Dynamic Symmetry in Multi-Cycle Open Chemical Reaction Networks

Abstract: This paper extends the Dynamic Symmetry programme from the minimal single-throughput open chemical reaction network to a richer multi-cycle setting and then pushes the same framework into a first time-local implementation. The aim is to test whether the proposed relation between informational diversity and boundary-maintained dissipation survives once the internal reaction topology contains competing loops, parallel pathways, and circulation currents rather than a single dominant chain, and whether the resulting index can begin to function as an early-warning indicator rather than only as a stationary diagnostic. The network is formulated as a finite-state, chemostatted stochastic chemical system with several internal species and multiple independent reaction cycles. As in the earlier open-network construction, microscopic variability is represented by the Shannon entropy of the internal probability distribution, while macroscopic organisation is represented by entropy production generated by externally maintained chemical affinities. These quantities are normalised to define diversity and order variables, which are then combined into a Dynamic Symmetry Index. The principal claim is modest but important. In a multi-cycle open network, the same architecture remains mathematically coherent and yields a more informative picture than the minimal case: weak forcing produces insufficient organised throughput, strong forcing canalises probability mass into dominant pathways, and an intermediate regime can again be identified in which internal diversity and dissipative structure are jointly sustained. A first time-local extension, implemented on the minimal chemostatted network as a proof of method, further shows that the same architecture can be monitored on sliding windows and paired with rolling variance and lag-one autocorrelation tests. The result does not establish universality, but it shows that the Dynamic Symmetry framework survives a nontrivial increase in topological and thermodynamic complexity while also admitting a first route toward early-warning analysis.

1. Introduction

The first open-network application of Dynamic Symmetry Theory considered a minimal chemostatted chemical reaction network in which a single throughput channel linked two boundary reservoirs through a small internal state space. That construction served its purpose. It showed that the bridge between Shannon entropy and boundary dissipation could be made mathematically explicit in an established thermodynamic framework, and it provided a numerical example in which the Dynamic Symmetry Index attained an interior maximum at intermediate forcing. Yet the same simplicity that made the model transparent also limited its reach. A single-chain open network contains little topological competition. Once the external driving is specified, there are few genuinely distinct pathways through which matter can circulate. The model therefore risks being read as a proof of concept for one privileged geometry rather than as evidence that the framework can survive richer nonequilibrium organisation.

The natural next step is to move from a single-throughput architecture to a multi-cycle one. Real open reaction networks often contain loops, branch points, competing channels, and internally circulating currents. These features matter because they enlarge the set of accessible mesoscopic states while simultaneously allowing several distinct dissipative structures to be maintained by the same or by coupled chemical affinities. A multi-cycle setting therefore provides a much stronger test of the Dynamic Symmetry proposal. If the index remains meaningful there, it is less likely to be a mere artefact of a one-dimensional throughput model.

This paper develops that extension. The goal is not yet to produce a general theorem for all open CRNs. The goal is to formulate a concrete multi-cycle network, define the relevant entropy and dissipation quantities on the same probabilistic state space, and show how the Dynamic Symmetry architecture carries over. The conceptual expectation is clear. At low forcing, multiple cycles exist in principle but are thermodynamically underused, so organised throughput is weak. At very high forcing, one or a small number of strongly favoured routes dominate the stationary distribution, reducing effective internal diversity. Between these limits lies the regime of interest, where several pathways remain active while nonequilibrium organisation is already substantial. That is the regime the Dynamic Symmetry Index is designed to detect.

2. Dynamic Symmetry Framework for Open CRNs

The open-network version of Dynamic Symmetry Theory already established the central formal ingredients. In a chemostatted stochastic reaction network, the internal chemical system is represented by a stationary probability distribution over copy-number states, and the nonequilibrium environment maintains persistent fluxes through fixed boundary concentrations or chemical potentials. The two central variables are then defined on the same stationary state. The diversity variable is derived from Shannon entropy over the internal distribution, and the order variable is derived from the entropy production rate sustained by the open, driven steady state. In the earlier minimal network, this led to an index that rose from near zero at weak forcing, peaked at intermediate driving, and declined again as the stationary measure became increasingly concentrated on high-throughput states.

That earlier construction provides the template for the present paper. The relevant quantities are not changed in kind, only extended to a more complex topology. Let p_s denote the stationary probability of internal mesoscopic state s in a finite-state stochastic CRN. The stationary Shannon entropy is

$$H = -\sum_s p_s \ln p_s. \quad (1)$$

The corresponding normalised diversity variable is

$$D = \frac{H}{H_{\max}}, 0 \leq D \leq 1, (2)$$

where $H_{\max}$ is chosen as the maximal stationary entropy across the sampled forcing range or, where analytically justified, across the accessible state support.

The order-like quantity is built from the stationary entropy production rate. For a continuous-time Markov jump process with stationary probabilities p_s and transition rates $W_{s\to s'}$, the standard stationary entropy production rate is

$$\sigma = \frac{1}{2}\sum_{s,s'} J_{s,s'} \ln \frac{p_s W_{s\to s'}}{p_{s'} W_{s'\to s}}, (3)$$

where $J_{s,s'} = p_s W_{s\to s'} - p_{s'} W_{s'\to s}$ is the stationary probability current. In an open network, σ is positive whenever the chemostats maintain broken detailed balance and persistent nonequilibrium circulation. It is then normalised to the bounded order variable

$$O = \frac{\sigma}{\sigma + K}, 0 \leq O \leq 1, (4)$$

with $K > 0$ a characteristic dissipation scale. The Dynamic Symmetry Index is taken in the same canonical form as before,

$$DSI = 4OD(1 - |O - D|). (5)$$

This formula was already adopted in the earlier open-network paper because it vanishes when either dissipation-based order or informational diversity vanishes, and because it penalises strong imbalance between the two variables. The present question is whether this same construction remains informative once the underlying network contains several internally coupled cycles.

3. Why Multi-Cycle Topology Matters

A single-throughput chain tests only a limited aspect of the theory. In such a model, increasing the external affinity mainly increases the current through one dominant route. The entropy decline at high forcing can then be interpreted as simple canalisation along that route. Although valuable, this does not address the more interesting question of how Dynamic Symmetry behaves when the system has a genuine choice among several recurrent pathways.

Multi-cycle topology changes the problem in three related ways. First, it increases structural degeneracy. The same net conversion between boundary reservoirs may be achieved through distinct internal cycles and sub-cycles, which creates multiple mesoscopic routes for the same macroscopic transformation. Second, it introduces competing or cooperating probability currents. In a nonequilibrium steady state, current need not pass monotonically from input to output; it may circulate internally through loops before contributing to net throughput. Third, it creates the possibility that increasing forcing does not simply amplify one current, but reorganises the distribution of current across the network. Such reorganisations are thermodynamically meaningful because entropy production is distributed over several cycle affinities rather than a single one.

For Dynamic Symmetry Theory, these features are not peripheral. They supply the very kind of internal richness that the theory is meant to describe. A multi-cycle network may retain high internal diversity even while strongly driven, provided several pathways remain statistically active. Conversely, a strongly driven network may undergo pathway selection, reducing diversity even as entropy production rises. The DSI should therefore not be read merely as a response to forcing magnitude. It is a response to how forcing interacts with topology to shape the coexistence of variability and organisation.

4. Multi-Cycle Network Construction

Consider an open network with two chemostatted boundary species, A and B, and three internal species, X, Y, and Z. The internal reactions are chosen so that the system contains two overlapping cycles and one bypass path:

$$A \rightleftharpoons X, X \rightleftharpoons Y, Y \rightleftharpoons B,$$

$$X \rightleftharpoons Z, Z \rightleftharpoons Y, Z \rightleftharpoons B. \quad (6)$$

The reactions $X \rightleftharpoons Y$, $X \rightleftharpoons Z$, and $Z \rightleftharpoons Y$ form an internal triangular loop, while the channels $A \to X \to Y \to B$ and $A \to X \to Z \to B$ provide two parallel boundary-to-boundary pathways. The link $Z \rightleftharpoons Y$ couples those pathways, so the network is not merely branched but genuinely multi-cycle.

The species A and B are maintained by large reservoirs. Their fixed concentrations or chemical potentials impose a nonequilibrium chemical affinity $\Delta\mu = \mu_A - \mu_B$, which drives the system away from detailed balance when nonzero, just as in the earlier minimal CRN application. The internal state is given by the copy-number vector

$$s = (n_X, n_Y, n_Z)$$

subject to a finite truncation

$$n_X + n_Y + n_Z \leq N,$$

chosen for numerical tractability. This produces a finite Markov state space on which the chemical master equation is defined. Because the network contains multiple loops, stationary currents may circulate through the internal triangle as well as through the net boundary conversion from A to B.

This geometry is the minimal extension that already introduces genuine thermodynamic competition. It is simple enough to simulate directly, but rich enough to distinguish between total dissipation and how that dissipation is distributed across alternative pathways.

5. Stochastic Formulation

At the mesoscopic level, the open network is described as a continuous-time Markov jump process over the truncated copy-number state space. Each elementary reaction contributes a birth-death type transition. For example, the boundary uptake and release step $A \rightleftharpoons X$ changes n_X by ± 1, while $X \rightleftharpoons Y$ moves one copy between the corresponding internal coordinates. Under mass-action scaling, each transition is assigned forward and backward rates determined by reaction constants and, for the boundary steps, by the fixed reservoir concentrations.

The generator $\mathcal{L}$ of the process is therefore a sparse matrix acting on probability vectors over the finite state space. The stationary distribution p is obtained by solving

$$\mathcal{L}^{\top} p = 0, \sum_s p_s = 1. \quad (7)$$

Once the stationary distribution is known, both the stationary Shannon entropy and the stationary current structure follow directly.

The presence of the internal triangle means that some stationary currents may contribute to net throughput from A to B, while others represent internal circulation among X, Y, and Z. Both types contribute to entropy production, but they need not affect informational diversity in the same way.

This distinction is precisely why the multi-cycle case is interesting. A network may dissipate strongly not only because it moves matter rapidly from one boundary to another, but because it sustains organised internal circulation. Such circulation can either broaden the stationary distribution by activating several regions of state space or narrow it by favouring a restricted current-carrying manifold. The DSI measures the outcome of that competition rather than presuming it.

6. Numerical Protocol

A practical implementation can closely follow the protocol already used in the minimal open-network paper. The finite state space is chosen by fixing a maximum total internal occupancy N, for example $N = 20$ or $N = 30$, and retaining all states satisfying $n_X + n_Y + n_Z \leq N$. Symmetric internal rate constants provide a neutral starting point, while the nonequilibrium driving is introduced by varying the ratio of effective reservoir activities associated with A and B. Equivalently, one may vary the chemical affinity $\Delta\mu$ over a geometric or linear range extending from near equilibrium to strong forcing.

For each forcing value, the numerical steps are as follows. First, construct the generator matrix from the reaction list and rate constants. Second, compute the stationary distribution as the normalised null vector of $\mathcal{L}^{\top}$. Third, evaluate the Shannon entropy using equation (1). Fourth, compute the stationary probability currents on all directed edges and from them the entropy production rate using equation (3). Fifth, normalise the results to obtain D, O, and the DSI using equations (2), (4), and (5). Sixth, examine not only the scalar values but also the distribution of currents across the two boundary pathways and the internal triangular loop.

This final step is particularly important in the multi-cycle case. The earlier single-channel model required only one dominant current measure to interpret its behaviour. Here, by contrast, one wants to know whether the DSI peak coincides with active participation of several cycles, with balanced use of the two parallel routes, or with some other current-redistribution phenomenon.

7. Expected Regime Structure

The theory predicts a three-regime structure closely analogous to, but richer than, the minimal case. Near equilibrium, the stationary distribution spreads relatively broadly over the accessible state space because no strong nonequilibrium preference selects one route over another. The diversity variable can therefore remain appreciable, but the entropy production rate stays small because persistent organised currents are weak. The resulting DSI is low because macroscopic dissipative order has not yet become substantial.

At strong forcing, the opposite problem emerges. The network is driven far from equilibrium, and entropy production rises as the imposed affinity generates strong stationary currents. Yet those currents need not remain topologically plural. One pathway may become dominant, or the internal loop may lock into a preferred circulation pattern. The stationary distribution then concentrates on a comparatively narrow family of current-carrying states. Diversity falls, and the DSI declines again despite high dissipation.

Between these extremes lies the regime of real interest. In this intermediate range, the network is driven strongly enough to sustain nontrivial boundary throughput and internal loop currents, but not so strongly

that one route suppresses all others. Several pathways remain statistically occupied, and several currents remain thermodynamically active. In that region the diversity and order variables can both be appreciable, so the DSI is expected to display an interior maximum, just as it did in the minimal single-throughput case. The difference is that the present peak has a stronger structural interpretation: it marks not simply intermediate forcing, but intermediate forcing in the presence of topological plurality.

8. Interpreting Multi-Cycle Dynamic Symmetry

The importance of a multi-cycle application is not merely numerical. It changes what the DSI peak means. In the minimal open network, a peak at intermediate forcing could be described as a balance between throughput and state-space spread. In the present model, the same peak may instead be read as evidence that several competing routes and an internal loop remain jointly active before extreme forcing induces pathway selection. The dynamically symmetric regime is therefore not just one of moderate flux. It is one of maintained pathway richness under sustained nonequilibrium organisation.

This interpretation aligns well with the broader ambition of Dynamic Symmetry Theory. The theory is not chiefly interested in systems that are simply half ordered and half disordered. It is interested in systems whose internal variability remains functionally available while large-scale organisation is maintained. A multi-cycle CRN supplies a much better analogue of that idea than a simple chain, because it contains both structure and choice. The stationary distribution measures how much of that choice remains alive; entropy production measures how strongly the structure is being driven and maintained.

Another advantage of the multi-cycle setting is that it suggests finer observables for later work. One may define cycle-resolved entropy production, pathway occupancy fractions, or route-participation entropies in addition to the global variables H and σ. These are not needed to define the present DSI, but they provide diagnostic tools for understanding why the DSI behaves as it does. They may also prove valuable if later applications seek early-warning signals of topological reorganisation rather than only scalar peaks.

9. Relationship to the Earlier Open-Network Paper

The earlier chemostatted CRN paper established three points. It located the theory within standard nonequilibrium thermodynamics, it defined an open-network DSI in terms of normalised Shannon entropy and entropy production rate, and it illustrated the resulting index numerically in a minimal network with a clear interior maximum at intermediate forcing. The present paper preserves all three commitments. What changes here is the burden of the example. In the earlier work, the main challenge was to show that the bridge between information entropy and boundary dissipation could be made explicit at all. In the present work, that bridge is taken as given. The challenge is instead to show that it remains informative when network topology becomes richer than a single dominant chain. That extension matters because it begins

to answer the obvious criticism that a minimal CRN is only a toy example. A multi-cycle network is still far from a biochemical cell, but it is no longer thermodynamically trivial.

This paper should therefore be read as an application note rather than as a replacement for the original open-network formulation. The architecture remains unchanged. What is tested is its robustness under topological enrichment.

10. Time-Local DSI Extension

The stationary analysis developed so far answers only one part of the larger programme. It shows how Dynamic Symmetry may be located across parameter regimes once a network has relaxed to a nonequilibrium steady state, but it does not yet describe how the same balance evolves in time as a system approaches or departs from such a regime. That temporal question is essential if the framework is to contribute to early-warning analysis. The broader stochastic formulation of Dynamic Symmetry already emphasised that a useful index should be definable on evolving probability laws rather than only on stationary ones. In the CRN setting, this suggests a natural next step: promote the stationary variables to sliding-window observables computed along stochastic trajectories.

To demonstrate the method, a first time-local implementation was carried out on the earlier minimal chemostatted CRN rather than on the multi-cycle network itself. This choice was deliberate. The minimal network already has a well-understood stationary DSI profile, so it provides a clean environment in which to test the machinery of local estimation before transferring that machinery to the topologically richer case. The underlying stochastic process was simulated by a Gillespie algorithm on the same finite triangular state space used in the stationary study, while the boundary forcing was slowly ramped from the vicinity of the stationary DSI optimum toward a more strongly driven regime. At each sliding time window, an empirical state-occupancy distribution was estimated from the trajectory, a windowed Shannon entropy was computed, and a local dissipation proxy was obtained from empirical jump-current asymmetries. These were normalised to define time-local variables $D_{\mathrm{loc}}(t)$, $O_{\mathrm{loc}}(t)$, and

$$DSI_{\mathrm{loc}}(t) = 4D_{\mathrm{loc}}(t)O_{\mathrm{loc}}(t)(1 - |D_{\mathrm{loc}}(t) - O_{\mathrm{loc}}(t)|).$$

The methodological point is more important than the specific numerical values. Once the index is written in this sliding-window form, standard early-warning statistics can be attached to it. In the present implementation, rolling variance and lag-one autocorrelation were computed on the local DSI series. These are familiar diagnostics in the critical-transitions literature and are natural first tests of whether dynamic symmetry becomes measurably unstable before a forced loss of balance. The present paper does not claim that these indicators are already calibrated or universally reliable. It claims only that the DSI architecture

can now be tracked as a temporal object, which is the precondition for any serious early-warning programme.

11. Numerical Early-Warning Test

The time-local test proceeded in two stages. First, the stationary parameter scan for the minimal chemostatted chain was reproduced to identify the forcing value at which the stationary DSI attains its interior maximum. In the present implementation, this reproduced the same qualitative structure reported in the earlier CRN paper: the information-based diversity variable remains high at weak forcing and then falls, the dissipation-based order variable rises with forcing, and the DSI reaches a distinct interior maximum rather than increasing monotonically. The numerical scan used in the code confirms that behaviour and locates a stationary optimum around $\log(A/B) \approx 1.91$, which was then used as the starting regime for the temporal experiment.

The second stage introduced a slow ramp in the chemostatted forcing, moving the system from the neighbourhood of that stationary optimum toward a more strongly driven regime. Along a Gillespie trajectory, local entropy, local dissipation, and local DSI were estimated on sliding windows. The resulting time series show the expected qualitative picture. The local DSI remains comparatively elevated before the forcing ramp, then becomes more irregular as the transition unfolds, while the rolling variance and lag-one autocorrelation increase in the transition window relative to the pre-transition baseline. This does not by itself prove a universal early-warning law, but it is exactly the kind of behaviour one would want to see in a first proof-of-method experiment.

The phase-averaged summary makes the same point more plainly. In the pre-transition regime, the mean local DSI is higher and its rolling variance and autocorrelation are lower. During the transition window, the mean local DSI declines while both variance and autocorrelation rise; in the post-transition regime, the index remains depressed relative to the pre-transition level and the transition statistics relax only partially. Interpreted cautiously, this suggests that a loss of dynamic symmetry can be registered not only by the fall of the index itself but also by increased temporal instability in the index before or during the onset of a new forced regime. That is enough to justify extending the same tests to the multi-cycle network, where current redistribution across multiple routes should make the interpretation of time-local DSI even richer.

12. Implications for the Multi-Cycle Case

Although the explicit time-local simulation reported here was carried out on the minimal open chain, its relevance to the multi-cycle paper is direct. The purpose of the stationary multi-cycle construction is to show that the Dynamic Symmetry framework survives greater topological complexity; the purpose of the time-local extension is to show that the same framework survives temporal localisation. Taken together, these two steps move the programme in the right direction. The first strengthens structural credibility, while the second opens the door to predictive use.

In a genuinely multi-cycle network, one expects the time-local picture to be more informative than in the minimal chain. A strongly driven transition need not simply collapse the distribution onto one throughput route; it may first redistribute current across parallel pathways or alter circulation through internal loops. A local DSI decline in that setting could therefore be decomposed into topologically meaningful causes: loss of route plurality, strengthening of a dominant loop, or collapse of occupancy across formerly active regions of state space. This is precisely why the time-local extension belongs inside the present paper rather than in a separate methodological appendix. The multi-cycle application is not only about richer stationarity; it is about setting up the next stage of dynamic testing.

13. Limitations

Several limitations should be stated clearly. First, the analysis remains finite-state and numerical. The truncation $n_X + n_Y + n_Z \leq N$ is mathematically convenient, but any quantitative conclusions will depend on checking that the observed regime structure is stable as N increases. Second, the present network is only one choice among many multi-cycle architectures. Different loop couplings, different stoichiometries, or asymmetric rate constants may shift the location or even the number of interior maxima.

Third, entropy production is used here as a global order-like variable exactly as in the earlier open-network paper. That choice is well motivated in chemostatted nonequilibrium systems, but it does not capture every aspect of chemical organisation. A network could in principle dissipate strongly while exhibiting qualitatively different forms of internal structure. For this reason, cycle-resolved observables should be treated as complementary diagnostics rather than competitors to the global DSI.

Fourth, the present paper does not yet present a theorem guaranteeing an interior maximiser for all multi-cycle CRNs. The strongest claim justified at this stage is more modest. The Dynamic Symmetry construction remains mathematically coherent in a topologically richer open network, and there are strong theoretical grounds to expect the same intermediate-force regime seen in the minimal case, now with a clearer pathway-based interpretation.

The time-local extension adds an equally important second dimension. By implementing sliding-window entropy, a local dissipation proxy, and rolling early-warning statistics on Gillespie trajectories of the minimal chemostatted network, the paper shows that Dynamic Symmetry can be followed as a temporal object rather than only as a stationary one. The first results are intentionally modest, but they are encouraging. The local DSI declines as the system is forced away from the stationary optimum, while rolling variance and lag-one autocorrelation rise through the transition window.

14. Conclusion

This paper has argued that the transition from a minimal chemostatted chain to a multi-cycle open chemical reaction network is the right next step for the Dynamic Symmetry programme, and that this structural extension should be paired from the outset with a time-local version of the index. The earlier open-network construction already showed that informational diversity and boundary dissipation can be placed on a common mathematical footing by combining stationary Shannon entropy with stationary entropy production in a bounded Dynamic Symmetry Index. The present paper preserves that architecture while placing it in a network whose topology is substantially richer. Instead of one dominant pathway, the system now contains overlapping loops, parallel channels, and the possibility of genuine internal circulation. That difference matters because it turns the DSI from a test of simple throughput balance into a test of whether organised dissipation and pathway plurality can coexist within the same nonequilibrium steady state.

The central theoretical picture is now sharper than before. Near equilibrium, the network may explore a broad region of state space, but it does not yet sustain strong organised currents, so the dissipation-based order variable remains small. Under very strong forcing, the opposite problem appears: entropy production rises, but the stationary distribution becomes canalised as one route or one circulation pattern begins to dominate. Between these limits lies the regime of interest, where several pathways remain statistically active while the network is already being thermodynamically organised by boundary driving. If the DSI peaks there, it does so for a stronger reason than in the minimal example. It indicates not merely moderate forcing, but the joint maintenance of nonequilibrium structure and internal route diversity.

The time-local extension adds a further reason to take the framework seriously. By implementing sliding-window entropy, a local dissipation proxy, and rolling early-warning statistics on Gillespie trajectories, the present paper shows that Dynamic Symmetry can be followed as a temporal object rather than only as a stationary one. The first results are intentionally modest, but they are encouraging. The local DSI declines as the system is forced away from the stationary optimum, while rolling variance and lag-one autocorrelation rise through the transition window. The phase-averaged summary points in the same direction. Dynamic Symmetry may therefore be detectable not only as a structural property of stationary

regimes, but as a measurable balance that becomes unstable in time before a new forced regime is fully established.

The present result is conceptually decisive. The Dynamic Symmetry framework is not confined to a single-chain toy model, and it is not confined to purely stationary analysis. It survives a meaningful increase in topological richness and admits a first temporal implementation. That is sufficient to justify the next phase of the programme.