

Article

A Drive towards Thermodynamic Efficiency for Dissipative Structures in Chemical Reaction Networks

Kai Ueltzhöffer ^{1,2,*}, Lancelot Da Costa ^{1,3} , Daniela Cialfi ⁴  and Karl Friston ¹ 

¹ Wellcome Centre for Human Neuroimaging, Queen Square Institute of Neurology, University College London, London WC1N 3AR, UK; l.da-costa@imperial.ac.uk (L.D.C.); k.friston@ucl.ac.uk (K.F.)

² Department of General Psychiatry, Center of Psychosocial Medicine, Heidelberg University, 69115 Heidelberg, Germany

³ Department of Mathematics, Imperial College London, London SW7 2AZ, UK

⁴ Department of Philosophical, Pedagogical and Economic-Quantitative Sciences, Economic and Quantitative Methods Section, University of Studies Gabriele d'Annunzio Chieti-Pescara, 65127 Pescara, Italy; daniela.cialfi@unich.it

* Correspondence: kai.uelthoeffer@med.uni-heidelberg.de

Abstract: Dissipative accounts of structure formation show that the self-organisation of complex structures is thermodynamically favoured, whenever these structures dissipate free energy that could not be accessed otherwise. These structures therefore open transition channels for the state of the universe to move from a frustrated, metastable state to another metastable state of higher entropy. However, these accounts apply as well to relatively simple, dissipative systems, such as convection cells, hurricanes, candle flames, lightning strikes, or mechanical cracks, as they do to complex biological systems. Conversely, interesting computational properties—that characterize complex biological systems, such as efficient, predictive representations of environmental dynamics—can be linked to the thermodynamic efficiency of underlying physical processes. However, the potential mechanisms that underwrite the selection of dissipative structures with thermodynamically efficient subprocesses is not completely understood. We address these mechanisms by explaining how bifurcation-based, work-harvesting processes—required to sustain complex dissipative structures—might be driven towards thermodynamic efficiency. We first demonstrate a simple mechanism that leads to self-selection of efficient dissipative structures in a stochastic chemical reaction network, when the dissipated driving chemical potential difference is decreased. We then discuss how such a drive can emerge naturally in a hierarchy of self-similar dissipative structures, each feeding on the dissipative structures of a previous level, when moving away from the initial, driving disequilibrium.

Keywords: stochastic thermodynamics; dissipative structures; thermodynamic efficiency; chemical reaction networks



Citation: Ueltzhöffer, K.; Da Costa, L.; Cialfi, D.; Friston, K. A Drive towards Thermodynamic Efficiency for Dissipative Structures in Chemical Reaction Networks. *Entropy* **2021**, *23*, 1115. <https://doi.org/10.3390/e23091115>

Academic Editors: David Wolpert and Jessica Flack

Received: 13 July 2021

Accepted: 25 August 2021

Published: 27 August 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

We start by briefly reviewing the role of dissipation in self-organisation on one hand, and the role of thermodynamic efficiency in the emergence of interesting computational properties—a hallmark of biological systems—on the other. We end by highlighting a small explanatory gap, namely how self-organising dissipative processes might be driven towards thermodynamic efficiency. We will then outline—through a heuristic argument substantiated by simulations of a stochastic chemical reaction network—a simple mechanism that pressures certain dissipative structures to become thermodynamically efficient.

Contrary to the—still widely held—belief that life is a struggle against the second law of thermodynamics, recent advances in nonequilibrium thermodynamics successfully recast biological systems as a subclass of dissipative structures. The formation of such dissipative structures is statistically favoured by (generalizations of) the second law of

thermodynamics, because their existence enables the dissipation of reservoirs of free energy, which could not be accessed otherwise. Therefore, their formation facilitates the irreversible relaxation of the associated disequilibria [1–5]. In other words, dissipative structures constitute channels in the universe’s highly structured state space, which enable transitions from one frustrated, metastable state to another metastable state of higher entropy [6]. This line of thinking dates back at least to the work of Lotka, who tried to relate natural selection to a physical principle of maximum energy transformation [7,8]. Dissipative structure formation has been well understood for many systems in the near-equilibrium, linear-response regime, due to the work of Prigogine and colleagues in the 1960s and 1970s [9], leading to the notion of biological systems as a class of self-organising free energy-conversion engines [10]. However, it took several decades until thermodynamic equalities were derived that hold for small systems arbitrarily far from equilibrium. These equalities take the form of fluctuation theorems, which relate the entropy produced by a microscopic forward trajectory with the probability ratio of observing the forward, versus the time-reversed backward trajectory [11,12]. These fluctuation theorems generalized the relationship between entropy increase and irreversibility from the regime of macroscopic, closed systems at equilibrium—i.e., the second law of thermodynamics—to microscopic, open systems arbitrarily far from equilibrium. For introductory surveys of the resulting field of stochastic thermodynamics see, for example, [13,14], for a comprehensive review see [15]. Recently, dissipative self-organisation of macroscopic structures far from equilibrium was formalized through a macroscopic coarse-graining of such a microscopic fluctuation theorem [16]. This work relates the finite-time transition probabilities from one macroscopic state to two possible outcome states not only to the energy of the final states—via an equilibrium-like Boltzmann-term—and the kinetic barriers separating the initial from the potential finite states, but crucially also to the amount of dissipated heat during the transitions. Thus, energy levels and kinetic accessibility of potential final states being equal, outcome states with a higher dissipative history are favoured, leading to the coining of the term dissipative adaptation for such a selection process [16].

Taken together, this line of research provides a powerful account of how highly ordered nonequilibrium systems can emerge from basic thermodynamic principles, namely fluctuation theorems in stochastic thermodynamics and their special limiting case, the second law of thermodynamics (combined with an initial state of low entropy and a highly structured state space). However, in principle these accounts apply as well to relatively simple, dissipative systems, such as convection cells, hurricanes, candle flames, lightning strikes, or mechanical cracks, as they do to complex biological systems.

One important difference between life and non-life is the role of information processing in the ongoing physical processes [17–19]. Although simple dissipative systems, such as convection cells, are governed by thermodynamic constraints and fluxes flowing through them, living systems contain large amounts of information (e.g., stored in a cell’s DNA and epigenome), which—in stark contrast to simpler systems—govern and structure the thermodynamic fluxes through them. What kind of drives or pressures could facilitate the transition from a simple thermodynamic to an information-governed regime in dissipative systems? Two lines of work show that thermodynamic efficiency might play a role in this transition: In stochastic thermodynamics, the dissipated heat in a thermodynamic system, which is driven by a time-varying potential, upper bounds the system’s non-predictive information about the time-dependent drive [20]. Thus, to minimize the dissipated heat during the work extraction process, the system must develop an efficient, predictive representation of the driving environmental dynamics. From the requisite predictive dynamics, which can be linked to an information bottleneck [21], agent-like features, such as action policies that balance exploration and exploitation [22], and curiosity driven reinforcement-learning [23] can be derived (c.f., [24] for a synthesis). In computational neuroscience, adding energetic costs to the objective function minimized by simulated neural networks leads to the emergence of computational properties such as distributed

neural codes, sparsity of representations, stochasticity, and the heterogeneity of neural populations [25–32] (c.f., [33] for a review).

It is still not fully understood, under which constraints thermodynamically efficient processes emerge during dissipative self-organisation. This work attempts to address this explanatory gap.

The results in this work are based on simulations and analyses of the dynamics and thermodynamics of stochastic chemical reaction networks. The application of stochastic thermodynamics to chemical reaction networks is a relatively young [34–37], and active field of research [38]. Recent works covered advances in energy-efficient dissipative chemical synthesis [39], and provided important extensions from homogeneous mixtures, ideal, and elementary chemical reactions, to spatially extended reaction-diffusion systems [40,41], non-ideal [42], and non-elementary chemical reaction networks [43]. Furthermore, thermodynamically consistent coarse-graining rules for chemical reaction networks were derived only recently [44]. These results might carry special importance for understanding large, hierarchical, and complex biochemical networks. Yet, in this work we call only on results concerning the nonequilibrium thermodynamics of spatially homogeneous mixtures, and ideal, elementary reactions in deterministic and stochastic chemical reaction networks [45,46].

The underlying idea is that in certain emergent, self-similar hierarchies of dissipative structures within a stochastic chemical reaction network, namely those without concentration processes, the available chemical driving potential decreases at each subsequent layer in the hierarchy, while the minimum work-rate required to sustain a new layer of dissipative structures might not. Thus, there might be a minimum viable efficiency of the thermodynamic processes, which must harvest the work required to maintain the associated dissipative structure from the conducted free energy fluxes. This pressure would increase with increasing hierarchical distance of the dissipative structures from the initial, driving disequilibrium.

To illustrate this, we construct chemical reaction networks, in which we explicitly implement instances of our proposed mechanisms: We start by demonstrating a simple mechanism that leads to a selection pressure on the thermodynamic efficiency of certain dissipative steady states in a stochastic chemical reaction network. The idea behind our simulations is simple: we construct a stochastic chemical reaction network, which may dissipate a large reservoir of chemical free energy via the maintenance of any one of a set of discrete, nonequilibrium steady states. Each of these nonequilibrium steady states is associated with the same, positive, finite minimum work-rate required to maintain it, and must maintain itself by harvesting the required work-rate from the conducted free energy flux via a bifurcation mechanism. Crucially, by making the thermodynamic efficiency of these work-harvesting processes different for each individual nonequilibrium steady state, a decrease of the dissipated driving chemical potential difference leads to a direct selection pressure on the nonequilibrium steady states. Indeed, only the corresponding bifurcation processes which are sufficiently efficient can sustain their associated nonequilibrium steady states.

We then discuss how this mechanism can lead to a selective pressure for thermodynamic efficiency in higher levels of a particular hierarchy of self-similar dissipative structures, based on the above network motif, where each dissipative structure is driven by (or feeding on) a dissipative structure at a lower level, i.e., closer to an initial, driving disequilibrium.

2. Materials and Methods

In this section, we explain the implementation details underlying our stochastic chemical reaction networks simulations. All calculations and simulations presented in this paper can be reproduced with code available at <https://github.com/kaiu85/CRNs> (accessed on 26 August 2021).

2.1. Stochastic Chemical Reaction Network Simulations

Most of our numerical experiments involving stochastic chemical reaction networks rest on sampling trajectories given a network architecture and initial network state. The state of a chemical reaction network, consisting of m chemical species S_1, \dots, S_m , is described by the number n of molecules of each species S present in the reaction volume V . This yields the state vector $\mathbf{n} = (n(S_1), n(S_2), \dots, n(S_m))$. The probability $p_{\mathbf{n}}(t)$ of finding the network in the state \mathbf{n} at a given time t evolves according to the chemical master equation

$$d_t p_{\mathbf{n}} = \sum_{\rho} w_{-\rho}(\mathbf{n} + \mathbf{S}_{\rho}) p_{\mathbf{n} + \mathbf{S}_{\rho}} - w_{\rho}(\mathbf{n}) p_{\mathbf{n}}.$$

The sum is taken over individual forward reactions ρ , and $-\rho$ denotes the corresponding backward reaction. \mathbf{S}_{ρ} represents the change in molecule numbers associated with the occurrence of the forward reaction ρ . Assuming elementary reactions and mass-action stochastic kinetics, we can write the reaction rates as

$$w_{\rho}(\mathbf{n}) = k_{\rho} \frac{V}{\prod_i V^{v_{\rho}^i}} \frac{\mathbf{n}!}{(\mathbf{n} - \mathbf{v}_{\rho})!}.$$

Here $\mathbf{v}_{\rho} = (v_{\rho}^1, v_{\rho}^2, \dots, v_{\rho}^m)$ is the vector of stoichiometric coefficients of the reaction ρ . This definition ensures that the corresponding chemical reaction rate constants k_{ρ} are the same as for the large volume limit considered in deterministic chemical kinetics [45,46].

To ensure that a closed stochastic chemical reaction network relaxes to a unique equilibrium steady state, therefore connecting dynamics to thermodynamics, we assume a local detailed balance relation between the rate constants of the forward and backward reaction directions

$$\ln \frac{k_{\rho}}{k_{-\rho}} = -\beta \boldsymbol{\mu}^0 \cdot \mathbf{S}_{\rho}.$$

Here $\boldsymbol{\mu}^0$ is the vector of standard-state chemical potentials, $\beta = \frac{1}{k_B T}$, k_B denotes the Boltzmann constant, and T denotes the temperature of a large heat bath, which is in contact with the reaction network. A detailed discussion of the dynamics and thermodynamics of stochastic chemical reaction networks can be found in [46].

For simplicity, we set the volume $V = 1$, the standard chemical potentials $\boldsymbol{\mu}^0 = 0$ for all species, corresponding to symmetric forward and backward rate constants $k_{\rho} = k_{-\rho}$, and use units of $k_B T = 1$.

The assumption of an equal standard chemical potential $\mu_S^0 = \mu^0 = 0$ for all species S effectively makes the equilibrium chemical potential a monotonous function of the species count $n(S)$, and renders forward and backward reaction rate constants symmetric. This assumption furnishes an intuitive understanding of the dynamics along individual reactions as being driven by concentration gradients, i.e., by the tendency to relax towards a homogeneous equilibrium state. This allows intuitions from other homogeneous relaxation processes (e.g., heat diffusion in homogeneous media, water levels in pipeline systems) to carry over, as it aligns the direction of the driving concentration gradient with the direction of the driving chemical potential gradient. This assumption can easily be relaxed by choosing individual values of μ_S^0 for each species. This will not affect our results or discussion on the self-selection of stable nonequilibrium steady states based on the efficiency of their work-harvesting processes, as long as one replaces “driving concentration gradient” with “driving chemical potential gradient”.

To sample individual trajectories for a given chemical reaction network and initial state, we implemented Gillespie’s stochastic simulation algorithm [47] based on Python 3 and PyTorch [48]. The latter allows us to use GPU-acceleration to parallelize over individual processes and chemical reactions.

2.2. Quantification of Minimum Work-Rate Required to Maintain Nonequilibrium Steady States

The minimum work-rate required to maintain a target nonequilibrium steady state equals the system's heat production rate in the instant, when the forcing maintaining the nonequilibrium steady state is stopped, i.e., when the system is closed and just begins to relax to equilibrium [49]. Therefore, we estimated the minimum work-rate required to maintain a given nonequilibrium steady state with the following approach: We simulated an ensemble of 10^5 stochastic trajectories at a target nonequilibrium steady state, by initializing the trajectories close to that particular nonequilibrium steady state, and stabilizing that state by opening the chemical reaction network via corresponding chemostatted species [46]—i.e., allowing for external nonequilibrium forces. We then closed the system, by releasing the clamps from the chemostatted species, treating them as any other species of the network henceforth. Therefore, the closed network immediately started to relax to its equilibrium steady state. Recall that the heat rate dissipated just at the beginning of this relaxation process equals the minimum work-rate required to maintain the corresponding nonequilibrium steady state [49]. Thus, we simulated ten additional time steps (i.e., reaction events) for each trajectory just after the system was closed. We used the time to Δt and free energy Δg released by the occurring reaction event to calculate the instantaneous heat dissipation rate via averaging the fraction $\frac{\Delta g}{\Delta t} = \frac{g(\mathbf{n}(t_{10})) - g(\mathbf{n}(t_0))}{t_{10} - t_0}$ over trajectories. Here t_0 is the time when the network was closed, t_{10} is the time of occurrence of the tenth reaction event thereafter, $g(\mathbf{n}(t_0))$ is the initial Gibbs free energy just when the network was closed, and $g(\mathbf{n}(t_{10}))$ is the Gibbs free energy of the closed network after the tenth reaction event. We quantified the Gibbs free energy g of the individual, closed chemical reaction network states using $g(\mathbf{n}) = kT \sum_S \ln(n(S)!)$, following [46]. The sum is taken over all chemical species S , and $n(S)$ is the number of molecules of chemical species S . This formula reflects our choice of the same chemical standard potential of $\mu_0 = 0$ for all species.

2.3. Quantification of the Thermodynamic Efficiency of Bifurcation-Based Work-Harvesting Processes

To quantify the average efficiency of different nonequilibrium steady states, we used the deterministic rate equations for the corresponding, large chemical reaction network (i.e., in the limit $n(S) \rightarrow \infty$, $V \rightarrow \infty$, $\frac{n(S)}{V} = c(S) = \text{const.}$ for species counts $n(S)$ and reaction volume V), and calculated the fluxes and the thermodynamic forces along each forward-backward pair of reactions, following [45]. This allowed us to calculate the heat dissipated along each forward-backward reaction pair via

$$\dot{Q}_\rho = -\Delta G_\rho J_\rho.$$

Here $J_\rho = J_\rho^+ - J_\rho^-$ is the net flux, i.e., the difference between the fluxes J_ρ^+ in the forward and J_ρ^- in the backward reaction direction. The thermodynamic forces ΔG_ρ are calculated as the chemical potential difference along the forward direction via

$$\Delta G_\rho = \sum_{p \in \text{products}, \rho^+} \mu_p - \sum_{e \in \text{educts}, \rho^+} \mu_e.$$

As before, we are using $\mu_0 = 0$ for all species, reducing the nonequilibrium chemical potentials to $\mu_\sigma = RT \ln \frac{Z_\sigma}{Z_\sigma^{\text{eq}}}$, where Z_σ is the concentration of species σ , and Z_σ^{eq} is the corresponding equilibrium concentration.

3. Results

In the following sections, we first generalize the bistable Schlögl network [50] to multiple, competing chemical species, leading to winner-take-all dynamics and multistability. We then show that to maintain any of the high-concentration steady states, a finite, minimum work-rate is required. In our simulations, this corresponds to a finite, positive driving chemical potential difference, which—due to our choice of $\mu^0 = 0$ for the

standard chemical potential of all chemical species—means just a finite, positive driving concentration gradient.

Using this winner-take-all motif as a chemical switch, we then integrate it into a larger network, where the dissipation of a concentration gradient is dependent on the presence of a high-concentration steady state of the winner-take-all network. At the same time, the winner-take-all network must harvest the work (or corresponding concentration gradient) to sustain itself directly from the conducted free energy flux via a bifurcation mechanism. By ensuring the efficiency of these bifurcation mechanisms differs for each of the high-concentration states, we see that a reduction of the driving concentration gradient leads to a strong selection pressure on the corresponding nonequilibrium steady states: Only the states with bifurcation processes of sufficient thermodynamic efficiency can retain their stability, when the driving concentration gradient is decreased.

In the discussion section, we will discuss how in a self-similar hierarchy of dissipative systems such a pressure might develop spontaneously, due to a decrease of the available, driving concentration gradient at each successive layer of dissipative structures.

3.1. Chemical Winner-Take-All Dynamics

The architecture of a chemical reaction network featuring winner-take-all dynamics is shown in Figure 1a. It is a straightforward extension of the nonlinear chemical reaction network introduced by Schlögl [50], which is the simplest bistable reaction network. The network is driven out of equilibrium by a fixed concentration difference between two chemostatted species, Hi and Lo, with $n(\text{Hi}) \gg n(\text{Lo})$. We chose the equilibrium potential $\mu^0 = 0$ for all species in our simulation, effectively rendering all reaction channels symmetric. Thus, the closed chemical reaction network would relax to an equilibrium distribution with equal counts n_{eq} for all species. To allow for nonlinear behavior, we introduce three dynamic chemical species, $X^{(1)}$, $X^{(2)}$, and $X^{(3)}$, which autocatalyze their own creation from the high-concentration species Hi. Furthermore, to introduce inhibitory competition between the dynamic species, we introduce decay reactions to the low-concentration species Lo for each dynamic species, which are catalysed by the competing species. Using the reaction rate constants $k_1 = k_2 = k_3 = 10^{-6}$ and $k_4 = k_5 = k_6 = 10^{-3}$, and the fixed concentrations $n(\text{Hi}) = 500$ and $n(\text{Lo}) = 5$ the resulting network features four stable attractors: One symmetric, low-concentration steady state, where $n(X^{(1)}) \approx n(X^{(2)}) \approx n(X^{(3)}) \approx n(\text{Lo})$, and three high-concentration steady states of species $X^{(1)}$, namely $n(X^{(1)}) \approx n(\text{Hi})$, $n(X^{(2)}) \approx n(X^{(3)}) \approx n(\text{Lo})$, species $X^{(2)}$, namely $n(X^{(2)}) \approx n(\text{Hi})$, $n(X^{(3)}) \approx n(X^{(1)}) \approx n(\text{Lo})$; and species $X^{(3)}$, namely $n(X^{(3)}) \approx n(\text{Hi})$, $n(X^{(1)}) \approx n(X^{(2)}) \approx n(\text{Lo})$, as shown in Figure 1b.

Using the same reaction rate constants but decreasing the concentration difference to $n(\text{Hi}) = 100$ and $n(\text{Lo}) = 5$, the high-concentration states lose their stability, and the only remaining stable attractor is the symmetric low-concentration state, as shown in Figure 1c.

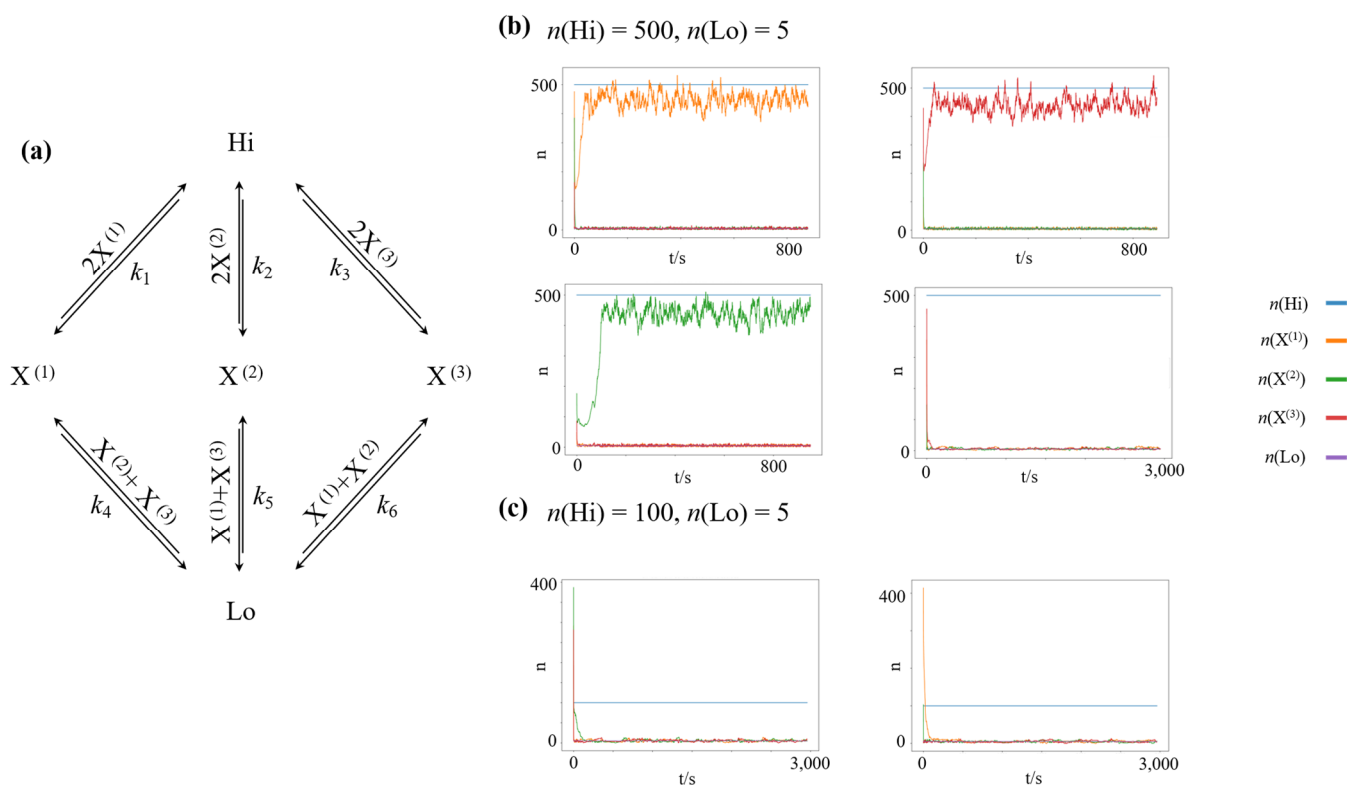


Figure 1. Nonlinear chemical reaction network, featuring winner-take-all attractor dynamics. Crucially, all three high-concentration nonequilibrium steady states are symmetric with respect to exchanging $X^{(1)}$, $X^{(2)}$, and $X^{(3)}$. Thus, the associated minimum work-rate required (and the associated minimum driving concentration gradient $n(\text{Hi}) - n(\text{Lo})$ to maintain each of these states is the same. (a) Layout of the chemical reaction network (b) Simulations of randomly initialized networks, with driving species clamped at $n(\text{Hi}) = 500$, $n(\text{Lo}) = 5$. Given this forcing, high-concentration states of an individual species $X^{(1)}$, $X^{(2)}$ or $X^{(3)}$, as well as a low-concentration state are stable attractors of the dynamics. (c) Simulations of randomly initialized networks, with driving species clamped at $n(\text{Hi}) = 100$, $n(\text{Lo}) = 5$. Given this forcing, only the low-concentration state constitutes an attractor of the dynamics.

3.2. Finite Minimum Work-Rate Required to Maintain High-Concentration States

To quantify the range of stability of high-concentration steady states, we run a series of simulations initializing the network close to the high-concentration steady state of species $X^{(1)}$, i.e., $n(X^{(1)}) = n(\text{Hi})$, $n(X^{(2)}) = n(X^{(3)}) = n(\text{Lo})$, and simulating 10^5 processes for 300,000 reaction steps. When the ensemble of trajectories has converged to a high-concentration steady state, we save the network state and calculate the minimum work-rate required to maintain the high-concentration steady state (c.f., methods section). As the network architecture is completely symmetric with respect to permutations of $X^{(1)}$, $X^{(2)}$, and $X^{(3)}$, the resulting minimum work-rate is the same also for high-concentration steady states of species $X^{(2)}$ and $X^{(3)}$.

We also performed a series of simulations, initializing the network close to the low-concentration steady state i.e., $n(X^{(1)}) = n(X^{(2)}) = n(X^{(3)}) = n(\text{Lo})$, using the same procedure and parameters. The resulting plot of the stability of the high- and low-concentration steady states, and the associated minimum work-rate required is shown in Figure 2a,b.

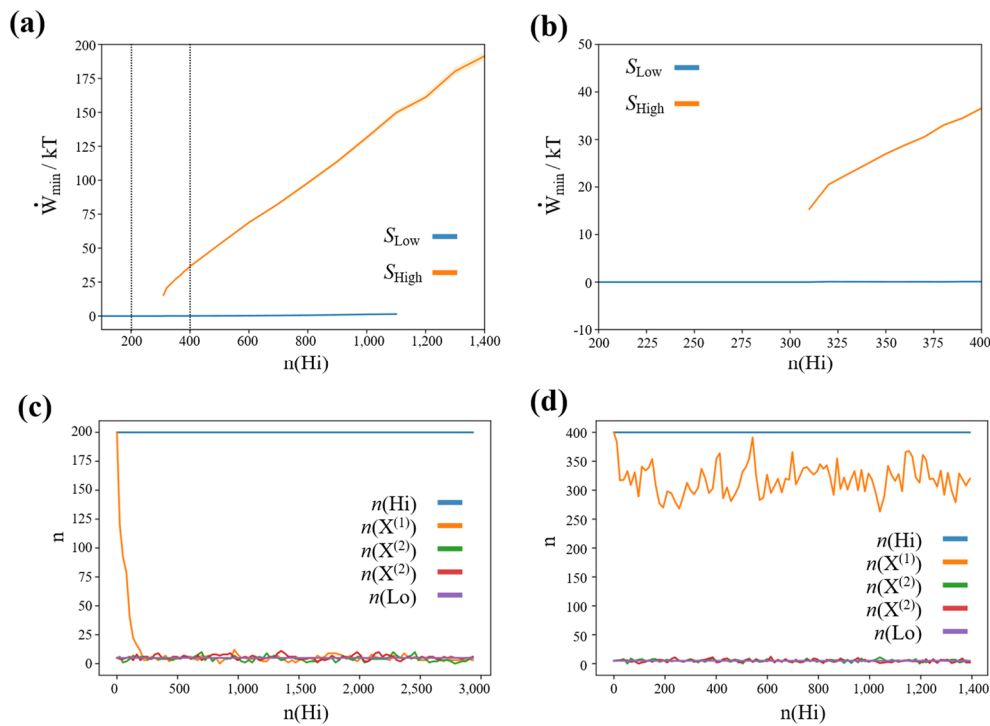


Figure 2. (a,b) Minimum work-rate (equivalent to a minimum concentration difference $n(\text{Hi}) - n(\text{Lo})$, where $n(\text{Lo}) = 5 = \text{const.}$) required to maintain stable nonequilibrium, high-concentration states in a stochastic chemical reaction network featuring winner-take-all dynamics. S_{Low} : Low-concentration steady state, S_{High} : High-concentration steady state. (c) Representative simulation trajectory for $n(\text{Hi}) = 200$. (d) Representative simulation trajectory for $n(\text{Hi}) = 400$.

We see that below a critical concentration of the high-concentration species (and therefore below a critical driving concentration difference) of $n_{\text{crit}}(\text{Hi}) \approx 310$ the high-concentration states lose their stability. Furthermore, we see that there is a positive, finite minimum work-rate associated with this concentration gradient.

Thus, in summary, the stability of high-concentration steady states in our winner-take-all chemical reaction network requires a positive, finite minimum work-rate (and associated concentration gradient). Furthermore, by construction this work-rate is the same for all high-concentration steady states.

3.3. Decreasing the Driving Chemical Potential Difference Leads to Self-Selection of Nonequilibrium Steady States with High-Efficiency Bifurcation Processes

We now can construct the chemical reaction network evincing the selection mechanism that we want to discuss in this paper: We use the same winner-take-all attractor motif; however, we do not directly connect $X^{(1)}$, $X^{(2)}$, and $X^{(3)}$ to the clamped high-concentration species Hi, but rather to a dynamic reservoir species Re, as shown in Figure 3a. Next, we introduce three potential relaxation channels, allowing the conversion of a molecule of the clamped, high-concentration species Hi to a molecule of the clamped low-concentration species Lo via three possible intermediate species, $B^{(1)}$, $B^{(2)}$, and $B^{(3)}$. Crucially, we make the synthesis of these intermediary species dependent on catalysis by two molecules of $X^{(1)}$, $X^{(2)}$, or $X^{(3)}$, respectively, as shown in Figure 3b. Finally, we couple the maintenance of the high-concentration winner-take-all attractor states to the conducted free energy flow, via a bifurcation process, turning a molecule of $B^{(1)}$, $B^{(2)}$, or $B^{(3)}$ into a molecule of the reservoir species Re. Again, we make these reactions depending on two molecules of $X^{(1)}$, $X^{(2)}$, or $X^{(3)}$, respectively, as shown in Figure 3c.

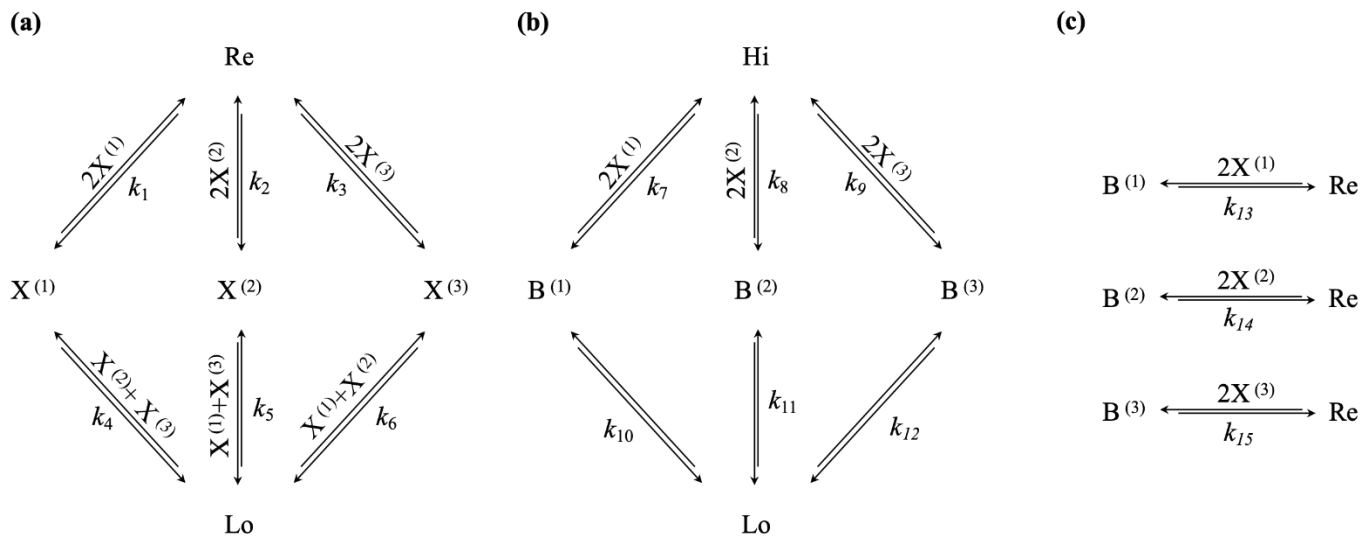


Figure 3. (a) Multi-stable winner-take-all network, featuring a single low- and three high-concentration nonequilibrium steady states, given a sufficiently high-concentration of reservoir species Re; (b) Disequilibrium/Free energy reservoir given by the concentration gradient $n(\text{Hi}) - n(\text{Lo})$, the dissipation of which is conditional on a high-concentration steady state of the winner-take-all network, by means of the catalysed reaction channels from Hi to $B^{(1)}$, $B^{(2)}$, or $B^{(3)}$; (c) Bifurcation mechanisms, coupling the dissipative channel to the maintenance of high-concentration states in the winner-take-all network, via the conversion of $B^{(1)}$, $B^{(2)}$, or $B^{(3)}$ to the reservoir species Re.

We run a series of simulations for a wide range of concentration gradients, by varying the concentration of the high-concentration species $n(\text{Hi})$, while keeping $n(\text{Lo}) = 5$ constant. We use the same rate constants as before for the winner-take-all subnetwork, namely $k_1 = k_2 = k_3 = 10^{-6}$ and $k_4 = k_5 = k_6 = 10^{-3}$. We set the other reaction rate constants to $k_7 = k_8 = k_9 = 10^{-5}$ and $k_{10} = k_{11} = k_{12} = 1.0$. We couple the intermediate species $B^{(1)}$, $B^{(2)}$, or $B^{(3)}$ to the reservoir species Re by different reaction rates, $k_{13} = 10^{-6}$, $k_{14} = 3 \cdot 10^{-7}$ and $k_{15} = 10^{-7}$, which leads to different efficiencies of the associated bifurcation process. We simulate 4000 processes, 1000 of which we initialize close to each of the four potentially stable states of the winner-take-all subnetwork. Specifically, 1000 close to the low-concentration state, $n(X^{(1)}) = n(X^{(2)}) = n(X^{(3)}) = n(B^{(1)}) = n(B^{(2)}) = n(B^{(3)}) = n(\text{Re}) = n(\text{Lo})$, 1000 close to the high-concentration state of species $X^{(1)}$, $n(X^{(1)}) = n(B^{(1)}) = n(\text{Re}) = n(\text{Hi})$, $n(X^{(2)}) = n(X^{(3)}) = n(B^{(2)}) = n(B^{(3)}) = n(\text{Lo})$, 1000 close to the high-concentration state of species $X^{(2)}$, $n(X^{(2)}) = n(B^{(2)}) = n(\text{Re}) = n(\text{Hi})$, $n(X^{(3)}) = n(X^{(1)}) = n(B^{(3)}) = n(B^{(1)}) = n(\text{Lo})$, and 1000 close to the high-concentration state of species $X^{(3)}$, $n(X^{(3)}) = n(B^{(3)}) = n(\text{Re}) = n(\text{Hi})$, $n(X^{(1)}) = n(X^{(2)}) = n(B^{(1)}) = n(B^{(2)}) = n(\text{Lo})$. We simulated 6,000,000 reaction steps for each process and checked if it had converged to the corresponding steady state. For each external forcing, we quantified the relative number of processes that converged to the steady state, which they were initialized close to, as a proxy for this state's stability. The resulting graphs are shown in Figure 4a. As expected, for a very strong driving concentration gradient, all possible high-concentration states are stable (and the low-concentration state actually loses its stability), in analogy to the Schlögl model [50]. However, as the driving concentration gradient is decreased, more and more high-concentration states lose their stability. The first state to lose its stability is the high-concentration state of species $X^{(3)}$, which features the bifurcation (i.e., work-harvesting) process of lowest efficiency. When decreasing the driving concentration gradient further, the next high-concentration state to lose its stability is the high-concentration state of species $X^{(2)}$, featuring the bifurcation process of second lowest efficiency. When we reduce the driving concentration gradient even further, the high-concentration state of species $X^{(1)}$, featuring the most efficient bifurcation process, also loses its stability.

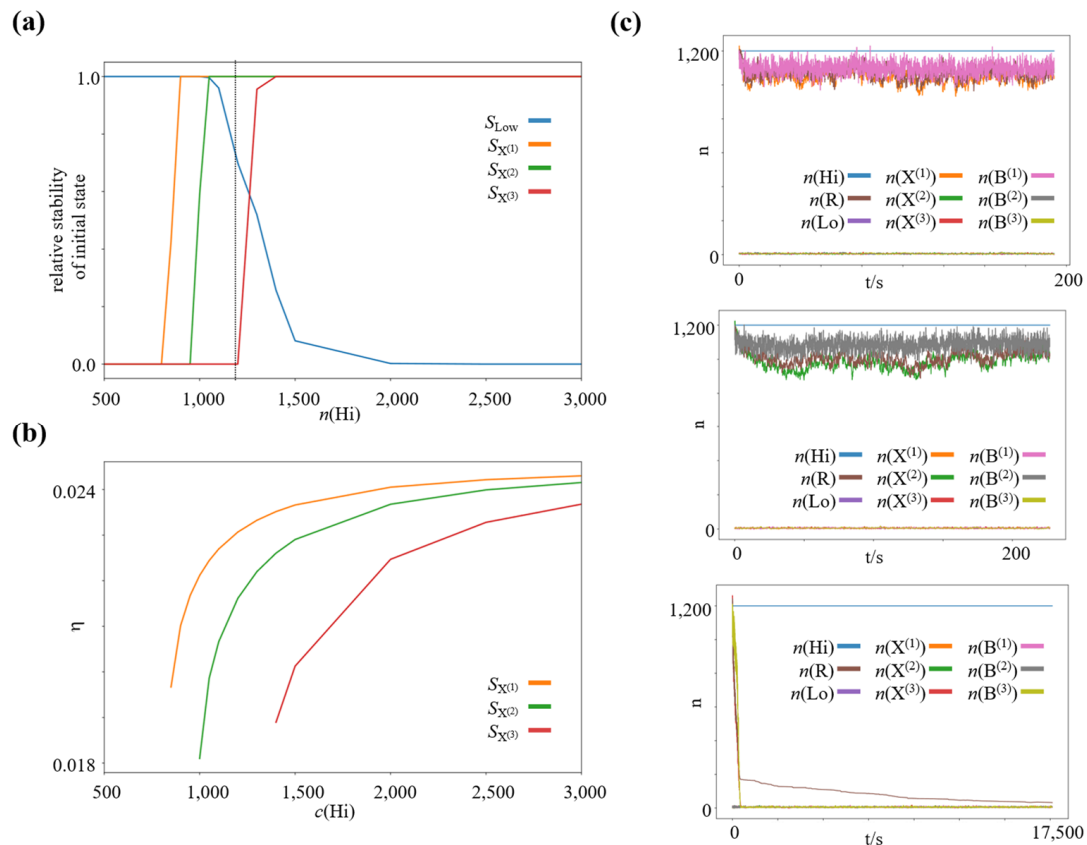


Figure 4. Increasing selection pressure on the efficiency of the work-harvesting processes maintaining the dissipative nonequilibrium steady states with decreasing available chemical potential difference. **(a)** Relative stability of the low and high-concentration steady states as function of $n(\text{Hi})$, where $n(\text{Lo}) = 5 = \text{const}$. **(b)** Efficiency η of the bifurcation processes at the individual high-concentration nonequilibrium steady states of species $X^{(1)}$, $X^{(2)}$, and $X^{(3)}$ as function of $c(\text{Hi})$, where $c(\text{Lo}) = 5 = \text{const}$. **(c)** Representative trajectories of a chemical reaction network initialized close to each high-concentration steady state for $n(\text{Hi}) = 1200$. $S_{X^{(1)}}$: high-concentration state of species $X^{(1)}$, $S_{X^{(2)}}$: high-concentration state of species $X^{(2)}$, $S_{X^{(3)}}$: high-concentration state of species $X^{(3)}$, S_{Low} : Low-concentration steady state.

The average thermodynamic efficiency of the bifurcation processes at the individual high-concentration steady states was computed using the deterministic limit of the stochastic chemical reaction network, i.e., taking $n(S) \rightarrow \infty$, $V \rightarrow \infty$, $\frac{n(S)}{V} = c(S) = \text{const}$. Using the established theory of thermodynamics in deterministic, open chemical reaction networks ([45]; c.f., Methods), we calculated the average total dissipated heat \dot{Q}_{tot} in the entire network, and the average heat \dot{Q}_{diss} dissipated along the direct decay pathways from Hi via $B^{(1)}$, $B^{(2)}$, and $B^{(3)}$ to Lo at each high-concentration steady state. At steady state, the total dissipated heat corresponds to the total chemical work put into the network by the driving potential gradient (see for example [45,46]), thus we can calculate $\dot{Q}_{\text{tot}} - \dot{Q}_{\text{diss}} = \dot{W}_{\text{chem,tot}} - \dot{Q}_{\text{diss}} = \dot{W}_{\text{chem,WTA}}$, which yields the chemical work, which is bifurcated to maintain the high-concentration state of the winner-take-all module (or equivalently, the chemical work required to pay the corresponding house-keeping heat). Now we can calculate the associated efficiency via $\eta = \frac{\dot{W}_{\text{chem,WTA}}}{\dot{W}_{\text{chem,tot}}}$. We plotted the corresponding efficiencies for the high-concentration steady states in Figure 4b. In this particular setup, all three processes are highly inefficient; however, the small absolute efficiency difference is sufficient to create strong, specific selection pressures on the corresponding nonequilibrium steady states when the driving concentration gradient is decreased.

4. Discussion

4.1. A Simple Selection Mechanism for Thermodynamic Efficiency in Dissipative Nonequilibrium Steady States of Chemical Reaction Networks

This work demonstrates a relatively simple mechanism, which leads to selective pressure on the efficiency of bifurcation-based work-harvesting processes in dissipative nonequilibrium steady states—following a decrease in the dissipated, driving disequilibrium. It rests on conditioning dissipative fluxes—along the driving concentration gradient in an open chemical reaction network—on a high-concentration state in a winner-take-all chemical reaction network. In turn, the winner-take-all network must harvest the work required to maintain this state from the very free energy fluxes that it enables, via a bifurcation mechanism.

Although it would be conceivable in principle that the dissipative structure was fueled by a different energy source than the one whose dissipation it facilitated, recent work showed that the molecular level free energy-conversion processes, which living systems use to harvest the work to maintain their own structure, must rely on such turnstile-, bifurcation- or escapement-like mechanisms [1,51]. Furthermore, the fact that at least some of the associated molecular engines—such as the F1-ATPase [52]—work very close to optimal efficiency, might be a hint that selection pressures towards thermodynamic efficiency might have played a role in the development of these structures.

Furthermore, while we only discuss a very specific special case here, based on simplifying assumptions such as symmetric reaction rate constants for forward and backward reaction directions, the general premise—namely that dissipative structures require a positive, finite minimum work-rate to persist, and that they have to harvest this work directly from the conducted dissipative free energy fluxes, therefore requiring a specific minimum efficiency of the associated work-harvesting processes—should be applicable to a wider class of dissipative structures.

4.2. Emergent Pressure towards Thermodynamic Efficiency in a Hierarchy of Dissipative Structures

Although we manually decreased the driving concentration difference in our simulations, one could easily imagine a similar pressure arising naturally, for example in an emergent, self-similar hierarchy of dissipative structures, each feeding on—i.e., dissipating—a lower level of dissipative structures (cf. [1,2]), in the following way:

The steady state concentration gradient $n(\text{Re}) - n(\text{Lo})$, which is required to sustain the dissipative fluxes between Hi and Lo in our simulated network, constitutes a disequilibrium by itself. This disequilibrium can drive the emergence of another layer of dissipative systems with a similar structure. For example, one could just add another winner-take-all dissipation motif in a hierarchical fashion, as shown in Figure 5. Due to the construction of our reaction network, the concentration gradient $n(\text{Re}) - n(\text{Lo})$ and the associated chemical potential difference will be smaller than that of the initial, driving disequilibrium $n(\text{Hi}) - n(\text{Lo})$. Furthermore, assuming a similar architecture of the “higher order” winner-take-all network, including the reaction rate constants, the minimum work-rate required to maintain the corresponding nonequilibrium steady state will be similar to that of the original dissipative structure. Thus, there is already a slight increase in the minimum efficiency of the new bifurcation process required to harvest enough work to maintain the additional layer of dissipative structures. Given a sufficiently complex chemical space, one might imagine another layer of dissipative processes, feeding on the structure afforded by the second order winner-take-all network, e.g., in terms of the concentration gradient $n(\text{Re}') - n(\text{Lo})$, where $n(\text{Re}') - n(\text{Lo}) < n(\text{Re}) - n(\text{Lo}) < n(\text{Hi}) - n(\text{Lo})$. Thus, by iterating this argument over additional, self-similar layers of higher order dissipative structures—each feeding on the structures afforded by the previous layer of dissipative structures—an increasing pressure towards efficient bifurcation processes would develop as one moved further away from the initial, driving disequilibrium in this dissipative hierarchy. In this setting, a new layer of dissipative structure could only form if the associated

work-harvesting bifurcation process were efficient enough to sustain it. As the available driving concentration gradient would monotonically decrease with each new layer of dissipative structures, but the minimum work required to sustain each successive layer would not, one could expect to find highly efficient, close to optimal dissipative structures in higher levels of this dissipative hierarchy—given a sufficiently complex chemical space and enough time for the chemical network to discover these highly efficient nonequilibrium steady states.

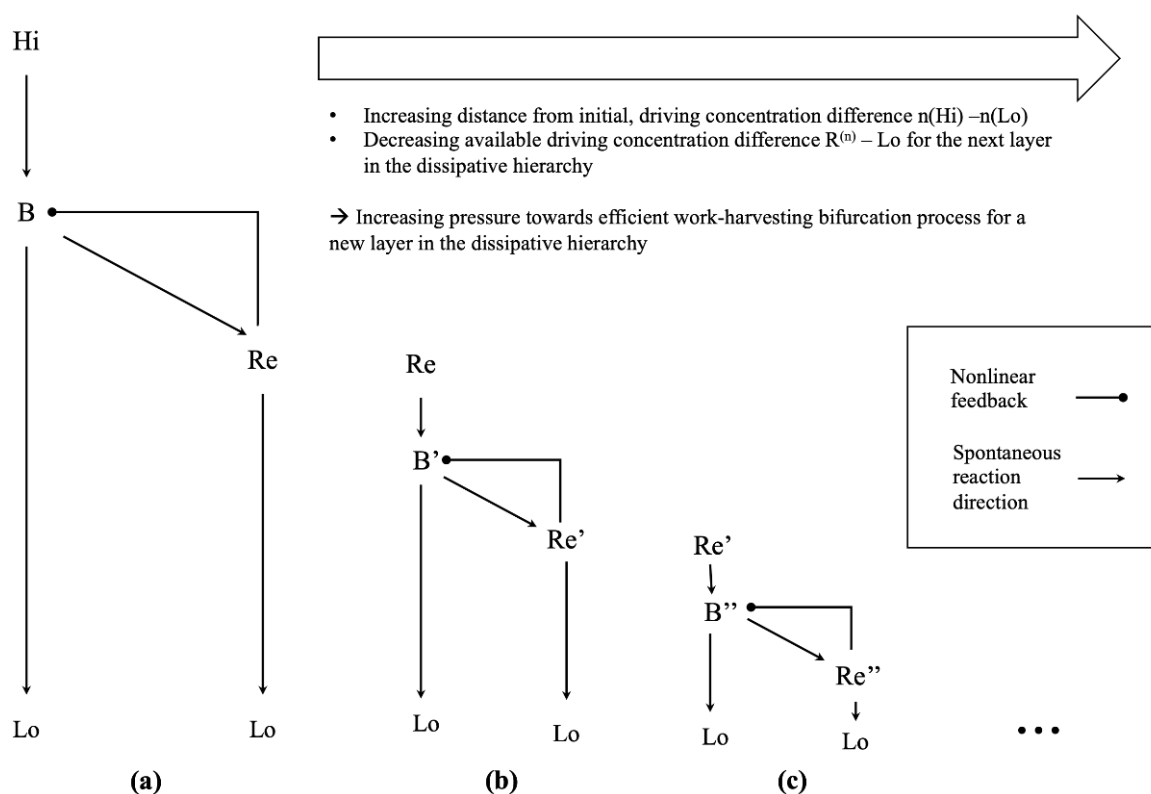


Figure 5. (a) Schematic depiction of the chemical reaction network discussed in this paper. The nonlinear feedback between the reservoir species Re and the intermediate species B in our simulation is mediated by the dynamics of the winner-take-all module. (b) It is conceivable that a similar dissipative structure develops, which does not feed on the initial driving concentration gradient $n(\text{Hi}) - n(\text{Lo})$, but on the structure realizing the dissipative process, in terms of the concentration gradient $n(\text{Re}) - n(\text{Lo})$ between the reservoir species of the initial dissipative structure and the low-concentration species. (c) This process could be iterated, by imagining a dissipative structure which now feeds on the reservoir species of the second order dissipative structure. One immediately sees that the concentration gradient available to drive the next layer in such a hierarchy of dissipative structures decreases the further one moves away from the initial driving concentration gradient.

In a more realistic setting, variations in the minimum work-rate, required to form a new layer of dissipative structures, might lead to a non-monotonic relationship between the distance of a layer of dissipative structure from the initial driving disequilibrium, and the minimum efficiency required by the associated work-harvesting process to sustain this new layer of dissipative structures. However, while the available driving chemical potential difference will in general decrease with the hierarchical level of dissipative processes in the absence of concentration processes (i.e., processes which can turn two or more molecules of low chemical potential into a molecule of higher chemical potential), the minimum work-rate required to maintain a new layer of dissipative structure does not have to: it might decrease, stay constant, or increase. Thus, it does not seem too implausible that—in many chemical systems—the required work-rate to sustain a new layer of dissipative structure might come close to the maximum available free energy flux, therefore leading to a correspondingly high minimum efficiency required by the associated work-harvesting

process. However, there are many processes in biology, for which this argument does not hold. Prominent counterexamples are food webs, in which plants concentrate the free energy of low entropy photons, e.g., in terms of vegetable fats and carbohydrates. These are then further concentrated by herbivores, e.g., in terms of animal fats and protein, which provide carnivores with a source of highly concentrated chemical free energy (for a discussion of adaptive dynamics and potential self-similarity in food webs, see [53]).

Although our work demonstrates self-selection of a subset of stable nonequilibrium steady states—namely those associated with sufficiently efficient work-harvesting processes to maintain themselves—we do not address the question of how probable different trajectories leading towards these states might be. This is an interesting question as recent work shows that the probability of arriving at a given nonequilibrium steady state is dominated by the dissipative histories of the trajectories leading there [54,55], which in general cannot be related directly to the dissipation at the steady state itself. However, we can recover a qualitative relationship between our results and the general notion of dissipative adaptation from the following, heuristic argument: In real systems, the dissipated concentration gradient would feed from a finite reservoir, i.e., disequilibrium, which must be dissipated at some point to allow for the whole—initially highly frustrated—system to relax to equilibrium. In our case, the relative probability of that reservoir being dissipated without realizing one of the stable high-concentration nonequilibrium steady states is, by construction, virtually zero. Using our simulation framework to quantify the relative probabilities of reaching equilibrium via the available, self-selected nonequilibrium steady states, might be an interesting avenue for future research.

5. Conclusions

Given the large and high-dimensional, yet highly structured (in the sense that not every chemical reaction is allowed) nature of chemical space, leading to intriguing—and only partly understood—structures and dynamics in the most complex chemical reaction networks we know, namely those in biochemistry (e.g., [17,56,57]), it seems reasonable to assume that similar dissipative hierarchies and associated pressures towards thermodynamic efficiency might have emerged in biological systems. This may be the case, as existing biogeochemical networks feature a universal scaling law, which is not a product of the underlying chemical space alone [56]. The mechanisms considered in this work might therefore constitute one of many drives leading to the transition of early biochemistry from a thermodynamic- to an information-determined regime, which is considered a hallmark of the living state [18,19].

Author Contributions: Conceptualization, K.U., L.D.C. and K.F.; methodology, K.U.; software, K.U.; writing—original draft preparation, K.U.; writing—review and editing, K.U., L.D.C., D.C. and K.F.; visualization, K.U.; All authors have read and agreed to the published version of the manuscript.

Funding: K.U. was supported by a DAAD PRIME fellowship. L.D.C. was supported by the Fonds National de la Recherche, Luxembourg (Project code: 13568875). D.C. was supported by the Department of Philosophical, Pedagogical, and Economic-Quantitative Sciences Fellowship. K.F. was a Wellcome Principal Research Fellow (Ref: 088130/Z/09/Z). This publication is based on work partially supported by the EPSRC Centre for Doctoral Training in Mathematics of Random Systems: Analysis, Modelling and Simulation (EP/S023925/1).

Data Availability Statement: Code reproducing the calculations, simulations and graphics shown in this paper can be accessed at: <https://github.com/kaiu85/CRNs> (accessed on 26 August 2021).

Acknowledgments: The authors thank Dimitrije Markovic, Conor Heins, Zahra Sheikhabaee, Danijar Hafner, Beren Millidge, and Alexander Tschantz for stimulating discussions. They furthermore thank two anonymous referees and the academic editor for their engaged, thoughtful, and constructive reviews, which substantially improved the clarity and quality of the presented work.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Branscomb, E.; Russell, M.J. Turnstiles and bifurcators: The disequilibrium converting engines that put metabolism on the road. *Biochim. Biophys. Acta (BBA) Bioenerg.* **2013**, *1827*, 62–78. [[CrossRef](#)]
2. Russell, M.J.; Nitschke, W.; Branscomb, E. The inevitable journey to being. *Philos. Trans. R. Soc. B Biol. Sci.* **2013**, *368*, 20120254. [[CrossRef](#)] [[PubMed](#)]
3. Morowitz, H.; Smith, E. Energy flow and the organization of life. *Complexity* **2007**, *13*, 51–59. [[CrossRef](#)]
4. Goldenfeld, N.; Biancalani, T.; Jafarpour, F. Universal biology and the statistical mechanics of early life. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **2017**, *375*, 20160341. [[CrossRef](#)] [[PubMed](#)]
5. Smith, E.; Morowitz, H.J. *The Origin and Nature of Life on Earth: The Emergence of the Fourth Geosphere*; Cambridge University Press: Cambridge, UK, 2016.
6. Jeffery, K.; Pollack, R.; Rovelli, C. On the Statistical Mechanics of Life: Schrödinger Revisited. *Entropy* **2019**, *21*, 1211. [[CrossRef](#)]
7. Lotka, A.J. Natural Selection as a Physical Principle. *Proc. Natl. Acad. Sci. USA* **1922**, *8*, 151. [[CrossRef](#)] [[PubMed](#)]
8. Lotka, A.J. Contribution to the Energetics of Evolution. *Proc. Natl. Acad. Sci. USA* **1922**, *8*, 147. [[CrossRef](#)] [[PubMed](#)]
9. Prigogine, I.; Nicolis, G. Biological order, structure and instabilities. *Q. Rev. Biophys.* **1971**, *4*, 107–148. [[CrossRef](#)] [[PubMed](#)]
10. Cottrell, A. The natural philosophy of engines. *Contemp. Phys.* **1979**, *20*, 1–10. [[CrossRef](#)]
11. Crooks, G.E. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Phys. Rev. E* **1999**, *60*, 2721–2726. [[CrossRef](#)]
12. Jarzynski, C. Nonequilibrium Equality for Free Energy Differences. *Phys. Rev. Lett.* **1997**, *78*, 2690–2693. [[CrossRef](#)]
13. Van den Broeck, C.; Esposito, M. Ensemble and trajectory thermodynamics: A brief introduction. *Phys. A Stat. Mech. Its Appl.* **2015**, *418*, 6–16. [[CrossRef](#)]
14. Seifert, U. Stochastic thermodynamics: From principles to the cost of precision. *Phys. A Stat. Mech. Its Appl.* **2018**, *504*, 176–191. [[CrossRef](#)]
15. Seifert, U. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Rep. Prog. Phys.* **2012**, *75*, 126001. [[CrossRef](#)] [[PubMed](#)]
16. Perunov, N.; Marsland, R.A.; England, J.L. Statistical Physics of Adaptation. *Phys. Rev. X* **2016**, *6*, 021036. [[CrossRef](#)]
17. Walker, S.I.; Kim, H.; Davies, P.C.W. The informational architecture of the cell. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **2016**, *374*, 20150057. [[CrossRef](#)] [[PubMed](#)]
18. Walker, S.I.; Davies, P.C.W. The algorithmic origins of life. *J. R. Soc. Interface* **2013**, *10*, 20120869. [[CrossRef](#)] [[PubMed](#)]
19. Davies, P.C.W.; Walker, S.I. The hidden simplicity of biology. *Rep. Prog. Phys.* **2016**, *79*, 102601. [[CrossRef](#)] [[PubMed](#)]
20. Still, S.; Sivak, D.A.; Bell, A.J.; Crooks, G.E. Thermodynamics of Prediction. *Phys. Rev. Lett.* **2012**, *109*, 120604. [[CrossRef](#)] [[PubMed](#)]
21. Tishby, N.; Pereira, F.C.; Bialek, W. The information bottleneck method. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing*; Hajek, B., Sreenivas, R.S., Eds.; University of Illinois: Urbana, IL, USA, 1999; pp. 368–377.
22. Still, S. Information-theoretic approach to interactive learning. *EPL* **2009**, *85*, 28005. [[CrossRef](#)]
23. Still, S.; Precup, D. An information-theoretic approach to curiosity-driven reinforcement learning. *Theory Biosci.* **2012**, *131*, 139–148. [[CrossRef](#)]
24. Still, S. Information Bottleneck Approach to Predictive Inference. *Entropy* **2014**, *16*, 968–989. [[CrossRef](#)]
25. Balasubramanian, V.; Kimber, D.; Li, M.J.B. Metabolically Efficient Information Processing. *Neural Comput.* **2001**, *13*, 799–815. [[CrossRef](#)]
26. Krieg, D.; Triesch, J. A unifying theory of synaptic long-term plasticity based on a sparse distribution of synaptic strength. *Front. Synaptic Neurosci.* **2014**, *6*, 3. [[CrossRef](#)]
27. Perge, J.A.; Koch, K.; Miller, R.; Sterling, P.; Balasubramanian, V. How the Optic Nerve Allocates Space, Energy Capacity, and Information. *J. Neurosci.* **2009**, *29*, 7917. [[CrossRef](#)] [[PubMed](#)]
28. Attwell, D.; Laughlin, S.B. An Energy Budget for Signaling in the Grey Matter of the Brain. *J. Cereb. Blood Flow Metab.* **2001**, *21*, 1133–1145. [[CrossRef](#)] [[PubMed](#)]
29. Koch, K.; McLean, J.; Segev, R.; Freed, M.A.; Berry, M.J., II; Balasubramanian, V.; Sterling, P. How Much the Eye Tells the Brain. *Curr. Biol.* **2006**, *16*, 1428–1434. [[CrossRef](#)] [[PubMed](#)]
30. Levy, W.B.; Baxter, R.A. Energy-Efficient Neuronal Computation via Quantal Synaptic Failures. *J. Neurosci.* **2002**, *22*, 4746. [[CrossRef](#)] [[PubMed](#)]
31. Levy, W.B.; Baxter, R.A. Energy Efficient Neural Codes. *Neural Comput.* **1996**, *8*, 531–543. [[CrossRef](#)] [[PubMed](#)]
32. Sarpeshkar, R. Analog Versus Digital: Extrapolating from Electronics to Neurobiology. *Neural Comput.* **1998**, *10*, 1601–1638. [[CrossRef](#)]
33. Balasubramanian, V. Heterogeneity and Efficiency in the Brain. *Proc. IEEE* **2015**, *103*, 1346–1358. [[CrossRef](#)]
34. Schmiedl, T.; Seifert, U. Stochastic thermodynamics of chemical reaction networks. *J. Chem. Phys.* **2007**, *126*, 044101. [[CrossRef](#)] [[PubMed](#)]
35. Schmiedl, T.; Speck, T.; Seifert, U. Entropy Production for Mechanically or Chemically Driven Biomolecules. *J. Stat. Phys.* **2007**, *128*, 77–93. [[CrossRef](#)]
36. Gaspard, P. Fluctuation theorem for nonequilibrium reactions. *J. Chem. Phys.* **2004**, *120*, 8898–8905. [[CrossRef](#)] [[PubMed](#)]

37. Andrieux, D.; Gaspard, P. Fluctuation theorem and Onsager reciprocity relations. *J. Chem. Phys.* **2004**, *121*, 6167–6174. [[CrossRef](#)] [[PubMed](#)]
38. Esposito, M. Open questions on nonequilibrium thermodynamics of chemical reaction networks. *Commun. Chem.* **2020**, *3*, 107. [[CrossRef](#)]
39. Penocchio, E.; Rao, R.; Esposito, M. Thermodynamic efficiency in dissipative chemistry. *Nat. Commun.* **2019**, *10*, 3865. [[CrossRef](#)] [[PubMed](#)]
40. Avanzini, F.; Falasco, G.; Esposito, M. Thermodynamics of chemical waves. *J. Chem. Phys.* **2019**, *151*, 234103. [[CrossRef](#)]
41. Falasco, G.; Rao, R.; Esposito, M. Information Thermodynamics of Turing Patterns. *Phys. Rev. Lett.* **2018**, *121*, 108301. [[CrossRef](#)] [[PubMed](#)]
42. Avanzini, F.; Penocchio, E.; Falasco, G.; Esposito, M. Nonequilibrium thermodynamics of non-ideal chemical reaction networks. *J. Chem. Phys.* **2021**, *154*, 094114. [[CrossRef](#)] [[PubMed](#)]
43. Avanzini, F.; Falasco, G.; Esposito, M. Thermodynamics of non-elementary chemical reaction networks. *New J. Phys.* **2020**, *22*, 093040. [[CrossRef](#)]
44. Wachtel, A.; Rao, R.; Esposito, M. Thermodynamically consistent coarse graining of biocatalysts beyond Michaelis–Menten. *New J. Phys.* **2018**, *20*, 042002. [[CrossRef](#)]
45. Rao, R.; Esposito, M. Nonequilibrium Thermodynamics of Chemical Reaction Networks: Wisdom from Stochastic Thermodynamics. *Phys. Rev. X* **2016**, *6*, 041064. [[CrossRef](#)]
46. Rao, R.; Esposito, M. Conservation laws and work fluctuation relations in chemical reaction networks. *J. Chem. Phys.* **2018**, *149*, 245101. [[CrossRef](#)]
47. Gillespie, D.T. Stochastic Simulation of Chemical Kinetics. *Annu. Rev. Phys. Chem.* **2007**, *58*, 35–55. [[CrossRef](#)]
48. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. *Annu. Conf. Neural Inf. Process. Syst.* **2019**, *32*, 8026–8037.
49. Horowitz, J.M.; Zhou, K.; England, J.L. Minimum energetic cost to maintain a target nonequilibrium state. *Phys. Rev. E* **2017**, *95*, 042102. [[CrossRef](#)] [[PubMed](#)]
50. Schlögl, F. On thermodynamics near a steady state. *Z. Für Phys. A Hadron. Nucl.* **1971**, *248*, 446–458. [[CrossRef](#)]
51. Branscomb, E.; Biancalani, T.; Goldenfeld, N.; Russell, M. Escapement mechanisms and the conversion of disequilibria; the engines of creation. *Phys. Rep.* **2017**, *677*, 1–60. [[CrossRef](#)]
52. Kinoshita, K.; Yasuda, R.; Noji, H.; Adachi, K. A rotary molecular motor that can work at near 100% efficiency. *Philos. Trans. R. Soc. London. Ser. B Biol. Sci.* **2000**, *355*, 473–489. [[CrossRef](#)] [[PubMed](#)]
53. McCann, K.S.; Rooney, N. The more food webs change, the more they stay the same. *Philos. Trans. R. Soc. B Biol. Sci.* **2009**, *364*, 1789–1801. [[CrossRef](#)]
54. Kachman, T.; Owen, J.A.; England, J.L. Self-Organized Resonance during Search of a Diverse Chemical Space. *Phys. Rev. Lett.* **2017**, *119*, 038001. [[CrossRef](#)] [[PubMed](#)]
55. Cook, J.; Endres, R.G. Thermodynamics of switching in multistable non-equilibrium systems. *J. Chem. Phys.* **2020**, *152*, 054108. [[CrossRef](#)] [[PubMed](#)]
56. Kim, H.; Smith, H.B.; Mathis, C.; Raymond, J.; Walker, S.I. Universal scaling across biochemical networks on Earth. *Sci. Adv.* **2019**, *5*, eaau0149. [[CrossRef](#)]
57. Marshall, S.M.; Mathis, C.; Carrick, E.; Keenan, G.; Cooper, G.J.T.; Graham, H.; Craven, M.; Gromski, P.S.; Moore, D.G.; Walker, S.I.; et al. Identifying molecules as biosignatures with assembly theory and mass spectrometry. *Nat. Commun.* **2021**, *12*, 3033. [[CrossRef](#)] [[PubMed](#)]