics at the
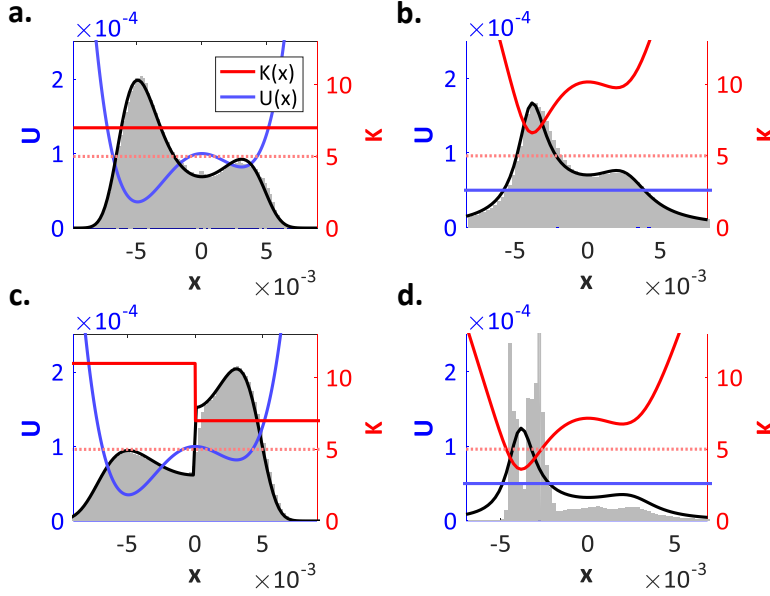
Figure 3-2: Histograms (grey) of the steady-state cart positions in simulation with:
a. constant $K(x) = K$ and in a potential landscape $U(x)$ plotted in blue (black
curve gives the expected Boltzmann distribution); b. $U(x) = const$ and $K(x)$
landscape in red (black curve shows $1/T_{eff}(x)$ – solution of the Fokker-Planck
equation) c. $U(x)$ as in a. and $K(x)$ step-function (analytical prediction plotted in
black is given by eq.3.5). d. $U(x) = const$ and $K(x)$ as in b., but shifted down to
dip below the critical $K_c \sim 5$ value – dotted red line (black curve again shows
$1/T_{eff}(x)$ outside of ordered region) – this shows that probability gets localized at
the two transition points at long times.

predicted temperature T_{eff}. Figure 3-2a shows the agreement between the histogram

produced by this simulation and the curve for the expected Boltzmann distribution.

To see that the $T_{eff}(x)$ landscape remains the appropriate description even for

non-uniform $K(x)$, we can calculate the steady-state distribution in a $K(x)$ landscape,

now letting $U(x) = 0$. The expected distribution for free diffusion in a temperature

landscape can easily be found using, e.g., Fokker-Planck equation, and gives $p(x) \propto$

$1/T(x)$ (note that this arises precisely because our effective slow dynamics have Itô

multiplicative noise – for Stratonovich it would be $1/\sqrt{T(x)}$). This is well confirmed

by simulations in fig. 3-2b, thus showing that at least in the steady-state, probability

density does indeed collect in low-temperature regions.

The last natural test that we mention here is to see how $T_{eff}(x)$ landscape can

counteract the forces of $U(x)$ – specifically changing the stability in a double-well

potential. This setup is shown in 3-2c, where the higher-energy potential well is stabilized by having a lower T_{eff}. The numerical result is correctly predicted by the steady-state solution of the Fokker-Planck equation with the expected effective temperatures in each well (labels L and R denote left and right wells respectively), as shown in fig. 3-2c:

$$p(x) = \frac{1}{Z} \begin{cases} \frac{1}{T_L} e^{-U(x)/T_L} & x < 0 \quad (L) \\ \frac{1}{T_R} e^{-U(x)/T_R + \Delta} & x > 0 \quad (R) \end{cases} \tag{3.5}$$

$$\Delta \equiv U(0) \left(\frac{1}{T_R} - \frac{1}{T_L} \right)$$

In the limit of a discrete jump process between the two wells (wells with equal internal entropy separated by a high barrier), this exact solution becomes well approximated by that obtained from current-matching with the expected jump rates: $r_{\rightarrow} = e^{-(U(0)-U_L)/T_L}$ and $r_{\leftarrow} = e^{-(U(0)-U_R)/T_R}$. The key non-equilibrium feature in these solutions is the dependence of the probabilities in either well on the barrier height $U(0)$ via Δ – for higher barriers the temperature difference becomes more important. This example gives the first non-trivial application of thermodynamic intuition from $T_{eff}(x)$ landscape to solution of our non-equilibrium system. Projecting this concept onto a broader context, we note that this setup is a particular realization in the class of problems of iterative annealing (e.g., used in chaperoned protein folding [68], etc).

3.4 Least rattling

Having confirmed the steady-state and thermal properties of the slow behaviors, we next want to look at the predictive power of our formalism for transient behaviors and currents, again in the presence of inhomogeneous fast dynamics. The first example we consider is transient cart motion in linearly varying $K(x) = \kappa\, x$. The simulation results are shown in fig. 3-3. As mentioned above and in Appendix 1.B, free diffusion in a temperature gradient results in a median drift to low T, as observed here.
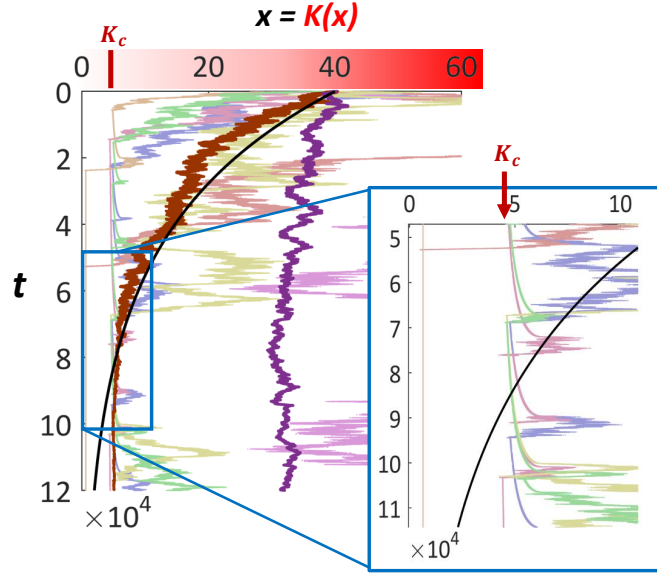
Figure 3-3: Typical cart trajectories in linear $K(x)$ landscape ($U(x) = const$) all starting from one point, along with their mean (purple) and median (brown). Black curve shows the analytical prediction for the median, while mean is expected to be constant at small times. Inset shows the regularization transition at $K_c \sim 5$, where effective temperature drops abruptly to 0, and median departs from the smooth decay. x-axis is labelled in units of K

Explicitly, the slow dynamics in this case $c\dot{x} = \frac{\kappa}{2\sqrt{2}} x \cdot \xi$ can be solved exactly to give $\ln x(t) = -\left(\frac{\kappa}{4c}\right)^2 t + \frac{\kappa}{2\sqrt{2}} \mathcal{N}(0, t)$ (with $\mathcal{N}$ giving the normal distribution with variance t), from which we can read off the mean $x(t) = x_0$ and median $x = x_0 \exp\left[-t\left(\kappa/4c\right)^2\right]$ behaviors. The latter is plotted in black in fig.3-3 and well reproduces the simulation result in brown. Note that for any finite ensemble of trajectories, or for a bounded system, the mean will eventually go to low temperatures as well, but not as cleanly or predictably – so the constant mean value is not practically realizable at long times. The inset focuses on the crossover into regular dynamics, where we see that the symmetry-broken drift-force $\langle F_x \rangle$ can either take the cart to the $K = 0$ absorbing state (as detailed in the further inset), or back out into the chaotic regime. In the latter case, the cart typically diffuses back down to the transition again. The resulting oscillations cause a (transient) accumulation of probability around the critical point, giving a peculiar realization of self-organized criticality. This critical region itself is also interesting as the correlations of the fast variables persist for long times,

and can thus break the time-scale separation assumption – but this will have little effect on the global system behavior. The overall takeaway here is the emergent "least rattling:" slow dynamics drift towards regions where fast ones are less stochastic.

To further illustrate the importance of the regularization transition on the slow dynamics, we consider the probability distribution $p(x)$ in the presence of $K(x)$ landscape (and no potential $U = 0$), as in fig. 3-2b., but shifted down such that it dips slightly below the regularization transition at its lowest point – fig. 3-2d. The resulting small zero-temperature region in x, corresponding to integrable fast dynamics, becomes absorbing, collecting most of the probability density over time (see fig. 3-2d). Note again that probability accumulates at the critical transition points, giving the two-pronged shape. We stress here the observed sharp localization transition of the steady-state distribution as soon as some regular regime of the fast dynamics becomes accessible – i.e., the slow variables find the regularized region even if it requires some fine-tuning. (This trade-off between least-rattling and entropic forces can be made quantitative.)

3.5   Anomalous diffusion

The last example we present shows that besides an effective temperature landscape, the regular dynamical phase accessible to this model can give rise to apparent anomalous diffusion. To begin, Fig. 3-4a shows an implementation of Buttiker-Landauer ratchet using our model: periodic $U(x)$ and $K(x)$ landscapes, with a relative phase-offset of $\pi/2$ create a steady-state current being pumped, in this case to the right. Intuitively, this happens because a higher effective temperature in the right half of the potential well makes it easier for the cart to overcome the right potential barrier than the left one. The interesting behavior appears when we shift the $K(x)$ wave downward to straddle the transition point at $K \sim 5$ (fig. 3-4b). In this case the pumped current reverses direction and becomes an order of magnitude *larger* – even if we reduce the amplitude of the $K(x)$ variation. To understand this, it helps to look at some typical realizations of barrier-crossing trajectories at the bottom of fig.
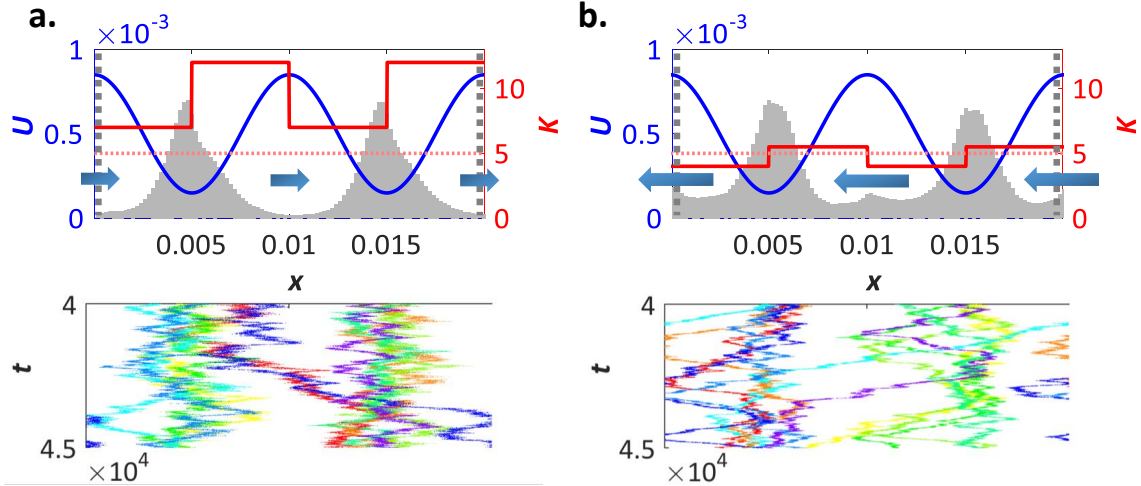
Figure 3-4: Simulated steady-state cart-position distributions for the shown $U(x)$ (blue) and $K(x)$ (red) landscapes (x is periodic). These result in a pumped steady-state current (block arrows), with the typical barrier-crossing trajectories shown at the bottom. Unlike in all the above simulations, thermal bath temperature $T_0 = 10^{-4} > 0$ was taken in these to smooth out the critical behavior. Straddling the critical point with $K(x)$ in panel b produces a ten-fold larger (and reversed) current, even for smaller absolute variation in $K(x)$

3-4. While in panel a. transitions are achieved by stochastic fluctuations that are exponentially suppressed by the Boltzmann factor, in panel b., these are achieved by a directed symmetry-broken drift force $\langle F \rangle > \partial_x U$, and thus the crossing probability is just the probability of the fast dynamics finding the appropriate regular attractors. These ballistic-like trajectories of the cart in the regular regime can be usefully thought of as anomalous super-diffusion with exponent $\alpha = 2$. Also, in so far as the barrier crossing becomes easier as we lower $K(x)$ through the critical value, we can say that the diffusion becomes stronger, thus showing non-monotonicity with K – reminiscent of the findings in [64].

We have thus confirmed by direct numerical simulations of this toy model that the framework derived in chapter 2 can give correct quantitative predictions in a context where no other analytical tools are available. We could analytically derive the effective temperature landscape $T_{eff}(x)$ for the slow variables based on energy-flow considerations and a qualitative knowledge of the character of fast dynamics. In the following chapters we will consider more complicated systems where there is no

way to analytically predict T_{eff} landscape, but it can be measured from studying only local short-time trajectories, and used to make global predictions about the system.

Appendix

3.A   Details of the analytical calculation

In this appendix we present the analytical calculations required to make the predictions for the Kicked Rotor on a cart toy model described in the main text. For convenience, we reproduce the dimensionless equations of motion here (this time including the direct effect of kicks on the cart in F_x):

$$\dot{\theta} = \underbrace{v}_{\equiv f_\theta}$$

$$\dot{v} = \underbrace{-K(x)\,\sin\theta\,\delta(t-n) - \ddot{x}\,\cos\theta - b\,v}_{\equiv f_v} + \sqrt{2\,T\,b}\,\xi_v$$

$$c\,\dot{x} = \underbrace{\sin\theta\,\left(v^2 + K(x)\,\cos\theta\,\delta(t-n) - \ddot{x}\,\sin\theta\right)}_{\equiv F_x}$$

$$- \partial_x U(x) + \sqrt{2\,T\,c}\,\xi_x \tag{3.6}$$

The units were chosen such that rotor arm length, mass, and kicking period are all $=1$. The part in black gives simple kicked rotor dynamics, red part weakly (for $b \ll 1$) couples it to a thermal bath at temperature T, and blue part gives the coupling and dynamics of the cart. We assume throughout that the bath temperature here is very low $T \sim 0$ to highlight the effect of the chaotic fluctuations of kicked rotor. Strong time-scale separation, which here is achieved by assuming $c \gg 1$, implies that terms $\propto \ddot{x}$ will be small. To see when precisely we can be justified in dropping these, we estimate their magnitude for the two regimes: the forced regime ("during the kick") and the free rotation. For the forced regime, since $\delta(t)$ is a distribution, we can only

talk about the integrals

$$\delta v_K = \lim_{\eta \to 0} \int_{-\eta}^{+\eta} dt \, \dot{v} = -K \, \sin \theta - \cos \theta \, [\dot{x}]_{-\eta}^{+\eta}$$

$$= - K \, \sin \theta - \frac{K}{c} \sin^2 \theta \, \cos \theta \, [K \sin \theta - 2 \, v]$$

$$\delta x_K = \lim_{\eta \to 0} \int_{-\eta}^{+\eta} dt \, \dot{x}$$

$$= K \, \sin \theta \, \cos \theta - \frac{K}{c} \sin^4 \theta \, [K \sin \theta - 2 \, v]$$

For the free rotation, we can simply differentiate the unforced part of the last line in eq. 3.6 with respect to time, which gives $c \, \ddot{x} = v^3 \cos \theta$ to leading order. Thus, to ignore the $\ddot{x}$ terms, we need $K \, v/c \ll K$ for the driven regime, and $v^3/c \ll b \, v$ for free rotation. While the former condition is easy to satisfy for a large c, the latter one competes with our additional assumption that $b \ll 1$ and can be difficult to satisfy numerically, especially as velocities v can sometimes get very large – thus we will keep the $\ddot{x} \sin \theta$ term as an additional perturbative correction to the free dynamics.

3.A.1   Chaotic Kicked Rotor steady-state

To proceed in evaluating the different terms in the expression for the effective slow dynamics (eq.2.3 in main text), we need to find the steady-state distribution over (θ, v) for a fixed cart position x. As mentioned in the main text, for strong driving $K \gtrsim 5$, the kicked rotor dynamics are fully chaotic, and thus the steady-state thermalizes all input energy. This immediately implies that the probability distribution is of form $p_{ss}(\theta, v \,|x) \propto \exp \left[-\frac{v^2}{2 \, T_R(x)} \right]$ – uniform over θ and Gaussian over v, parametrised by a single number $T_R(x)$ – effective rotor temperature. The symmetries of this distribution guarantee that $\langle F_x \rangle_{ss} = 0$.

To find this temperature, we can use the argument from eq. 3.2 of the main text,

which tells us that this steady-state will have a dissipation rate:

$$\delta Q = \int dt\, v \circ \left(b\, v - \sqrt{2\, T_0\, b}\; \xi \right) = b\left(\langle v^2 \rangle - \sqrt{2\, T_0\, b}\; \left\langle \frac{v_{t+\delta t} + v_t}{2}\, \xi_t \right\rangle \right) \delta t$$

where for underdamped, forced Langevin dynamics, we have in general:

$v_{t+\delta t} = v_t + \delta t\left(-b\, v_t + F(x, v, t) + \sqrt{2\, T_0\, b}\; \xi \right)$. Thus, while v_t is completely un-correlated with ξ_t: $\langle v_t\, \xi_t \rangle = 0$, $v_{t+\delta t}$ is correlated only via the thermal noise term: $\langle v_{t+\delta t}\, \xi_t \rangle = \sqrt{2\, T_0\, b}$, and is independent of any interaction or driving forces $F(x, v, t)$. This gives, for mass=1:

$$\mathcal{P} \equiv \frac{\partial Q}{\partial t} = b\, (T_{eff} - T_0). \tag{3.7}$$

(note that if v were a vector in d-dimensions, we would multiply this expression by d). With this, and neglecting the bath temperature T_0, we can balance the work flow in and heat flow out per kick, to get:

$$
\begin{aligned}
0 =& \delta W - \delta Q \\
=& \left\langle \lim_{\eta \to 0} \int_{-\eta}^{+\eta} dt\, v \circ (-K\, \sin\theta) \right\rangle_{ss} - b\, \langle v^2 \rangle_{ss} \\
=& \left\langle \left(v_{pre} - \frac{K}{2}\, \sin\theta \right) (-K\, \sin\theta) \right\rangle_{ss} - b\, T_R \\
\Rightarrow\quad & T_R = \frac{K^2}{4\, b}
\end{aligned}
$$

where v_{pre} is the pre-kick velocity, which is uncorrelated with θ. Since this gives the variance of typical rotor velocities, we can use it to simplify the time-scale separation condition derived above $v^3/c \ll b\, v$ down to $K/c \ll b^2$ – which is quite difficult to satisfy in addition to $b \ll 1$. Thus, while this result is correct, it turns out that $O\,[1/c]$ correction coming from the cart coupling term $\ddot{x}\, \cos\theta$ is very important here. Note that eq. 2.3 in main text tells us to include x-dependence of the steady-state in correcting the effective damping coefficient γ, however as this term depends only

on $\ddot{x}$ and not x itself, the steady-state does not gain any x-dependence from it, but rather an x-uniform correction which must be included directly. In this case, as the dynamics are still chaotic and distribution thermal, while the work extracted from the drive δW is not affected (since any work done on the overdamped cart is immediately dissipated and so can be ignored), so the coupling to cart simply adds another channel for heat dissipation:

$$
\begin{aligned}
\delta Q =& b \left\langle v^2 \right\rangle_{ss} + \left\langle \int_0^1 dt \, c \, \dot{x}^2 \right\rangle_{ss} \\
=& b \, T_R + \frac{1}{c} \int_0^1 dt \, \left\langle F_x \big|_t F_x \big|_t \right\rangle_{ss} \\
=& b \, T_R + \frac{1}{c} \Bigg[\int_0^1 dt \, \left\langle \left(v^2 \sin\theta \right)^2 \right\rangle_{ss} \\
& \quad + \left\langle 2 \, K \, v^2 \, \sin^2\theta \, \cos\theta \big|_{t=0} \right\rangle_{ss} \Bigg] \\
=& b \, T_R + \frac{1}{c} \left[\frac{3}{2} T_R^2 + 0 \right] \\
\Rightarrow \quad & T_R = \frac{b \, c}{3} \left(\sqrt{1 + \frac{3 \, K^2}{2 \, b^2 \, c}} - 1 \right) < \frac{K^2}{4 \, b}
\end{aligned}
$$

where we assumed that v and θ are uncorrelated during most of the time $0 < t < 1$, as justified below. This correction significantly lowers the typical velocities and well reproduced in the results of the simulations for large, but practical values of c (e.g., for $b = 10^{-2}$, $c = 10^4$, $K = 10 \Rightarrow T_R \approx 375 < K^2/4b = 2500$, or for $b = 10^{-1}$, $c = 5 \times 10^4$, $K = 10 \Rightarrow T_R \approx 230 < K^2/4b = 250$).

3.A.2 Cart Damping and Noise correction

With the above understanding of the steady-state, we now proceed to compute the two-time correlations functions needed to get the corrections on the slow dynamics given by eq.2.3 of the main text. We begin by noting their general structure here: each kick introduces correlations between θ and v, after which, while v remain approximately constant until the next kick (for $b \ll 1$), θ spins around and correlations decay. The typical decay time-scale in this system can be estimated by looking at the

decay:

$$\langle \sin \theta(0) \, \sin \theta(t) \rangle_{ss} = \langle \sin \theta(0) \, \sin (\theta(0) + v \, t) \rangle_{ss} = \frac{1}{2} - \frac{T_R}{4} t^2 + O\left[t^4\right]$$

where we assume θ and v to be uncorrelated over the time-window. This gives decay $\tau \sim 1/\sqrt{T_R}$, which for typical values of parameters (e.g., for $b = 10^{-1}$, $c = 5 \times 10^4$, $K = 10$) could be around $1/20$. The key here is that in most cases the decay time is much shorter than 1 (the kicking period). Figure 3-5 shows numerical results for $\langle F_x(t), F_x(t') \rangle$ correlation in the chaotic phase to give a sense of how these quantities typically look for the given system. Since the $\theta - v$ correlations are only generated by kicks, this implies that most of the time they are uncorrelated – as we have assumed a few times above. Moreover, this means that our noise and damping corrections should always be interpreted as Itô for this system, as discussed in the Appendix 2.B.2 of the main text.

Using this result we can immediately see that the damping correction vanishes (see eq.3.6 for definitions of f_θ, f_v):

$$\delta\gamma = \frac{1}{c} \int dt' \, (t - t') \left\langle i \left(\tilde{\theta} \, \partial_x f_\theta + \tilde{v} \, \partial_x f_v \right) \big|_{t'}, F_x \big|_t \right\rangle_{(\theta,v)|fix \, x}$$

$$= \frac{1}{c} \int dt' \, (t - t') \left\langle -i \, \tilde{v} \, K'(x) \sin \theta \sum_n \delta(t' - n) \big|_{t'}, F_x \big|_t \right\rangle_{(\theta,v)|fix \, x}$$

$$= \frac{1}{c} \sum_n (t - n) \left\langle -i \, \tilde{v} \, K'(x) \sin \theta \big|_n, F_x \big|_t \right\rangle_{(\theta,v)|fix \, x} = 0$$

as the correlator vanishes by the symmetries of the thermal steady-state.

With that, to find the effective temperature experienced by the cart, we need only compute the noise correction, as: $T_{eff} = \frac{1}{2c} \int dt' \left\langle F_x \big|_{t'}, F_x \big|_t \right\rangle_{(\theta,v)|fix \, x}$. From eq.3.6, we get $F_x = v^2 \sin \theta + K(x) \sin \theta \cos \theta \, \delta(t - n) - \ddot{x} \sin^2 \theta = $ (centripetal F_c) + (direct kick coupling F_k) – (inertia F_i). Unlike in f_v, where the term $b \, v$ was comparable magnitude to $\ddot{x} \cos \theta$, here F_c and F_k are both $> O[1]$, and thus the inertia F_i is distinctly sub-leading and can be dropped. We now proceed to individually compute the $\langle F_c, F_c \rangle$, $\langle F_k, F_k \rangle$, $\langle F_c, F_k \rangle = \langle F_k, F_c \rangle$ contributions.
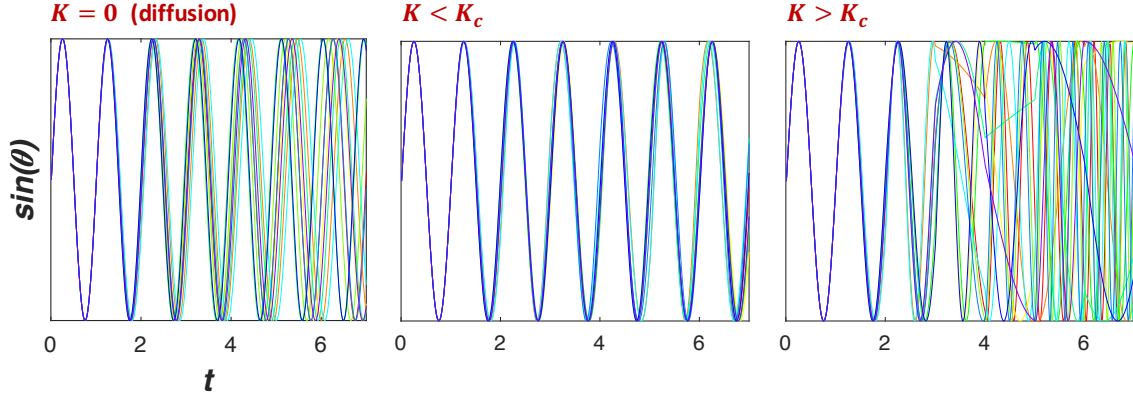
Figure 3-5: Different realizations of kicked-rotor trajectories starting from same initial conditions in the regular and chaotic driven phases, as well as the freely-diffusing undriven case (for reference). While the steady-state in the chaotic phase looks like large random fluctuations about a constant mean (thus giving a high T_{eff}), in the regular phase, the drive serves to synchronize all trajectories, thus yielding small fluctuations about an oscillating mean in the steady-state (and so a low T_{eff}).

For the $\langle F_c, F_c \rangle$ term, we see that far from the kicks, where θ and v are uncorrelated, we get:

$$\int dt' \; \left\langle v^2 \sin\left(\theta + t'\,v\right), \; v^2 \sin\theta \right\rangle_{(\theta,v)|fix\,x} = 0$$

which can be evaluated analytically in Mathematica. The leading correction to this quantity then comes from the $\theta - v$ correlations generated by the kicks. To capture these, we write all velocities and angles in terms of their values before the last kick – at a time when they were guaranteed to be uncorrelated. Thus

$$v(t) = \begin{cases} v_{pre} \\ v_{pre} - K \sin\theta_{pre} \end{cases}$$

$$\text{and} \quad \theta(t) = \begin{cases} \theta_{pre} + t\,v_{pre} & t < 0 \\ \theta_{pre} + t\,(v_{pre} - K \sin\theta_{pre}) & t > 0 \end{cases}$$

Using these expressions, we can thus evaluate $\iint dt\,dt' \left\langle F_c \big|_t , F_c \big|_{t'} \right\rangle$ piecewise (where a second integral must be included since the time-translation-invariance is now broken).

Dropping the subscript *pre*, we have:

$$\int_{-1/2}^{1/2} dt \int_{-\infty}^{\infty} dt' \left\langle v(t)^2 \sin\theta(t), \ v(t')^2 \sin\theta(t') \right\rangle_{(\theta,v)|fix\ x}$$

$$= \begin{cases} \left\langle \int_0^{1/2} dt \int_0^{\infty} dt' \ (v - K\sin\theta)^4 \sin(\theta + t(v - K\sin\theta)) \sin(\theta + t'(v - K\sin\theta)) \right\rangle_{(\theta,v)|fix\ x} \\[2mm] +2 \left\langle \int_{-1/2}^{0} dt \int_0^{\infty} dt' \ v^2 \sin(\theta + t\,v) \ (v - K\sin\theta)^2 \sin(\theta + t'(v - K\sin\theta)) \right\rangle_{(\theta,v)|fix\ x} \\[2mm] + \left\langle \int_{-1/2}^{0} dt \int_{-\infty}^{0} dt' \ v^4 \sin(\theta + t\,v) \sin(\theta + t'\,v) \right\rangle_{(\theta,v)|fix\ x} \end{cases}$$

$$= \begin{cases} K^2/8 + T_R/2 \\[2mm] -T_R \qquad\qquad = K^2/8 \\[2mm] +T_R/2 \end{cases}$$

where all integrals can be evaluated analytically (in Mathematica) if we take the time integrals first, and then average over the (uncorrelated) θ and v. Similarly for the other terms:

$$\int dt' \left\langle F_k \big|_t, F_k \big|_{t'} \right\rangle_{(\theta,v)|fix\ x} =$$

$$= \left\langle \int_{-\infty}^{\infty} dt' \int_{-1/2}^{1/2} dt \sum_{n,m} \frac{K}{2} \sin 2\theta(t)\, \delta(t - n)\, \frac{K}{2} \sin 2\theta(t')\, \delta(t' - m) \right\rangle_{(\theta,v)|fix\ x}$$

$$= \sum_m \frac{K^2}{4} \left\langle \sin 2\theta(0) \ \sin 2\theta(m) \right\rangle_{(\theta,v)|fix\ x} = K^2/8$$

$$\int_{-\infty}^{\infty} dt' \left\langle F_k \big|_t, F_c \big|_{t'} \right\rangle_{(\theta,v)|fix\ x} =$$

$$= \int_0^{\infty} dt' \left\langle \frac{K}{2} \sin 2\theta \ (v - K\sin\theta)^2 \sin(\theta + t'(v - K\sin\theta)) \right\rangle_{(\theta,v)|fix\ x}$$

$$= -K^2/8 = \int_{-\infty}^{\infty} dt' \left\langle F_c \big|_t, F_k \big|_{t'} \right\rangle_{(\theta,v)|fix\ x}$$

Adding up all four terms, we thus get $T_{eff} = \frac{1}{2c} \int dt' \left\langle F_x \big|_{t'}, F_x \big|_t \right\rangle_{(\theta,v)|fix\ x} = 0$. Here we clearly see that the cancellation comes up due to the anti-correlations between the centripetal force and the kick coupling. In this particular case, the cancellation

is somewhat accidental, and is a consequence of the simplicity of the system – the functional form of couplings is quite restricted. In general, we expect such cancellations to be unlikely in higher-dimensional systems. As discussed in the main text, to make the system interesting and get a finite T_{eff}, we can simply eliminate the direct kick-cart coupling F_k from the dynamics altogether, with the physical interpretation of "pinning" down the cart at the instant of the kick. This leaves only F_c, thus giving $T_{eff} = K^2/16c$, as desired.

Note also that the rotor temperature T_R calculated above ends up dropping out and does not affect any of the time-integrated correlators, but only the particulars of their time-dependence as $\langle F_x(t), F_x(t') \rangle_{(\theta, v)|fix\, x}$. Thus the only really key role it played for us was to show that these correlators decay faster than kicking period.

3.A.3 Ordered Kicked Rotor steady-state

On the other hand for weak driving $K \lesssim 5$, the kicked rotor undergoes dynamic regularization, and in steady-state is found in one of the integrable attractors in its phase space. Thus, none of the above arguments apply here. Instead, the main regular regions correspond to the rotor completing n full revolutions per kick, with $n = ..., -2, -1, 0, 1, 2....$ As it does, there are no stochastic fluctuations, other than those from the thermal bath, and as the steady-state lacks the symmetries of the thermal state, $\langle F_x \rangle_{ss} \neq 0$ except at $n = 0$. In fact, depending on the attractor that the rotor falls into, it will exert a persistent force on the cart, causing constant directed drift. We can easily estimate this drift force for the n'th attractor as (here we again assume that $b \sim O\left[1/c\right] \ll 1$, and let $v_n \equiv 2\pi n$):

$$\langle F_x \rangle_{ss} = \int_0^1 dt\, v^2(t)\, \sin\theta(t) = b\, v_n + \frac{v_n^3}{2\,c} + O\left[b^2\right]$$

$$\text{where} \quad v(t) = 2\pi n - b\, v\, t - \frac{1}{c} \int_0^t d\tau\, v^3 \cos^2(v\,\tau) + O\left[b^2\right]$$

$$\text{and} \quad \theta(t) = \int_0^t d\tau\, v(\tau)$$

where $v(t)$ dynamics are found by directly integrating eq.3.6 to first order in small

parameters. However, as it is impossible to predict which of the attractors will be chosen, we can't a-priori tell the direction or speed that the cart will be moving at – though the options are restricted to the above small discrete set of possibilities parametrized by n.

Chapter 4

Defining rattling

While the intuition that things end up in "colder" regions seems simple, it is not immediately clear how we should define a scalar quantity that would be quantitatively predictive of the steady-state probability density for a range of values – not just as a binary "high-low." For systems with time-scale separation, we were able to derive a concrete form for this T_{eff} based on the fast dynamics, but without the scale-separation assumption, it isn't immediately clear how that formulation should be generalized. Consider the general diffusion process:

$$\dot{x}_i = D_{ij}(x) \cdot \xi_j \quad [\text{Itô}] \tag{4.1}$$

where $\xi_i(t)$ - univariate Gaussian white noise: $\langle \xi_i \rangle = 0$ and $\langle \xi_i(t), \xi_j(s) \rangle = \delta_{ij}\,\delta(t-s)$, indices $\{i, j, k, n, m\}$ run over the d dimensions of the configuration space, and temperature tensor is defined as $T(x) \equiv \frac{1}{2}D\,D^T$. By analogy with these dynamics, we could define effective temperature generally as $T_{ij}(x) = \frac{1}{2}\int ds\ \langle \dot{x}_i(t), \dot{x}_j(s) \rangle$ (also motivated by discussion in the last chapters). However, without clear scale-separation, it isn't a-priori clear what the range of integration should be, if any, especially if correlations don't really drop to zero. We may also ask whether we should be taking only the connected part of this correlator, and what's the right way to construct a scalar from this so we can talk about $p_{s.s.}$. Additionally, we want to somehow filter out, at least partly, the periodic or regular part of the motion which does not produce

any probability drop.

4.1 As inference problem

One perspective we can take on addressing these issues is by solving the following concrete problem. Let's assume that for some observed dynamical process $\vec{x}(t)$ (generated by a black box), the underlying motion is indeed Brownian, governed by eq.4.1, and we want to infer the inhomogeneous temperature tensor $T(x) \equiv \frac{1}{2}D\,D^T$ from observations of trajectories. This amounts to a Bayesian inference problem: $P\left[T(x)\,|\,x(t)\right] = P\left[x(t)\,|\,T(x)\right]\,P\left[T(x)\right]/Z$ ($Z = P[x(t)]$ – normalization). For our prior $P\left[T(x)\right]$, we assume that temperature landscape varies smoothly over x – otherwise it would have infinite degrees of freedom and we could never hope to estimate it from a finite observation of trajectory. Explicitly, this can be implemented via: $P\left[T(x)\right] \propto \exp\left[-\epsilon \int d^d x\; \partial_k T_{ij}\partial_k T_{ij}\right]$, among other ways. Additionally, from the distribution for Gaussian white noise, we get $P\left[x(t)\,|\,T(x)\right] = \prod_t \left[\sqrt{\frac{1}{\det(4\pi T)}}\exp\left[-\frac{1}{4}T_{ij}^{-1}\dot{x}_i\dot{x}_j\right]\right]$ this way:

$$P\left[T(x)\,|\,x(t)\right] = \frac{1}{Z[x(t)]}\exp\left[-\int dt\left(\frac{1}{2}\operatorname{tr}\log T(x(t)) + \frac{1}{4}T_{ij}^{-1}(x(t))\;\dot{x}_i(t)\,\dot{x}_j(t)\right)\right.$$
$$\left. -\epsilon\int d^d x\;\partial_k T_{ij}\partial_k T_{ij}\right]$$

Note that, interestingly, we have an integral over time and another integral over configuration space in the exponent. We then want to choose some estimator for $T(x)$. The easiest to evaluate here is the maximum-likelihood, which we can get by setting the variation $\frac{\delta}{\delta T(x)}$ of the exponent to 0, which gives:

$$4\epsilon\; T_{in}T_{jm}\partial_k^2 T_{mn}\big|_X = T_{ij}(X) - \frac{1}{2}\left\langle\dot{x}_i(t)\,\dot{x}_j(t)\right\rangle_{t\,|\,x(t)=X}$$

where the average is over all the times when trajectory $x(t)$ passes through X. So if we had a uniform prior $\epsilon \to 0$, then we get that $T_{ij}(X)$ is simply the equal-time second-moment of velocity that trajectories have at that point. For $\epsilon > 0$, any deviations

from this rule become sources for a diffusions equation, smoothing out the landscape – and thus for small ϵ, we simply get that the average $\langle \cdot \rangle$ above is not only at X, but over its neighborhood, weighted by a Gaussian kernel:

$$T_{ij}(X) = \frac{1}{2} \langle \dot{x}_i(t) \, \dot{x}_j(t) \rangle_{t \,|\, (x(t)-X)^2 < \epsilon} \,. \tag{4.2}$$

Now, we can use this result to take our trajectories $x_i(t)$, which came from some non-diffusive dynamical system, and infer the temperature landscape that might have created such trajectories. This will always produce some result, regardless of how poor of a model diffusive behavior is for our dynamics. But if we do want to use the diffusive approximation, then this gives the optimal expression for T_{eff}, showing that we need only consider the equal-time correlator, rather than the integrated one.

4.2   As max-entropy modelling

To motivate this diffusive approximation better, we can view it as an example of maximum entropy modelling (or more precisely, "maximum caliber," as in [18]). We ask, if we knew nothing about our dynamics besides the amplitude of fluctuations (2-point correlators) throughout the configuration space, what would be our most unbiased guess at what the dynamics really are? We can frame this question precisely by finding the maximum-entropy probability distribution over the space of possible trajectories $\{x(t)\}$, under the constraint that two-point correlators match the data. As with all maximum-entropy modeling the result depends crucially on what we choose to constrain – even with two-point functions, we have the freedom to choose between equal-time correlator, integrated correlator, different choices of spatial averaging, etc. Motivated by the above inference result, eq.4.2, we will first try to constrain the equal-time correlator at every point in configuration space. Thus the entropy functional we want to maximize w.r.t. $P[x(t)]$ is:

$$\mathcal{S} = P[x] \, \log\left(P[x]\right) + \lambda_0 \left(1 - \int \mathcal{D}x \, P[x]\right) + \int dX \, \lambda_{ij}(X) \left(\frac{1}{2} \langle \dot{x}_i \, \dot{x}_j \rangle_X - T_{ij}(X)\right)$$

where $\langle \dot{x}_i \dot{x}_j \rangle_X \equiv \int \mathcal{D}x \ P[x] \int dt \ \dot{x}_i(t)\dot{x}_j(t) \, \delta(x(t) - X)$, and $T_{ij}(X)$ is the corresponding quantity measured from the data. λ_0 and $\lambda_{ij}(X)$ are Lagrange multipliers for the normalization and fluctuation constraints respectively. As in the inference section above, we have an interesting mix of integrals over configuration space and over trajectory duration. Setting the variation $\frac{\delta}{\delta P[x]}$ to 0, we get:

$$1 + \log(P[x]) + \lambda_0 + \int dX \ \lambda_{ij}(X) \int dt \ \dot{x}_i(t)\dot{x}_j(t) \, \delta(x(t) - X) = 0$$

$$\Rightarrow \quad P[x] = \exp\left[-1 - \lambda_0 - \int dt \ \lambda_{ij}(x(t)) \, \dot{x}_i(t)\dot{x}_j(t)\right] \quad (4.3)$$

We must then set the Lagrange multipliers such that our constraints are satisfied. This means that λ_0 is chosen to give proper normalization, and $\langle \dot{x}_i \dot{x}_j \rangle_X = \frac{1}{2}\lambda_{ij}^{-1}(X) = 2\,T_{ij}(X)$, giving us $P[x] = \frac{1}{Z} \exp\left[-\frac{1}{4} \int dt \ T_{ij}^{-1}(x(t)) \, \dot{x}_i(t)\dot{x}_j(t)\right]$. But this is precisely the probability distribution for the diffusive dynamics we had in eq.4.1!

We can thus view our diffusive dynamics assumption as the natural least-informative approximation reproducing the 2-point correlators of the actual dynamics. It is interesting to note here that this derivation is surprisingly fragile to the choice of constraint. If instead of the equal-time correlator (eq.4.2) we chose to constrain almost anything else, such as integrated correlator, we would have no way to easily express $\lambda_{ij}(X)$ in terms of $T_{ij}(X)$ – and in particular, the dependence would be non-local! This is because the distribution in eq.4.3 is non-Gaussian, and so the two-point correlator is not in general simple.

4.3   Filtering out regular motion

Clearly, the diffusive approximation badly fails if you have motion around a limit-cycle. In this case, you exit a particular region of configuration space, only to come back there every time, and so without suppressing the steady-state probability of that region. Statistically, this break-down between exit rates and $p_{s.s.}$ happens because entrance rates become strongly correlated with exit rates (see sec.5.1.1) – similar to a break-down of molecular chaos assumption. Therefore, we need some way to filter

out periodic motion from our estimates of effective temperature. To proceed, we can express our velocity, whose correlators we've been studying, in frequency space as $\dot{x}_i = \text{Im}\left[\int d\omega\,\omega\,x_i(\omega)\right]$. The evolution of this spectrum over time is basically what's plotted in fig.6-2b for the example system considered there. We can then try to filter it in different ways to remove any regular part of the motion.

For systems with periodic driving, we can simply suppress all signal at the driving frequency – this does not guarantee that other emergent periodicity cannot happen, but it will remove the dominant source of regularities. In general, there is no objective way to tease apart regular from chaotic behaviors – since to some extent it depends subjectively on the time-scale one which you observe your system. This is famously illustrated by the Fermi-Pasta-Ulam paradox, where dynamics that seemed chaotic at first, turned out to be exactly periodic on a longer time-scale. However, such effects are, in some sense, highly unlikely and idealized – we never expect to find them in real systems with noise and imperfect interaction potentials. This is the key point where we rely on our assumption of "complexity:" we assume that typically, for real many-body interacting systems, the probability of a trajectory returning to an earlier configuration is indistinguishable from chance, after a sufficiently long delay time (akin to molecular-chaos assumption). Note that this is a strictly weaker assumption than saying that complex systems lose memory of where they were over time – in fact, the entire point of this thesis is to discuss what sort of long-term memory effects we can expect in such systems.

This thinking does motivate another way to estimate effective temperature from velocity spectrum: since we believe that most periodicity or organization in such systems should happen on faster time-scale, we can simply use a low-pass filter. E.g., in the wavelet transform of fig.6-2b, we see the noise-component of the dynamics showing up as the broad-band spectral component in the lower-frequency range. Similarly, we can directly define the noise component as the broad-band part of the power-spectrum, after removing any sharply-peaked (narrow-band) features. This is also motivated by the intuition that white-noise is distinguished by having power at all frequencies.

This way, while at this point, we cannot make a precise argument for how to reliably separate out the truly chaotic part of the dynamics, there are a number of reasonable guesses we can make. Most definitively, we can say that at the end of the day, our chosen definition of effective temperature does seem to be predictive of the steady-state in our experiments. Moreover, this correlation remains largely unchanged for the different choices listed in this section, showing that the result is robust.

4.4   Constructing a scalar

With the above methods, we motivated how to measure effective temperature tensor $T_{ij}(x)$ from the trajectory data. To predict the steady-state distribution, however, we must first convert this to a scalar $\mathcal{R}(x)$, which we call rattling, with the hypothesis that $p_{s.s.}(x) \sim 1/\mathcal{R}(x)$. Note that no result like this can hold exactly in general, as the diffusion process in eq.4.1 does not generally admit a closed-form analytical solution for the steady-state. This can be understood by recognizing that this stochastic process is generally not detailed-balanced, and so $p_{s.s.}(x)$ can depend non-locally on values of $D_{ij}(y)$ anywhere.

Thus, all we can hope to get is an approximation of $p_{s.s}$, and we can thus be motivated in the choice of an approximation by studying solvable limiting behaviors. The two leading candidates for a scalar that could work here are the trace of T_{ij} and its determinant. As we have not yet been able to conclusively decide which of these two is more generally appropriate, following we present the arguments in favor of each choice.

The first argument to use the trace comes from the fact that for the isotropic limit of diffusion in eq.4.1:

$$\dot{x}_i = \sqrt{2\,T(x)} \cdot \xi_i \quad [\text{Itô}] \tag{4.4}$$

(i.e., if the temperature tensor $T_{ij}(x) = \delta_{ij}\,T(x)$), we get the Fokker-Planck equation

$\dot{p} = \partial_i^2 (T(x)\, p)$, whose steady-state is $p_{s.s.} \propto 1/T(x) \propto 1/\mathrm{tr}[T_{ij}(x)]$. Note that determinant here would $= T^d$, with d – dimensionality of configuration space, and so would not be inverse-linear with $p_{s.s.}$. The numerical tests in sec.5.2 similarly support that $p_{s.s.} \propto 1/\mathrm{tr}[T_{ij}(x)]$.

Another argument in favor of trace comes from noting that if the inference in section 4.1 were carried through from the assumption of isotropic temperature, the maximum-likelihood estimator for that scalar temperature would have been $T(X) = \frac{1}{2} \sum_i \langle \dot{x}_i(t)\, \dot{x}_i(t) \rangle_{t \,|\, (x(t)-X)^2 < \epsilon} = \mathrm{tr}\, [T_{ij}(X)]$. Similarly, in section 4.2, if instead of constraining the tensor $\langle \dot{x}_i \dot{x}_j \rangle_X$, we chose to constrain its trace, we would have landed again on eq.4.4, but if we instead chose to constrain its determinant, we could not arrive at any local closed-form result due to the additional non-linearity of determinant.

The determinant, on the other hand, is nice in that it is consistent with composite systems. I.e., if we have two independent subsystems with $p_{s.s.}^{(1)}(x) \propto 1/\det\left[T^{(1)}\right]$ and $p_{s.s.}^{(2)}(y) \propto 1/\det\left[T^{(2)}\right]$, then the total $p_{s.s.}(x, y) = p_{s.s.}^{(1)}(x)\, p_{s.s.}^{(2)}(y) \propto 1/\det\left[T^{(1)}\, T^{(2)}\right]$. In particular, if the temperature tensor takes the form $T_{ij}(\vec{x}) = \delta_{ij}\, T_i(x_i)$, where its i-th component is independent of all $\{x_j\}_{j \neq i}$, then the Fokker-Planck equation $\dot{p} = \sum_i \partial_i^2 (T_i(x_i)\, p)$ has the steady-state $p_{s.s.} \propto \prod_i 1/T_i(x) = 1/\det\left[T_{ij}(x)\right]$. We thus see that for different x-dependencies of the temperature tensor, $1/p_{s.s.}(x)$ can be proportional to either the trace or the determinant – we explore this cross-over directly in section 5.2 below.

Another evidence in favor of the determinant is its behavior under coordinate-changes. This is of crucial importance, since the estimation of T_{ij} in eq. 4.2 is entirely dependent on how we choose to parametrize our data. Concretely, if we change variables $x_i \to y_i(\vec{x})$, then the probability density picks up the Jacobian of the transformation $J(y) \equiv \det\left[\partial^i y_j\right]$ (where partial derivative is w.r.t. x-coordinates), such that $p_{s.s.}(x) \to \tilde{p}_{s.s.}(y) = p_{s.s.}(x(y))/J(y)$. Temperature tensor, on the other hand, transforms $T_{ij}(x) \to \tilde{T}_{ij}(y) = T_{kl}(x(y))\partial^k y_i \partial^l y_j$, and so $\det\left[\tilde{T}_{ij}\right] = J^2 \det\left[T_{ij}\right]$. This way, the expression $p_{s.s.}(x) \propto 1/\sqrt{\det\left[T_{ij}(x)\right]}$ is reparametrization-invariant. But what about this square root?

We must also remember that by Itô's lemma, a change of coordinates also intro-duces drift terms, such that the correct y-dynamics for eq.4.1 are $\dot{y}_i = T_{jk}\,\partial^j\partial^k y_i + \partial^k y_i D_k{}^j \cdot \xi_j$. Thus, for 1D, or for 1D systems composed together as above, this drift ensures the correct transformation of the expression $p_{s.s.}(x) \propto 1/T(x)$. In high di-mensions, we could argue that unless the coordinate change is fine-tuned to the profile of $T_{ij}(x)$, these drift terms may be small – but on the other hand, they are correlated with magnitude of T_{ij}, and so may have a substantial effect on $p_{s.s.} - T$ correlation.

In all, it seems that the correct way to construct a scalar for rattling (such that $p_{s.s.} \sim 1/\mathcal{R}$) may depend on the statistics of disorder in our dynamical system. Em-pirically, $\mathcal{R} = \det[T_{ij}]$ seems to be the most robust choice for the smarticle system in chapter 6 – though we see some cases where the power $p_{s.s.} \sim 1/\mathcal{R}^p$ deviates from 1, e.g., fig.6-4a. In Chapter 5 we explore other examples of dynamical systems where different choices for $\mathcal{R}$ may be more reliable.

Chapter 5

Random dynamical systems

In this section, we attempt to make more precise statements about "typical" dynamical systems and their steady-states by randomly constructing and analysing various types of such systems. With some of these we can make analytical progress, with others we must mostly resort to numerics. The general theme we are focusing on here is that while the exact steady-state at a point $p_{s.s.}(x)$ will depend non-locally on system properties throughout configuration space, we can often construct a local observable (i.e., rattling $\mathcal{R}(x)$), which is approximately predictive of the steady-state. Some related systems have been explored in the 80s, with a particularly clear review in [4]. The main development of this research was an understanding of anomalous diffusion processes, and subsequent applications of it to all kinds of systems. Here, we focus on a different facet: the steady-state in a bounded domain.

5.1   Random Markov process

5.1.1   Continuous time

While we have up to this point only talked about dynamical systems with continuous state-space, the easiest example of random dynamics we can solve is a Markov process

on N discrete states. Consider the master equation

$$\dot{p}_i = \sum_j R_{ij}\, p_j - \sum_j R_{ji}\, p_i \tag{5.1}$$

with $\{i, j, ...\} \in \{1, ..., N\}$ labeling discrete states. In general the steady-state distribution $p_{s.s.}(i)$ is the null eigenvector of the transition-graph Laplacian matrix $L_{ij} = -R_{ij} + \delta_{ij} \sum_k R_{ki}$, and thus will depend in a complicated way on all the N^2 transition rates.

Now, as we want to understand large disordered dynamical systems, we let the transition rates R_{ji} be independent identically distributed random variables, with some mean $\bar{R}$ and standard deviation σ, and let $N \to \infty$. We now want to try finding the steady-state as an asymptotic series in powers of $\frac{1}{N}$: $p_{s.s.}(i) = p_i^{(0)} + p_i^{(1)} + ...$ First, by the central limit theorem, we can write the mean entrance and exit rates for state i as:

$$\frac{1}{N} \sum_j R_{ij} = \bar{R} + \frac{\sigma\, \xi_i}{\sqrt{N}}$$

$$\frac{1}{N} \sum_j R_{ji} = \bar{R} + \frac{\sigma\, \zeta_i}{\sqrt{N}}$$

respectively, where ξ_i and ζ_i are univariate Gaussian random variables specifying the deviation of i-th entrance and exit rate from the mean. We immediately see that at leading order, constant $p_i^{(0)} = \frac{1}{N}$ is a self-consistent choice:

$$\dot{p} = \sum_j R_{ij} \frac{1}{N} - \sum_j R_{ji} \frac{1}{N} = \bar{R} + \frac{\sigma\, \xi_i}{\sqrt{N}} - \bar{R} - \frac{\sigma\, \zeta_i}{\sqrt{N}}$$

which vanishes for large N. The first correction to the steady-state must further reduce $\dot{p}$:

$$\frac{\sigma\, (\xi_i - \zeta_i)}{\sqrt{N}} + \sum_j R_{ij}\, p_j^{(1)} - \sum_j R_{ji}\, p_i^{(1)} \sim o\left(\frac{1}{\sqrt{N}}\right)$$

(using the little-o notation). We can check that letting $p_i^{(1)} = \frac{\sigma\, (\xi_i - \zeta_i)}{\bar{R}\, N^{3/2}}$ accomplishes

this:

$$\sum_j R_{ij}\, p_j^{(1)} - \sigma\, \zeta_i\, p_i^{(1)}\, \sqrt{N} \sim o\left(\frac{1}{\sqrt{N}}\right)$$

where the second term is explicitly of $O\left(\frac{1}{N}\right)$, and the first term averages out to be of that order if R_{ij} is uncorrelated with $p_j^{(1)}$.

This last assumption is the crucial step of the derivation, and plays a similar role to that of the molecular chaos assumption in equilibrium thermodynamics. In particular, it breaks time-reversal symmetry by distinguishing between exit and entrance currents for a given state i: the former are all correlated as they are all $\propto p_i^{(1)}$, while the latter average out as they are proportional to the independent random variables $p_j^{(1)}$. In reality, this assumption does not exactly hold, as $p_j^{(1)}$ depends on jth exit rate ζ_j, which correlates with R_{ij} – but as all the rates are assumed independent, the effect of this correlation is suppressed by an additional factor of $\frac{1}{N}$. Thus we see that if we introduce correlations in the rates R_{ji}, this assumption will be the first failure mode of the derivation – but by the same token, as long as it holds, so should our result.

Putting the pieces together, we thus get the approximate expression for the steady-state

$$p_{s.s.}(i) = \frac{1}{N} + \frac{\sigma\,(\xi_i - \zeta_i)}{\bar{R}\, N^{3/2}} + O\left(\frac{1}{N^2}\right) = \frac{1}{N} + \sum_j (R_{ij} - R_{ji})/(N^2\,\bar{R}) + O\left(N^{-2}\right)$$

$$(5.2)$$

We can verify this result numerically, as illustrated by the red points in fig.5-1. The first thing we notice about this expression is that our random transitions rates produce only small fluctuations in the steady-state distribution on top of a dominant uniform background. This is in sharp contrast to what we see for the smarticles, where steady-state probability piles up almost entirely in certain stable states. We may wonder if extreme values of the $p_{s.s.}(i)$ variation here may be comparable to the uniform background – but from asymptotic expansion of the error function, we can check that for large N, the maximum amplitude of ζ_i we can expect only reaches $O\left(\sqrt{\log N}\right)$,
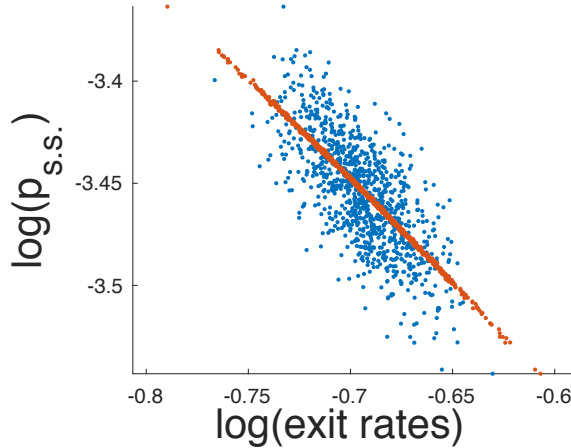
83

Figure 5-1: Steady state for a Markov process with 1000×1000 random transition matrix. x-axis for blue points gives exit rates, while for red points it also include the correction from entrance rates.

short of the $O\left(\sqrt{N}\right)$ we would need. This way even the largest amplitude $p_{s.s.}(i)$ variations are small compared to the background.

Next, we see that steady-state probability at a state depends in an equal measure on both the exit, and the entrance rates of that state. In practice, however, the total entrance rate may be hard to measure, as it would require initializing the system in all possible configurations and seeing how often it enters i. The exit rate, on the other hand, is akin to our measure of rattling – it is basically a measure of how stable the state i is and only requires local measurements initialized in that state. While not enough to predict the steady-state exactly, it can already tell us a lot, and will be strongly correlated with it (see blue points in fig.5-1).

This way, we see that while in general we need all N^2 rates to predict the steady-state, the randomness in rates and large system size allow us to express it in terms of just $2N$ local measurements of the exit and entrance rates.

5.1.2 Discrete time

As we considered periodic driving in some of our examples, which could be thought of stroboscopically as discrete-time dynamics, it is natural to extend the current discus-

sion to Markov process with discrete time-steps. Starting from our setup in the last section, $\dot{p}_i = -L_{ij} p_j$ ($\{i, j, ...\}$ again label discrete states), we could naturally define a discrete-time process given by one-time-step evolution of this equation, governed by the transfer matrix $\mathcal{T}_{ij} = \exp[-\tau L_{ij}]$. However, as we are looking to relate the $p_{s.s.}(i)$ to the actual entries of the given transfer matrix (which we might hope to measure experimentally), this will not be useful, as expressions in eq.5.1.1 are in terms of L_{ij} matrix elements, which have a non-local complicated relationship to the entries of $\mathcal{T}_{ij}$. Of course, for small times τ, we can Taylor expand to get a simple local relation $\tau L_{ij} = \delta_{ij} - \mathcal{T}_{ij}$, but this will not be true more generally.

Instead, let's carry out a similar derivation to that of sec.5.1.1. The only restriction on the transfer matrix is that its columns sum to 1 to preserve normalization, and so the simplest starting point is thus to take $\mathcal{T}_{ij} = t_{ij}/\sum_k t_{kj}$, with t_{ij} i.i.d. random variables ≥ 0. As before, we can define the pre-normalization entrance and exit probabilities for state i respectively via

$$\frac{1}{N}\sum_j t_{ij} = \bar{t} + \frac{\sigma \xi_i}{\sqrt{N}}$$

$$\frac{1}{N}\sum_j t_{ji} = \bar{t} + \frac{\sigma \zeta_i}{\sqrt{N}}$$

so that $\mathcal{T}_{ij} = \frac{t_{ij}}{N\bar{t} + \sqrt{N}\sigma\zeta_j}$, $\sum_i \mathcal{T}_{ij} = 1$ by construction, and $\sum_j \mathcal{T}_{ij} = 1 + \frac{\sigma \xi_i}{\bar{t}\sqrt{N}} - \frac{\sigma}{\bar{t}^2 N^{3/2}}\sum_j t_{ij}\zeta_j$, where the last term turns out to be of $O\left(\frac{1}{N}\right)$ as t_{ij} and ζ_j are mostly uncorrelated.

Proceeding with the asymptotic expansion $p_{s.s.}(i) = p_i^{(0)} + p_i^{(1)} + ...$ for the steady-state $p_{s.s.}(i) = \mathcal{T}_{ij} p_{s.s.}(j)$, we can once again check that at leading order we have a constant background $p_i^{(0)} = 1/N$, followed by:

$$p^{(0)} + p_i^{(1)} = p^{(0)}\left(1 + \frac{\sigma \xi_i}{\bar{t}\sqrt{N}}\right) + \sum_j \mathcal{T}_{ij} p_j^{(1)}$$

from which we can check that $p_i^{(1)} = \frac{\sigma \xi_i}{\bar{t} N^{3/2}}$ is a self-consistent choice as $\sum_j \mathcal{T}_{ij} p_j^{(1)}$

averages out to be of higher order, just as in sec.5.1.1. This way we get the result:

$$p_{s.s.}(i) = \frac{1}{N} + \frac{\sigma \xi_i}{\bar{t} \, N^{3/2}} = \frac{1}{N} \sum_j \mathcal{T}_{ij} \qquad (5.3)$$

which are just the entrance probabilities. The most obvious difference from previous section's result in eq.5.1.1 is that this only depends on the entrance probabilities – exit probabilities, which we argued were our main point of interest there, do not correlate with $p_{s.s.}(i)$ here! This is, quite trivially, because the normalization here enforces all exit probabilities to be $1 - O\left(\frac{1}{N}\right)$ – i.e., since we have N states we could transition into, it is very unlikely that we will stay in the same place over one time-step.

More generally, if there is a continuous-time process underlying our dynamics, such that $\mathcal{T}_{ij} = \exp\left[-\tau L_{ij}\right]$, then we could argue that for a sufficiently short time-step τ, the chances of leaving a state are low and can be informative – however if we increase τ, exit probabilities grow until they are capped at 1, and thus there is very little room for them to be informative.

Another way to motivate this result (eq.5.3) more broadly (beyond random rates) is to see that by raising the transfer matrix to a sufficiently high power τ (increasing the length of our chosen time-step) we can get $p_{s.s.}(i) = \lim_{\tau \to \infty} \mathcal{T}^\tau_{ij} 1_j$, where we let $1_i \equiv 1/N$ for all i – uniform initial probability, giving $p_{s.s.}(i) = \lim_{\tau \to \infty} \frac{1}{N} \sum_j \mathcal{T}^\tau_{ij}$, and so the row sums of $\mathcal{T}^\tau_{ij}$ should generally become progressively more correlated with $p_{s.s.}(i)$ as we increase the power (time-step) τ. E.g., starting from the result of the last section (eq.5.1.1), we see that for small τ, Taylor expanding $\mathcal{T}_{ij} = \exp\left[-\tau L_{ij}\right]$ allows us to write $p_{s.s.}(i) = \frac{1}{N} + \sum_j (\mathcal{T}_{ij} - \mathcal{T}_{ji})/(N^2 \bar{R} \tau)$, but as we increase τ, exit probabilities become less important, and eventually we just get $p_{s.s.}(i) = \frac{1}{N} \sum_j \mathcal{T}_{ij}$.

5.2 Diffusion in random medium

Here we look more carefully at the generic multi-dimensional diffusion process of eq.4.1, reproduced here:

$$\dot{x}_i = D_{ij}(x) \cdot \xi_j \quad [\text{Itô}] \tag{5.4}$$

with $T \equiv \frac{1}{2} D\, D^T$, and $\{i, j, ...\} \in \{1, ..., d\}$ index dimensions of the configuration space. As we cannot find the steady-state analytically, we study this numerically. Again, as we want to study some "typical" regime, we choose $D_{ij}(x)$ landscape randomly, but so that it varies smoothly in x. Explicitly, for each of d^2 entries of this matrix, we independently generate a random d-dim grid of numbers $\in (-1, 1)$ (e.g., for $d = 4$, this may be $5 \times 5 \times 5 \times 5$), and then use cubic interpolation to fill in the values in-between, thus making a smooth landscape (e.g., making a $15 \times 15 \times 15 \times 15$ grid). To then find the steady-state distribution, we explicitly construct the discretized Fokker-Planck operator $\partial_i \partial_j \left(T_{ij}(x) \cdot \right)$ as a matrix, implementing periodic boundary conditions, and find its null eigenvector. Note that if we now view this matrix as the Laplacian matrix of a Markov jump process, as in section 5.1.1 above, then the exit rates found on the diagonal, will be simply $\text{tr}\,[T_{ij}]$ (though this could be slightly different depending on the choice of discretization). Plotting the computed exact steady-state vs. trace or determinant of T_{ij} (fig. 5-2), we see that for this setup, trace correlates much better, and unlike determinant, scales linearly with $p_{s.s.}^{-1}$. The correlation becomes more pronounced if, after constructing the diffusion tensor $D_{ij}(x)$ as described above, we raise it to the 4th power to amplify its variation over x (fig.5-2b).

We can run a couple other tests with this setup. First, we claim that this correlation crucially relies on the disorder and "typicality" of our chosen dynamics. Thus, we should be able to break it by introducing some fine-tuning or correlations into the diffusion landscape. In particular, if we have strong diffusion along a closed ring, then while $\text{tr}\,[T_{ij}]$ will be large there, it will not contribute to any suppression of $p_{s.s.}$ as trajectories don't leave that ring any more frequently. We implement this scenario in
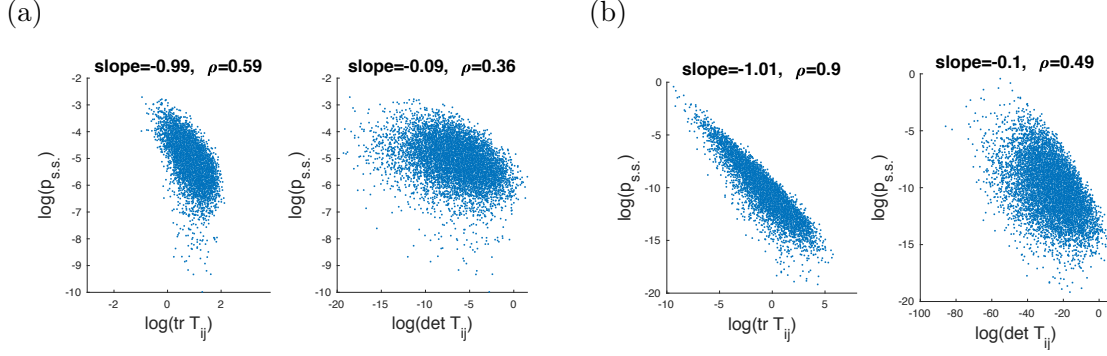
Figure 5-2: Numerical results for 4D diffusion process in eq.5.4 with randomly varying diffusion tensor $D_{ij}(x)$. (b) shows results when $D(x) \to D^4(x)$ to allow larger amplitude variation. Linear regression slopes and Pearson correlation coefficients are listed above.

2D (for better visualization – works same way in higher dimensions), as shown in fig. 5-3a, second panel. We also see in the third panel that if the direction of enhanced diffusion is not aligned along the ring, then the correlation is restored – illustrating that breaking it really requires strong fine-tuning.

Another correlation that we can introduce to the diffusion landscape to change our result is to make it a product of d 1D diffusion processes – in which limit we expect $p_{s.s.} \propto \det[T_{ij}]$, as discussed in section 4.4. Since thus far we've shown the trace is the one that is inversely linear with $p_{s.s.}$, this would mean that we should see a crossover between the two regimes. To achieve that limit, let's first restrict D_{ij} to be diagonal, and then ensure that D_{ii} does not vary along all $x_{j \neq i}$ dimensions. We can then smoothly tune up this cross variation $var_{j \neq i}(D_{ii})$ up to the full dependence $D_{ii}(\vec{x})$ to study the cross-over. Numerically, we implement this by first confining all interpolation query points along $j \neq i$ dimension to a tiny range of the generated random grid (see description above), and then spreading them out to the max range (illustrated by point grids in fig.5-3b). As seen in fig.5-3b for 4D diffusion, we indeed see this crossover in regression slopes of our distributions. As shown by the correlation coefficient plot, on the other hand, despite changing slope, trace remains well-correlated with $p_{s.s.}$ throughout, while determinant becomes uncorrelated as we approach full $\vec{x}$-dependence (the regime in fig.5-2).
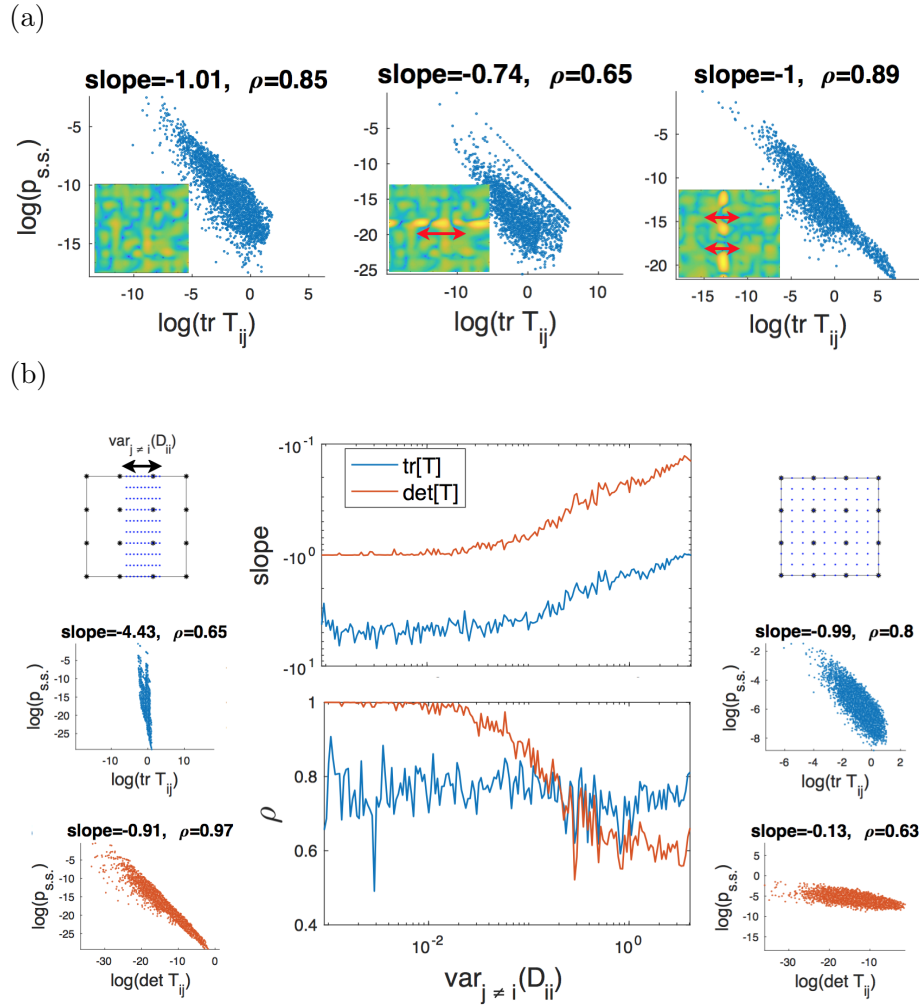
Figure 5-3: (a) Breaking the correlation by enhancing diffusion along a closed ring (in 2D). Insets show the $\log \mathrm{tr}\,[T_{ij}(x)]$ landscape (boundaries are periodic, so the bright bands are actually closed). Three panels show respectively the setups with diffusion unperturbed random, enhanced along the ring, and orthogonal to the ring – only the second case breaks the correlation. (b) Crossover from determinant to trace being predictive of $p_{s.s.}$ as variation of D_{ii} along $x_{j\neq i}$ dimensions is scaled up, showing the change of regression slope and correlation coefficient. Point grids illustrate how D_{yy} landscape is generated: random numbers are generated at black stars, and then cubic interpolation is taken at the blue query points, which form the landscape.

5.3 Random force field

Most continuous dynamical systems can be written as a first-order system

$$\dot{x}_i = F_i(x, t) + \sqrt{2\, T_0}\, \xi_i \tag{5.5}$$

If the system is sufficiently complex and high-dimensional, then we can try to gain insight about "typical" behaviors by letting the force-field $F_i(x)$ be random. Turns out we can gain some analytical insights on this problem – though the full result will still require numerics.

5.3.1 Analytics

We consider a smoothly-varying force-landscape, which we can analytically construct by letting $F_i(x)$ be Gaussian random field with $\langle F_i(x) \rangle = 0$ and $\langle F_i(x) F_j(y) \rangle = G_{ij}(x, y)$ – some correlation function that decays with $|x - y|$ (for cleaner notation, we focus on time-independent systems for now). As we are focusing on the effect of forces here, we take temperature to be a uniform scalar. This setup is nice in that if G decays very quickly in $|x - y|$, it approximates diffusive behavior, as any trajectory $x(t)$ experiences nearly independent forces over time. On the other hand, if there are system-spanning correlations in $F(x)$, then the behavior is more like deterministic chaos, where trajectories separate faster the further apart they already are. We can thus hope to see these limiting behaviors emerge in our results, and try to understand the crossover.

To evaluate a general observable $\mathcal{O}[x(t)]$ under disorder averaging, we must compute the partition function

$$Z[h_{\mathcal{O}}] = \left\langle \int \mathcal{D}x \, \exp\left[-\frac{1}{4T} \int dt \, (\dot{x}_i - F_i(x))^2 + h_{\mathcal{O}}\mathcal{O} \right] \right\rangle_{F_i(x)}$$

such that $\langle \mathcal{O} \rangle = \left. \frac{\partial Z[h_{\mathcal{O}}]}{\partial h_{\mathcal{O}}} \right|_{h_{\mathcal{O}}=0}$. Note that we can take the disorder average directly of $\langle Z \rangle$, rather that of $\langle \log Z \rangle$ because by construction our probabilities are already

normalized, as $Z[h_\mathcal{O} = 0] = \text{const.}$ independent of $F_i(x)$. In a sense, time takes on the role of a replica-index, and just as we get replica-coupling, we will see that here we get global time-coupling.

Evaluating the disorder average using cumulant expansion we thus get:

$$Z[h_\mathcal{O}] = \int \mathcal{D}x \, \exp\left[-\frac{1}{4T} \int dt \, \left(\dot{x}_i^2 + G_{ii}(x, x)\right) + \right.$$

$$\left. \frac{1}{(4T)^2} \iint dt \, ds \, \left(G_{ij}^2\left(x(t), x(s)\right) + 2\,\dot{x}_i(t)\dot{x}_j(s)G_{ij}\left(x(t), x(s)\right)\right) + O\left(\frac{1}{T^3}\right) + h_\mathcal{O}\mathcal{O}\right]$$

While this looks bad, it turns out we can build substantial intuition for the various terms if we think of it as a partition function for a polymer. From this perspective, the first integral just gives the energy of a Gaussian coil in an energy landscape given by $\text{tr}\,[G(x)]$ – i.e., the polymer avoids regions with large amplitude F_i, as the trajectory tends to leave these. The double integral then gives polymer self-interaction energy: the first term is an attractive force with range set by $F(x)$ correlation length (i.e., if a particle spent time in some region before, then forces are probably weak there, and it will spend time there again); while the second term is a collimating-interaction (i.e., if a particle was moving one way in a region, it will probably move the same way next time it is in that region). Note that in a high-dimensional configuration space, orientation-dependent forces will tend to be small unless the polymer is already tightly packed, and so can be ignored when we study how our trajectory first deviates from random exploration behavior. The self-attraction, thought, may already cause the polymer to form a globule and localize – corresponding to localized regular dynamics. Finally, for time-dependent forces, the intuition generally persists, except now the polymer will self-interact only among certain segments.

The exciting thing about this result is that while the global time-couplings that disorder averaging generates are entirely unnatural for a causal dynamical system, they are completely normal interaction for a polymer to have. We thus get a qualitative shift of description, and behaviors that were hard to describe from one perspective can become simple in the other. Speculatively, as polymers have no notion of causal-

ity, we can thus hope to see apparently a-causal behaviors, such as prediction or anticipation, emerge in the dynamics. However, as we have not yet explored these directions further, we now shift gears to show numerical results.

5.3.2 Numerics

Numerically, the simulation proceeds very much along the same lines as in section 5.2: we construct a smoothly varying random force field, from which we build a discretized Fokker-Planck operator $\partial_i(F_i(x)\cdot)+\partial_i^2(T\cdot)$ and find its null eigenvector. There are a few small changes we make compared to random diffusion runs. First, in generating the force landscape, we again start with the random d-dim grid, but instead of interpolating between it, we just smooth it out by a d-dim running average (set each point to the mean of its neighborhood). Again we do this independently for each component of $F_i(x)$. This generates a more disordered field, as the other method introduces some complicated correlations, in particular along our cubic grid. While for the random diffusion simulation the two methods gave very similar results, and interpolation was more convenient for studying the trace-determinant crossover of fig.5-3b, here the correlations in fig.5-4 turn out much weaker if we use interpolation. This may signify that in this setup, the relationship $p_{s.s.} \sim 1/\mathcal{R}$ is more sensitive to correlations in $F_i(x)$.

Second, as we don't have an explicit temperature landscape here, we have to talk about the effective temperature, which we define, according to the discussion in sec.4.1, as $T_{ij}(x) \equiv F_i(x)F_j(x) + T_0$, where T_0 is the inherent temperature from eq. 5.5, which we assume to be negligible in comparison. This way $\operatorname{tr}[T_{ij}] = F_i^2$ is just the force magnitude, while determinant doesn't really have a natural interpretation in terms of the force. Finally, to make force fluctuations more pronounced, we scale $F_i(x) \to F_i(x)\,|\vec{F}(x)|^3$ after generating it, giving us the results in fig.5-4.

The first thing we note from these results is that both trace and determinant here are similarly well-correlated with $p_{s.s.}$. We also note that $p_{s.s.}(x)$ is no longer inverse linear with the trace, as it was in section 5.2, but now $p_{s.s.}(x) \sim 1/\sqrt{\operatorname{tr}[T_{ij}(x)]}$. This could be understood by seeing that in 1D, the Fokker-Planck dynamics $\dot{p} =$
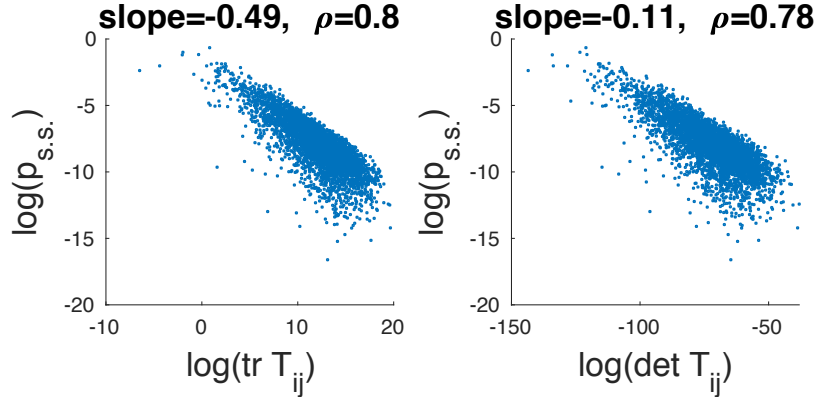
92

Figure 5-4: Numerical results for motion in 4D random force field, per eq.5.5. Effective temperature tensor here $T_{ij}(x) \equiv F_i(x)F_j(x)$. Regression slope and correlation coefficient are labeled above.

$\partial_x(F(x)\,p(x))$ are stationary when $p(x) \propto 1/F(x) = 1/\sqrt{T(x)}$, for our definition of effective temperature.

Finally, as in section 5.2, we can break this correlation by introducing fine-tuned patterns into the force-field. E.g., enhancing the force field along a closed ring, similar as was done for fig. 5-3a, produces qualitatively the same sort of results as shown in that figure.

5.4   Random discrete map

For periodically driven systems, it may be useful to view the system stroboscopically, making the effective dynamics discrete in time – but unlike sec.5.1.2, here we consider having a continuous configuration space. We can generally express this as:

$$x_i(t+1) = f_i(\vec{x}(t)) + \sqrt{2\,T}\,\xi_i \qquad (5.6)$$

This then gives the probability evolution, which can be expressed in terms of a transfer matrix:

$$p_{t+1}(\vec{y}) = \int d^d\vec{x} \, \mathcal{T}(\vec{y}, \vec{x}) \, p_t(\vec{x})$$

$$\mathcal{T}(\vec{y}, \vec{x}) = \frac{1}{\sqrt{4\pi T}} \exp\left[-\frac{1}{4T}(y_i - f_i(\vec{x}))^2\right]$$

The steady-state $p_{s.s.}$ is thus the dominant eigenvector of $\mathcal{T}(\vec{y}, \vec{x})$ (with the eigenvalue $= 1$). As in the other sections, we want to see what we can say about this eigenvector just from looking at some local properties of the transfer matrix, under the assumption that our map $f_i(x)$ is high-dimensional and disordered.

Again, analytical progress turns out to be unlikely here, and we proceed numerically. We construct the map $f_i(x)$ much as we constructed the force-field in the last section 5.3.2 – generating a d-dimensional grid of random numbers, and smoothing it out by averaging over local neighborhoods. As anticipated in sec.5.1.2, exit probabilities, as given by $\mathcal{T}(\vec{y}, \vec{x})$, are all ~ 1, as long as T is smaller than typical step-size $(x_i - f_i(\vec{x}))^2$, and so don't help us to predict the steady-state. Entrance probabilities (row sums of $\mathcal{T}(\vec{y}, \vec{x})$), on the other hand, are very well correlated with $p_{s.s.}(\vec{x})$ (fig.5-5), consistent with the prediction in eq.5.3. These are, however, hard to get from any local measurements of our dynamics – as they require knowing all the locations that map to $\vec{x}$ under f.

If we look more carefully at the exit probabilities, we see that they deviate from 1 by the diagonal elements $\mathcal{T}(\vec{x}, \vec{x})$ – staying probability. Plotting $p_{s.s.}$ vs. $\mathcal{T}(\vec{x}, \vec{x})$ on a log-log plot (fig.5-5), we see that these do indeed correlate, albeit nowhere near as much as entrance rates. Moreover, we see that the x-axis of this plot, $\log \mathcal{T}(\vec{x}, \vec{x}) \propto -(x_i - f_i(\vec{x}))^2$ is just the map's step-sizes, which is a very natural local measure of state-stability that can be readily measured experimentally. Additionally, this step-size is quite similar to the effective temperature measures we have looked at in other sections (e.g., sec.5.3.2), showing that the general intuition we've built is still relevant for discrete-time processes.
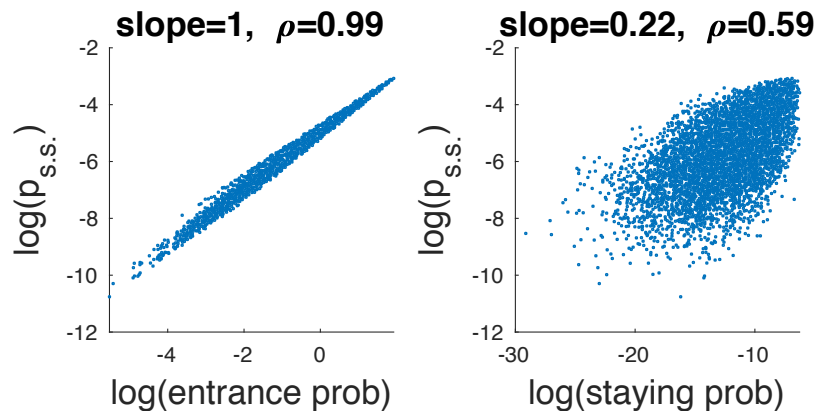
Figure 5-5: Numerical results for random 4D map, per eq.5.6. Entrance prob. are row sums of the transfer matrix $\mathcal{T}(\vec{x}, \vec{y})$, and staying prob. are its diagonal elements, whose log (x-axis of the right plot) is $\propto -(x_i - f_i(\vec{x}))^2$ = map step size.

Chapter 6

Smarticles experiment

In this chapter we describe another example system we used to test our proposal –
a robot swarm developed in Dan Goldman's lab at GA Tech (see, e.g., [7] for a past
report of these). We carried out a number of different experiments in collaboration
with his group, and ran more extensive tests in simulations of their system. Using
ideas presented above, we were able to design an experiment that exhibited the sort
of least-rattling self-organization we've argued for, and have then characterized and
checked this phenomenon quantitatively against some of our predictions. Such spon-
taneous ordering in swarm-robotic context can take us a long way towards engineering
complex collective behaviors – as robust collaborative dynamics are generally difficult
to orchestrate by hand [65, 56]. This chapter describes our setup, experiments and
results in detail.

6.1   Experiment design

We conducted our experiments with a collection of small simple robots called "smar-
ticles" (smart active particles) (fig. 6-1). Each one is comprised of three connected
links, with the two hinges being controlled by a programmable micro-controller via
servo motors. When the smarticle is sitting on its base, the arms do not touch the
ground, and so an individual smarticle cannot move. A group of them, on the other
hand, can achieve complex motion by pushing off each other. Here we will focus on

the regime where their arms are executing a motion pattern, which we call "gait," with no feedback from the environment (note that the gait timing is kept by the controller, and hence is not affected by forces on the arms). For most of the paper, we will consider the scenario where an ensemble of smarticles executes a gait in perfect synchrony with each other (e.g., fig.6-2 a). The ensemble is then placed on a stage inside a plastic ring that keeps the group confined as they push each other by moving their arms. The ring is in turn glued to the aluminum stage-plate for simplicity (fig.6-1 b). To collect data, we use infrared video recording suite "OptiTrack," which helps track the time-evolution of the 2D coordinates and body-angle (x, y, θ) of each robot at 120 frames per second (we do not track the arm positions).

Most of the technical details of the smarticle construction and experimental realization of this setup were described in [7]. One key difference from those experiments is that in this work, the gaits of the different smarticles in the ensemble were kept synchronized throughout the runs, which is what allowed them to fall into regular dynamical patterns not previously observed. For periodic gaits, this synchronization could be achieved by connecting a microphone sensor to an interrupt pin of the microcontroller in each smarticle, and then triggering every gait cycle with a beep from external speaker placed next to the setup. This meant that all the smarticles stopped momentarily between consecutive gait cycles to listen for the trigger – but this stop was kept down to a fraction of a second by carefully timing the triggering pulse-train with the gait period. The resulting setup was robust enough to keep the experiment running and synchronized indefinitely (though runs presented here were limited to 10 min each). It also ensured that the arm motion was precisely periodic throughout the experiment, with no possibility for any phase drifts of fluctuations.

For experiment prototyping and better sampling statistics, we have also built a simulation of this setup, described in Appendix 6.A, and benchmarked it against experiments. Both experiment and simulation generates time-series $(x(t), y(t), \theta(t))$ data for each smarticle, which we then run through the same data processing pipeline, detailed in Appendix 6.B. The main steps in this pipeline are as follows: First, we convert the $3N$ (where N is number of smarticles) raw coordinates into $2N$ rotation-
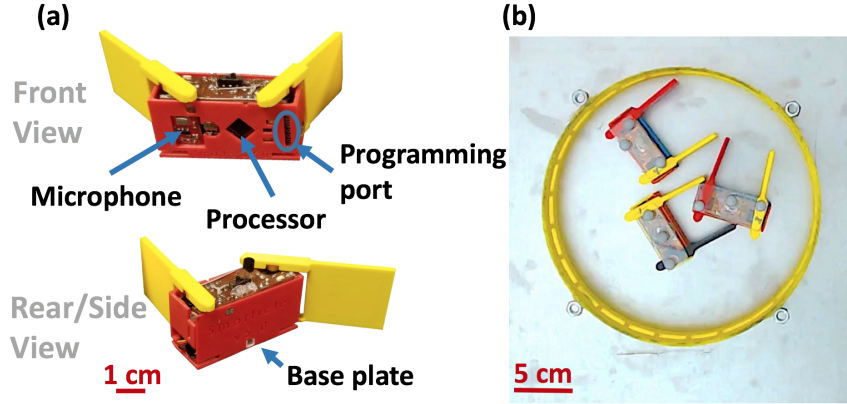
Figure 6-1: (a) Diagram of a single Smarticle robot, both body and arms are 5.2 cm long. (b) One simple setup to look for collective self-organizing dynamics (view from above, the 19.2 cm diameter ring confinement is glued to the aluminum stage leveled to < 0.1°).

and permutation-invariant observables we call $x_i(t)$ $(i \in 1, ..., 2N)$ that parametrize our space of distinct ensemble configurations. Some 3D projection of these coordinates, sampled stroboscopically whenever the two arm angles $= (\frac{\pi}{2}, \frac{\pi}{2})$, is shown in fig.6-2 c scatter plot. Second, we use short snippets of runs (3-gaits long) initialized in a given neighborhood of the configuration space, to measure the apparent noise, or effective temperature, there. We define the tensor $T_{eff} \equiv \langle \dot{x}_i \dot{x}_j \rangle$ (averaged over the neighborhood), where $\dot{x}_i$ is defined as the "noise-only" velocity, based on the broadband component of the spectrum of the trajectory snippet: $\dot{x}_i \equiv \text{Im} \left[\sum_{\omega \in b.b.} \omega \, \tilde{x}_i(\omega) \right]$ (see fig.6-2 b, which shows a related way to extract this quantity). We then define a scalar quantity "rattling" as $\mathcal{R} \equiv \det T_{eff}$, which we show along the color axis (with log scale) in fig.6-2 d. Chapter 4 explores these choice for our definition of $\mathcal{R}$ in more detail. Third, for each neighborhood, we measure how often a long trajectory $\vec{x}(t)$ visits it, and scatter plot the resulting steady-state density $p_{s.s.}(\vec{x})$ vs. $\mathcal{R}(\vec{x})$ we found – fig.6-2 e.

6.2 Self-organization

In general, the response properties of a dynamical system to an external drive will depend on the system's configuration. In particular, a given drive can cause chaotic
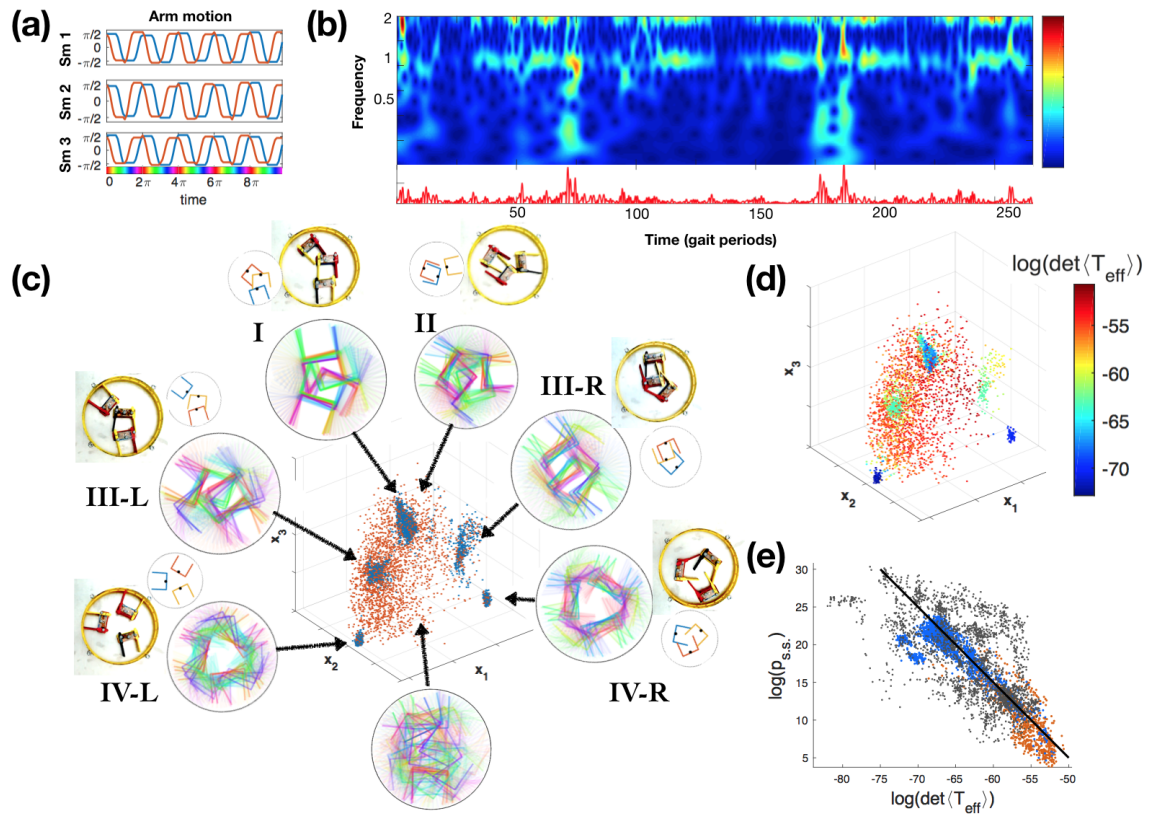
Figure 6-2: a) Motion pattern for the two arms of each of the three smarticles for this setup. b) Wavelet transform of one Smarticle's x-coordinate over a 10-min experimental run (gait frequency = 1). Amplitude of the broad-band component is plotted below, and is indicative of how organized the dynamics are at different times. c) Dynamical regular states that emerge in this setup: initial conditions for each [exp't and sim]; time evolution along the color-axis (as labeled in (a)) for 3 gait periods [sim]; ensemble position in its configuration space (stroboscopic at time-points colored red in (a), projected to 3D) - blue points sampled from long-time evolution, red sampled by randomly dropping smarticles in the ring [sim] – illustrating the spontaneous fine-tuning (note: states I and II are indistinguishable in this projection; states III and IV each have two chiral compliments). d) Configuration space points colored by $\log(\det \langle T_{eff} \rangle)$ of their neighborhood [sim] – long trajectories spend most of their time in low-T_{eff} regions. e) correlation illustrating steady-state probability $p_{s.s.} \propto 1/\det \langle T_{eff} \rangle$ (black line): experimental data for long trajectories is in grey, simulation data is in blue/red (for different regions of configuration space, as in (c))

.

motion in some parts of the configuration space, while being perfectly organized elsewhere. To check if this applies to our swarm dynamics, we set an ensemble of three smarticles to execute a synchronized periodic gait, as shown in fig.6-2 a, inside a glued-down ring, track the $(x(t), y(t), \theta(t))$ trajectories for each smarticle, and plot out the wavelet transform of the orientation $\theta(t)$ of any one smarticle – fig. 6-2 b. As energy from the periodic drive (arm-motion) is injected at a particular frequency (normalized to =1 here), we see a peak in the power-spectrum there for all times. In contrast, the amount to which that energy gets dispersed to other frequencies before dissipating out varies as the system explores different configurations over time. In the regions where a large broad-band component emerges, we have very noisy dynamics. The lack of a broad-band component, on the other hand, indicates that the dynamics are orderly, which we may be surprised to find in such a messy system – yet here we see that this is the case for most times.

Taking a closer look at configurations found during these regular dynamics, we see that there the three smarticles organize into one of several periodic dynamical patterns – a sort of collectively organized dance. Figure 6-2 c attempts to convey some illustration of these dynamical patterns by overlaying the configurations visited over 3 gait cycles, colored (periodically) by time – we clearly see that the six patterns are qualitatively distinct from each other, as well as from chaotic motion shown below. The scatter plot in that figure illustrates the fine-tuning associated with these regular dynamical states (blue), on contrast with a sampling of random 3-smarticle arrangements (red).

Such self-organization is perplexing, and its interpretation depends on the perspective we take – that of dynamical systems theory, or thermodynamics. In the context of dynamical systems, perhaps it is not surprising that a periodically-driven system is governed by several limit-cycle attractors. On the other hand, not all periodic systems relax to such attractors – even here we can make slight changes to our setup, such as giving Smarticles inertia, so as to make the system behave chaotically everywhere (fig.6-3). Furthermore, this system cannot really be seen as deterministic, as we know that it is strongly affected by the stochasticity coming from the hard-

boundary interactions between non-smooth shapes, and the highly surface-sensitive effects of sliding friction.

From the perspective of equilibrium thermodynamics, where entropy and disorder tend to a maximum, this sort of self-organization may seem extremely surprising. Of course, this system is far from equilibrium, and so entropy need not generally grow. Indeed, for a Brownian particle diffusing in a compact domain of inhomogeneous temperature $T(\vec{x})$ (and no other forces), we know that the steady-state probability distribution will be $p_{s.s.} \propto 1/T(\vec{x})$ – which can be far from the maximum-entropy uniform distribution (even though energy landscape is flat). We want to use this example to motivate our explanation of the fine-tuning we observe.

Indeed, the intuition is the same: probability concentrates in regular states, characterized by having an orderly response to the drive, which are contrasted with the rest of the configuration space, where response is chaotic. We can make this intuition precise by defining an effective temperature $T_{eff}(\vec{x})$ at every point in the configuration space, and checking if it correlates with $p_{s.s.}(\vec{x})$. While this system is too complicated to calculate $T_{eff}(\vec{x})$ from first principles, we can measure it from local information only, by looking at short-term trajectories starting in the neighborhood of $\vec{x}$. Our hypothesis, therefore, is that we can predict the long-term behavior from looking only at short trajectories – which need not be possible for far-from-equilibrium systems in general, as discussed in [38] (commonly referred to as Landauer's blowtorch theorem).

Chapter 4 goes through the considerations of defining this T_{eff} in detail, but intuitively, it is based on the broad-band component of the spectrum of these short trajectories (see fig.6-2 b). Coloring the points in fig. 6-2, panel d, by the scalar $\mathcal{R} \equiv \det T_{eff}$, which we call rattling, we see that the regions of low $\mathcal{R}$ correspond precisely to the concentration of steady-state probability in panel c. More precisely, panel e shows the scatter plot of $p_{s.s.}(\vec{x})$ vs. $\mathcal{R}(\vec{x})$, showing that rattling is highly predictive of the steady-state throughout the configuration space, both for in experiment and simulation.

We have thus identified a local quantity (akin to exit rates from a region), which empirically allows us to approximately predict $p_{s.s.}$, which may seem at odds with

the blowtorch theorem [38]. At the same time, the statement that at long times the system tends to be found in places from where it does not leave, seems almost tautologically true. These considerations are made more precise in random dynamical systems of Chapter 5. We argue that for a high-dimensional complicated dynamical system, we can assume some sort of "typicality," which provides the structure we need to preclude the fine-tuning required for the blowtorch effect. We thus argue that our claim is both, highly non-trivial in a precise mathematical sense, and also profound for understanding real-world strongly-driven systems.

6.3   Breaking regular states

To further establish the applicability of this framework, we want to try making predictions about new experimental conditions. In particular, we can dramatically reduce friction between the smarticles and the stage by filling the ring with a loose single layer of small plastic beads, and putting the smarticles on top of these, thus allowing them to skate around with very low friction. In simulation, we can run a series of experiments interpolating between these two regimes by gradually reducing damping (fig.6-3). While the exact damping force in experiment is complicated, and besides the rolling friction also depends on collisions among smarticles and the ring, we simply approximate it by Stokes drag in simulation. This may be unrealistic in detail, but guarantees stability as velocities are capped – which may not be the case if we used a constant friction force. Empirically, the grey cluster in fig.6-3 a illustrates that experimental data fits nicely into the pattern found in simulation.

This way our velocities get a finite decay curve: $\ddot{x} + \dot{x}/\tau = \delta(t) \quad \Rightarrow \quad v(t) = v(0) \exp\left[-t/\tau\right]$. We then predict that the associated increase in the mean free path acts to raise the effective noise level everywhere by the same τ-dependent isotropic value $\delta T(\tau)$ (since $v(0)$ would be roughly set by motor speeds, independent of configuration). This way, we expect $p_{s.s.}(\vec{x}) \propto 1/\det\left(T_0(\vec{x}) + \delta T(\tau)\right)$. Thus, as we tune up τ, configurations that used to be low-rattling – i.e., the regular states – get substantially suppressed, while chaotic configurations are only weakly affected (and may even

become slightly more probable as normalization changes). Perhaps the key falsifiable prediction we make here is that the stochasticity coming from dynamical chaos inherent to our overdamped system, and stochasticity coming from the enhanced mean-free path, while of very different origins and dynamical properties, are indistinguishable in their effect on the steady-state. This is reminiscent of universality in statistical mechanics, where most microscopic details of interactions and dynamics become irrelevant for determining some global property. All these predictions are then confirmed in fig.6-3 a: colored points lie along the black ones, indicating that T_{eff} is predictive of $p_{s.s.}$, whether the noise comes from inertia or from chaotic interactions; increasing τ suppresses the likelihood of the regular states, and collapses the range of $p_{s.s.}$ variation; all colored points shift progressively further to the right as τ increases (though only weakly due to the log-scale).

We can do even more by analytically estimating the dependence $\delta T(\tau)$. Intuitively, for larger τ, any energy input from the drive stays in the system longer, thus making the velocity amplitude larger at the steady-state. We can thus estimate the change in $\dot{x}^2$, and hence δT, by balancing the rate of heat dissipation $dQ \sim \delta \dot{x}^2/\tau$ with the rate of input work dW, which will be related to motor speeds, and so independent of both τ and configuration $\vec{x}$. This way, noting that $T_0(\vec{x}) + \delta T(\tau) = \langle(\dot{x}_0 + \delta\dot{x})^2\rangle = \langle\dot{x}_0^2\rangle + \langle\delta\dot{x}^2\rangle$ (as inertial steady-state velocity contribution is uncorrelated with our definition of filtered $\dot{x}$), we can therefore write $p_{s.s.}(\vec{x}) \propto 1/\det(T_0(\vec{x}) + dW\,\tau)$. In particular, we know that $T_0(\vec{x})$ is positive definite everywhere, approaching 0 in regular states, so that $\log p_{s.s.}(\vec{x}_{reg}) = -\log(Z\,dW^d) - d\log\tau$ will be a bound, where the power d (configuration-space dimension) comes from the determinant, and Z is the normalization. For small τ, such that $\delta T \lesssim \langle T_0(\vec{x})\rangle$, Z is independent of τ, and so we get a concrete analytical prediction for the suppression of regular state probabilities, with Z as a fitting parameter. This is shown as the black curve in fig.6-3 b, with dotted lines showing departure from this bound outside of regular states where $T_0 > 0$, and along with simulation measurements of $p_{s.s.}$ at 1000 points in configuration space. Note that when τ is large and δT dominates $T_0(\vec{x})$ in most regions, we get a flat total temperature landscape, and thus a flat $p_{s.s.}$. Fig.6-3 b shows the crossover
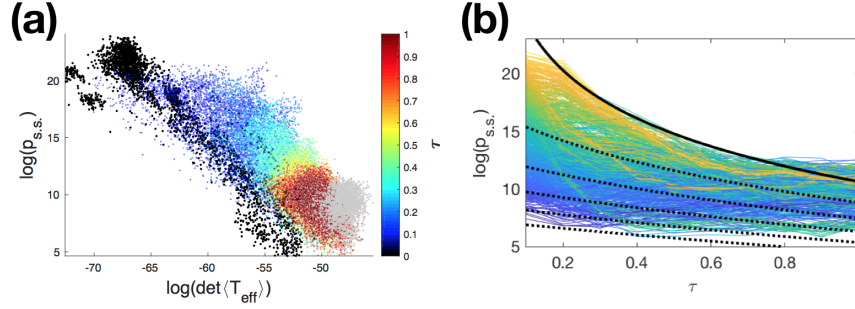
Figure 6-3: Destroying regularization by reducing Smarticle damping. a) The overdamped run (black points, from fig. 6-2e) is predictive of all simulations with lower damping (color = velocity decay time-scale τ, in units such that gait period $= 2\pi$), which fall along the same line; light grey – experimental data for weakly-damped Smarticles. b) Effect of damping on $p_{s.s.}$ at 1000 different points in configuration space [sim] (color by plot's left edge); black curve is a theoretical prediction for the upper edge (with one fitting parameter = vertical shift), with dashed lines illustrating the expected relaxation away from that bound.

to this regime setting in around $\tau \gtrsim 0.8$, as our assumption of constant Z begins to break-down.

6.4 Drive-specificity

Why should a driven dynamical system get an inhomogeneous $T_{eff}(x)$ landscape? For a given form of the drive signal, different system configurations will have a different response. In particular, some configurations will respond in a more predictable way than others. In our Smarticle setup, for some initial conditions the sources of stochasticity (such as corner collisions) will be more prevalent than elsewhere. This way, a system settles into drive-specific configurations that give more orderly and stable response properties, and changing the drive will change these selected states. For smarticles, this would imply that the regular states found at long times are specific to the gait. To explore this, we run simulations with 3 smarticles executing a gait that is random, but still identical and synchronized across the ensemble – see fig.6-4 a. While this system has no periodicity, we again see regular states emerge, as shown by the distribution of $\mathcal{R} = \det T_{eff}$ in the configuration space shown (as before, only configurations with U-shaped arms are plotted). On the other hand, we

see that this distribution is distinct from that for periodic gait in fig.6-2 d, illustrating the drive-specificity of $\mathcal{R}(\vec{x})$ – here we only have 3 regular states as opposed to the 6 from before, and even these are less sharply localized. As shown by the scatter plot, our measure of rattling is still predictive of the steady-state distribution, though as the slope is different from -1, we do not have the inverse-linear relation that was clear before. As a control, we can repeat this experiment with each of the 3 smarticles running its own, independent, random gait. As expected, we see that this produces no correlation (grey points in the scatter plot), as all configurations have similar T_{eff}: there is no longer any structure in the drive that some configurations could be better or worse adapted to.

Finally, to illustrate further generality of our framework, we repeat our experiment (again in simulation) with 5 smarticles inside the ring, all executing synchronized periodic gait, as shown in fig.6-4 b. As the volume of configuration-space grows exponentially with the number of smarticles N, so does the number of possible regular states (at least for this simple gait). Two representative regular dynamical patters are shown. The scatter plot, produced the same way as before, again shows a very clear predictive power of $\mathcal{R}$ for the steady-state. Note also that as the red points were sampled by randomly dropping smarticles in a ring, the fact that they almost never sample low-rattling regions indicates that regular states are much more fine-tuned in this setup – and so the fact that they are typical at the steady-state is more surprising.
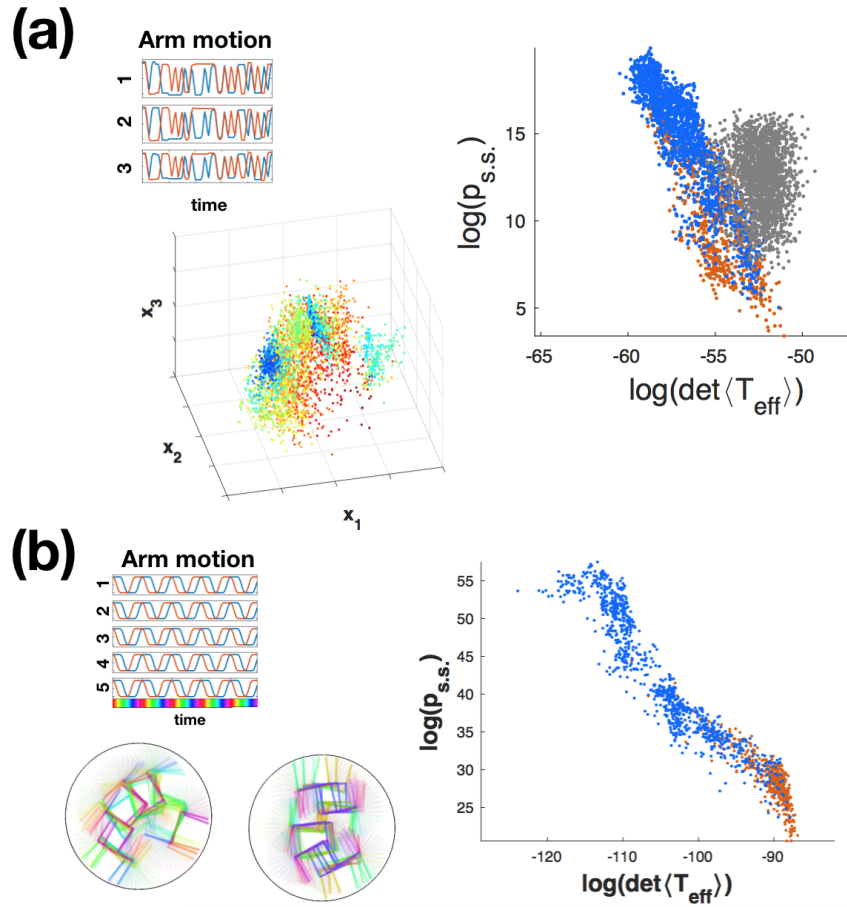
Figure 6-4: a) Aperiodic drive: arms move in a random, but synchronized fashion. T_{eff} distribution is different from fig.6-2d – different regular states emerge. $p_{s.s.}$ is still strongly correlated with T_{eff}, though the exponent is no longer -1. Grey – random and independent arm motion, leaving no structure in the drive for adaptation. b) 5 Smarticle group executing synchronized square gait shown similarly produces fine-tuned regular behaviors (two sample ones are shown). Note that most of the red points in the scatter plot are confined to high-T_{eff} region, indicating a higher degree of fine-tuning in regular states than for 3 Smarticles. [all data from simulation].

Appendix

6.A   Simulation details

For easier exploration of hypothesis space of this system, we have constructed a numerical simulation with the following algorithm. We approximate smarticles by three-segment lines (insets in fig. 6-2 c). At each time-step, the algorithm moves the arms slightly according to the chosen gait, and then iteratively cycles through smarticle pairs in random order, checking for collisions, and moving one in each pair slightly according to the net interaction force. If there are multiple points of contact for a given pair, the move is a translation in the direction of the total force, otherwise, it is a rotation about a pivot point chosen so as to balance the forces and torques. Choosing which in the pair moves is random, weighted by their relative friction coefficients (as motivated by difficulty of predicting static friction). Note that since a move can create new collisions with other smarticles, it is important to take small steps and iterate. The algorithm continues looping through pairs until all collisions are resolved, then proceeding to the next time-step. While this describes the core of the algorithm, there are a number of bells and whistles necessary to improve its stability and reliability:

- If two arms are near-parallel when they approach each other, they can pass through each other between ticks without ever intersecting. To prevent this, along with collision detection, we must explicitly test for this in each pair. We then store the order of the arms for a few ticks into the future to prevent them passing through each other in any of those times.

- In case a smarticle with small friction gets trapped between two others, it might rattle back and forth on each iteration of collision-resolution, with no net effect. To prevent this, we temporarily (until next tick) increase its friction each time a smarticle moves.

- In experiment, when resolving collision is too hard, the motor simply does not move (jams up). To allow for that possibility in simulation, we add an exit condition from the collision-resolving loop for when any one smarticle's temporary friction (from last bullet) becomes very large - as this serves as a proxy for how much force a motor must provide. We then move the most colliding arm back to its last time-step, and try collision-resolving again. If everything resolves, that arm will then have a chance to catch up to where it needed to be over the following ticks (its speed being capped at some ω_{max}).

- The resulting simulation turns out to be too clean, despite the multiple stochastic elements of the algorithm, and so dynamical phases much more stable than in experiments. Since we want to study transition statistics, we must find ways to destabilize them. There are a number of places we can add more noise: slight fluctuations of smarticles' positions and angles at each tick, or proportional to each move, randomly varying gait amplitude, varying arm velocity, etc. Adding inertia or inter-Smarticle friction forces (see below) also helps to destabilize the dynamics.

- The ring boundary is implemented similar to other smarticles, and collisions with it are resolved in the same loop. Alternatively, we can implement weakly confining potential to keep smarticles together in a more smooth way (for different experiments).

- It is easy to adjust the simulation to give the Smarticle inertia: at each tick, simply move them according to last step's velocity before resolving collisions.

- If inter-Smarticle friction is 0, then each interaction force is directed normal to one Smarticle's surface. To include effects of such friction, we can add a small

lateral component to these forces that depends on the interaction angle.

Even with all these additions, many differences remain between simulation and experiment: smarticles have non-zero thickness in experiments, there are relief features on smarticle body not present in simulation that can get caught, the precise force-response profile of the motors is not captured, etc. The consequence in this system is that while qualitative features can be recovered in the simulation, we don't generally expect precise quantitative agreement. We guess that these differences may be less important for ensembles with more smarticles – as universal collective properties start to dominate. We also guess that analytical predictive power of least-rattling, or other similar statistical techniques, could be more applicable in that regime.

6.B   Data analysis

6.B.1   Constructing the configuration space

Once we have the tracked coordinates and body orientations of the individual smarticles over time $(x_s(t), y_s(t), \theta_s(t))$, either from experiment or simulation (fig.6-5 a), we want to construct new coordinates that are optimally informative about collective ensemble configuration. As all experiments presented in this work were done in a circular ring boundary and treated all the smarticles as identical, we want configurations that differ only by global rotation or smarticle permutations to be seen as identical. The easiest way to accomplish this is to construct a complete basis of inherently rotation and permutation invariant observables. Moreover, as we want to study probability density and fluctuation amplitudes in this space, both of these will be sensitive to the distance metric we get – and so our chosen observables should give a faithful notion of distance between configurations. Doing this optimally is highly non-trivial (as it is related to parametrizing a quotient space, which isn't even smooth, see [29]), and so we instead settle for a heuristic guess of observables that give 0-distance between identified configurations (related by rotation or permutation), while somehow distinguishing between the rest. Note that while any choice of such coordinates will

necessarily represent a huge distortion to our original space, as it folds in on itself in complicated ways, the high-dimensionality can help to average out these distortions such that we can hope to get a distance metric faithful to the distinguishability of configurations, even for sub-optimal choice of observables. Either way, as we will see, our analysis gives results that are largely reparametrization-invariant, and so is robust to coordinate-choice (see fig.6-6 b).

To take care of the global rotation symmetry, we find the coordinates of the c.o.m. $\vec{X} = \frac{1}{3}\sum_s \vec{x}_s$ (note: permutation invariant) in the reference frame of each smarticle: $\vec{\chi}_s \equiv \mathcal{R}(\theta_s) \cdot (\vec{X} - \vec{x}_s)$, where $\mathcal{R}(\theta_s)$ is the rotation matrix for s-th smarticle. Note that this is not a perfect choice: it is entirely invariant under c.o.m. translations (which is relevant, but we assume it to be small for a confined system), and smarticle orientation matters more the further it is from c.o.m. (so we must assume that their distance to c.o.m. stays relatively constant) – but these are the inevitable distortions that we discussed above. Also, this choice projects our number of tracked variables from $3N$ down to $2N$ (where N is number of smarticles), which, for $N = 3$ can be attributed to the rotation and translation invariance of the new coordinates, but for larger N will constitute a loss of some information about the configuration as well – and so we may seek out better choices. For the one experiment with $N = 5$ that we present in this work, we still use this analysis for consistency (again, counting on the robustness of our analysis illustrated by fig.6-6 b).

Next, the permutation-invariant functions we choose are just the first three moments of the distribution of the three $\{\chi_s\}$ at any fixed time. In order for these functions to have relatively similar sensitivity to changes in raw data, it helps to ensure that they have the same units – by raising them to the appropriate fractional powers. I.e., $\mu_1 = \frac{1}{3}\sum_s \chi_s$, $\mu_2 = \sqrt{\frac{1}{3}\sum_s (\chi_s - \mu_1)^2}$, $\mu_3 = \sqrt[3]{\frac{1}{3}\sum_s (\chi_s - \mu_1)^3}$ (fig. 6-5 b).

Finally, we recall that the full description of smarticle configuration also crucially depends on all the arm angles $(\alpha_s^{(1)}, \alpha_s^{(2)})$. However, for the experiments in this work, arm motion does not depend on any of the other variables, and is uniquely (or statistically, for the random arm motion) determined by the arms' current state. This way,

we need not consider arm angles as additional dynamical degrees of freedom, and it is sufficient to simply ensure that we are only comparing configurations at the same arm angles. For all the plots in this work, we choose these to be $(\alpha_s^{(1)}, \alpha_s^{(2)}) = (\frac{\pi}{2}, \frac{\pi}{2})$ (U-shape, for all smarticles). Note that while for periodic gaits, this simply means that we sample the dynamics at stroboscopic time-intervals (fig.6-5 b), for the random arm motion, we post-select only those configurations where arms were at the specified angles. Also note that as we cannot track the arms directly in experiment, we do the runs with random arm motion in simulation only, and even for periodic gait we must make do with inferring the stroboscopic time-points indirectly from motion pattern of the bodies and the knowledge that arms are synchronized and exactly periodic for the entire run.

This gives us a set of discrete points in $2N$ dimensional configuration space corresponding to distinct collective configurations visited along the run. The 3D scatter plots in fig.6-2 c,d, fig.6-4 a, and fig.6-5 c directly show these sets of points, projected down to $(x_1, x_2, x_3) = (\mu_1^{(1)}, \mu_1^{(2)}, \mu_2^{(1)})$ subspace (chosen arbitrarily, just so it distinguishes the regular states clearly). Some of the points in those plots are sampled from long runs, while others come from randomly dropping down the 3 smarticles inside a ring, all in simulation, to have a faithful sampling of the entire configuration space.

6.B.2 Estimating steady-state distribution and effective temperature

The challenge with extracting the steady-state $p_{s.s.}(x)$ and effective temperature tensor $T_{ij}(x)$ from the data is that since the configuration-space is high-dimensional, we can't sample it very well, and so our chosen algorithm must be robust to poor and non-uniform sampling. While there are many possible ways to proceed with this, we choose a method that seems more reliable, unbiased, and informationally-efficient than some of the other possible choices. Qualitatively though, other simple algorithms that we tested for this always reproduce the same patterns. We start by outlining our entire algorithm, after which we will explain how we actually implement each step.

Parametrizing our $d = 2N$-dimensional configuration space by observables μ defined in the last section (which we now call x_i, indexed by $i, j \in \{1, ...d\}$) allows us to compute distances using Euclidean metric, as these observables were already chosen to be faithful to the distinguishability of smarticle configurations – though we try to rely on this metric as little as we can.

1. Begin by defining two sets of points on the configuration space, $\mathcal{D}_p$, which samples the steady-state distribution, and $\mathcal{D}_T$, sampling all regions on interest in this configuration space.

2. Choose a subset $\mathcal{S}_S \subseteq \mathcal{D}_T$ of points where to evaluate $p_{s.s.}$ and T_{ij}

3. For each point $\vec{x}^* \in \mathcal{S}_S$, find the subset $\mathcal{S}_n(\vec{x}^*) \subset \mathcal{D}_T$ of its n nearest-neighbors. We will estimate $p_{s.s.}(\vec{x}^*)$ and $T_{ij}(\vec{x}^*)$ from the region covered by this set, see fig.6-5 c.

4. $p_{s.s.}(\vec{x}^*)$

   (a) Compute the variance tensor $V_{ij}(\vec{x}^*) = \langle (x - x^*)_i (x - x^*)_j \rangle_{\vec{x} \in \mathcal{S}_n(\vec{x}^*)}$ and the volume $V(\vec{x}^*) = \sqrt{\det V_{ij}(\vec{x}^*)}$ of the region occupied by $\mathcal{S}_n(\vec{x}^*)$ (note: this requires $n \gg d$ to give sensible estimates).

   (b) Find the steady-state probability of this region $P[\mathcal{S}_n(\vec{x}^*)]$ as the fraction of all points $\nu \in \mathcal{D}_p$ that satisfy $(\nu - x^*)_i (\nu - x^*)_j V^{-1}_{ij}(\vec{x}^*) < 5$ (are within 2.2σ of $\vec{x}^*$).

   (c) The steady-state probability density is thus estimated as $p_{s.s.}(\vec{x}^*) = P[\mathcal{S}_n(\vec{x}^*)]/V(\vec{x}^*)$.

5. $T_{ij}(\vec{x}^*)$

   (a) For each point $x_i \in \mathcal{S}_n(\vec{x}^*)$, take a short run $x_i(t)$ (e.g., 3 gait-periods long if periodic), initialized at that point $x_i(0) = x_i$.

   (b) Fourier transform to get $x_i(t) \to \tilde{x}_i(\omega)$

   (c) Take the average broad-band (b.b.) signal: $\dot{x}_i \equiv \text{Im}\left[\sum_{\omega \in b.b.} \omega \, \tilde{x}_i(\omega) \right]$. For periodic drives, broad-band would typically be all frequencies not propor-

114

tional to the drive frequency. This gives a sort of "noise-only" velocity, see fig.6-5 d.

(d) Thus define $T_{ij}(\vec{x}^*) = \langle \dot{x}_i \dot{x}_j \rangle_{\vec{x}(0) \in \mathcal{S}_n(\vec{x}^*)}$

In simulation, we have the flexibility to carry out these steps directly, so we begin by going through that. To generate $\mathcal{D}_p$ in step 1, we directly take the set of ensemble configurations $\{\vec{x}\}$ measured stroboscopically from a long run, as described in the last section. Note that we view these points as just a set, ignoring any time-ordering information. For $\mathcal{D}_T$, on the other hand, we take the union of two sets: one generated by randomly dropping smarticles inside the ring (resolving any collisions to produce valid configurations), and another chosen as a random subset of $\mathcal{D}_p$ (the points corresponding to the two sets are colored red and blue respectively in our correlation plots, fig.6-2e, 6-4). This allows us to explore the full configuration space, while also getting better sampling on the interesting regions near regular states (see blue points in fig.6-5 c). The algorithm is designed so that the non-uniform sampling this creates does not introduce any bias into our results.

Since to estimate $p_{s.s.}(\vec{x}^*)$ we must choose a neighborhood of $\vec{x}^*$, we can do this optimally according to the locally available sampling by using the set $\mathcal{S}_n(\vec{x}^*)$ to specify a variable-size region of the configuration space. Most importantly, this ensures that our results are robust to strange coordinate choices, as averaging regions are determined by sampling densities rather than by distance metric directly. Additionally, this ensures that all estimates of $p_{s.s.}$ and T_{ij} are based on a comparable number of samples, and thus have similar uncertainties. This allows getting higher-resolution in the densely-sampled neighborhoods of regular states, while optimally extracting the available information from our data. The variance tensor $V_{ij}(\vec{x}^*)$ for the point-set thus estimates the shape of this region, and its inverse specifies a local distance metric we can then use in step 4b to count the fraction of the steady-state samples that fall within this region. Normalizing this by the region's volume gives us the density estimate we wanted.

The choice of how to define our effective temperature T_{ij} based only on local infor-
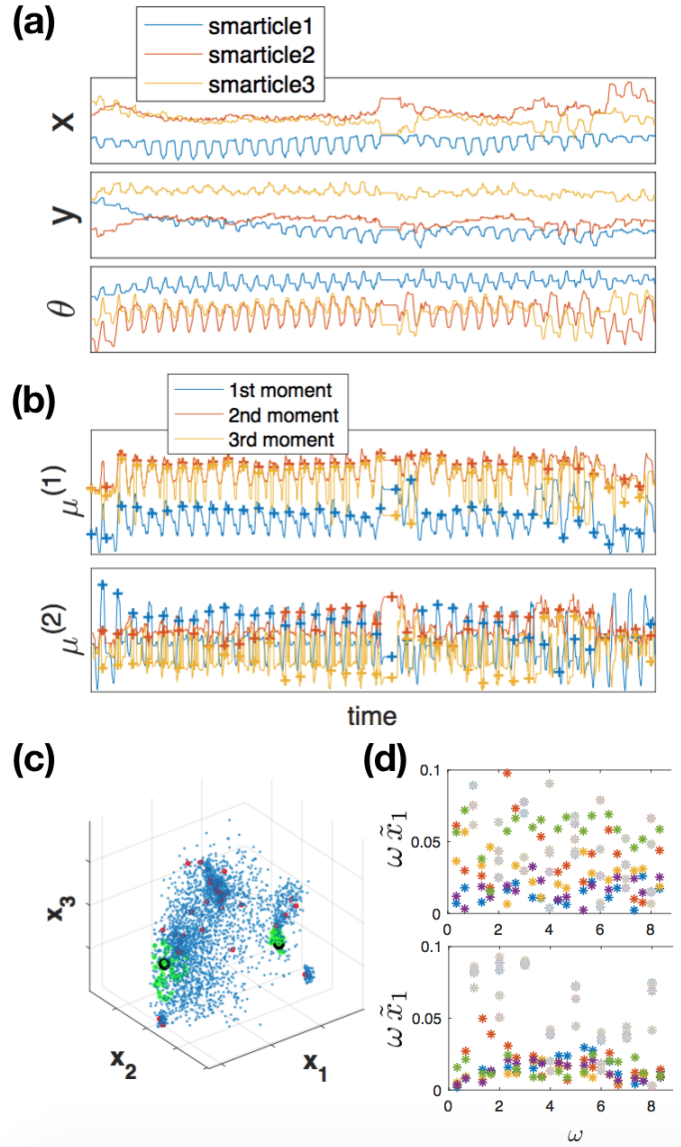
Figure 6-5: Data analysis steps: (a) time-traces of smarticle coordinates and angles as tracked from experiments $(x(t), y(t), \theta(t))$ for 3-smartice ensemble. (b) time-traces of rotation- and permutation-invariant observables constructed from (x, y, θ), which define our configuration space. (c) $(\mu_1^{(1)}, \mu_1^{(2)}, \mu_2^{(1)})$ projection of points marked by crosses in (b) (configurations when arm angles $(\alpha_1, \alpha_2) = (\frac{\pi}{2}, \frac{\pi}{2})$). Blue also includes configurations sampled by random initialization (set $\mathcal{D}_T$), red illustrates sub-sampling $\mathcal{S}_S$, for two representatives of which (black), we highlight the 50 nearest-neighbors $\mathcal{S}_n(\vec{x}^*)$ (green) – showing that neighborhoods are different size and shape depending on the local sampling. (d) shows spectra of 3-gait-long trajectories, realized 5 times from the same initial configuration. Top plot is for chaotic-state, bottom – regular state, showing that after notch-filtering harmonics of the gait frequency (light grey), the broad-band spectral amplitude is reproducible and distinct for the two cases.

mation and so that it can do a good job in predicting $p_{s.s.}$, is much less straightforward, and is, in some sense, the key question of this work. This topic is discussed in detail in Chapter 4. Here, we simply show results for a few possible choices. The best choice we found (eq.4.2), and one used for all the figures above, just relies on the velocities of trajectories passing through our averaging region: $T_{ij}(\vec{x}^*) = \langle \dot{x}_i \dot{x}_j \rangle_{\mathcal{S}_n(\vec{x}^*)}$. A key caveat to this, however, is that any motion that is periodic or regular should not contribute to a measure of effective temperature, which should capture only the chaotic or disordered part of the motion. This is why instead of taking the velocities directly for this measure in step 5c, we take an entire 3-period segment of trajectory initialized near $\vec{x}^*$, and notch-filter its spectrum to remove any contributions to $\dot{x}$ coming from the periodic motion. Fig.6-5 d, illustrates that this choice is well-founded: in regular states, we can have significant motion at the gait frequency, which nonetheless does not contribute to the flat white-noise-like background. Note that in our simulations, we actually ran 5 realizations of these short trajectories from each point of $\mathcal{S}_n(\vec{x}^*)$ for better averaging, and the spectra of these different realization shown in fig.6-5 d attest to the reproducibility of our chosen measure for $\dot{x}$. More generally, for aperiodic drives, we can simply average the part of the spectrum that we believe to be representative of the broad-band component: this could often be the IR frequencies – slower than any regular motion we may expect – or just the background amplitude after having removed any outliers.

Specifically for periodic drives, another natural quantity we could use to estimate noise-level is how much you move over one full period: $\delta x_i \equiv (x(\tau) - x(0))_i$, where τ is the drive period. We can see that this is clearly distinct from the quantity $\dot{x}$ above: if we Fourier transform, we get $\delta x_i = \text{Re}\left[\int d\omega \, (\mathrm{e}^{i\omega\tau} - 1)\, \tilde{x}(\omega)\right] = \int d\omega \, (\cos(\omega\tau) - 1)\, \text{Re}\left[\tilde{x}(\omega)\right] - \sin(\omega\tau)\, \text{Im}\left[\tilde{x}(\omega)\right]$ – which also notch-filters the drive frequency harmonics, but otherwise looks quite different. The result from using $T_{ij}(\vec{x}^*) = \langle \delta x_i \delta x_j \rangle_{\mathcal{S}_n(\vec{x}^*)}$ is plotted in fig.6-6 a, and turns out to be almost identical to that coming from $\dot{x}$ – though is only defined for periodically driven systems. This measure δx_i and time-discrete dynamics is further developed in the sec.5.4.

After computing the tensor T_{ij} from our data, we then still need to construct
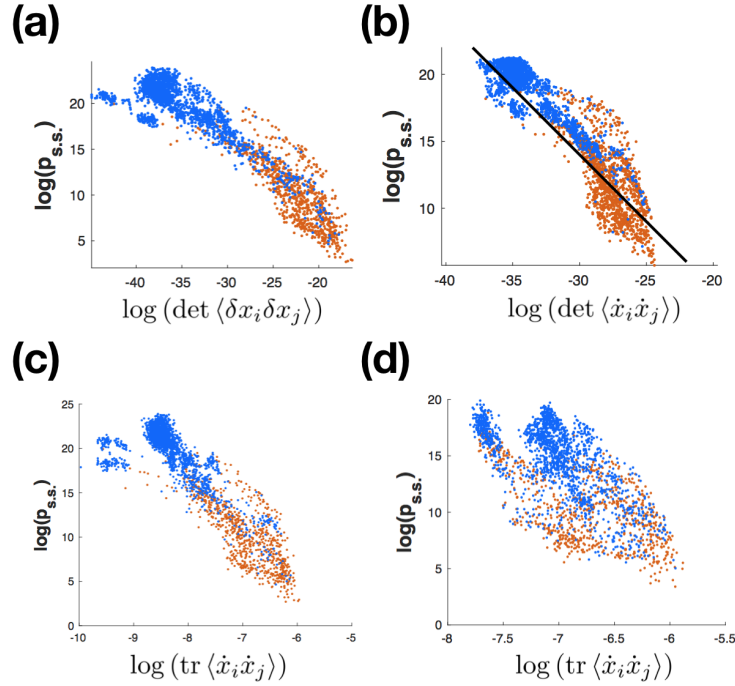
Figure 6-6: Correlations for alternative setups: (a) uses the displacement over one gait period δx_i to get effective temperature T_{ij}. (b) uses our usual choice $\mathcal{R} = \det \langle \dot{x}_i \dot{x}_j \rangle$, but the entire analysis is ran on a 3D projection $(\mu_1^{(1)}, \mu_1^{(2)}, \mu_2^{(1)})$ of the full 6D configuration space, illustrating the robustness of this metric. (c) and (d) use trace T_{ij} instead of determinant, for the runs with periodic (fig.6-2) and random (fig.6-4a) gaits respectively. Unlike in (a) and (b), plots (c) and (d) don't have equal spacing for the $x - y$ axes and so show slopes very different from -1, while also begin less correlated and prone to outliers. $\log (\det \langle \dot{x}_i \dot{x}_j \rangle)$ $\log (\operatorname{tr} \langle \dot{x}_i \dot{x}_j \rangle)$ $\log (\det \langle \delta x_i \delta x_j \rangle)$

a scalar out of it, which we can then correlate with $p_{s.s.}$. This choice of scalar is discussed in sec.4.4, and we use determinant throughout the paper above. Here we also show results for the other natural choice: trace, see fig. 6-6 c,d. The key point we want to illustrate is that unlike the trace, determinant comes out to be both, inversely linear to the $p_{s.s.}$, and robust across systems and changes of variables (even as dramatic as throwing away half the configuration-space dimensions in fig.6-6 b). Trace, on the other hand, generally correlates with $p_{s.s.}^{-n}$, where the power n changes under reparametrizations, and is more sensitive to outliers and noise.

Finally, to do the analysis for experimental data, where all we can get are several long-term time-traces of $(x_s(t), y_s(t), \theta_s(t))$, we simply let $\mathcal{D}_T = \mathcal{D}_p$ in step 1, letting

both be sampled from these long runs. While this does a worse job at sampling the chaotic regions of the configuration space, as the trajectory tends to stay around regular regions, the analysis can be done with no changes, and is sufficiently robust to still produce good results. The only obvious effect of under-sampling chaotic configurations is that in fig.6-2 e, the experimental data only covers the part of the simulation data corresponding to regular-state neighborhoods (blue points), and does not extend into the region we sample in simulation by random initializations (red points) – just as we would expect. One other caveat is that for step 5a, while in simulation we could independently run short trajectories starting from desired initializations, here we must get these by snipping them out of the long runs starting from the corresponding points, and then proceeding in the same way. Again, this should not introduce any biases or additional dependencies into our data.

Chapter 7

Conclusions

The equilibrium partition function that is computed for the Boltzmann distribution is a powerful formal tool for making predictive calculations in thermally fluctuating systems. Its success stems from two key simplifying assumptions: first, that energy only enters or leaves the system of interest in the form of heat exchanged at a single temperature, and second, that the system and surrounding heat bath uniformly sample joint states of constant energy. This latter ergodic assumption essentially amounts to eliminating time from the picture, so that energy and probability become interchangeable.

The nonequilibrium scenario is generally less tractable than its equilibrium counterpart both because time has not been eliminated from our description of the system, and also because energy is permitted to enter and leave the system via different couplings to the external environment. Thus, the specific approach to modeling some nonequilibrium systems we have described here seeks to recover some of the desirable advantages of the equilibrium description by exploiting an assumption of typicality: considering dynamical systems that have so little structure or correlation in their dynamical rules that their global behavior is indistinguishable from that produced by certain random dynamics. There are, of course, many ways to construct random dynamical systems, as we go through some cases in chapter 5. To get a good approximation, we must thus constrain some of the random choices so as to reproduce the key structures of the original dynamics. One such structure may often be locality:

dynamics only move locally in configuration space. With no other constraints, the random approximation turns out to be a diffusion process (chapter 4).

We have established that this effective picture is quantitatively predictive of the transient and stationary behavior of distributions for slow variables in systems with clear time-scale separation (chapters 2 and 3), and in a more complex example of an interacting robot swarm (chapter 6). The tendency of such systems to gravitate to configurations that reduce the effective temperature suggests a interesting relationship between dissipation and kinetic stability in overdamped driven systems. Although nonequilibrium steady-states are not in general required to be extrema of the average dissipation rate, it is true that the minimum required dissipation to maintain an effective temperature scales with T_{eff}. This suggests a lower bound on dissipation rate for the unstable configurations in a system. Stable configurations, on the other hand, can be achieved not only by reducing dissipation and decoupling from the drive, but also by maintaining dynamically regular motions, which may be highly dissipative.

We have therefore argued that such low-T_{eff} dynamical regularities may be typical in complex systems at long times, especially if decoupling from the drive is hard or impossible. This is especially exciting when such behavior requires substantial fine-tuning and adaptation to the driving environment. We have confirmed that this adaptation is indeed observed in the two specific toy-systems we have studied, both in simulation and in experiment. In the swarm-robotic system we looked at, such spontaneous dynamical ordering could give significant headway on engineering collaborative task-oriented behaviors.

Upon further development, we can take our framework in two general directions. One progression would be to try to refine it as a method of analysing otherwise intractable dynamical systems. This would entail identifying the key non-generic structures of the system (such as specific symmetries, regularities, conserved quantities, etc.), and by letting all other elements of the dynamics be random, constructing a tractable approximation. The other direction is to use least-rattling as a guiding principle to designing dynamical systems that are capable of fine-tuning and adapting

to complicated driving environments. We may hope to thus construct goal-oriented dynamical behaviors (either for self-assembly or task execution) only by tuning environmental selection pressures. The key advantage here would be that we need not know a-priori the successful design or process for the prescribed task, but instead let dynamics themselves search the configuration space via something like a genetic algorithm.

Bibliography

[1] Urna Basu, Christian Maes, and Karel Netočný. Statistical forces from close-to-equilibrium media. *New Journal of Physics*, 17(11):115006, 2015.

[2] MV Berry and JM Robbins. Chaotic classical and half-classical adiabatic reactions: geometric magnetism and deterministic friction. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 442, pages 659–672. The Royal Society, 1993.

[3] Stefano Bo and Antonio Celani. Multiple-scale stochastic processes: Decimation, averaging and beyond. *Physics reports*, 2016.

[4] Jean-Philippe Bouchaud and Antoine Georges. Anomalous diffusion in disordered media: statistical mechanisms, models and physical applications. *Physics reports*, 195(4-5):127–293, 1990.

[5] Barbara Bravi and Peter Sollich. Statistical physics approaches to subnetwork dynamics in biochemical systems. *Physical Biology*, 2017.

[6] Stijn Bruers, Christian Maes, and Karel Netočný. On the validity of entropy production principles for linear electrical circuits. *Journal of Statistical Physics*, 129(4):725–740, 2007.

[7] Sarah Cannon, Joshua J Daymude, William Savoie, Ross Warkentin, Shengkai Li, Daniel I Goldman, Dana Randall, and Andréa W Richa. Phototactic supersmarticles. *arXiv preprint arXiv:1711.01327*, 2017.

[8] J. Cardy, G. Falkovich, K. Gawedzki, S. Nazarenko, and O.V. Zaboronski. *Non-equilibrium Statistical Mechanics and Turbulence*. London Mathematical Society Le. Cambridge University Press, 2008.

[9] Michael E. Cates and Julien Tailleur. Motility-induced phase separation. *Annual Review of Condensed Matter Physics*, 6(1):219–244, 2015.

[10] Antonio Celani, Stefano Bo, Ralf Eichhorn, and Erik Aurell. Anomalous thermodynamics at the microscale. *Phys. Rev. Lett.*, 109:260603, Dec 2012.

[11] Raghunath Chelakkot, Arvind Gopinath, Lakshminarayanan Mahadevan, and Michael F Hagan. Flagellar dynamics of a connected chain of active, polar,

brownian particles. *Journal of The Royal Society Interface*, 11(92):20130884, 2014.

[12] Pavel Chvykov and Jeremy England. Least-rattling feedback from strong time-scale separation. *Phys. Rev. E*, 97:032115, Mar 2018.

[13] Laurent Corte, PM Chaikin, Jerry P Gollub, and DJ Pine. Random organization in periodically driven systems. *Nature Physics*, 4(5):420–424, 2008.

[14] Gavin E Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E*, 60(3):2721, 1999.

[15] Mark C Cross and Pierre C Hohenberg. Pattern formation outside of equilibrium. *Reviews of modern physics*, 65(3):851, 1993.

[16] Leticia F Cugliandolo, Jorge Kurchan, and Luca Peliti. Energy flow, partial equilibration, and effective temperatures in systems with slow dynamics. *Physical Review E*, 55(4):3898, 1997.

[17] Wenping Cui, Robert Marsland III, and Pankaj Mehta. Diverse communities behave like typical random ecosystems. *arXiv preprint arXiv:1904.02610*, 2019.

[18] Purushottam D Dixit, Jason Wagoner, Corey Weistuch, Steve Pressé, Kingshuk Ghosh, and Ken A Dill. Perspective: Maximum caliber is a general variational principle for dynamical systems. *The Journal of chemical physics*, 148(1):010901, 2018.

[19] Luca DâĂŹAlessio, Yariv Kafri, and Anatoli Polkovnikov. Negative mass corrections in a dissipative stochastic environment. *Journal of Statistical Mechanics: Theory and Experiment*, 2016(2):023105, 2016.

[20] Jeremy L England. Statistical physics of self-replication. *The Journal of chemical physics*, 139(12):09B623_1, 2013.

[21] Mohammadjavad Faraji, Kerstin Preuschoff, and Wulfram Gerstner. Balancing new against old information: The role of surprise in learning. *arXiv preprint arXiv:1606.05642*, 2016.

[22] Ulrike Feudel and Celso Grebogi. Multistability and the control of complexity. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 7(4):597–604, 1997.

[23] Ulrike Feudel, Celso Grebogi, Brian R Hunt, and James A Yorke. Map with more than 100 coexisting low-period periodic attractors. *Physical Review E*, 54(1):71, 1996.

[24] Richard Phillips Feynman and Frank Lee Vernon. The theory of a general quantum system interacting with a linear dissipative system. *Annals of physics*, 24:118–173, 1963.

[25] Karl Friston. Life as we know it. *Journal of the Royal Society Interface*, 10(86):20130475, 2013.

[26] Crispin W Gardiner. *Handbook of stochastic methods for physics, chemistry and the natural sciences, vol. 13 of.* 1985.

[27] Todd R Gingrich, Jordan M Horowitz, Nikolay Perunov, and Jeremy L England. Dissipation bounds all steady-state current fluctuations. *Physical review letters*, 116(12):120601, 2016.

[28] AY Grosberg and J-F Joanny. Nonequilibrium statistical mechanics of mixtures of particles in contact with different thermostats. *Physical Review E*, 92(3):032118, 2015.

[29] Li Han and Lee Rudolph. Explicit parametrizations of the configuration spaces of anthropomorphic multi-linkage systems. In *Robotics: Science and Systems*, 2009.

[30] Takahiro Hatano and Shin-ichi Sasa. Steady-state thermodynamics of langevin systems. *Physical review letters*, 86(16):3463, 2001.

[31] Jordan M Horowitz, Kevin Zhou, and Jeremy L England. Minimum energetic cost to maintain a target nonequilibrium state. *Physical Review E*, 95(4):042102, 2017.

[32] Christopher Jarzynski. Thermalization of a brownian particle via coupling to low-dimensional chaos. *Physical review letters*, 74(15):2937, 1995.

[33] Christopher Jarzynski. Nonequilibrium equality for free energy differences. *Physical Review Letters*, 78(14):2690, 1997.

[34] Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.

[35] Mehran Kardar. *Statistical physics of fields*. Cambridge University Press, 2007.

[36] Axel Kleidon and Ralph D Lorenz. *Non-equilibrium thermodynamics and the production of entropy: life, earth, and beyond*. Springer Science & Business Media, 2004.

[37] Suso Kraut, Ulrike Feudel, and Celso Grebogi. Preference of attractors in noisy multistable systems. *Physical Review E*, 59(5):5253, 1999.

[38] Rolf Landauer. Inadequacy of entropy and entropy derivatives in characterizing the steady state. *Physical Review A*, 12(2):636, 1975.

[39] Eric Lutz. Fractional langevin equation. *Physical Review E*, 64(5):051106, 2001.

[40] Benjamin B Machta. Dissipation bound for thermodynamic control. *Physical review letters*, 115(26):260603, 2015.

[41] Christian Maes and Stefano Steffenoni. Friction and noise for a probe in a nonequilibrium fluid. *Physical Review E*, 91(2):022128, 2015.

[42] Martin P Magiera and Lothar Brendel. Trapping of interacting propelled colloidal particles in inhomogeneous media. *Physical Review E*, 92(1):012304, 2015.

[43] M Cristina Marchetti, Jean-François Joanny, Sriram Ramaswamy, Tanniemola B Liverpool, Jacques Prost, Madan Rao, and R Aditi Simha. Hydrodynamics of soft active matter. *Reviews of Modern Physics*, 85(3):1143, 2013.

[44] Samuel R. McCandlish, Aparna Baskaran, and Michael F. Hagan. Spontaneous segregation of self-propelled particles with different motilities. *Soft Matter*, 8:2527, Jan 2012.

[45] Leenoy Meshulam, Jeffrey L Gauthier, Carlos D Brody, David W Tank, and William Bialek. Collective behavior of place and non-place neurons in the hippocampal network. *Neuron*, 96(5):1178–1191, 2017.

[46] Thierry Mora, Aleksandra M Walczak, William Bialek, and Curtis G Callan. Maximum entropy models for antibody diversity. *Proceedings of the National Academy of Sciences*, 107(12):5405–5410, 2010.

[47] Harold Morowitz and Eric Smith. Energy flow and the organization of life. *Complexity*, 13(1):51–59, 2007.

[48] Mark EJ Newman, Duncan J Watts, and Steven H Strogatz. Random graph models of social networks. *Proceedings of the National Academy of Sciences*, 99(suppl 1):2566–2572, 2002.

[49] Nikolai Nikola, Alexandre P Solon, Yariv Kafri, Mehran Kardar, Julien Tailleur, and Raphaël Voituriez. Active particles with soft and curved walls: Equation of state, ratchets, and instabilities. *Physical Review Letters*, 117(9):098001, 2016.

[50] Ole Peters. Optimal leverage from non-ergodicity. *Quantitative Finance*, 11(11):1593–1602, 2011.

[51] Ole Peters and Alexander Adamou. The evolutionary advantage of cooperation. *arXiv preprint arXiv:1506.03414*, 2015.

[52] Ole Peters and William Klein. Ergodicity breaking in geometric brownian motion. *Physical review letters*, 110(10):100603, 2013.

[53] Alexander N Pisarchik and Ulrike Feudel. Control of multistability. *Physics Reports*, 540(4):167–218, 2014.

[54] Sriram Ramaswamy. The mechanics and statistics of active matter. *Annu. Rev. Condens. Matter Phys.*, 1(1):323–345, 2010.

[55] Gabriel S Redner, Michael F Hagan, and Aparna Baskaran. Structure and dynamics of a phase-separating active colloidal fluid. *Physical review letters*, 110(5):055701, 2013.

[56] Michael Rubenstein, Alejandro Cornejo, and Radhika Nagpal. Programmable self-assembly in a thousand-robot swarm. *Science*, 345(6198):795–799, 2014.

[57] Shin-ichi Sasa. Collective dynamics from stochastic thermodynamics. *New Journal of Physics*, 17(4):045024, 2015.

[58] Volker Schaller, Christoph Weber, Christine Semmrich, Erwin Frey, and Andreas R Bausch. Polar patterns of driven filaments. *Nature*, 467(7311):73–77, 2010.

[59] Mark J. Schnitzer. Theory of continuum random walks and application to chemotaxis. *Phys. Rev. E*, 48:2553–2568, Oct 1993.

[60] Philipp Schwartenbeck, Thomas H. B. FitzGerald, Christoph Mathys, Ray Dolan, Martin Kronbichler, and Karl Friston. Evidence for surprise minimization over value maximization in choice behavior. *Scientific Reports*, 5:16575 EP –, 11 2015.

[61] Ya G Sinai. The limiting behavior of a one-dimensional random walk in a random medium. *Theory of Probability & Its Applications*, 27(2):256–268, 1983.

[62] AP Solon and Julien Tailleur. Revisiting the flocking transition using active spins. *Physical review letters*, 111(7):078101, 2013.

[63] Thomas Speck. Stochastic thermodynamics for active matter. *EPL (Europhysics Letters)*, 114(3):30006, may 2016.

[64] Jakub Spiechowicz, Marcin Kostur, and Jerzy Łuczka. Brownian ratchets: How stronger thermal noise can reduce diffusion. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(2):023111, 2017.

[65] Luc Steels. Cooperation between distributed agents through self-organisation. In *EEE International Workshop on Intelligent Robots and Systems, Towards a New Frontier of Applications*, pages 8–14. IEEE, 1990.

[66] Igor Tikhonenkov, Amichay Vardi, James R Anglin, and Doron Cohen. Minimal fokker-planck theory for the thermalization of mesoscopic subsystems. *Physical review letters*, 110(5):050401, 2013.

[67] Mikhail Tikhonov and Remi Monasson. Collective phase in resource competition in a highly diverse ecosystem. *Phys. Rev. Lett.*, 118:048103, Jan 2017.

[68] Matthew J Todd, George H Lorimer, and D Thirumalai. Chaperonin-facilitated protein folding: optimization of rate and yield by an iterative annealing mechanism. *Proceedings of the National Academy of Sciences*, 93(9):4030–4035, 1996.

[69] Mingcheng Yang and Marisol Ripoll. Brownian motion in inhomogeneous suspensions. *Phys. Rev. E*, 87:062110, Jun 2013.

[70] DN Zubarev and AG Bashkirov. Statistical theory of brownian motion in a moving fluid in the presence of a temperature gradient. *Physica*, 39(3):334–340, 1968.

[71] Patrick R Zulkowski, David A Sivak, Gavin E Crooks, and Michael R DeWeese. Geometry of thermodynamic control. *Physical Review E*, 86(4):041148, 2012.

[72] Patrick R Zulkowski, David A Sivak, and Michael R DeWeese. Optimal control of transitions between nonequilibrium steady states. *PloS one*, 8(12):e82754, 2013.