Entropy-production-driven oscillators in simple nonequilibrium networks

Jeffrey K. Weber and Vijay S. Pande

Phys. Rev. E **91**, 032136 — Published 24 March 2015

DOI: 10.1103/PhysRevE.91.032136

Entropy Production-Driven Oscillators in Simple Non-Equilibrium Networks

Jeffrey K. Weber and Vijay S. Pande*

Department of Chemistry, Stanford University, Stanford, California 94305
*Corresponding author. Contact: pande@stanford.edu

Abstract: The development of tractable non-equilibrium simulation methods represents a bottleneck for efforts to describe the functional dynamics that occur within living cells. In this Article, we employ a non-equilibrium approach called the λ-ensemble to characterize the dissipative dynamics of a simple Markovian network driven by an external potential. In the highly dissipative regime brought about by the λ-bias, we observe a dynamical structure characteristic of cellular architectures: the entropy production drives a damped oscillator over state populations in the network. We illustrate the properties of such oscillations in weakly and strongly driven regimes, and we discuss how control structures associated with the "dynamical phase transition" in the system can be related to switches and oscillators in cellular dynamics.

INTRODUCTION

The processes that govern life share one common precept: living things are coupled to and driven by an external input of energy, and therefore operate outside the constraints of equilibrium. In performing the physical work required to grow and reproduce, organisms inevitably waste a portion of the energy they consume. This "dissipative" work, which takes the form of heat evolved into the surrounding bath, is intimately connected to the irreversibility of the system's corresponding functional trajectories and the production of entropy [1]. Theoretical results from Evans, Searles, Gallavotti, Cohen, and their successors over the past two decades describe deep and powerful connections between entropy production and time reversibility and dissipative work and equilibrium free energies [2–9].

Despite such theoretical advances, methods for conducting rigorous non-equilibrium simulations lag behind the scope and impact of potential applications. The ability to reliably modulate and engineer cellular processes would revolutionize molecular and medical biology [10]. Proficiency with non-equilibrium control systems would also prove highly useful in designing synthetic constructs of mechanical or self-assembly processes [11]. Before mastering the design and control of

cellular and other non-equilibrium pathways, however, one first needs to acquire an acute understanding of the nanoscale dissipative processes that drive molecular components involved in the dynamics. Even simulating such functionally relevant non-equilibrium systems, however, presents an enormous challenge. Biological machines are coupled to an exceptionally complicated field of external energy (manifested in chemical potential and interactions with other cellular components); in artificial systems, no such inscrutable blueprint for a driving potential even exists. Connecting mechanical components to their full functional environments requires the continuous manipulation of energy input and dissipation, a task for which standard simulation methods are poorly suited.

A promising means for facilitating such manipulations, however, has emerged from the theory of large deviations: the so-called "λ-ensemble" provides a method by which non-equilibrium trajectories can be sampled according to their degree of entropy production [4–8, 12–15]. By constraining entropy production like temperature constrains energy in canonical statistical mechanics, the λ-parameter brings rare fluctuations that give rise to dissipative processes to the forefront. Coupled to the theory of Markov chains, such an approach provides a means for interrogating driven processes in a general framework. Given the advent of techniques like Markov state models (MSMs), which provide network descriptions of biomolecular dynamics on a grand scale [16–19] the λ-ensemble promises to become a powerful non-equilibrium simulation tool.

Recently, Vaikuntanathan, et al. analyzed entropy production in a selection of simple driven networks using the λ-ensemble approach [14, 15]. Principally, the authors demonstrate a propensity for so-called "entropy production phase transitions" to occur in such systems, and characterize the steady states assumed at low and high entropy production rates. A "localization" to particular dynamical bottlenecks is observed in the low entropy production regime.

In the realm of dynamical systems, the fluctuations and relaxation around steady states are also of great interest. Small impulses to well-designed parameters can force a system to jump to a disparate fixed point (thus flipping a "switch"); perturbations to basic rates can induce a system into cycles around a once-stable fixed point (creating an "oscillator") [20, 21]. Such rudimentary notions of switches and oscillators are ubiquitous in biology, and the simple dynamical principles described here govern the operation of switches and oscillators in gene transcription [22], neuronal dynamics [23], the control of cycles like the cell cycle and circadian rhythm [24, 25], and a myriad of other processes. In a sense, dynamical phase transitions (like that seen in entropy production) represent switches in driven Markovian networks: an impulse in a control parameter (like λ) can be used to flip the system between distinct fixed points [14, 26, 27]. Can conditions be set in driven networks to give rise to oscillating behavior? By manipulating parameters in a simple driven system, we provide a hint at the answer.

In this Article, we construct a triangle network nearly identical to that invoked by Vaikuntanathan, et. al, and we employ tilted transfer operator theory to study the dynamics of entropy

production in the low and high production rate phases. We show that, under the correct conditions, entropy is produced via a damped oscillation mechanism, and populations of states in the system cycle back and forth at a regular period.

THEORY, MODELS, AND METHODS

Taken from large deviation theory, the λ-ensemble is treated like a standard thermodynamic ensemble: relative probability weights for states are calculated and summed into a partition function, and cumulant properties are derived from a free energy analogues. While traditional ensembles place constraints on bulk properties like the energy and volume, the λ-ensemble constrains entropy production rates in trajectories sampled from a system of interest. For an arbitrary trajectory, $x(t_{\text{obs}})$, observed for a given time t_{obs}, the entropy production rate, Ω, is given by [2, 12, 13]

$$\Omega(x(t_{\text{obs}})) = \frac{1}{t_{\text{obs}}}\left[\frac{p_{\text{forward}}}{p_{\text{reverse}}}\right] \tag{1}$$

where p_{forward} and p_{reverse} are the forward and reverse probabilities for the microscopic trajectory of interest. In the framework of Markov chains, one defines an $N \times N$ matrix of transition probabilities, $\mathbb{T}$, among a network of discrete states. With this matrix in hand, forward and reverse trajectory probabilities are calculated trivially as products of microscopic transition probabilities. In conducting a λ-ensemble analysis on a Markovian system, one can thus derive all needed information from a "tilted" transition matrix, $\mathbb{T}_\lambda$ that obeys the entropy-producing constraints of the scalar field, λ [12, 13]. Exponentiating the λ-multiplied entropy production rate expression, one finds that equilibrium transition probabilities need to be modified as [13, 14]

$$(\mathbb{T}_\lambda)_{ij} = (\mathbb{T})_{ij}^{\lambda} \times (\mathbb{T})_{ji}^{(1-\lambda)} \tag{2}$$

Trajectories sampled from the tilted matrix produce entropy (on average) according to the value to which λ is tuned. The thermodynamic properties of the ensemble are computed from its so-called dynamical partition function, Z_λ: [28, 29]

$$Z_\lambda(t_{\text{obs}}) = \left\langle \mathrm{e}^{-\lambda\Omega} \right\rangle_{t_{\text{obs}}} = \langle \pi | \mathbb{T}_\lambda^{t_{\text{obs}}} | o \rangle \tag{3}$$

where π is the stationary distribution for the equilibrium model and o is a vector containing all unit entries. In principle, any discrete probability distribution can be inserted in place of π to simulate arbitrary initial conditions for the transfer operator $\mathbb{T}_\lambda$.

The corresponding dynamical free energy for entropy production can be calculated as

$$F_{\text{ep}}(\lambda, t_{\text{obs}}) = -\frac{1}{t_{\text{obs}}} \log Z_\lambda(t_{\text{obs}}) \tag{4}$$

The mean entropy production rate as a function of λ is obtained by differentiating the free energy with respect to λ. In the long-time, ergodic limit, one can derive stationary probabilities over Markov states by taking the inner product of left and right eigenvectors corresponding to the largest eigenvalue of the tilted matrix [29]. To ascertain the dynamics of entropy-producing trajectories, however, one must compute time-dependent state probabilities. One employs a transfer matrix theory, detailed in the literature [28, 29], to do so. At a given time t in a trajectory of length t_{obs}, the probability of residing in a given Markov state X is given by

$$p(X, \lambda, t, t_{\text{obs}}) = \frac{1}{Z_\lambda(t_{\text{obs}})} \langle \pi | \mathbb{T}_\lambda^{t_{\text{obs}}-t} \mathbb{M}(X) \mathbb{T}_\lambda^t | o \rangle \tag{5}$$

where $\mathbb{M}$, called a projection matrix, contains only one non-zero entry: $\mathbb{M}(X, X) = 1$. The evolution of probabilities in the driven system can thus be traced over time.

The driven network on which we focus this analysis is very similar to the "triangle network" designed and coined by Vaikuntanathan, et al [14, 15]. Illustrated in Fig. 1, the triangle network consists of an amalgam of three-state cycles connected periodically with a slow, symmetry-breaking link. Trajectories are driven from left to right by an imbalance of rates (exemplified by the parameters x and y) for traversing the small triangular cycles. The presence of directed cycles in the network, of course, gives rise to non-equilibrium dynamics: entropy cannot be produced at steady state without the input of external energy, here transduced through an arbitrary external force that circumvents detailed balance. The symmetry-breaking link is thought to break time translational invariance in the system, giving rise to a dynamical phase transition in entropy production as a function of λ.

To augment the methodology of Vaikuntanathan, et al., we first construct the triangle network from the perspective of transition probabilities, rather than transition rates. To build a transition probability matrix, we assign transition counts between states according to the connectivity and rough edge weights shown in Fig. 1, and we normalize the rows of the resulting count matrix to unity. In particular, we study the dynamics within low and high entropy producing phases for weakly driven and strongly driven networks defined by different parameter sets. Common between the weakly and strongly driven networks, the parameters h, x, and b are set as follows: $h = 0.02$, $x = 1$, and $b = 0.1$. The parameter y is switched between 2 to 20 to define the two distinct driven regimes. To promote numerical stability, unit values were added to the all diagonal entries in the count matrices, and probabilities are computed at a time $t = t_{\text{obs}}/2$ within ensembles of increasingly long trajectories. The reasoning behind this mode of time evolution is founded on considerations related to fluctuation theorems and time non-locality within trajectory ensembles, as is explicated in the Results and Discussion.

In general, it is also possible to define an auxiliary propagator that provides a stochastic model for the dynamics that occur at arbitrary values of the λ-bias. To do so, it is often useful to first

transform the unbiased transition matrix, $\mathbb{T}$, into its rate matrix analogue, $\mathbb{W}$, using simple matrix diagonalization (the eigenvalues of the rate matrix can be computed via the logarithm of the transition matrix eigenvalues). Elements of the tilted rate matrix, $\mathbb{W}_\lambda$, of the λ-ensemble are then computed as [29]

$$(\mathbb{W}_\lambda)_{ij} = (1-\delta_{ij})\left[(\mathbb{W})_{ij}^{\lambda} \times (\mathbb{W})_{ji}^{(1-\lambda)}\right] + \delta_{ij}(\mathbb{W})_{ij} \tag{6}$$

The aforementioned auxiliary stochastic model, $\mathbb{W}_\lambda^{\text{aux}}$, is generated from the dominant eigenvalue of $\mathbb{W}_\lambda$, $\omega_{l,\lambda} \equiv \omega$, and its respective right (transfer operator) and left (propagator) eigenvectors, $\mathbf{u}_{l,\lambda} \equiv u$ and $\mathbf{v}_{l,\lambda} \equiv v$. Applying normalization constraints such that $\langle v|u\rangle = 1 = \langle v|o\rangle$, the auxiliary rate matrix for dynamics at a given value of λ, $\mathbb{W}_\lambda^{\text{aux}}$, is calculated using the diagonal operator $\mathbb{U}_{ii} = u_i$: [29]

$$\mathbb{W}_\lambda^{\text{aux}} = \mathbb{U}\mathbb{W}_\lambda\mathbb{U}^{-1} - \omega \tag{7}$$

The off-diagonal elements of $\mathbb{W}_\lambda^{\text{aux}}$ are thus given by

$$(\mathbb{W}_\lambda^{\text{aux}})_{ij} = u(j)(\mathbb{W}_\lambda)_{ij}u(i) \tag{8}$$

and the complementary diagonal elements of the auxiliary matrix are subject to the constraint of stochasticity:

$$(\mathbb{W}_\lambda^{\text{aux}})_{ii} = -\sum_{j\neq i}(\mathbb{W}_\lambda^{\text{aux}})_{ij} \tag{9}$$

It can be formally shown that $\mathbb{W}_\lambda^{\text{aux}}$ reproduces the dynamics of the λ-ensemble exactly in the steady state regime [29]. In practice, it is also informative to study the slow dynamical processes (embedded in eigenvectors with negative rate matrix eigenvalues) that are generated by the auxiliary stochastic model. The eigenproperties of the auxiliary propagator for the simple networks studied here can be obtained through direct diagonalization of the matrix $\mathbb{W}_\lambda^{\text{aux}}$.

In summary, the above methodology allows one to characterize the time evolution of state probabilities within a stochastic network – in particular when dynamics are subjected to a constraint of enhanced (or suppressed) entropy production. The application of such constraints yields rich behavior in the dynamics of driven triangle networks, as demonstrated below.

RESULTS AND DISCUSSION

The driven network illustrated in Fig. 1 can, of course, be extended to an arbitrary length along its horizontal axis. In the following analysis, we choose to highlight the properties of a relatively short network (containing 50 triangular cycles and 101 discrete states) because the probability structure induced by the triangular cycles is readily evident at small scales. However, for completeness,

aspects of the finite-size scaling of network properties are also noted in the subsequent text and illustrated in supplementary figures.

Fig. 2 illustrates the mean entropy production rate observed in a strongly driven, 50-cycle network, as a function of the control parameter λ. As the abrupt jump in entropy production indicates, we observe a dynamical phase transition as λ is tuned in the neighborhood of its unbiased (unit) value, in a manner consistent with Vaikuntanathan, et. al [14]. Unsurprisingly (and also in agreement with [14]), the phase transition becomes sharper as the number of cycles in the network is increased; the qualitative features of the crossover do not change, though, in networks containing a greater number of states. Fig. 3 shows the scaling of the phase transition over a range of network sizes. Notably, a well known symmetry in the dynamical free energy about $\lambda = 0.5$ implies that the properties of entropy production need only be calculated above or below this λ-value [12–14].

Fig. 4 shows steady state probabilities in an unbiased ($\lambda = 1$), strongly driven model containing 50 cycles. Population is delocalized across states in the middle of the network (in a sawtooth pattern generated by the triangular structure), and a peak of probability density appears as flow is impeded by the symmetry-breaking link. As illustrated in Fig. 5, the absolute magnitude of this peak decays as more and more cycles are added to the network; when scaled by the number of states in the system, however, the relative peak height tends toward a plateau in larger and larger networks.

As λ is modulated away from a unit value, dynamics induced by the entropy production constraints become immediately apparent. The steady state distributions corresponding to the high and low entropy production phases are presented for both the strongly and weakly driven networks in Fig. 6; the rightward peak discussed above is greatly enhanced at low values of dissipation and largely decimated at high levels of dissipation.

In parallel to previous findings, we thus indeed find that probability density becomes highly localized in the ergodic limit of low entropy production phase [14]. Fig. 7 demonstrates the dynamics involved in reaching this steady state on the low side of the dynamical phase transition. Under both weak and strong driving conditions, density in the middle portion of the network quickly evaporates, and populations spike to the left (and to a minute extent, to the right) of the symmetry-breaking connection. With the system localized on both sides of the slow link, little entropy is produced in traversing the triangular cycles from the left to right in the network.

This picture of gentle relaxation to a dynamical steady state, however, changes radically within the high entropy production phase. Fig. 8 catalogues the high entropy-producing population dynamics in the strongly driven network as a function of time. Immediately, such dynamics diverge from the localization paradigm: the plateau in the middle of the network pulls density away from the slow link, resulting in a central peak in the probability distribution after about 75 time steps. This peak, however, is quickly annihilated, and a probability deficit takes its place, leading to population enhancement on both sides of the slow link boundary. The central peak is reestablished after about

250 time steps with a somewhat damped maximum amplitude. This oscillating behavior continues at a regular period of approximately 200 time steps, with substantial peaks appearing the middle of the network for many cycles. Fig. 9 illustrates that, in the weakly driven system, more highly damped oscillatory behavior (with a longer period of approximately 500 time steps) is observed. A longer period is perhaps expected in the weakly driven network, given that left-to-right traversal of the system is, in the absence of bias, slower under weak driving conditions. The dynamics in both cases reach reasonable approximations of a steady state (shown in Fig. 6) before 5000 time steps have been evaluated. The absolute magnitude of the first central peak in the oscillations decays as more cycles are added to the network, as one might expect (see Fig. 10 for an exhibition of finite size effects); when scaled by the state space cardinality, however, the magnitude of the peak again plateaus in larger and larger systems.

On a side note, the reasoning behind our mechanism of time evolution (i.e., varying t_{obs} and fixing $t = t_{\mathrm{obs}}/2$) is worthy of some exposition. The specific choice of the value $t = t_{\mathrm{obs}}/2$ is related to the aforementioned symmetry of observables about $\lambda = 0.5$, a property rooted in the Crooks fluctuation theorem [12]. Based on the structures of Equations 2 and 5, one should observe that extending the time parameter t beyond $t_{\mathrm{obs}}/2$ is equivalent to evaluating the transposition of the transfer operator $\mathbb{T}_\lambda$ at short times. This transposed operator simply represents a momentum-reversed analogue to the original propagator, implying that a time- and state index-reversal of dynamics after $t_{\mathrm{obs}}/2$ will reproduce dynamics seen at $t < t_{\mathrm{obs}}/2$. Accordingly, unique information within the λ-ensemble can only be obtained up to a time of $t = t_{\mathrm{obs}}/2$ within a given trajectory ensemble.

One should next note that, rather than measuring various time points up to $t = t_{\mathrm{obs}}/2$ in an ensemble of very long trajectories, we evaluate different ensembles comprised of successively longer trajectories. To justify this choice, one needs to call upon simple physical arguments and examine their consequences within space-time statistical mechanics. While the physics of entropy production are not time reversible, the processes that result in energy dissipation in real systems are certainly time local. However, the physics of large deviation trajectory ensembles are decidedly not time local: from the perspective of any single, non-boundary point in time, trajectories are sampled based on their past and future states. This dependence on future events, of course, violates our traditional sense of causality. Causality violations in the context of large deviations have been comprehensively discussed in the past [29]; in general, it's agreed that considerations of causality can be safely ignored in the long-time, time-translationally invariant (TTI) regime of dynamics. The limited success, to date, of large deviation fields in biasing physical simulations on the fly can likely be attributed to issues of time non-locality.

In the present work, we do analyze the time evolution of trajectories in the sub-TTI regime, meaning that preservation of causality must be taken into account. From a physical viewpoint, one

would expect short-time dynamics to be agnostic to events in the near and distant future. To satisfy this expectation in large deviation theory, one must analyze ensembles of trajectories of minimal length relative to a given time point of interest. Based on our above discussion of symmetries in the λ-ensemble, that minimal length is $t_{\text{obs}} = 2t$, where t is the desired time for probability evaluation.

Under the paradigm of scanning various time points within a single, long-trajectory ensemble, one still in practice observes oscillations at high levels of dissipation; these alternative dynamics (for the strongly driven network) are now illustrated in the Fig. 11. The observed oscillator has a similar period to that presented in Fig. 8 but exhibits a somewhat different waveform. Since these long-trajectory probability distributions are based on future information about entropy-producing trajectories, however, one should consider them to be less physical than those collected under the $t_{\text{obs}} = 2t$ lens.

The genesis of the observed dissipative oscillations can be manifested in the striking, near-harmonic eigenvectors that appear in the auxiliary model above the dynamical phase transition. Fig. 12 shows the real and imaginary components of the four most dominant conjugate sets of such eigenvectors derived from the strongly-driven form of $\mathbb{W}_{\lambda}^{\text{aux}}$. Each set of eigenvectors is coupled to a complex eigenvalue, which flips in sign as the matrix is propagated to yield cyclic positive and negative contributions to the evolving density. As one might expect, the eigenvectors with the slowest associated time scales (indicated, roughly, by the real component of the related eigenvalue) exhibit the longest spatial wavelengths. The high frequency modes contribute most to the dynamics at short times, when the probability density waveform is flat and broad; the slowest, single-peaked mode comes to dominate the oscillations as time progresses, yielding the smooth probability peaks seen in later panels of Fig. 8. Below the dynamical phase transition, these hallmarks of oscillatory behavior disappear: as demonstrated by Fig. 13, the slowest eigenvectors in the low entropy production regime simply serve to transfer density toward the network's rightward bottleneck.

Throughout the course of oscillations that occur above the phase transition, a nearly constant mean entropy production rate is maintained. In order for this production to be sustained, trajectories must be driven more prominently rightward (on average) in the network over the total duration of an observation window. Trajectories which cause density to shift from right to left, of course, are counterproductive to a high level of dissipation in the dynamics. At large values of the λ-parameter, such motions must be counteracted by more dramatic shifts to the right. An additional, asymmetric probability enhancement (correlated with an eigenvector asymmetry) indeed appears on the left side of the slow link boundary, accounting for the high level of continuous entropy production in the network.

Fig. 14 provides further comparison of the oscillators that emerge under strong and weak driving conditions. One clearly sees the asymmetry that emerges on the far right side of the network: the maximum amplitudes for the extreme rightward states are larger than those in the middle. Oscil-

lations are effectively damped out under weak conditions after about five cycles, while analogous oscillations are far more persistent in the strongly driven case. Both systems evolve at very regular periods, and the shapes of the resulting waveforms are highly correlated with the characteristic periodic time scales.

We thus observe that, upon addition of a symmetry-breaking link to a simple driven network, high entropy production drives an oscillator that transfers population back and forth between extreme regions of the state space. The slow link introduces a separation of time scales in the network dynamics, and coupling this adiabaticity to a strong driving potential (manifested by large values of λ) yields oscillating behavior centered around the boundary. The scheme of adding an adiabatic element to induce oscillations in a driven system is not uncommon [20]; it is intriguing and encouraging, however, that this behavior emerges naturally from the λ-ensemble approach. One could imagine coupling this oscillator among states to a particular operation: if the states in the center of the network and its extremities were connected separate but dependent tasks, the system could (during the course of oscillations) complete these undertakings in sequence and at regular intervals.

In biology, oscillators related to processes like the gene regulation [22], neuronal dynamics [23], the cell cycle [24], and circadian rhythms [25] are perhaps the most prominent. To be successful, however, these types of oscillators need to be "relaxation" oscillators. The period in a relaxation oscillator is largely independent of the waveform's morphology; in the absence of this constraint, a system's residence in certain states (like the arrested state of cell division or the wakeful state of an organism) can thus be long and stable without sacrificing the expediency of transitions between different regimes [20]. The oscillators observed here decidedly do not share this characteristic, as the populations are deformed nearly continuously throughout the course of a cycle. The construction of a reliable relaxation oscillator using the λ-ensemble framework presents an interesting exercise. As a rough comparison, the oscillators observed in the triangle network are closer in character to some seen in developmental biology (perhaps first, and most elegantly, described by Turing [21, 30, 31]), which are responsible for cyclic morphogenesis in organisms. The central and extreme populations in the high entropy production regime of the triangle network are phase locked and decay with time [21]. Though grossly oversimplified for the purposes of this comparison, a similar phase-locked and decaying motif is needed to generate repeating and terminating patterns like stripes and vertebrae in developing creatures [30, 31]. The triangular structure of the network also evokes simple two/three-state models of molecular motors, wherein the driven link in each miniature cycle could serve to represent the power stroke of the machine, while the passive links might describe fluctuations between inactive and active states. Under this interpretation, the network's symmetry-breaking link could simulate an interaction with a spatially localized inhibitory complex, from which the motor might disengage and operate in reverse. At this juncture, such connections to biological systems are admittedly hypothetical; dissipative analyses of more physically rigorous biomolecular models will

certainly be of interest in the future.

CONCLUSION

Above, we have provided evidence that the λ-ensemble approach can capture the dissipative dynamics related to important non-equilibrium control structures like switches and oscillators. By employing simple dynamical principles in the construction of driven networks, we can thus use a λ-bias simulate two of the most archetypical constructions biology has evolved to help living things function and propagate. A converse analysis awaits: given a detailed network model (for example, a Markov state model) of a biological system of interest, can the λ-ensemble be used to characterize the control structures present in the dissipative regime? Though one would need to start small (perhaps on the level of single protein molecules), the results presented here provide a blueprint for accessing the non-equilibrium physics of arbitrary cellular architectures.

We thank the NSF (MCB-0954714) and NIH (R01-GM062868) for their support of this work. J.K.W. was supported by the Fannie and John Hertz Foundation on the endowed Professor Yaser S. Abu-Mostafa Fellowship.

[1] J. L. England, The Journal of Chemical Physics **139**, 121923 (2013).
[2] D. J. Evans and D. J. Searles, Physical Review E **50**, 1645 (1994).
[3] G. Gallavotti and E. Cohen, Physical Review Letters **74**, 2694 (1995).
[4] D. Ruelle, Journal of Statistical Physics **85**, 1 (1996).
[5] C. Maes, Journal of Statistical Physics **95**, 367 (1999).
[6] G. E. Crooks, Physical Review E **61**, 2361 (2000).
[7] J. Kurchan, Journal of Physics A: Mathematical and General **31**, 3719 (1998).
[8] B. Derrida, J. Lebowitz, and E. Speer, Physical Review Letters **87**, 150601 (2001).
[9] C. Jarzynski, Physical Review Letters **78**, 2690 (1997).
[10] D. E. Cameron, C. J. Bashor, and J. J. Collins, Nature Reviews Microbiology **12**, 381 (2014).
[11] B. H. Northrop, Y.-R. Zheng, K.-W. Chi, and P. J. Stang, Accounts of chemical research **42**, 1554 (2009).
[12] H. Touchette, Physics Reports **478**, 1 (2009).
[13] J. L. Lebowitz and H. Spohn, Journal of Statistical Physics **95**, 333 (1999).
[14] S. Vaikuntanathan, T. R. Gingrich, and P. L. Geissler, Physical Review E **89**, 062108 (2014).
[15] T. R. Gingrich, S. Vaikuntanathan, and P. L. Geissler, arXiv preprint arXiv:1406.3311 (2014).
[16] J.-H. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. D. Chodera, C. Schütte, and F. Noé, The Journal of Chemical Physics **134**, 174105 (2011).
[17] V. S. Pande, K. Beauchamp, and G. R. Bowman, Methods **52**, 99 (2010).

[18] K. A. Beauchamp, R. McGibbon, Y.-S. Lin, and V. S. Pande, Proceedings of the National Academy of Sciences **109**, 17807 (2012).
[19] D. Shukla, Y. Meng, B. Roux, and V. S. Pande, Nat. Commun. **5** (2014).
[20] A. Pikovsky, M. Rosenblum, and J. Kurths, Self **2**, 3 (2001).
[21] A. Koseska, E. Volkov, and J. Kurths, Physics Reports **531**, 173 (2013).
[22] J. Hasty, D. McMillen, and J. J. Collins, Nature **420**, 224 (2002).
[23] G. Buzsáki and A. Draguhn, Science **304**, 1926 (2004).
[24] J. R. Pomerening, E. D. Sontag, and J. E. Ferrell, Nature Cell Biology **5**, 346 (2003).
[25] M. J. Rust, J. S. Markson, W. S. Lane, D. S. Fisher, and E. K. O'Shea, Science **318**, 809 (2007).
[26] L. O. Hedges, R. L. Jack, J. P. Garrahan, and D. Chandler, Science **323**, 1309 (2009).
[27] J. K. Weber, R. L. Jack, and V. S. Pande, Journal of the American Chemical Society **135**, 5501 (2013).
[28] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, Physical review letters **98**, 195702 (2007).
[29] R. L. Jack and P. Sollich, Progress of Theoretical Physics Supplement **184**, 304 (2010).
[30] A. M. Turing, Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences **237**, 37 (1952).
[31] C. J. Tomlin and J. D. Axelrod, Nature Reviews Genetics **8**, 331 (2007).

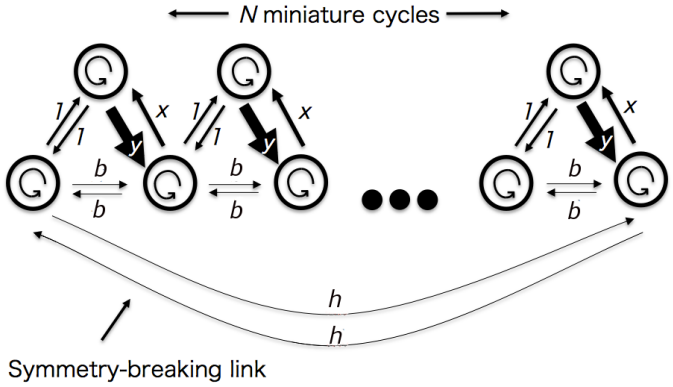

FIG. 1. Schematic illustration of the non-equilibrium triangle network. Trajectories are driven around the total cycle according to an imbalance between rates x and y, which govern the traversal of N miniature, triangular cycles. The symmetry-breaking link (parameterized by a relatively slow rate, σ) between the right and gives rise to interesting behavior in the observed entropy production. Rates are translated into transition counts for the transition probability framework; several self-transition counts are added to each node to improve numerical stability.

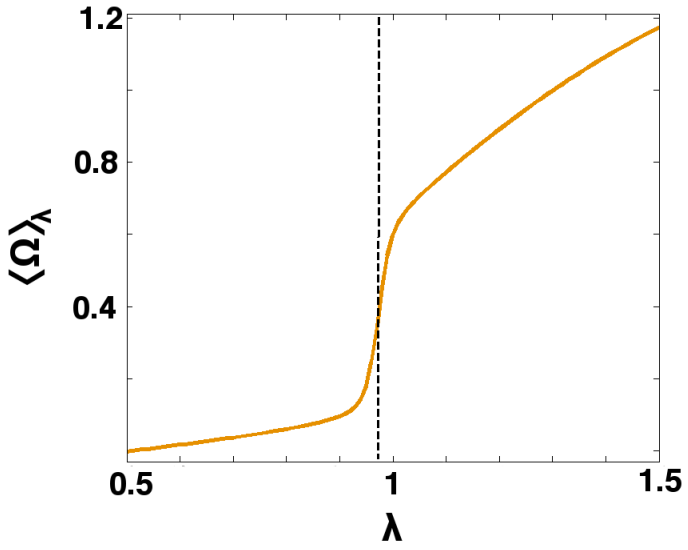

FIG. 2. (Color online) Mean entropy production rate, Ω, as a function of the biasing parameter λ in a 50-cycle triangle network. A dynamical phase transition (characterized by the sharp jump in the entropy production rate near $\lambda = 1$) separates two distinct regimes of dissipative behavior.

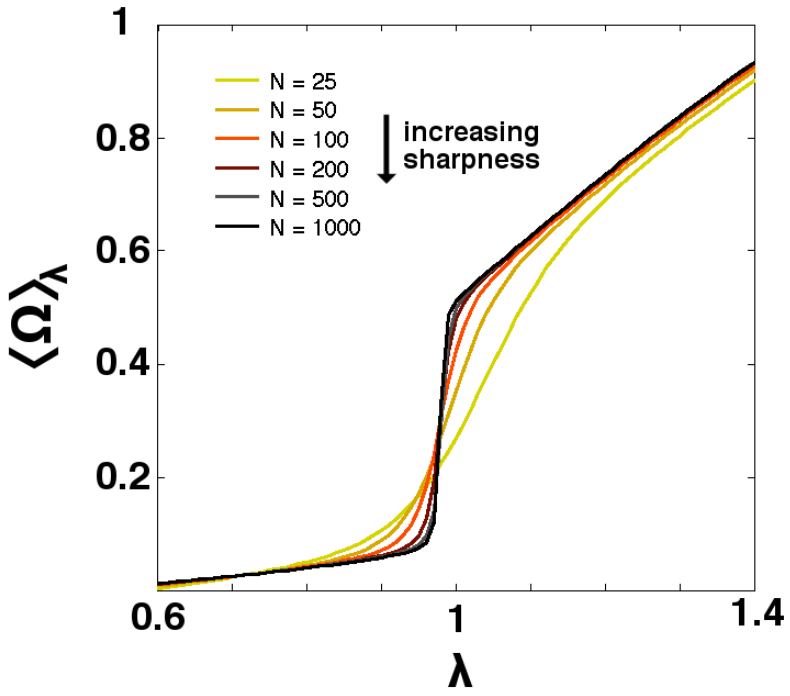

FIG. 3. (Color online) Finite size scaling of the mean entropy production rate as a function of λ in triangle networks. As the number of cycles in the network, N, is increased, the dynamical phase transition grows uniformly sharper but maintains its qualitative characteristics. Since mean entropy production rates vary slightly as $t_{\rm obs}$ is modulated, $\langle\Omega\rangle$ is calculated at $t_{\rm obs} = 2N$ for purposes of comparison.

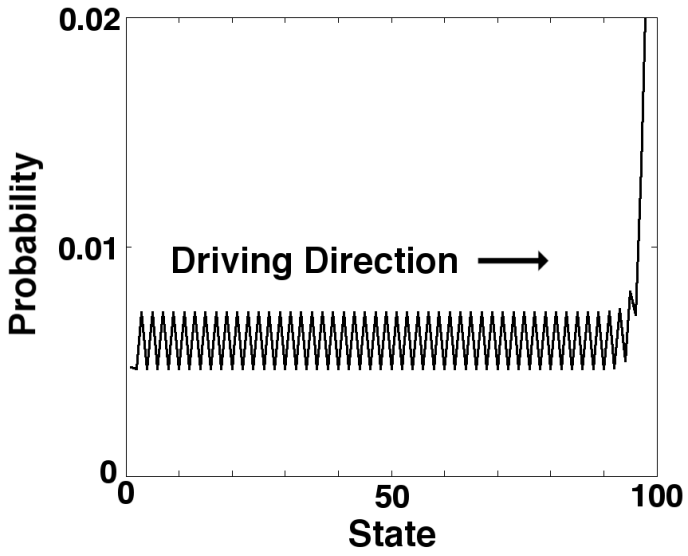

FIG. 4. Steady state probability distribution for the unbiased, strongly-driven triangle network. The peak in probability at the extreme right (near the symmetry-breaking link) is truncated to preserve scale. States are indexed according to a left-to-right orientation of Figure 1, where the slow link connects states 1 and 101. No dynamics as a function of trajectory length are observed in this unbiased case.

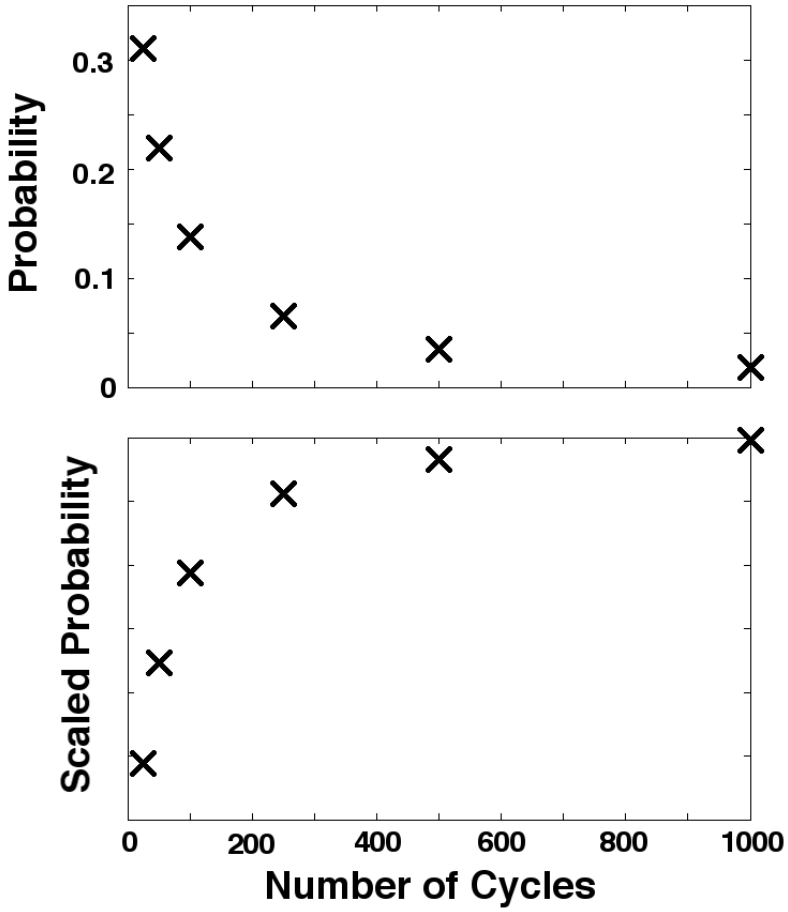

FIG. 5. Finite size scaling of the steady-state, "bottleneck" probability peak height (defined as the probability of the right-most state) in unbiased triangle networks. In the lower plot, probabilities are scaled by the number of states in each respective network; all values shown correspond to strong driving conditions.

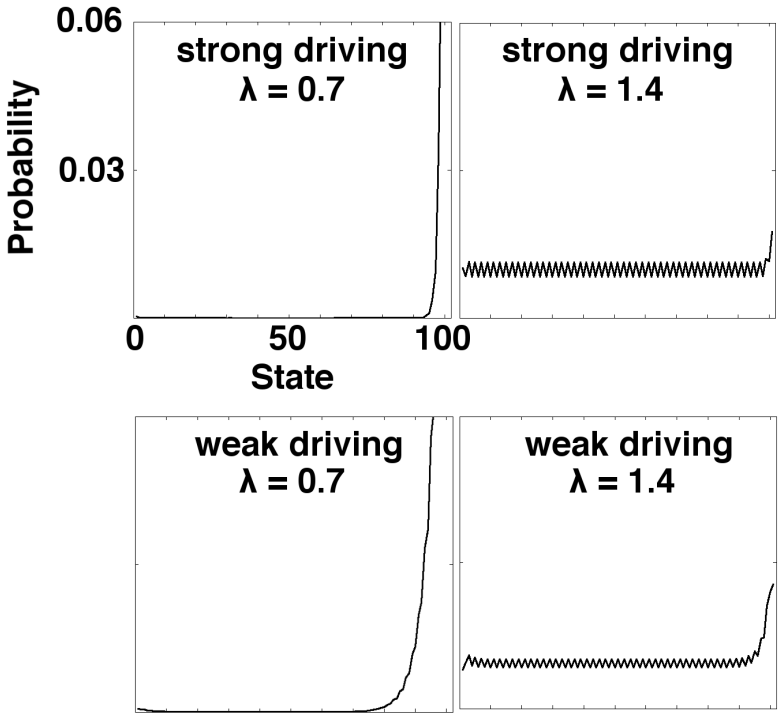

FIG. 6. Steady state probability distributions for the four λ-biased networks analyzed in the main text. As always, the driving direction from left-to-right in each of the subplots. The curves shown here describe probabilities in the long-time, ergodic limit; the dynamical processes illustrated in Figs. 4 - 6 converge to these steady states at respective low and high levels of entropy production.

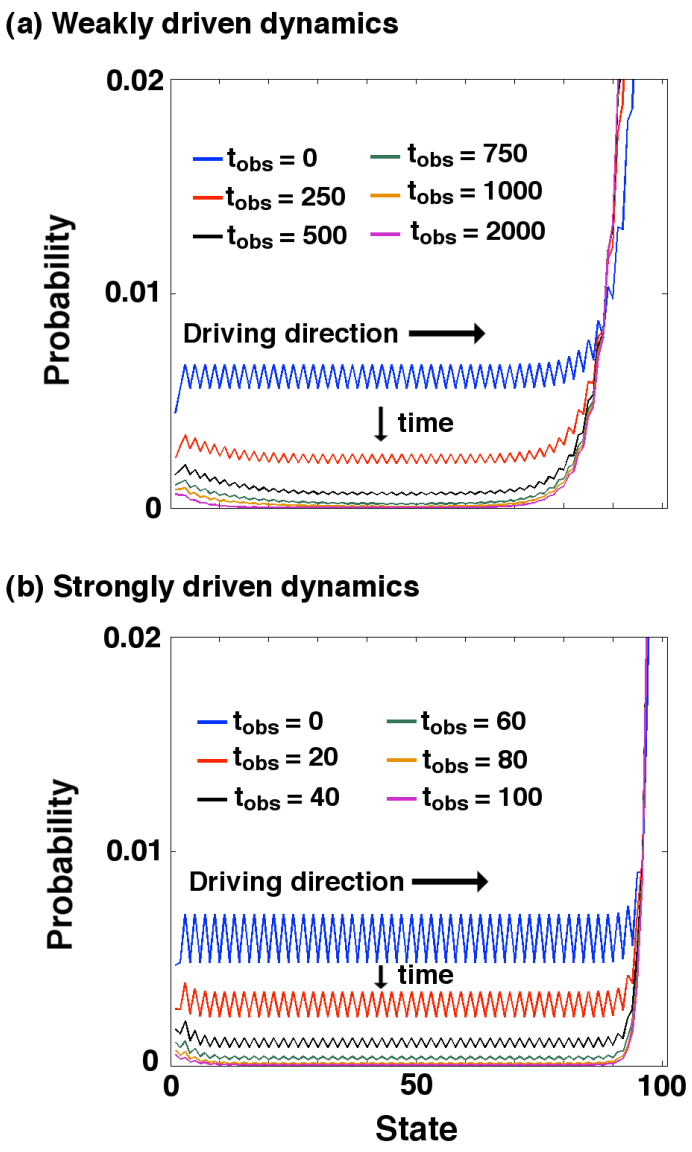

FIG. 7. (Color online) Dynamics within the low entropy production phases (at $\lambda = 0.7$) observed for the **(a)** weakly and **(b)** strongly driven networks. In both cases, the dynamics describe a "localization transition" onto states near the slow-link boundary. Probabilities are calculated for times t according to $t = t_{\rm obs}/2$, and the probability spike at the right side of the network is truncated to preserve scale. Probabilities in the central states decay monotonically as time advances.

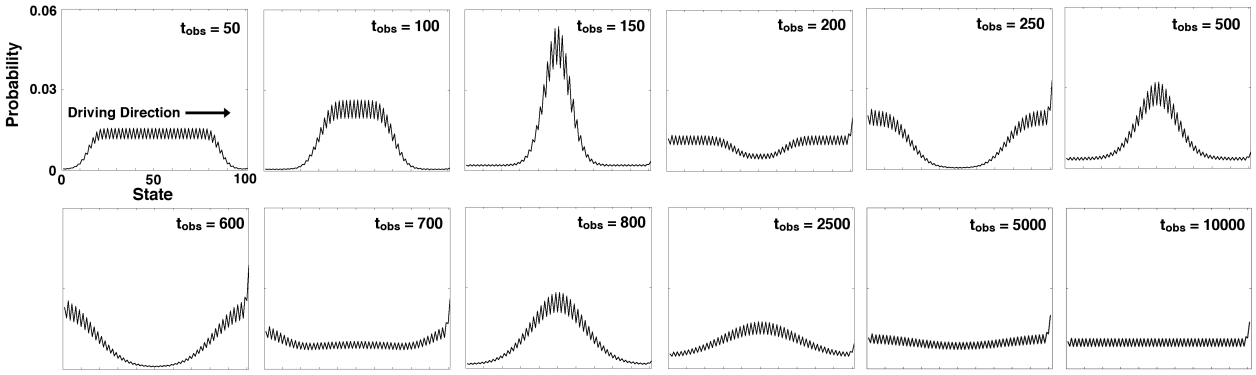

FIG. 8. Dissipative dynamics in the high entropy production phase ($\lambda = 1.4$) for the strongly driven triangle network. Once more, probabilities are calculated for times $t = t_{\rm obs}/2$. Oscillations persist over the course of the twelve frames; significant central peaks in probability still appear after seven cycles, as illustrated in the final frame.

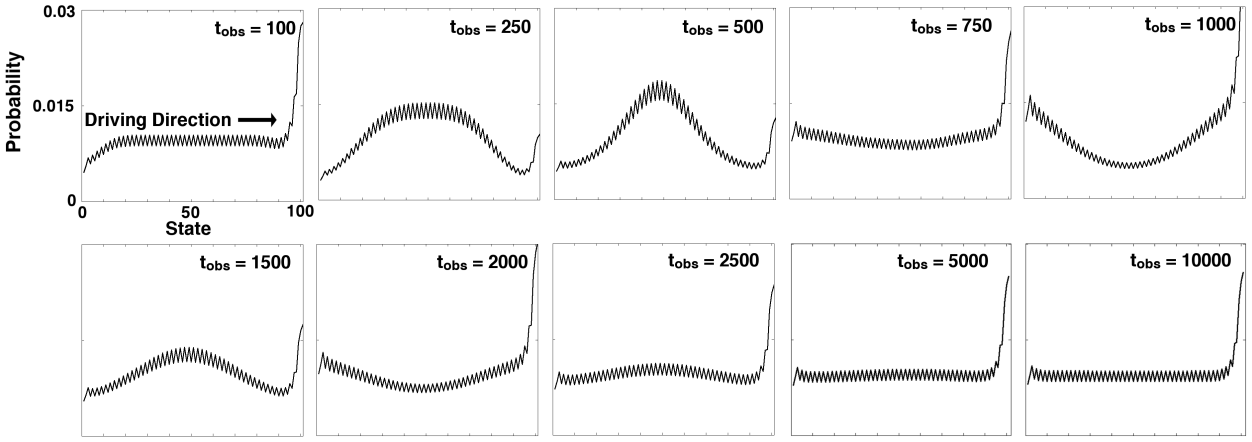

FIG. 9. Dissipative dynamics in the high entropy production phase ($\lambda = 1.4$) for the weakly driven triangle network. Once more, probabilities are calculated for times $t = t_{\rm obs}/2$. Oscillations are damped more quickly and occur over a longer period than in the strongly driven case.

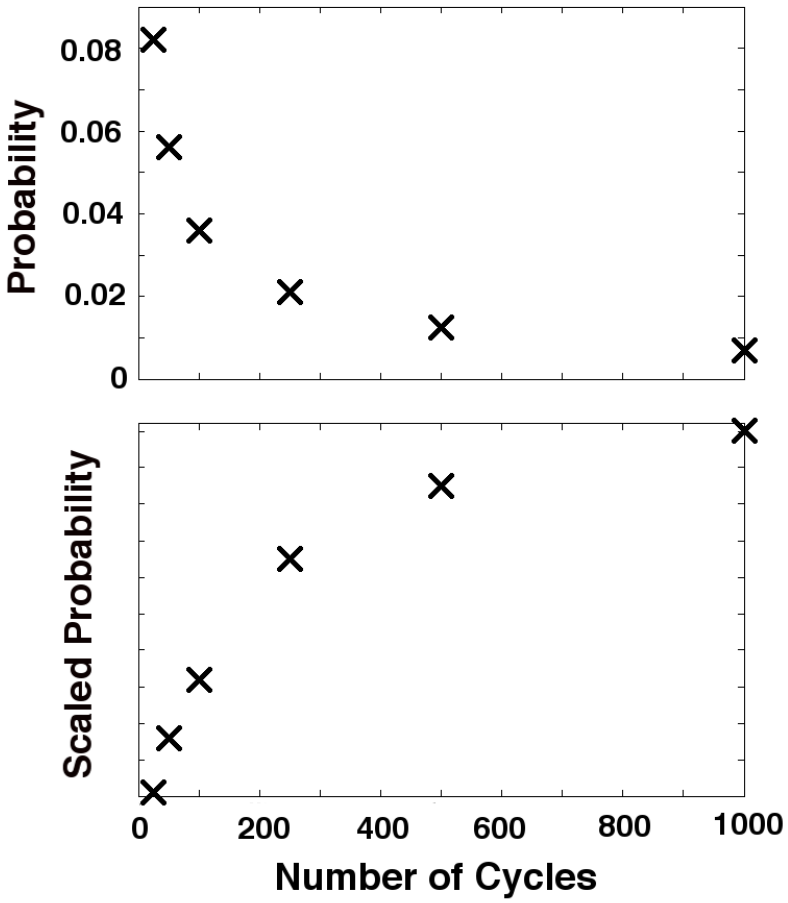

FIG. 10. Finite size scaling of the first central probability peak height in oscillations observed under the constraint of high entropy production ($\lambda = 1.4$) and under strong driving conditiions. In the lower plot, probabilities are once more scaled by state cardinalities.

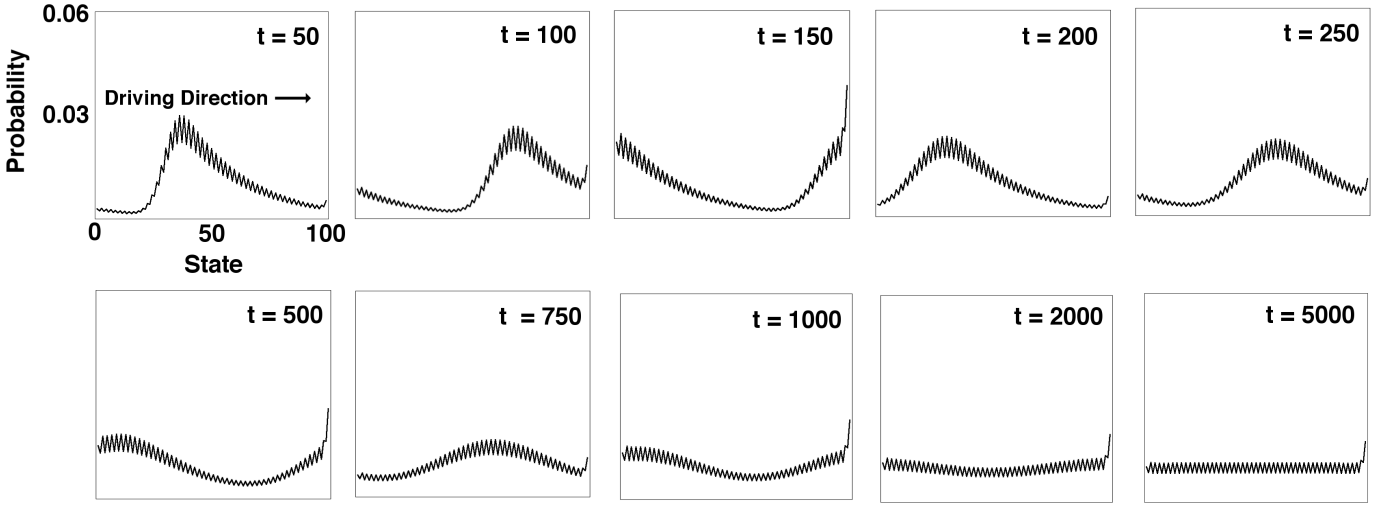

FIG. 11. Probability oscillations measured explicitly as a function of time within a single ensemble of long ($t_{\rm obs} = 10000$) trajectories at $\lambda = 1.4$ and under strong driving conditions. Deviating from the $t = t_{\rm obs}/2$ convention in the λ-ensemble yields results biased by future outcomes; in this case, oscillation waveforms are rightward-leaning compared to the curves presented in the main text. The period and peak amplitude of this oscillator (approximately 200 time steps and 0.03, respectively), however, are very similar to values extracted from time points at $t_{\rm obs}/2$.

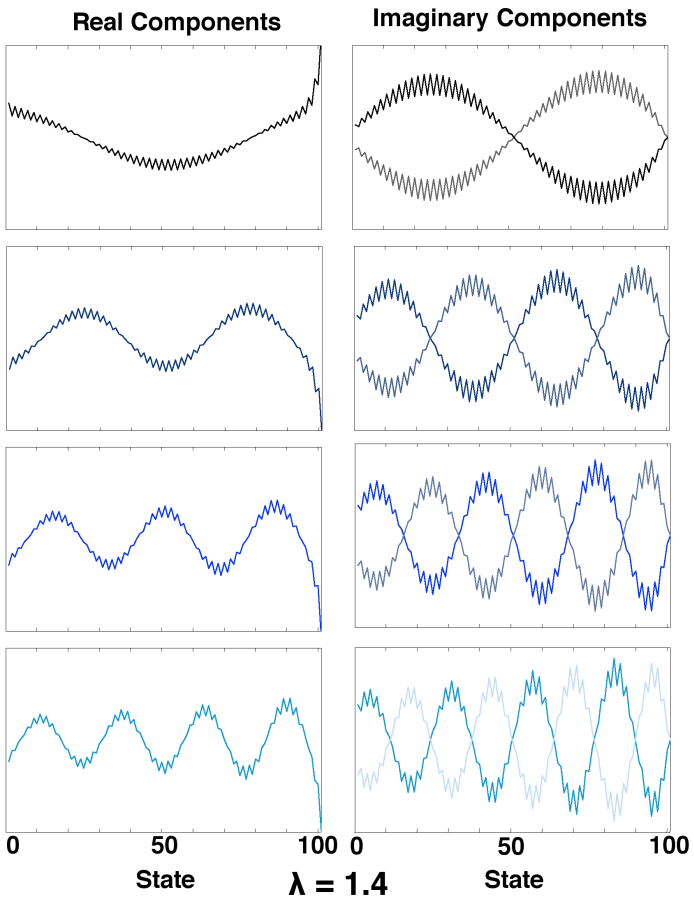

FIG. 12. (Color online) Dominant eigenvectors corresponding to the auxiliary operator, $\mathbb{W}_{\lambda}^{\mathrm{aux}}$, in the high entropy production phase. Real eigenvector components are illustrated at left, and imaginary eigenvector components are shown at right; each couple of vectors coincides with a complement of complex eigenvalues, appearing in complex conjugate pairs. The eigenvectors presented here exhibit the four largest real eigenvalue components that contribute to the auxiliary dynamics. These eigenvectors collectively form a set of near-harmonic modes that describe the oscillatory behavior seen at high levels of mean entropy production. In particular, the slowest mode (top) can be seen in the evolution of dynamics in Figure 4.

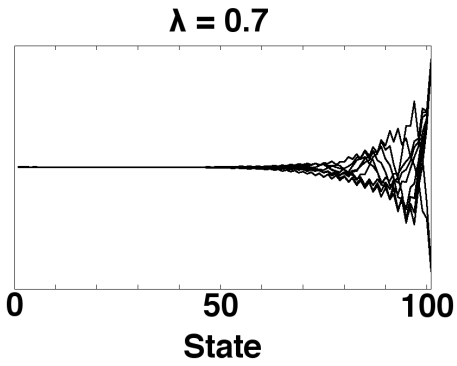

FIG. 13. Real components of dominant eigenvectors corresponding to the auxiliary operator, $\mathbb{W}_{\lambda}^{\mathrm{aux}}$, in the low entropy production phase. While some fine structure appears near the right side of the cycle, most oscillatory character observed above the phase transition disappears in the low entropy production phase. Below the phase transition, the dominant eigenvectors serve to transfer density toward the network's bottleneck in a nearly uniform fashion. Imaginary eigenvector components are also localized to the left side of the bottleneck.

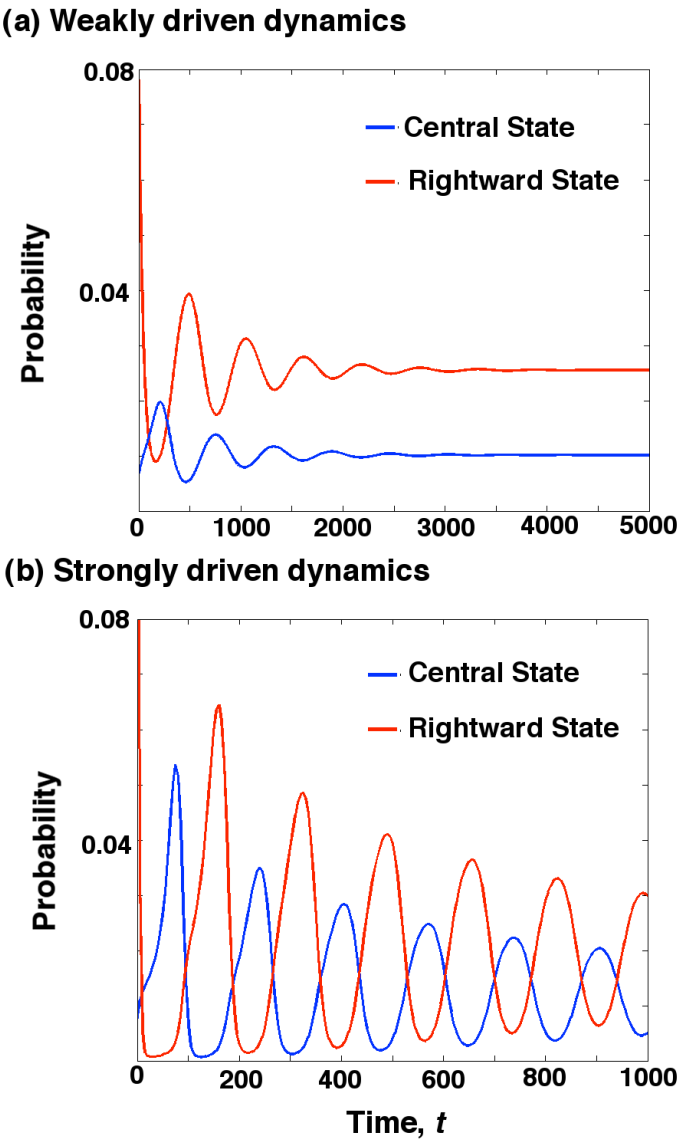

FIG. 14. (Color online) Illustration of probability oscillations over time in the middle and extreme right portions of the network. In both the **(a)** weakly driven and **(b)** strongly driven cases, oscillations follow a regular period, with peak probability amplitudes for extreme rightward states (red, higher amplitude) exceeding the peak amplitudes for central state (blue, lower amplitude) probabilities. In both cases, the two probabilities are phase-locked in opposition and evolve over time according to a smooth waveform.