

Review



Cite this article: Nam K-M, Martinez-Corral R, Gunawardena J. 2022 The linear framework: using graph theory to reveal the algebra and thermodynamics of biomolecular systems.

Interface Focus **12**: 20220013.

<https://doi.org/10.1098/rsfs.2022.0013>

Received: 14 March 2022

Accepted: 25 April 2022

One contribution of 6 to a theme issue 'Time-keeping and decision-making in living cells; Cell fate decisions (Part II)'.

Subject Areas:

systems biology

Keywords:

linear framework, graph theory, rational parametrization, Hopfield barrier, Hill functions, post-translational modification systems, gene regulation

Author for correspondence:

Jeremy Gunawardena

e-mail: jeremy@hms.harvard.edu

[†]These authors contributed equally.

The linear framework: using graph theory to reveal the algebra and thermodynamics of biomolecular systems

Kee-Myoung Nam[†], Rosa Martinez-Corral[†] and Jeremy Gunawardena

Department of Systems Biology, Harvard Medical School, Boston, MA 02115, USA

K-MN, 0000-0002-3594-6141; RM-C, 0000-0003-3600-3601; JG, 0000-0002-7280-1152

The linear framework uses finite, directed graphs with labelled edges to model biomolecular systems. Graph vertices represent biochemical species or molecular states, edges represent reactions or transitions and labels represent rates. The graph yields a linear dynamics for molecular concentrations or state probabilities, with the graph Laplacian as the operator, and the labels encode the nonlinear interactions between system and environment. The labels can be specified by vertices of other graphs or by conservation laws or, when the environment consists of thermodynamic reservoirs, they may be constants. In the latter case, the graphs correspond to infinitesimal generators of Markov processes. The key advantage of the framework has been that steady states are determined as rational algebraic functions of the labels by the Matrix-Tree theorems of graph theory. When the system is at thermodynamic equilibrium, this prescription recovers equilibrium statistical mechanics but it continues to hold for non-equilibrium steady states. The framework goes beyond other graph-based approaches in treating the graph as a mathematical object, for which general theorems can be formulated that accommodate biomolecular complexity. It has been particularly effective at analysing enzyme-catalysed modification systems and input–output responses.

1. Introduction

The linear framework is a graph-based approach to biochemical reaction networks [1–3] that arose from studying post-translational modification systems at steady state [4,5]; for reviews, see [6,7]. When it can be deployed, the framework enables a nonlinear biochemical reaction network based on mass-action kinetics to be decomposed into a coupled set of graphs, each of which has a linear dynamics, from which the name 'linear framework' is derived. The linear operator is given by the Laplacian matrix of the graph. The main value of the framework has been in giving mathematical access to the steady states of a network as rational algebraic functions of the parameters. This rational algebraic approach has been particularly useful in analysing enzyme-catalysed biochemical networks, such as post-translational modification systems, [4,8–11], and also input–output responses, which arise in several domains, including biochemistry, gene regulation and pharmacology, [7,12–18].

The language of graphs allows essential biological requirements to be specified while leaving other details implicit in the structure of the graph. The framework thereby offers a way to derive general theorems that can rise above the overwhelming molecular complexity that confronts us in modern biology [4,13,16,19]. An example is the rational parametrization theorem for multisite post-translational modification systems, described in §3, which shows that the steady state of such a system is rationally determined just by the enzymes, independently of the number of sites of modification. Despite such successes, the range of applicability of the framework to biochemical

reaction networks remains unclear. We will return to this question in the Discussion once we have seen the framework in action.

The case of a single graph coupled to reservoirs of chemical potential corresponds to a continuous-time, finite-state Markov process for which the linear Laplacian dynamics is the master equation [2]. The linear framework thereby offers a graph-based approach to Markov processes. This Markovian interpretation allows thermodynamic concepts like entropy to be specified within the framework, which reduces to equilibrium statistical mechanics for systems at thermodynamic equilibrium. However, importantly, it also provides a setting in which non-equilibrium statistical mechanics is exactly solvable in rational algebraic form (§4). Despite important progress in non-equilibrium physics, the subject remains much less developed than its equilibrium counterpart. The role of energy expenditure in information processing has been particularly elusive and the linear framework offers ways to approach this problem [7,12–14,20–22]. Among the insights to have emerged are that of the ‘Hopfield barrier’ for an information processing task and the identification of the widely used Hill function as the Hopfield barrier for the sharpness of an input–output response [14] (§5).

The mechanisms discussed here provide some of the underpinnings for the study of biological decision-making and time keeping, the subjects of the theme issue in which this paper appears. Post-translational modification systems are found in most biological processes. The pioneering study by Goldbeter and Koshland of a cycle of modification and demodification (§3) provided one of the earliest examples of a molecular switch that can participate in decision-making [23]. The methods described here show how the characteristics of such switches can be understood in terms of the properties of the constituent enzymes (§3). As for input–output responses, they are often convenient ways to summarize the function of mechanisms for which some degree of sharpness is required. Sharpness is often a necessary requirement for the dynamical bistability that may underlie a decision or for the instability that gives rise to a limit cycle oscillation. Hill functions are frequently used to introduce such sharpness in mathematical models (e.g. [24,25]). However, the Hill functions have always lacked rigorous justification, having been originally introduced only as a convenient mathematical family for fitting data [26]. Our results give them, for the first time, a biophysical justification and clarify the conditions under which they can arise biologically (§5).

The sections that follow describe these developments in more detail, drawing in part on recent unpublished work. We follow here the perspective set out in [27], that mathematical models in biology are not descriptions of reality but, rather, accurate descriptions of our (pathetic) assumptions. Hopefully, the assumptions are reasonable in the light of current biological knowledge and of the questions being asked. Having made those assumptions transparent, we try to show how the resulting conclusions are useful for understanding biology. We will also point out some of the open questions and unexplored directions that arise.

The results described here are distributed across papers on a variety of topics, with many of the details buried in supplementary material. We are most grateful to John Tyson and Tim Holt for the opportunity to bring this material together.

We hope that what follows will be useful in introducing readers to the capabilities of the linear framework and to some of the open questions that arise from this approach.

2. Laplacian dynamics and the Matrix-Tree theorems

We introduce the framework here and outline some of the basic results that provide the foundation for the applications described in subsequent sections. The framework is based on finite, simple, directed graphs with labelled edges. (A simple graph has no more than one directed edge from one vertex to another and has no self loops.) An example graph is shown in figure 1*a*. Let G denote a linear framework graph. The vertices of G correspond to molecular entities or states and are usually indexed $1, \dots, N$, where N is the number of vertices. The edges correspond to reactions or transitions and are denoted $i \rightarrow j$. The labels correspond to rates, with dimensions of $(\text{time})^{-1}$, and are denoted $\ell(i \rightarrow j)$. The labels are particularly significant: they can be algebraic expressions which encode the potentially time-dependent interactions between the system described by the graph and its surrounding environment. It is the labels which thereby incorporate the nonlinearity that is usually present. In this section, however, the labels will be treated as abstract parameters, as in k_1, \dots, k_9 in figure 1*a*. We will explain how labels are interpreted in §3.

A linear framework graph gives rise to a dynamics which is most simply explained as one-dimensional chemistry: each edge is treated as a chemical reaction under mass-action kinetics with the edge label as the rate constant. Since each edge has only one source vertex, the dynamics must be linear and is therefore described by a matrix equation,

$$\frac{du(t)}{dt} = \mathcal{L}(G) \cdot u(t). \quad (2.1)$$

Here, $\mathcal{L}(G)$ is the so-called *Laplacian matrix* of G , of size $N \times N$, and $u(t)$ is the time-dependent $N \times 1$ column vector of vertex ‘concentrations’. These may be actual concentrations if the vertices represent molecular entities or they may be probabilities if the vertices represent molecular states; we will take advantage of both interpretations below. The 4×4 Laplacian matrix of the example graph in figure 1*a* is shown in figure 1*b*. Since the chemistry simply moves matter around the graph, without creating or destroying it, there is an obvious conservation law,

$$u_1(t) + \dots + u_N(t) = u_{\text{tot}}. \quad (2.2)$$

Equation (2.2) corresponds in matrix terms to the column sums of $\mathcal{L}(G)$ being zero (figure 1*b*), or $1^T \cdot \mathcal{L}(G) = 0$, where 1 denotes the all-ones column vector and A^T denotes the transpose of matrix A . $\mathcal{L}(G)$ is a singular matrix and has 0 as an eigenvalue. (We try to keep the notation as light as possible, leaving it to the context to disambiguate the meaning where necessary. If we need to, we will use brackets, as in $u_{\text{tot}}(G)$, or a subscript, as in $i \rightarrow_G j$, to specify which graph is being discussed.) Synthesis and degradation of matter can be accommodated within the graph [1] and some of the consequences are explored in [3].

Laplacian matrices for various kinds of graphs have been widely studied, often with different orientations, signs and scalings [28]. They can be seen from one perspective as




Figure 1. Linear framework graph, Laplacian matrix and spanning trees. (a) A strongly connected linear framework graph on four vertices, indexed by $1, \dots, 4$, with symbolic edge labels, k_1, \dots, k_9 . (b) The Laplacian matrix corresponding to the graph in (a). (c) Example spanning forests of the graph in (a), demarcated by magenta edges, whose source and target vertices belong to the forest, with root sets (cyan) $\{4\}$, $\{1, 4\}$, $\{1, 3, 4\}$ and $\{1, 2, 3, 4\}$, reading from left to right. (d) All eight spanning trees (magenta edges) rooted at vertex 1 (cyan) for the graph in (a), with the corresponding monomial from the MTT in equation (2.4) given below.

discrete approximations to the Laplace–Beltrami operator on a Riemannian manifold [29], so that equation (2.1) resembles a discrete-space diffusion equation. Another interpretation is that equation (2.1) is the master equation of a continuous-time, finite-state Markov process. In this case, the linearity is to be expected, rather than being the consequence of what looks like a rather trivial chemistry. We will explore this stochastic interpretation of equation (2.1) and its thermodynamic implications in §4, after we have looked into the labels in §3.

The framework was introduced to study timescale separation, in which the system under study is assumed to have reached a steady state and one wants to simplify its behaviour accordingly [1]. This is an old technique in physics and it enters biology in the work of Michaelis & Menten [30]. Since equation (2.1) is linear, it is seemingly easy to solve analytically but this is not straightforward to do in terms of the edge labels, for instance by using eigenvalues. Instead, we

can exploit graph theory in the form of the *Matrix-Tree Theorems*. First, note that a steady state of equation (2.1), u^* , must be in the kernel of the Laplacian, $u^* \in \ker \mathcal{L}(G)$. If G is *strongly connected* then it is not hard to show that this kernel is one dimensional [1],

$$\dim \ker \mathcal{L}(G) = 1. \quad (2.3)$$

Recall that a graph is strongly connected if, given any ordered pair of distinct vertices, (i, j) , there is a path of directed edges from i to j ,

$$i = i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_{k-1} \rightarrow i_k = j.$$

The example in figure 1a is strongly connected but ceases to be if the edges $1 \rightarrow 2$ and $1 \rightarrow 3$ are removed. Strong connectivity depends only on the *structure* of the graph, which is to say its vertices and edges alone; it is independent of the edge labels. The non-strongly connected case is well understood [2] but we will not need it here.

Despite its simplicity, equation (2.3) is the nub of what follows here. To set up an initial condition, there are as many degrees of freedom for distributing matter among the vertices of the graph as there are vertices. However, the one-dimensional chemistry of the Laplacian dynamics (equation (2.1)) ensures that, when the graph is strongly connected, there is only a single degree of freedom left in the steady state (equation (2.3)). The steady state forgets the initial conditions but remembers the graph. The essential subsequent step is to determine the steady state in terms of the edge labels, which is where the Matrix-Tree Theorems come into play.

The Matrix-Tree Theorems reveal a remarkable property of Laplacian matrices: their minors—the determinants of the square submatrices obtained by removing equal-sized subsets of rows and columns—can be expressed in terms of *spanning forests* of G . A spanning forest, F , is a subgraph of G which contains all vertices in G (spanning), has no cycles when edge directions are ignored (forest) and for which each vertex has at most one outgoing edge (which orients the forest). Examples are shown in figure 1c. A spanning forest depends only on the structure of a graph. The vertices of F with no outgoing edges are called roots. If a forest has only one root, it is a *tree*, so that forests are disjoint unions of trees! Let $v(G) = \{1, \dots, N\}$ denote the vertex set of G . If $\emptyset \neq U \subseteq v(G)$, let $\Phi_U(G)$ be the set of spanning forests of G that are rooted at U . If G is strongly connected, then it is easy to see that $\Phi_U(G) \neq \emptyset$. The classical *First-Minors MTT* goes back to Gustav Kirchhoff [31], although the version for labelled directed graphs first appears in Bill Tutte's PhD thesis [32]; for a statement and proof, see [2]. The *All-Minors MTT* first appeared in the Czech mathematical literature [33] but independent proofs have been given in English (e.g. [34]). These sources give the full statements. We will focus here on using the First-Minors MTT (hereafter, MTT) but the All-Minors MTT comes into its own for some of the new directions mentioned in the Discussion.

Because the first minors give the cofactors of $\mathcal{L}(G)$, and hence its adjugate matrix, it is not difficult to see that the MTT leads to a canonical basis element $\rho(G) \in \ker \mathcal{L}(G)$, which is simply a right eigenvector for the zero eigenvalue. If H is any graph, let $\lambda(H)$ denote the product of the labels over all the edges in H : $\lambda(H) = \prod_{j \rightarrow k \in H} \ell(j \rightarrow k)$. The MTT implies that

$$\rho_i(G) = \sum_{T \in \Phi_{\{i\}}(G)} \lambda(T). \quad (2.4)$$

The spanning forests in equation (2.4) have only one root and are therefore *spanning trees*. In words, equation (2.4) says that, to obtain a basis element in $\ker \mathcal{L}(G)$, we take the product of the labels over the edges of each spanning tree rooted at i and add these products up over all such trees. This can be done by enumerating all the trees rooted at a vertex, as in figure 1d, which becomes the source of the complexity that we will discuss later (§4). The force of the MTT lies in the terms on the right-hand side of equation (2.4) all being positive. Determinants of sub-matrices typically have both positive and negative terms but in the case of Laplacian matrices, there are massive cancellations that result in the sum of positive monomials in equation (2.4). The MTT starkly reveals the underlying positive polynomial dependence on the parameters.

Because of equation (2.3), the steady-state vector, $u^*(G)$, must be a scalar multiple of $\rho(G)$, which can be written more concisely as $u^*(G) \propto \rho(G)$. The proportionality constant can be dispensed with by division,

$$\frac{u_i^*(G)}{u_j^*(G)} = \frac{\rho_i(G)}{\rho_j(G)}, \quad (2.5)$$

for any $i, j \in v(G)$, or by normalization to the total,

$$u_i^*(G) = \left(\frac{\rho_i(G)}{\rho_1(G) + \dots + \rho_N(G)} \right) u_{\text{tot}}, \quad (2.6)$$


which exposes the last remaining degree of freedom at steady state in the total concentration of material, u_{tot} , from equation (2.2).

The positive polynomial dependence arising from the MTT leads to steady-state concentrations being rational algebraic functions of the parameters that are always positive for positive parameter values. We see in equation (2.6) the origin of the rational functions that pervade molecular biology. The Michaelis–Menten and King–Altman formulae in enzyme kinetics [35,36], the Monod–Wyman–Changeux and Koshland–Némethy–Filmer formulae in allostery [37,38] and the Ackers–Johnson–Shea formula in gene regulation [39] can all be derived from equation (2.6) using appropriate graphs [1]. The linear framework not only offers a systematic procedure in place of ad hoc methods, it also reveals the hidden linearity within biochemical reaction networks. (The linearity occurs in the dynamical variables and is captured in equation (2.1); by contrast, the parametric dependence of the dynamics is highly nonlinear, as equation (2.4) shows.) To appreciate how this method works in practice, we need to understand how the edge labels provide an interface between the graph and its environment that brokers linearity out of nonlinear biochemistry.

3. Enzyme kinetics and multisite modification systems

Figure 2a depicts a modification–demodification cycle, sometimes called a Goldbeter–Koshland loop after the pioneering study that first clarified the switch-like behaviour of this molecular circuit [23]. The modification is shown as a phosphorylation to make it familiar but it could be any small-molecule modification, such as methylation or acetylation. (Ubiquitination and other polypeptide modifications on proteins have a different biochemistry that is not covered here.) The substrate, S , could be a protein, so that the process would be one of post-translational modification [40,41], but it could also be some other modifiable molecule such as a carbohydrate or a lipid. S is interconverted between its unmodified state, S_0 , and its modified state, S_1 , by enzymes; for phosphorylation, the forward enzyme, E , is a kinase and the reverse enzyme, F , is a phosphatase. We will explain here how modification systems like this can be decomposed into linear framework graphs.

We will start by showing how enzymes give rise to graphs which describe their mechanisms. In the literature, it is customary for the mechanisms of E and F to be described by the Michaelis–Menten reaction scheme,







Figure 2. A modification cycle in the linear framework. (a) Schematic of a modification–demodification cycle or Goldbeter–Koshland loop with phosphorylation as the modification. (b) Hypothetical reaction network for the kinase mechanism in the grammar of equation (3.2). (c) Hypothetical reaction network for the phosphatase mechanism with a dead-end complex Y_7 , also in the grammar of equation (3.2). (d–f) Linear framework graphs for the kinase, phosphatase and substrate, respectively (see text). The labels which are simple rates are omitted for clarity; the others, which involve entities from the environment of the graph, are in blue font. (g) Metagraph for the modification cycle, where the edges show the labelling relationships between the linear framework graphs represented by the vertices (see text).

Michaelis and Menten undertook their *in vitro* studies of enzyme kinetics by measuring initial rates of reaction, with essentially no product present [35,42]. They were therefore entirely justified in treating product release as irreversible and adopting equation (3.1). However, equation (3.1) has continued to be used in circumstances, like modification cycles, where there can be substantial amounts of product. *Strongly irreversible* mechanisms like equation (3.1), in which product is unable to rebind to enzyme, are biophysically unrealistic in contexts like this [6,9,43], as Michaelis and Menten well understood [30]. The dangers of assuming strong irreversibility have been repeatedly pointed out [43–45]: it is a limiting assumption that is highly singular [9] and can lead to serious overestimation of the parametric robustness of different dynamical behaviours [11,44]. It is sometimes claimed that equation (3.1) is necessary because it exhibits irreversibility. However, as a matter of thermodynamics, all biochemical transitions are reversible, even if the reverse transition has a relatively much lower rate. Irreversibility is another limiting assumption, which may be appropriate in some physiological contexts. For example, protein kinases and phospho-protein phosphatases appear to operate nearly irreversibly in most cellular contexts. A thermodynamically acceptable way to model this would be to use a *weakly irreversible* mechanism, in which product may rebind but cannot be converted back to substrate. This can be readily accommodated within the linear framework.

The framework allows the use of any enzyme mechanism that can be built up out of the following ‘grammar’

of reactions:



Here, S_* can be substrate or product and the Y 's are arbitrary intermediate enzyme–substrate complexes [9,43]. (There is a mild condition: the linear framework graph for the enzyme mechanism, which we will introduce below, needs to be strongly connected. Reaction networks that do not satisfy this requirement are implausible.) The graph-theoretic machinery described in §2 allows any such mechanism to be described by four aggregated parameters at steady state, as we will see below. This allows us to pay some attention to the details of enzyme mechanisms, which biochemists have worked so hard to disentangle [46–48], rather than persisting with the convenient delusion that all enzymes work the same way as in equation (3.1).

Figure 2b shows a hypothetical reaction network for the kinase E in the grammar of equation (3.2). In its forward direction, when it converts S_0 to S_1 , it has two routes for binding S_0 . After the catalytic step between the intermediate complexes Y_1 and Y_3 , it releases its product, S_1 , also by two routes. The two routes can be thought of as an approximation to what actually happens in a kinase, which has to bind two substrates, S_0 and ATP, in some order and release two products, S_1 and ADP, in some order. Here, ATP and ADP are not explicitly represented as substrates and their effects are assumed to be absorbed into the reaction rates [43]. Enzyme biochemists describe figure 2b as a random-order

bi-bi mechanism [49]. Phosphatases are much simpler single-substrate hydrolases and a hypothetical mechanism for F in the grammar of equation (3.2) is shown in figure 2c, in which we have included a dead-end complex. Both mechanisms are fully reversible but they can be made weakly irreversible by removing the reaction from Y_3 to Y_1 for E and from Y_6 to Y_5 for F .

The approximation in describing the kinase, which is explained in more detail in [43, §2.1], arises because the grammar in equation (3.2) excludes the possibility of a substrate binding to an intermediate complex. This exclusion is important for what follows. It is an interesting question as to whether this constraint can be lifted and a more elaborate grammar developed that would allow multiple-substrate reactions to be more accurately accommodated.

Figure 2d,e shows the corresponding linear framework graphs G_E and G_F for E and F , respectively. The vertices are the states of the respective enzyme, namely the free enzyme and whatever intermediate complexes are involved. Vertices have been indexed for convenience by the corresponding enzyme states. The edges correspond to the biochemical reactions in figure 2b,c. Note that both G_E and G_F are strongly connected. The labels encode how the enzyme states interact with their environment of substrate and product. In particular, an edge label can be an algebraic expression involving concentrations, such as $\ell(E \rightarrow Y_2) = k_2[S_0]$. Here, $[X]$ denotes the time-dependent concentration of the biochemical entity X and k_2 is the second-order rate constant for the bimolecular binding reaction. In this way, the entities that are not represented by vertices, such as S_0 and S_1 , are encoded in the edge labels. It is not difficult to see that the linear Laplacian dynamics in equation (2.1) for the graph of any mechanism built up from the grammar in equation (3.2) is just a rewriting of the nonlinear biochemistry of the mechanism, as given by mass-action kinetics [4]. The key requirement here is the *uncoupling condition*: a biochemical entity that appears in an edge label cannot also be represented by a vertex. If this condition is not met, then equation (2.1) would cease to be linear. Uncoupling implies a sharp separation between the vertices and the graph's environment, for which the labels act as the interface.

Now that we have understood how enzymes can be described in the linear framework, we can turn our attention to the more interesting problem of analysing the modification cycle in figure 2a as a whole. We will show that this cycle can be decomposed into three linear framework graphs, two of which are the enzyme graphs that we have already defined. The missing ingredient is a graph for the substrate, G_S (figure 2f).

A word on the terminology that we will use from now on: we will write 'parameters' or 'rates' for rate constants like k_2 in G_E and distinguish them from 'edge labels', which may be algebraic expressions involving both parameters and concentrations; 'aggregated parameters' are rational algebraic functions of the parameters.

The graphs G_E and G_F are easily derived from the reaction mechanisms in figure 2b,c, respectively, and correspond to a rewriting of the dynamics; their edge labels involve the time-dependent concentrations of entities in their environment. The situation is more subtle for G_S . Here, the vertices are the substrate forms, S_0 and S_1 , and the edges represent the action of the enzymes. But what should the labels be? It is hard to imagine labels that can capture the time-dependent

dynamics. However, labels can be found which reproduce the steady state. The details go back to [4]; they are described in linear framework terms in [9,43] and summarized in [11].

The edge labels in G_S are defined in terms of aggregated parameters that come from the enzyme mechanisms which implement the edge in question. It is easier to see how these aggregated parameters are defined by stepping back from the particular enzyme examples in figure 2 and working through the construction more generally. Let us consider an arbitrary enzyme X that converts between S_0 and S_1 using any mechanism based on equation (3.2) with generic intermediate complexes, Y_i . Let G_X be the corresponding enzyme graph, which we assume to be strongly connected. It follows from equation (2.5) that

$$[Y_i]^* = \left(\frac{\rho_{Y_i}(G_X)}{\rho_X(G_X)} \right) [X]^*, \quad (3.3)$$

where we have used $[-]^*$ to denote concentrations at steady state, so that $[X]^* = u_X^*(G_X)$. Recalling the MTT prescription in equation (2.4), and following the examples in figure 2d,e, we see that $\rho(G_X)$ may involve $[S_0]^*$ and $[S_1]^*$. Specifically, any spanning tree rooted at Y_i will have an edge label involving either $[S_0]^*$ or $[S_1]^*$ but not both, since such labels only occur on outgoing edges from X . Hence, $\rho_{Y_i}(G_X)$ is linear in $[S_0]^*$ and $[S_1]^*$. By the same token, a spanning tree rooted at X has no labels involving either $[S_0]^*$ or $[S_1]^*$. Hence, $\rho_X(G_X)$ is a constant. Accordingly, we can write, at steady state,

$$[Y_i]^* = (\kappa_{0,Y_i}^X [S_0]^* + \kappa_{1,Y_i}^X [S_1]^*) [X]^*. \quad (3.4)$$

We call the aggregated parameters, κ_{0,Y_i}^X and κ_{1,Y_i}^X , *reciprocal generalized Michaelis–Menten constants* (rgMMCs). They tell us through equation (3.4) how much of each substrate is locked up in each intermediate complex at steady state. By equation (2.4), they are given as sums of products of parameters over spanning trees of G_X rooted at Y_i . These expressions may be very complicated, reflecting the complexity of X 's reaction mechanism, but it is only the aggregate rgMMCs from equation (3.4) that are needed for what follows. The constraints imposed in the grammar of equation (3.2) are essential for the linearity of equation (3.4), from which these aggregated parameters emerge.

We can now define *total generalized catalytic efficiencies* (tgCEs), $c_{0,1}^X$ and $c_{1,0}^X$, which capture the contribution of the enzyme, X , at steady state, to the production rate of S_1 from S_0 , which is $c_{0,1}^X$, or of S_0 from S_1 , which is $c_{1,0}^X$. It is easier to describe how this works by returning to the modification cycle in figure 2. For example, for G_E in figure 2d,

$$c_{0,1}^E = \ell(Y_3 \rightarrow E) \kappa_{0,Y_3}^E + \ell(Y_4 \rightarrow E) \kappa_{0,Y_4}^E \quad \text{and} \\ c_{1,0}^E = \ell(Y_1 \rightarrow E) \kappa_{1,Y_1}^E + \ell(Y_2 \rightarrow E) \kappa_{1,Y_2}^E.$$

The edge labels on the substrate graph, G_S , are then given in terms of these tgCEs and the corresponding steady-state enzyme concentrations, as shown in figure 2f. Hopefully, it should now be reasonably clear how this works for any enzyme reaction mechanism. The key point is that, with this labelling, the steady state of G_S , as given by equation (2.6), exactly reproduces the steady state of the underlying biochemical reaction network.

We have now decomposed the modification cycle in figure 2 into three linear framework graphs, G_E , G_F and G_S . This graph decomposition is very useful because it allows

us to undertake *variable elimination*: we can use the graphs to iteratively eliminate the steady-state dynamical variables in terms of the free enzyme concentrations, $[E]^*$ and $[F]^*$. The latter can then be determined by the enzyme conservation laws. This elimination procedure works as follows.

First, using the substrate graph, G_S , in figure 2f, it follows from equation (2.5) that

$$\frac{[S_1]^*}{[S_0]^*} = \frac{\rho_{S_1}(G_S)}{\rho_{S_0}(G_S)} = \frac{c_{0,1}^E[E]^* + c_{0,1}^F[F]^*}{c_{1,0}^F[F]^* + c_{1,0}^E[E]^*}. \quad (3.5)$$

Using equation (3.5) together with equation (3.4) in the substrate conservation law at steady state allows $[S_0]^*$ and $[S_1]^*$ to be expressed as rational functions of $[E]^*$ and $[F]^*$, with S_{tot} as an additional parameter. This procedure introduces two further aggregated parameters, the *total trgMMCs* (trgMMCs), κ_0^X and κ_1^X :

$$\kappa_0^X = \sum_{Y_i \in \nu(G_X)} \kappa_{0,Y_i}^X \quad \text{and} \quad \kappa_1^X = \sum_{Y_i \in \nu(G_X)} \kappa_{1,Y_i}^X,$$

which determine how much of the substrate is locked up in the intermediate complexes at steady state:

$$\sum_{Y_i \in \nu(G_X)} [Y_i]^* = (\kappa_0^X [S_0]^* + \kappa_1^X [S_1]^*) [X]^*.$$

We can substitute the expressions for $[S_0]^*$ and $[S_1]^*$ into equation (3.4) to obtain expressions for the $[Y_i]^*$ as rational functions of $[E]^*$ and $[F]^*$. We have now eliminated all of the variables at steady state in favour of $[E]^*$ and $[F]^*$. Finally, we can substitute the resulting expressions into the conservation laws for the enzymes, to give a pair of simultaneous equations for $[E]^*$ and $[F]^*$,

$$\text{and} \quad \left. \begin{aligned} E_{\text{tot}} &= R_E([E]^*, [F]^*) \\ F_{\text{tot}} &= R_F([E]^*, [F]^*). \end{aligned} \right\} \quad (3.6)$$

Here, $R_E(-, -)$ and $R_F(-, -)$ are both polynomials in their two arguments and the tgCEs and trgMMCs become the effective parameters along with S_{tot} . The modification system in figure 2 has been reduced at steady state to the solution of a pair of simultaneous polynomial equations for the enzyme concentrations (equation (3.6)), with all the other steady-state variables being determined as rational algebraic functions of the enzyme concentrations.

Where does all this machinery get us? First, at the enzyme level, it gives us a set of four aggregated parameters for describing any enzyme mechanism in the grammar in equation (3.2). There is a tgCE and a trgMMC for both the forward and the reverse direction of the enzyme. These parameters determine the nature of the reaction mechanism: when converting S_0 to S_1 , X is strongly irreversible if, and only if, $\kappa_1^X = 0$ and X is weakly irreversible if, and only if, $\kappa_1^X > 0$ and $c_{1,0}^X = 0$. There is no longer any mathematical reason for persisting with strongly irreversible enzyme mechanisms when they are not appropriate.

Second, and as a consequence of the first point, the Goldbeter–Koshland loop can be analysed without making unreasonable assumptions about the enzyme mechanisms [9,43]. The switch that the loop implements, between S_0 and S_1 , cannot be infinitely sharp, as it appears to be when the enzymes are strongly irreversible [23]. Formulae can be given for both the sharpness and the dynamic range of the switch, in terms of the aggregated parameters for arbitrary enzyme mechanisms based on the grammar in equation

(3.2) [9]. This brings unexpected insights into the phenomenon of enzyme bifunctionality, in which, for example, a single enzyme acts as both a kinase and a phosphatase. A case in point is the enzyme 6-phosphofructo-2-kinase/fructose-2,6-bisphosphatase, which regulates the switch between glycolysis and gluconeogenesis in the mammalian liver [50]. Why an enzyme should both modify and demodify a substrate has seemed quite mysterious. The methods above show how this architecture enables sharp switching which is also coherent across heterogeneous individual cells, which seems to be particularly relevant for zonation in the liver [9].

Third, and most importantly, the example in figure 2 works in much greater generality for any modification system with multiple sites. Suppose that we have a substrate with any number of sites of modification. Proteins can have an extraordinary number of such sites. The tumour-suppressor p53, which is the protein most frequently mutated in cancers, has over 100 sites of PTM [51]. Even if those sites were all simple binary modifications, which they are far from being, this would imply over 10^{30} possible global patterns of modification, or *modforms*, on a single molecule of p53 [40]. Of course, only a tiny fraction of these modforms can be present at any time but the amount of molecular state that is potentially available is jaw-dropping. (What on Earth is it used for [40]?) Suppose that each of the enzymes responsible for making and unmaking these modifications follows some mechanism based on the grammar in equation (3.2). A given enzyme could have several modforms as substrates and could use a different mechanism on each one. Following the same procedure as described above for the simple modification cycle in figure 2a, the multisite modification system can be decomposed at steady state into a collection of enzyme graphs and a substrate graph; the steady-state enzyme concentrations are determined as the solution to a system of simultaneous polynomial equations arising from the enzyme conservation laws, as in equation (3.6); and all the other steady-state dynamical variables are expressed as rational algebraic functions of the enzyme concentrations [4]. The mass-action differential equations describing such a multisite modification system are polynomial and the steady state is therefore the positive part of a real algebraic variety [52]. The main result can be informally stated in the language of algebraic geometry as follows.

Rational parametrization theorem (RPT) [4]. *The steady-state variety of such a multisite modification system is birationally equivalent to the variety determined by the conservation laws for the enzymes (equation (3.6)).*

This result was unexpected. Despite the apparent intractability of multisite modification systems, with the number of variables scaling exponentially with the number of sites, the steady-state variety suffers an immense reduction to requiring only the enzyme concentrations to specify it completely, no matter how many sites are present or the complexity of the enzymology. The steady states of all the other variables are given by rational algebraic functions of the steady-state enzyme concentrations and these functions can be explicitly constructed from the graphs. This is not a reduction in the dimension of the variety but, rather, in the dimension of its ambient space. The steady-state variety of a multisite modification system has dimension zero, as does the variety defined by the enzyme conservation laws: both varieties consist of isolated points. The salient issue is that the former variety

lives in a potentially extremely high-dimensional ambient space while the latter lives in an ambient space whose dimension is the number of enzymes. It is far more straightforward to deal with the latter than with the former.

The RPT has several consequences. The enzyme conservation laws in equation (3.6) capture the essential nonlinearity of a multisite modification system. These equations can have multiple solutions, which accounts for the multistability that has been found in such systems [5,53]. Everything else in the system is determined as rational functions of these nonlinear solutions. For a system with just two enzymes, equation (3.6) offers ‘pseudo-nullclines’ for determining steady states as the intersections of two curves in \mathbb{R}^2 [5]. By combining the ambient dimension reduction provided by the RPT with homotopy continuation methods for solving polynomial equations, available in software tools like Bertini [54], it becomes possible to map out the ‘parameter geography’ of multistability in multisite modification systems [11]. Several conjectures and problems arise from this, which we lack the space to discuss further here.

The implications of the RPT seem not to have been widely appreciated. It is not unusual to see steady-state varieties of multisite modification systems being algebraically solved by ad hoc methods or, what is worse, being determined by numerical simulation, which requires making specific choices about mechanisms and parameter values rather than exploiting the generality which the RPT offers. The RPT demonstrates how the graph-based approach of the linear framework allows us to rise above some of the molecular complexity within biology, and our lack of knowledge about the details of mechanisms, while still drawing useful conclusions.

The RPT raises several questions that have yet to be explored. For example, although the statement of the RPT captures what has been most useful—that the enzymes alone determine the steady state—it glosses over what happens to the parameters. These are also greatly reduced, from the many original parameters required for the multisite modification system to the few aggregated parameters required for the enzyme equations. Here too, the latter are rational algebraic functions of the former. A more complete statement of the RPT would formalize the relationship between parametrized families of varieties that are rationally related in this way. The language to do this may already exist within algebraic geometry. This may seem an unfamiliar strategy in today’s systems biology but the benefit of such formalization is that it can place the result in a broader mathematical context and clarify the kinds of theorems that we may want to look for elsewhere [55–57].

Further questions arise in asking how far the RPT can be pushed. The possibility was mentioned above of expanding the grammar in equation (3.2) to accommodate reaction mechanisms in which multiple substrates participate. More broadly, does the general form of the RPT, in which a small subset of variables carries the nonlinearity and rationally determines all other variables at steady state, apply to biochemical reaction networks beyond those for single-substrate multisite modification systems?

A potential way to formulate this may be in terms of the ‘metagraph’ determined by the labelling relationship. A preliminary definition could be given as follows. The vertices of the metagraph are linear framework graphs and there is a directed edge from graph G_i to graph G_j if the steady states

of G_i enter into the labels of G_j . This is a coarse description that takes no account of the details of the labelling, which may need to be specified more carefully, as in [8], but it suffices to pose the problem. The metagraph for the RPT is a star, with the substrate graph at its centre and the enzyme graphs leading to it (figure 2g). What happens when the metagraph is more complicated? It seems plausible that, when the metagraph is a tree, a similar result to the RPT might hold, with the reduced set of variables, akin to the ‘enzymes’ in the RPT, corresponding to those metagraph vertices with no incoming edges. When the metagraph is not a tree and has cycles, the situation seems fundamentally different. In this case, the iterative construction of rational functions that is possible in a tree must be replaced by something like the fixed point of a recursive functional equation. Can there still be proper subsets of the variables in terms of which all others are rationally determined? The scope of the metagraph idea has not yet been explored but it seems likely that it could accommodate a broad class of enzyme-catalysed biochemical networks, at least at steady state. If results similar to the RPT could be found in this setting, it would show that biochemical reaction networks have a far richer algebraic structure at steady state than has been apparent up to now.

Mercedes Pérez Millán and Alicia Dickenstein have introduced MESSI (modifications of type enzyme–substrate or swap with intermediates) systems, which substantially generalize the multisite modification systems described here [58], albeit without the graph infrastructure. Other approaches have also been introduced for variable elimination in biochemical reaction networks [59,60]. These developments may provide the tools for exploring the graph-based questions raised above. One of the potential advantages of casting these questions in terms of linear framework graphs is that they can be more readily related to the underlying biochemical architectures that are found in the cell (figure 2) and the aggregated parameters that arise can also be given biochemical meaning. Metatrees of Goldbeter–Koshland loops were introduced in [8], where they were shown to have unique steady states, but they have not yet been placed in the setting of the RPT.

4. Markov processes and stochastic thermodynamics

The edge labels in a linear framework graph, G , and specifically those which involve entities outside G , may be dealt with in different ways. We have seen one of the more elaborate ways in §3, where the entities correspond to vertices in some other graph. A simpler possibility is that they just occupy a buffer, from which they bind and unbind to the system described by G . If the ligand L is such an entity binding to the vertices of G , then the edge labels involving L through binding would carry the concentration $[L]$ of free ligand. In the absence of synthesis and degradation of L , there would be a conservation law for total ligand,

$$L_{\text{tot}} = [L] + \sum_{i \in \nu(G), L \in i} u_i(G), \quad (4.1)$$

where $L \in i$ denotes those vertices of G in which L is bound. Equation (4.1), when considered at steady state, would play a similar role to equation (3.6) in determining

$[L]^*$. If L were also bound to other graphs, there would be similar contributions to equation (4.1) from those graphs, and if there were multiple ligands present, a system of simultaneous equations would emerge that is similar to equation (3.6).

A simplification of equation (4.1) is to assume that there is so much ligand present that binding and unbinding to G does not change its free concentration, so that $[L] \approx L_{\text{tot}}$. We call such a buffer a *reservoir*, by analogy with thermodynamic reservoirs such as heat baths. In classical thermodynamics, exchanging heat with a reservoir does not change the reservoir's temperature. When chemistry is incorporated, exchanging molecular particles with a reservoir does not change its chemical potential, or, in our context, the concentration. If all the entities interacting with a graph are in reservoirs, then their concentrations become constants, like the parameters, which remain unchanged over the timescale of the graph dynamics.

The advantage of this approximation is that the linear framework graph then corresponds to a Markov process [2]. By the latter term, we mean a finite-state, continuous-time, time-homogeneous Markov process, $X(t)$, specified by a conditional probability distribution, $\Pr(X(t+h)=j|X(t)=i)$ for $h \geq 0$, which is independent of t . Such a process gives rise to a graph whose vertices are the states of the Markov process and for which there is an edge $i \rightarrow j$ if, and only if, the infinitesimal transition rate from i to j is positive, in which case this rate becomes the edge label,

$$\ell(i \rightarrow j) = \lim_{h \rightarrow 0} \frac{\Pr(X(t+h)=j|X(t)=i)}{h} > 0. \quad (4.2)$$

Provided the Markov process is sufficiently well behaved for the infinitesimal rates to exist, linear framework graphs and Markov processes are equivalent. In effect, the graph specifies the infinitesimal generator of the corresponding process. From this perspective, the linear Laplacian dynamics in equation (2.1), with $u(t)$ now the vector of vertex probabilities and $u_{\text{tot}} = 1$ in equation (2.2), becomes the master equation, or Kolmogorov forward equation, for the time evolution of probabilities [2, theorem 4]. We have moved from the deterministic world of concentrations in §3 to the stochastic world of probabilities with the same mathematics.

The graph rarely makes an appearance in the theory of Markov processes, perhaps because the latter has been more concerned with behaviour in the transient regime. (The developments mentioned in the Discussion may be of interest in this respect.) It is also not unusual to see master equations described in a balance of flux form,

$$\frac{du_i}{dt} = \sum_{j \neq i} (r(i, j)u_j - r(j, i)u_i),$$

where $r(i, j)$ is what we would call $\ell(j \rightarrow i)$, again without the underlying graph or the Laplacian operator being made explicit. One of the messages of this paper is to point out how useful it can be to take a graph-theoretic approach, even to objects like Markov processes that seem very familiar from a different perspective.

The Markovian interpretation allows us to bring thermodynamics into the picture. For this, we make the standing assumption that whatever graph, G , we are dealing with is 'reversible', so that if $i \rightarrow_G j$, then also $j \rightarrow_G i$. In the interpretation of the graph, it is important that $j \rightarrow i$ represents the

reverse process to $i \rightarrow j$ and not merely some other means to get from j to i [61]. If $i \rightleftharpoons j$ is any pair of such reversible edges, then the ratio of their edge labels may be given a thermodynamic interpretation in terms of the entropy change during the transition from i to j ,

$$\frac{\ell(i \rightarrow j)}{\ell(j \rightarrow i)} = \exp\left(\frac{\Delta S^{\text{env}} + (S_j^{\text{sys}} - S_i^{\text{sys}})}{k_B}\right). \quad (4.3)$$

Here, ΔS^{env} is the total change in the entropy of the environmental reservoirs during the transition $i \rightarrow j$, S_i^{sys} is the internal entropy of vertex i and k_B is Boltzmann's constant. Equation (4.3), which asserts that the label ratio is the exponential of the total entropy change, is known as 'local detailed balance', where 'local' signifies that this is an assertion about a pair of reversible edges. Equation (4.3) goes back to the pioneering work of Terrell Hill [62] and Jürgen Schnakenberg [63], who began using graph-based descriptions to study biophysical systems. It has now been justified in many different settings [64,65], for example, when the vertices of the Markov process describe the mesoscopic states of some biomolecular system, such as an enzyme or a molecular motor, operating within the cellular environment, with which the system exchanges particles and energy.

Equation (4.3) allows us to interpret the central concept of thermodynamic equilibrium. In mathematics, the word 'equilibrium' is often used as a synonym for 'steady state'. This is not so in physics. Thermodynamic equilibrium is a very special kind of steady state with remarkable properties. In linear framework terms, a graph is at thermodynamic equilibrium if there is no entropy change along any cycle of reversible edges. To formalize this, let $P: i_1 \rightleftharpoons i_2 \rightleftharpoons \dots \rightleftharpoons i_{k-1} \rightleftharpoons i_k$ be any path of reversible edges in G and let $\mu(P)$ denote the product of label ratios along the path,

$$\mu(P) = \left(\frac{\ell(i_1 \rightarrow i_2)}{\ell(i_2 \rightarrow i_1)}\right) \dots \left(\frac{\ell(i_{k-1} \rightarrow i_k)}{\ell(i_k \rightarrow i_{k-1})}\right). \quad (4.4)$$

If P is a cycle, so that $i_1 = i_k$, then it follows from equation (4.3) that

$$\mu(P) = \exp\left(\frac{\Delta S^{\text{env}}}{k_B}\right),$$

where ΔS^{env} is the total entropy change in the reservoirs over one traversal of the cycle. It follows that the graph can be at thermodynamic equilibrium if, and only if, $\mu(P) = 1$ for all cycles of reversible edges, P . This is the 'cycle condition' for G . It is sometimes known in the reaction network literature as Wegscheider's condition but its thermodynamic significance was first appreciated by Gilbert Lewis [66]. It says that the product of the edge labels going clockwise around any cycle equals the product going counterclockwise. It is sufficient to check this condition only on the finitely many cycles in any basis of cycles (for details of cycle bases, see [21]).

An equivalent formulation of thermodynamic equilibrium is that it is a steady state, u^* , in which any pair of reversible edges, $i \rightleftharpoons j$, is in flux balance,

$$u_i^* \ell(i \rightarrow j) = u_j^* \ell(j \rightarrow i), \quad (4.5)$$

irrespective of any other transitions in which the vertices i and j are engaged. It is not hard to check that if the cycle condition holds, then any steady state satisfies equation (4.5) and,

conversely, if equation (4.5) holds in some steady state, then the cycle condition also holds. Equation (4.5) is ‘detailed balance’ or ‘microscopic reversibility’. It brings out the remarkable nature of thermodynamic equilibrium, in which pairs of reversible transitions become effectively decoupled from each other.

At thermodynamic equilibrium, the prescription for the steady state that comes from the MTT in equation (2.4) can be greatly simplified. Choose any vertex of G as a reference, which we take to be 1 by convention. Let P_i be any path of reversible edges from 1 to vertex i and let $\mu_i(G) = \mu(P_i)$, which is well defined no matter which P_i is chosen because of the cycle condition. It is not hard to show, again exploiting the cycle condition, that, in vector terms, $\mu(G) = \rho(G)/\rho_1(G)$, where $\rho(G)$ is the canonical basis element in $\ker \mathcal{L}(G)$ that comes from the MTT in equation (2.4). Hence, $\mu(G)$ is an alternative canonical basis element for $\ker \mathcal{L}(G)$. Using equation (2.6), we can write

$$u_i^*(G) = \frac{\mu_i(G)}{\mu_1(G) + \mu_2(G) + \dots + \mu_N(G)}. \quad (4.6)$$

Note that $\mu_1(G) = 1$. Using equation (4.3) and a little thermodynamics, the terms $\mu_i(G)$ yield the Boltzmann factors for the vertices, referred to vertex 1 as the zero for the free energy. Equation (4.6) is exactly the prescription for the steady state that comes from equilibrium statistical mechanics, with the denominator being the partition function for the grand canonical ensemble [67]. It follows from equations (4.4) and (4.6) that the only parameters needed at thermodynamic equilibrium are the label ratios, $\ell(i \rightarrow j)/\ell(j \rightarrow i)$, not the individual labels.

Linear framework graphs often appear unnecessary when working at thermodynamic equilibrium. This is especially so for physicists. They consider the free energies of vertices to be the fundamental parameters at equilibrium, so that graph edges are irrelevant. What the edges provide, however, are representations of the underlying molecular mechanisms, which are often crucial for interpreting the biology. For example, the widely used hypercube graph structure \mathcal{C}_n can represent the binding and unbinding of ligands to n sites on a biomolecule [14]. Figure 3a shows the structure that we denote \mathcal{C}_{2+1} , reflecting two ligands, one of which binds at 2 sites and the other at 1 site. Hypercube graphs can be used to model a receptor [68], an allosteric system [19] or a gene-regulation system [14]. Such graphs enable *higher-order cooperativities* (HOCs) to be defined at thermodynamic equilibrium, in which binding of a ligand at one site can be modulated by the ligand binding status of multiple other sites [7,14,19]. Prior to this, cooperativity had been largely studied as a pairwise effect involving just two binding sites but there is now compelling evidence that higher-order effects are biologically important [19]. HOCs provide an alternative parametrization of the free-energy landscape and a rigorous language in which to express cooperative interactions involving multiple ligands and sites [7,19]. (HOCs are associated with edges, not vertices, and the cycle condition at thermodynamic equilibrium implies that they are not independent quantities. However, it is straightforward to choose independent subsets of HOCs, in terms of which all other HOCs are rationally expressible [14].) An interesting question is how HOCs arise at a molecular level. It has been known since the pioneering work of Monod,

Wyman and Changeux that biomolecular systems with multiple interchanging conformations can create pairwise cooperativity through *allostery* [37,38]. The linear framework can be used to calculate the HOCs that arise from any allosteric conformational ensemble—using the technique of *coarse graining* described in §5—and to prove that such ensembles can implement any pattern of HOCs that is achievable at thermodynamic equilibrium [19].

Despite the algebraic similarity in the formulas for the steady state in equations (4.6) and (2.6), there is a profound distinction between $\mu(G)$ at equilibrium and $\rho(G)$ away from equilibrium, whose significance becomes clear from the examples calculated below in equations (4.7) and (4.8). At equilibrium, $\mu(G)$ shows that it is only necessary to pick any one path to a given vertex to calculate the steady-state probability of that vertex (up to normalization by the denominator in equation (4.6), which is the partition function). Only the edge label ratios along the path are required. By contrast, away from equilibrium, $\rho(G)$ shows that every path to the vertex is needed and the MTT in equation (2.4) provides the bookkeeping for this calculation. This poses a major difficulty because the enumeration of spanning trees becomes intractable even for simple graphs. The hypercube graph \mathcal{C}_2 has four spanning trees rooted at each vertex; \mathcal{C}_3 has 384 spanning trees; but \mathcal{C}_4 has 42 467 328 spanning trees [69]! If the combinatorics were not bad enough, $\rho(G)$ exhibits global parametric dependency: every edge label and parameter in the graph influences the steady-state probability of a vertex. This is startling in its singularity. We can imagine a graph at thermodynamic equilibrium in which the label on a single edge is modified so that the cycle condition is broken. The steady-state probabilities abruptly jump from being locally determined by $\mu(G)$ to being globally determined by $\rho(G)$, and all parameters, even those arbitrarily far from the edge whose label was changed, now have an effect on the probabilities. The extraordinary path-dependent complexity of calculating non-equilibrium steady states has been a formidable roadblock to progress.

The simple graph in figure 3b serves to illustrate the problem. Such a graph, whose structure is the hypercube \mathcal{C}_2 , could represent an allosteric system with two conformations, interchanging horizontally, and a single ligand binding site, with binding and unbinding in the vertical transitions. The concentration of the ligand, x , appears in the labels for the binding edges. There is a single minimal cycle. Let $f(x)$ be the steady-state probability of the ligand being bound, considered as a function of x ; with the notation used in §2, $f(x) = u_3^*(G) + u_4^*(G)$. Equivalently, for this simple example, $f(x)$ is also the average number of bound sites at steady state. $f(x)$ is an input–output function of the kind that we will discuss in more detail in §5. This function looks very different depending on whether the steady state is one of thermodynamic equilibrium or not. If the graph can reach thermodynamic equilibrium, the cycle condition that must be satisfied is

$$k_1 k_2 k_3 k_4 = k_5 k_6 k_7 k_8.$$

With this assumption, we can use equations (4.4) and (4.6) to show that the equilibrium input–output function is of Michaelis–Menten type,

$$f_{\text{eq}}(x) = \frac{Ax}{B + Ax}, \quad (4.7)$$

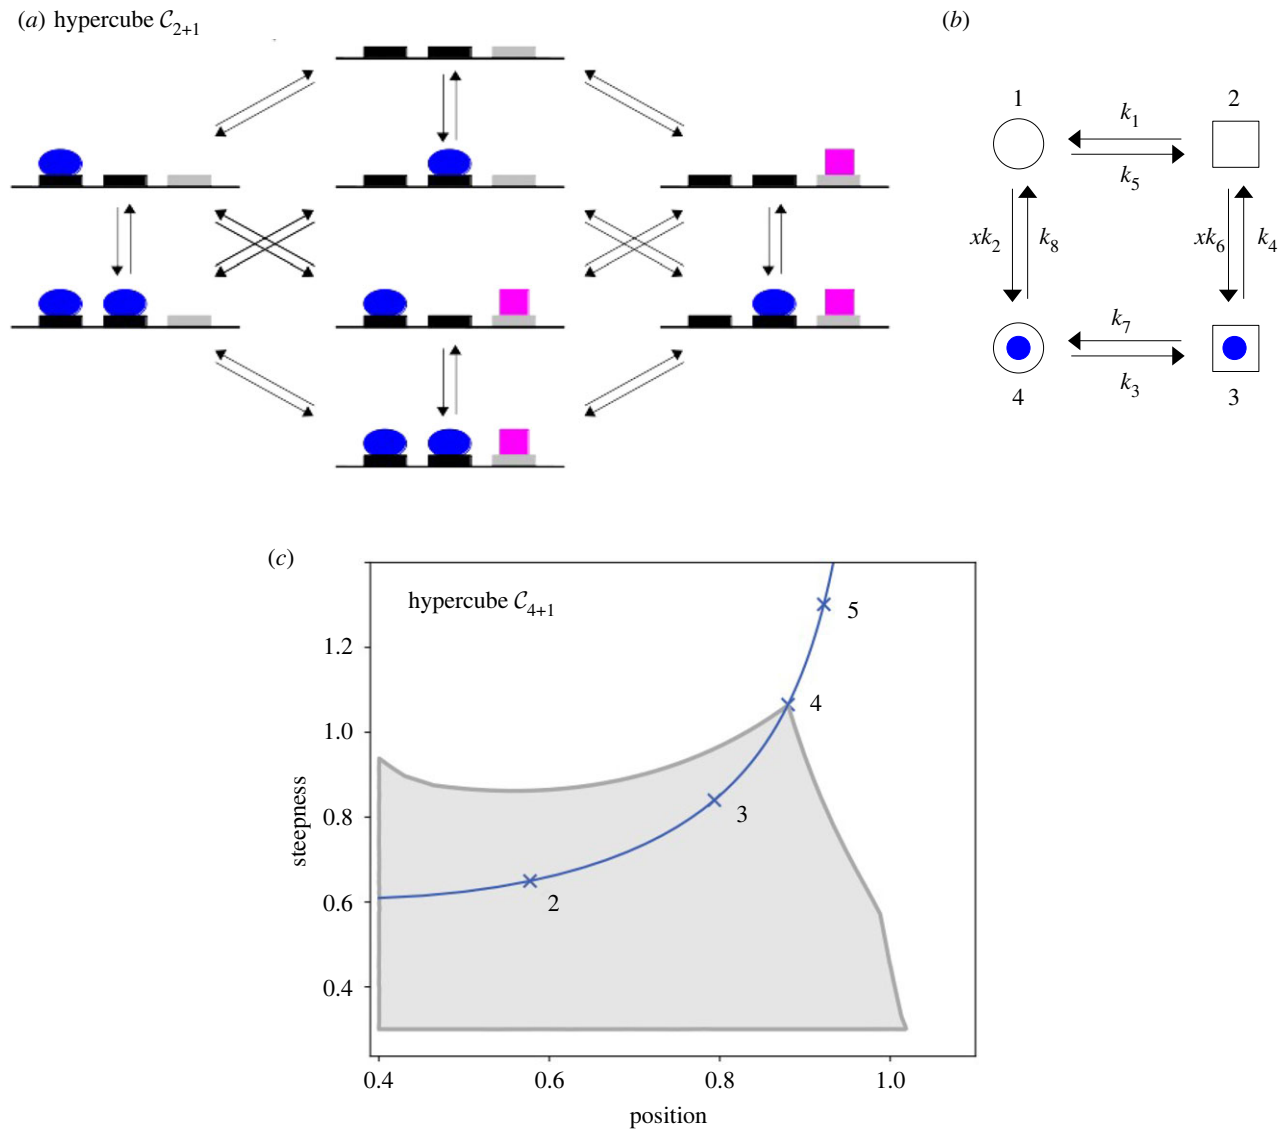


Figure 3. Graphs and position-steepness region. (a) Hypercube graph structure C_3 , in the form C_{2+1} , with three sites for binding of two ligands, a blue oval to 2 sites and a magenta square to 1 site. One interpretation could be a model of genetic regulated recruitment with the blue oval as a transcription factor and the magenta square as RNA Polymerase. Details of the labels and notation used with such models are given in [7]. (b) Linear framework graph to illustrate non-equilibrium complexity, representing a biomolecule in two conformations (circle, square) with a single site for binding of a ligand (blue disc). Vertices are indexed $1, \dots, 4$, parameters are k_1, \dots, k_8 and x is the concentration of ligand. (c) The grey region shows the thermodynamic equilibrium (p, s) region for C_{4+1} , the regulated recruitment example from (a). The input is the concentration of a transcription factor which binds to 4 sites and the output is the steady-state probability of RNA Polymerase bound to the 5th site. The outer curve in darker grey shows the asymptotic boundary of the (p, s) region, for arbitrary parameter values, with the region being truncated to the left and below. The blue curve is the Hill line, the locus of (p, s) points for the Hill functions, with integer Hill coefficients marked.

where the aggregated parameters are given by the appropriate ratios of rates,

$$A = \left(\frac{k_2}{k_8}\right) \left(1 + \frac{k_3}{k_7}\right) \quad \text{and} \quad B = 1 + \frac{k_5}{k_1}.$$

If the cycle condition is not satisfied, we can use equations (2.4) and (2.6) to calculate the input-output function. There are four spanning trees rooted at each vertex, which may be enumerated in a similar way to those in figure 1d. We find that the non-equilibrium input-output function has degree 2,

$$f_{ne}(x) = \frac{A^*x + B^*x^2}{C^* + D^*x + B^*x^2}, \quad (4.8)$$

where the aggregated parameters are given in terms of the

individual rates,

$$\begin{aligned} A^* &= k_1k_2k_3 + k_3k_5k_6 + k_5k_6k_8 + k_1k_2k_7 + k_5k_6k_7 + k_1k_2k_4, \\ B^* &= k_2k_3k_6 + k_2k_6k_7, \\ C^* &= k_1k_7k_8 + k_1k_3k_4 + k_1k_4k_8 + k_5k_7k_8 + k_3k_4k_5 + k_4k_5k_8 \\ \text{and } D^* &= A^* + k_2k_3k_4 + k_6k_7k_8. \end{aligned}$$

Even for the very simple example in figure 3b, the difference between equations (4.7) and (4.8) is striking in both rational structure and parametric complexity.

Physicists have largely dealt with the problem of path-dependent complexity by astute approximations, in which most of the spanning trees are ignored. But this tactic may be particularly problematic in biology in which phenomena often depend on the collective influence of many small effects. Mathematically, the complexity cannot be avoided:

equation (2.6) is an exact formula. It is necessary to somehow reorganize the complexity.

Recently, two approaches have emerged which have finally shed some light on this challenging problem. From the mathematical perspective, Pavel Chebotarev and Rafiq Agaev have previously interpreted the Faddeev–LeVerrier matrix inversion method for the Laplacian matrix of a graph [70]. Their method allows steady-state probabilities to be recursively and symbolically calculated without explicitly enumerating all the spanning trees. From the physical perspective, non-equilibrium steady-state probabilities can be interpreted as an average of the thermodynamic entropy production along paths [21]. The number of distinct path entropies that are required, and hence the complexity of the average, depends on the *entropy production index* of the graph [21]. If this index is low, calculations can be undertaken away from thermodynamic equilibrium which were previously infeasible [21]. These approaches have provided the first breakthroughs in dealing with path-dependent complexity away from equilibrium.

5. Hopfield barriers, coarse graining and Hill functions

The remarkable mathematical difference between the algebraic structure of steady-state probabilities at equilibrium (equation (4.6)) and away from equilibrium (equation (2.6)) strongly suggests that it must have significant biological implications. Surprisingly, these have only slowly emerged. At one level, the physics has always been obvious: we are only at thermodynamic equilibrium when we are dead! The role of energy expenditure in force generation or in pattern formation has been widely studied [71,72]. What has been much less clear, however, has been the implications of energy expenditure for cellular information processing. It was John Hopfield who provided the first insight into this in his pioneering work on kinetic proofreading [73]. The central idea behind his analysis can be restated and updated as follows: *given any information processing task, there is an upper bound to how well it can be undertaken by any mechanism that operates at thermodynamic equilibrium*. We call this bound the *Hopfield barrier* for that task [14]. The only way to exceed this barrier is for the underlying mechanism to operate away from thermodynamic equilibrium.

We have been interested in identifying the Hopfield barriers for various cellular information processing tasks, of which amplification has become the best understood. The problem can be formulated for steady-state input–output responses of biomolecular systems: by how much does the output change for a given change in an input? Here, the system is represented by a linear framework graph, G ; the input is the concentration, x , of some interacting ligand; and the output is any non-negative linear combination of steady-state probabilities,

$$f(x) = \sum_{1 \leq i \leq N} \lambda_i u_i^*(G), \quad \lambda_i \geq 0. \quad (5.1)$$

Such outputs are often measures of the overall state of the system that are appropriate for the specific biological context, as in the treatment of figure 3b in §4. Examples include receptors or allosteric systems responding to ligands, where the output could be the average number of bound sites [19],

and gene-regulation systems responding to transcription factors, where the output could be the probability of RNA Polymerase being bound to the gene promoter. The latter kind of input–output response has been widely used to study gene regulation [7], where it represents Mark Ptashne’s concept of regulated recruitment [74]. The hypercube graph C_n (figure 3a) can model such systems with n being the total number of binding sites. The internal complexities of the system are ignored in such a hypercube model and they are considered to exert their effect through the parameter values. The biomolecular system is assumed, under a time-scale separation, to have reached a steady state, with changes to the inputs taking place quasi-statically, with sufficiently small and slow increments that the system is not being forced but, rather, relaxes back to steady state after each change. Inputs have typically been examined singly in isolation, although there is no theoretical barrier to greater generality. Such steady-state, single-input, single-output responses have been widely measured and analysed and linear framework graphs are well suited to describing them [14,16,19,75].

The extent to which an input–output response is capable of amplification is usually quantified by some measure of sharpness. The hyperbolic Michaelis–Menten response, $x/(1+x)$, has typically been regarded as the baseline, with the Hill functions, $\mathcal{H}_h(x) = x^h/(1+x^h)$, with Hill coefficient, $h > 0$, regarded as displaying sharpness, or ultrasensitivity, which is positive when $h > 1$ and negative when $h < 1$. The Hill coefficient itself is often quoted as a measure of sharpness. The Hill functions can provide surprisingly good empirical fits to input–output response data, going back to their introduction by Archibald Hill to fit the data on oxygen binding to haemoglobin [26]. However, Hill functions, at least those which are not Michaelis–Menten, have no plausible justification [76,77]; they are merely an algebraically convenient functional family for fitting data. That is not the case for the Michaelis–Menten response, for which the linear framework yields theorems that explain its remarkably wide occurrence at the molecular level [16]. The empirical nature of Hill functions has frequently been overlooked, perhaps because they do their job of fitting so well, and they are often used in mathematical models when a sharp response is required without having to justify how the sharpness arises. It is commonplace to see them in models of gene-regulation networks, such as those used to model the transcription–translation circadian clock in metazoa [24]. Given the widespread use of Hill functions in the literature, the justification that we provide for them here may be of some interest.

To analyse sharpness, we use two non-dimensional, intrinsic measures, not just one, which is important for what follows. The definitions build on those previously given in [14] and are taken from the new findings described in [78]. In the generality considered here, input–output functions $f(x)$ are bounded rational functions for $x \in [0, \infty)$ and may not be monotonic. As a concrete example, consider the input–output function for a regulated-recruitment model of gene regulation, as described above. If the transcription factor binds at m sites, the corresponding graph structure is the hypercube C_{m+1} , shown in figure 3a for $m = 2$. The remaining site is bound by RNA Polymerase. The input, x , is the concentration of the transcription factor and the output is the steady-state probability of RNA Polymerase being

bound, which, in figure 3a, is the total probability coming from equation (4.6) for the vertices with the magenta square.

We will assume that the output value of f is naturally non-dimensional, as it is for the example just described, or that it has been otherwise normalized in a way appropriate to its context. Normalizing the input takes more care because an intrinsic normalization depends on f . Let $M(f)$ and $m(f)$ be the global maximum and minimum, respectively, of f : $M(f) = \max_{x \in [0, \infty)} f(x)$ and $m(f) = \min_{x \in [0, \infty)} f(x)$. We normalize the input to $x_{0.5}$, which is the minimum positive value for which $f(x_{0.5})$ is halfway between $m(f)$ and $M(f)$:

$$x_{0.5} = \min_{x \in (0, \infty)} \left\{ x \mid f(x) = \frac{M(f) + m(f)}{2} \right\}.$$

This yields the function $g(y) = f(yx_{0.5})$, where y is non-dimensional. The input normalization is important for what follows, although the specific definition of $x_{0.5}$ does not seem to qualitatively influence the final results. We then define the two measures of sharpness to be the ‘steepness’ of f , $s(f)$, which is the maximum of the absolute value of the derivative of g , and the ‘position’ of f , $p(f)$, which is the y value at which that maximum is attained,

$$s(f) = \max_{y \in (0, \infty)} \left| \frac{dg}{dy} \right| \quad \text{and} \quad p(f) = \operatorname{argmax}_{y \in (0, \infty)} \left| \frac{dg}{dy} \right|.$$

We can now plot *position–steepness regions*, abbreviated to (p, s) regions, assuming that the underlying graphs are at thermodynamic equilibrium, so that the cycle condition holds (§4). Equilibrium parameter values are randomly chosen in some range, the corresponding (p, s) points calculated and the resulting point cloud is incrementally grown until a boundary is reached. The algorithms are elaborated in [14,15,78]. Part of the equilibrium position–steepness region for the regulated-recruitment gene-regulation model described above, based on the hypercube structure \mathcal{C}_{4+1} , is shown in figure 3c.

Three consistent findings emerge from such plots. First, the (p, s) region becomes effectively asymptotically bounded as the parametric range is increased. Some care is needed to state this precisely because the boundary at zero position, $p = 0$, is anomalous: when position becomes very low, steepness can become very high. However, this arises from degenerate functions, as discussed in [14]. The more interesting behaviour occurs away from this boundary. To specify it, choose any $a > 0$, no matter how small. No matter what parameter values are chosen, that part of the (p, s) region within the quadrant $[a, \infty) \times [a, \infty)$ is bounded. Specifically, for all sufficiently large $A > a$, the region lies within the finite box $[a, A] \times [a, A]$. Figure 3c shows the (p, s) region lying within the finite box $[0.4, 1.2] \times [0.3, 1.2]$. The boundedness then leads to the second finding. The (p, s) region exhibits a cusp which lies on the *Hill line*, the locus of (p, s) points for the Hill functions (figure 3c). The tip of the cusp is always below the (p, s) point of the Hill function whose coefficient is the number of sites at which the input binds and the cusp approaches closer to this Hill point as the parametric range increases. So, if the input ligand binds at m sites of the biomolecular system, the cusp lies below, and asymptotically approaches, the (p, s) of \mathcal{H}_m . Third, if the assumption of thermodynamic equilibrium is dropped, (p, s) points can be found which lie above and to the right of the (p, s) of \mathcal{H}_m .

The asymptotic boundedness of the (p, s) region allows us to draw strong conclusions when experimental data fall outside this region [15]. We can claim that the underlying model is unable to account for the data without having to fit any data. Of particular interest is the possibility that energy expenditure is required to explain the data. This goes to the heart of the striking difference between prokaryotic and eukaryotic gene regulation [7]. The former appears to take place without any expenditure of energy, while the latter uses many sources of energy expenditure, including chromosome and nucleosome reorganization, post-translational modification of histones and co-regulators, and DNA methylation. We believe that this reflects a greatly increased capability for regulatory information processing in eukaryotic genomes. The methods described here provide the theoretical tools for addressing this question [7,13–15].

The three findings described above imply that the Hill function \mathcal{H}_m is the Hopfield barrier for the sharpness of input–output functions with m binding sites for the input ligand. This has been confirmed by numerical plotting of (p, s) regions for several models based on hypercube graphs with different output functions [14,15], including the regulated recruitment model whose (p, s) region is shown in figure 3c. However, models could be much more complex than those represented by hypercube graphs and could incorporate many details of the internal mechanisms of the underlying biomolecular system. It seems impossible, on the face of it, to make a definitive statement about the (p, s) regions for all of these models. It turns out, however, that a much stronger statement is possible of the universality of the Hill function as a Hopfield barrier. This relies on the technique of *coarse graining*, which was introduced in [19] to analyse allosteric conformational ensembles.

Suppose G is any strongly connected, reversible linear framework graph, which need not satisfy the cycle condition. Given any partition of the vertices of G , $\nu(G) = G_1 \cup \dots \cup G_s$ with $G_w \cap G_z = \emptyset$ when $w \neq z$, the coarse-grained graph $C(G)$ is defined as follows. $C(G)$ has vertices $1, \dots, s$ corresponding to the subsets of the partition. It has an edge $w \rightarrow_{C(G)} z$ whenever there exists $i \in G_w$ and $j \in G_z$ such that $i \rightarrow_G j$. $C(G)$ thereby inherits reversibility from G . Finally, $C(G)$ has labels given by

$$\ell(w \rightarrow_{C(G)} z) = Q \left(\sum_{j \in G_z} p_j(G) \right). \quad (5.2)$$

Here, the quantity Q is chosen arbitrarily to ensure the dimension of the label is $(\text{time})^{-1}$ but it plays no essential role, as we will see. It is shown in [19] that, with the labelling given by equation (5.2), $C(G)$ satisfies the cycle condition, even when G does not, and, furthermore, that the following coarse-graining equation is satisfied:

$$u_w^*(C(G)) = \sum_{i \in G_w} u_i^*(G). \quad (5.3)$$

In other words, the steady-state probability of being at w in $C(G)$ is the same as the steady-state probability of being at any of the vertices of G_w in G . This is exactly what we would expect from a coarse graining. Although equation (5.2) is not intuitive, it is the unique specification that allows equation (5.3) to hold, with $C(G)$ satisfying the cycle condition [19].

It is important to note that the prescription above is not a coarse graining of the dynamics. There is no reason to expect that $C(G)$ will approximate the dynamics of G , only that it reaches the expected steady state under the coarse graining. Because $C(G)$ satisfies the cycle condition and can therefore reach thermodynamic equilibrium, the only parameters that are required at steady state are the label ratios, so that the quantity Q in equation (5.2) cancels out and plays no role.

Coarse graining appears to be a powerful method for making general statements about steady-state behaviour, especially at thermodynamic equilibrium [19], and we expect that it will be broadly useful. To reveal the universality of the sharpness Hopfield barrier, we can consider any strongly connected, reversible graph, G , involving the binding of some ligand, L , to m sites. The graph may involve many other features, such as other ligands, allosteric conformations, co-regulators, patterns of post-translational modifications, scaffold configurations involving locations, and, in the context of gene regulation, chromatin state, nucleosomes, DNA methylation states, etc. (e.g. [12]). No matter how complex the graph, we can always coarse-grain it by bringing together all those vertices of G which exhibit the same pattern of binding of L . The resulting coarse-grained graph, $C(G)$, will have the same structure as the hypercube \mathcal{C}_m . (In fact, $C(G)$ may be a proper subgraph of \mathcal{C}_m but we will ignore this possibility in the interests of keeping the exposition simple; see [78] for a full discussion.)

For example, the hypercube structure \mathcal{C}_{2+1} in figure 3a, with the blue ligand as the input, would be coarse-grained into the hypercube structure \mathcal{C}_2 by making the subsets of the partition contain those vertices with the blue ligand bound at exactly the same sites. Since there are two binding sites for the blue ligand, there are four subsets in the partition, each of which contains two vertices, depending on whether or not the magenta ligand is bound.

Of course, the edge labels of $C(G)$, which arise from equation (5.2), could potentially be very complicated but, nevertheless, we know that $C(G)$ satisfies the cycle condition with this labelling. What happens to an output function on G ? It can be shown that, provided G itself is at thermodynamic equilibrium, any output function on G can be rewritten as a non-negative linear combination of steady-state probabilities of $C(G)$,

$$\sum_{i \in \nu(C(G))} \lambda_i w_i^*(C(G)), \quad \lambda_i \geq 0, \quad (5.4)$$

where, crucially, the coefficients λ_i do not depend on x , as long as G is at thermodynamic equilibrium [78]. Comparing with equation (5.1), we see that equation (5.4) defines an input–output function on $C(G)$, which, as noted above, has the hypercube structure \mathcal{C}_m . It follows that, by varying the coefficients λ_i in equation (5.4) as well as the edge labels imposed on \mathcal{C}_m , while ensuring the cycle condition holds, any input–output function on any graph G can be recovered, provided G also satisfies the cycle condition. In this way, a universal (p, s) region can be plotted, which looks similar in shape to but is larger in size than that shown in figure 3c [78]. In particular, it is bounded in $[a, \infty) \times [a, \infty)$ for any $a > 0$ and has the same kind of cusp that falls on the Hill line. Although we have not provided the details here, it should at least be plausible that the following result holds.

Universal Hopfield barrier for sharpness [78]. *The Hill function \mathcal{H}_m is the universal Hopfield barrier for the (p, s) sharpness of any thermodynamic equilibrium input–output function with m binding sites for the input ligand, no matter how complicated the underlying Markov process.*

We see that, despite its empirical origin, the Hill function has a deep biophysical justification, as the limit of sharpness that can be achieved at thermodynamic equilibrium. This finally provides a rigorous basis for its widespread use in mathematical models. However, that freedom brings with it a new obligation because we are now in a position to ask what kinds of HOCs are needed to approach \mathcal{H}_m . While many different patterns of HOCs can do so, one requirement that seems essential is that HOCs up to the highest possible order are present, with binding being modulated by all other available binding sites; see fig. 9 in [19]. Allosterism offers one possible mechanism for implementing such HOCs and allosteric conformational ensembles are now widely seen in all cellular contexts [79]. Further experimental study is needed to understand if this is the principal mechanism that is at work in creating sharp input–output responses.

The universal (p, s) region also provides a surprisingly strong claim when experimental data fall outside the region. No equilibrium Markov-process model, no matter how complicated, can account for such data, thereby offering a powerful incentive to explore non-equilibrium explanations.

The universal Hopfield barrier for sharpness is not yet a theorem, since (p, s) regions have so far only been plotted numerically, but we have no doubt as to its veracity. It is a tantalizing open problem to formulate the appropriate mathematical statement, which requires a better understanding of the shape of the cusp that approaches \mathcal{H}_m asymptotically. The (p, s) region is the image of a map from a space of rational input–output functions to \mathbb{R}^2 , which encodes the shape of each function in terms of extrema of its normalized derivative. The shape of the (p, s) region is telling us that the components of this map are highly correlated and constrained in the vicinity of the cusp. We do not yet understand the mathematical form of this relationship.

As mentioned above, if the graph does not satisfy the cycle condition, so that it reaches a non-equilibrium steady state, then it is not difficult to find (p, s) points that lie outside the equilibrium region. The situation away from thermodynamic equilibrium is strikingly different from the universality found at equilibrium: simple families of graphs can be found whose non-equilibrium (p, s) regions cover the entire positive quadrant [80]. Plotting non-equilibrium regions for the hypercube graph structures (figure 3a) remains challenging because of the path-dependent complexity described in §4 [14]. Little is known about the shape of these regions, whether or not they also have cusps and, if so, where those cusps lie with respect to the Hill line. The general problem of characterizing non-equilibrium (p, s) regions remains wide open.

6. Discussion

Graphs similar to those used here are found elsewhere in biology. The usage that is closest to the linear framework goes back to the non-equilibrium physics of Terrell Hill [62] and Jürgen Schnakenberg [63]. The Matrix-Tree Theorem, specifically the First-Minors MTT which underlies equation

(2.4), has been independently rediscovered multiple times in different areas of science, by Hill himself [62], by Raoul Bott when he was working on economics [81], from where it finds its way into quantum field theory [82], and by Mark Freidlin and Alexander Wentzell in their work on random dynamical systems and large deviations [83]; for more of the history, see [2]. Indeed, biochemists should recognize the MTT as the King–Altman procedure in enzyme kinetics [36]. These independent manifestations are usually disconnected from graph theory, where the result had already been established (§2). Many of the statements use their own specialized terms in place of the spanning trees that are fundamental concepts of graph theory. We have tried to avoid such parochial tendencies in the linear framework and to deliberately bring out the connections to areas like graph theory, enzyme kinetics, algebraic geometry, Markov processes and non-equilibrium physics. It is this nexus of overlapping perspectives that gives the framework much of its power and appeal.

The main difference between previous uses of graphs and the linear framework is that the graph is regarded not merely as a picture but as a mathematical object in its own right, in terms of which general theorems can be formulated [4,13,16,19]. Graphs are particularly flexible in allowing certain features, such as a type of structure, or a subgraph, or particular edges or labels, to be specified while leaving other aspects of the graph to vary arbitrarily. A case in point is the grammar of enzyme mechanisms in equation (3.2). This flexibility allows theorems to be proved about processes of interest to biology, such as post-translational modification systems and input–output responses, while accommodating some of the extraordinary complexity and diversity that is present at the molecular level. We can summarize the principal insights from the work described here in the following four points.

- There is a surprising linearity that underlies many biochemical networks, which manifests itself when they can be uncoupled into linear framework graphs (§3). This insight also arises from chemical reaction network theory [1,84] and it hints at a deeper simplicity behind the biochemical complexity that we confront.
- A rational algebraic structure, derived from the Matrix-Tree Theorem in equation (2.4), accompanies the linearity at steady state. The rational algebraic formulae of Michaelis–Menten, King–Altman, Monod–Wyman–Changeux, Koshland–Némethy–Filmer and Ackers–Johnson–Shea, familiar to many biologists, are all instances of equation (2.6) applied to some graph (§2). This gives the quantitative foundation of molecular biology a unity which it has previously lacked. The rationality has been central to the analysis of both post-translational modification systems (§3) and the sharpness of input–output responses (§5).
- There is an equivalence between linear framework graphs, under reservoir assumptions, and finite-state, continuous-time Markov processes (§4). The Laplacian matrix of the graph is the linear operator that defines the master equation. Up to now, this relationship has been exploited at steady state but this is no longer a restriction, as we will mention below. The framework again gives access through equation (2.4) to the rational algebraic structure of steady states, which has not been widely explored within Markov process theory.

- The graph theory also provides an algebraic reformulation of stochastic thermodynamics, not only at equilibrium, where the conventional physics is recovered through a new parametrization based on HOCs, but also, more importantly, away from equilibrium (§4). It thereby enables the analysis of energy dissipation in cellular information processing, from which a new interpretation emerges for that most pervasive, but least justified, of rational functions, the Hill function, as the Hopfield barrier for sharpness (§5).

We noted in the Introduction that the range of applicability of the linear framework remains unclear. As far as Markov processes are concerned, the framework offers an alternative approach based on graph theory (§4). As for biochemical reaction networks, the material in §3 suggests that the framework can be deployed when the metagraph that describes how enzymes act on their substrates is acyclic. We conjectured that a generalized version of the RPT may be true under such circumstances. This restriction rules out systems in which enzymatic action occurs along pathways that feed back, which would yield cyclic metagraphs, but it remains an interesting question whether a recursive structure with fixed points can be uncovered here (§3). The framework is also suited to analysing networks that reach steady states, as opposed to those with more complex dynamics, such as limit cycles.

Finally, two recent developments of the linear framework seem worth mentioning. We had long thought that the rational algebraic approach based on equation (2.4) was limited to the steady state and that the transient regime could only be approached by very different methods. In fact, the linear framework can be extended to the transient regime, at least within the Markov process interpretation (§4). We have shown that the All-Minors MTT (§2) allows first-passage times to be calculated as rational algebraic functions of the parameters [85]. This development comes at an appropriate time in biology, where new experimental methods are giving access to faster dynamics at the molecular scale and raising important questions about the range of validity of steady-state assumptions [86]. Another restriction of the framework, which also seemed unavoidable, has been to finite graphs. However, here too, we have shown that a particular class of semi-infinite ‘cylinder’ graphs can be treated recursively, again exploiting the All-Minors MTT [87]. This development is also timely because it applies, in particular, to integrating models of gene regulation with models of gene expression, for which the number of transcribed mRNA molecules becomes part of the system state [7]. These numbers can now be estimated by single-molecule methods and offer a more informative measure of gene output than the regulated recruitment of RNA Polymerase used in the models described above (figure 3*a*). In the light of these exciting developments, many new questions arise for theoretical exploration, which we hope to report on in subsequent work.

Data accessibility. This article has no additional data.

Authors’ contributions. K.-M.N.: formal analysis, investigation, software, writing—review and editing; R.M.-C.: formal analysis, investigation, software, writing—review and editing; J.G.: conceptualization, writing—original draft, writing—review and editing.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. We declare that we have no competing interests.

Funding. K.-M.N., R.M.-C. and J.G. were supported by award no. GM122928 and J.G. was supported by award no. GM105375 from

the National Institutes of Health, USA; R.M.-C. was supported by EMBO Fellowship no. ALTF 683-2019.

Acknowledgements. We thank members of the Gunawardena laboratory and two anonymous reviewers for their helpful comments.

References

- Gunawardena J. 2012 A linear framework for time-scale separation in nonlinear biochemical systems. *PLoS ONE* **7**, e36321. (doi:10.1371/journal.pone.0036321)
- Mirzaev I, Gunawardena J. 2013 Laplacian dynamics on general graphs. *Bull. Math. Biol.* **75**, 2118–2149. (doi:10.1007/s11538-013-9884-8)
- Mirzaev I, Bortz DM. 2015 Laplacian dynamics with synthesis and degradation. *Bull. Math. Biol.* **77**, 1013–1045. (doi:10.1007/s11538-015-0075-7)
- Thomson M, Gunawardena J. 2009 The rational parameterisation theorem for multisite post-translational modification systems. *J. Theor. Biol.* **261**, 626–636. (doi:10.1016/j.jtbi.2009.09.003)
- Thomson M, Gunawardena J. 2009 Unlimited multistability in multisite phosphorylation systems. *Nature* **460**, 274–277. (doi:10.1038/nature08102)
- Gunawardena J. 2014 Time-scale separation: Michaelis and Menten's old idea, still bearing fruit. *FEBS J.* **281**, 473–488. (doi:10.1111/febs.12532)
- Wong F, Gunawardena J. 2020 Gene regulation in and out of equilibrium. *Annu. Rev. Biophys.* **49**, 199–226. (doi:10.1146/annurev-biophys-121219-081542)
- Barabanshikov A, Gunawardena J. 2019 Trees of Goldbeter-Koshland loops. *Bull. Math. Biol.* **81**, 2463–2509. (doi:10.1007/s11538-019-00615-y)
- Dasgupta T, Croll DH, Owen JA, Vander Heiden MG, Locasale JW, Alon U, Cantley LC, Gunawardena J. 2014 A fundamental trade off in covalent switching and its circumvention by enzyme bifunctionality in glucose homeostasis. *J. Biol. Chem.* **289**, 13 010–13 025. (doi:10.1074/jbc.M113.546515)
- Dexter JP, Dasgupta T, Gunawardena J. 2015 Invariants reveal multiple forms of robustness in bifunctional enzyme systems. *Integr. Biol.* **7**, 883–894. (doi:10.1039/c5ib00009b)
- Nam K-M, Gyori BM, Amethyst SV, Bates DJ, Gunawardena J. 2020 Robustness and parameter geography in post-translational modification systems. *PLoS Comput. Biol.* **16**, e1007573. (doi:10.1371/journal.pcbi.1007573)
- Ahsendorf T, Wong F, Eils R, Gunawardena J. 2014 A framework for modelling gene regulation which accommodates non-equilibrium mechanisms. *BMC Biol.* **12**, 102. (doi:10.1186/s12915-014-0102-4)
- Biddle JW, Nguyen M, Gunawardena J. 2019 Negative reciprocity, not ordered assembly, underlies the interaction of Sox2 and Oct4 on DNA. *eLife* **8**, e410172018. (doi:10.7554/eLife.41017)
- Estrada J, Wong F, DePace A, Gunawardena J. 2016 Information integration and energy expenditure in gene regulation. *Cell* **166**, 234–244. (doi:10.1016/j.cell.2016.06.012)
- Park J, Estrada J, Johnson G, Ricci-Tam C, Bragdon M, Shulgina Y, Cha A, Gunawardena J, DePace AH. 2019 Dissecting the sharp response of a canonical developmental enhancer reveals multiple sources of cooperativity. *eLife* **8**, e41266. (doi:10.7554/eLife.41266)
- Wong F, Dutta A, Chowdhury D, Gunawardena J. 2018 Structural conditions on complex networks for the Michaelis-Menten input-output response. *Proc. Natl Acad. Sci. USA* **115**, 9738–9743. (doi:10.1073/pnas.1808053115)
- Yordanov P, Stelling J. 2018 Steady-state differential dose response in biological systems. *Biophys. J.* **114**, 723–736. (doi:10.1016/j.bpj.2017.11.3780)
- Yordanov P, Stelling J. 2020 Efficient manipulation and generation of Kirchhoff polynomials for the analysis of non-equilibrium biochemical reaction networks. *J. R. Soc. Interface* **17**, 20190828. (doi:10.1098/rsif.2019.0828)
- Biddle JW, Martinez-Corral R, Wong F, Gunawardena J. 2021 Allosteric conformational ensembles have unlimited capacity for integrating information. *eLife* **10**, e65498. (doi:10.7554/eLife.65498)
- Biddle JW, Gunawardena J. 2020 Reversal symmetries for cyclic paths away from thermodynamic equilibrium. *Phys. Rev. E* **101**, 062165. (doi:10.1103/PhysRevE.101.062125)
- Cetiner U, Gunawardena J. 2020 Reformulating non-equilibrium steady states. *bioRxiv*. (doi:10.1101/456640)
- Wong F, Amir A, Gunawardena J. 2018 Energy-speed-accuracy relation in complex networks for biological discrimination. *Phys. Rev. E* **98**, 012420. (doi:10.1103/PhysRevE.98.012420)
- Goldbeter A, Koshland DE. 1981 An amplified sensitivity arising from covalent modification in biological systems. *Proc. Natl Acad. Sci. USA* **78**, 6840–6844. (doi:10.1073/pnas.78.11.6840)
- Forger DB, Peskin CS. 2003 A detailed predictive model of the mammalian circadian clock. *Proc. Natl Acad. Sci. USA* **100**, 14 806–14 811. (doi:10.1073/pnas.2036281100)
- Wannig CT, Kulasiri D, Samarasinghe S. 2014 A nutrient dependant switch explains mutually exclusive existence of meiosis and mitosis initiation in budding yeast. *J. Theor. Biol.* **341**, 88–101. (doi:10.1016/j.jtbi.2013.09.030)
- Hill AV. 1910 The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves. *J. Physiol.* **40**(iv–vii).
- Gunawardena J. 2014 Models in biology: 'accurate descriptions of our pathetic thinking'. *BMC Biol.* **12**, 29. (doi:10.1186/1741-7007-12-29)
- Chung FRK. 1997 *Spectral graph theory*. Regional Conference Series in Mathematics, vol. 92. Providence, RI: American Mathematical Society.
- Belkin M, Niyogi P. 2008 Towards a theoretical foundation for Laplacian-based manifold methods. *J. Comput. Syst. Sci.* **74**, 1289–1308. (doi:10.1016/j.jcss.2007.08.006)
- Gunawardena J. 2012 Some lessons about models from Michaelis and Menten. *Mol. Biol. Cell* **23**, 517–519. (doi:10.1091/mbc.e11-07-0643)
- Kirchhoff G. 1847 Ueber der Auflösung der Gleichungen, auf welche man bei Untersuchung der linearen Vertheilung galvanischer Ströme geführt wird. *Ann. Phys. Chem.* **72**, 497–508. (doi:10.1002/andp.18471481202)
- Tutte WT. 1948 The dissection of equilateral triangles into equilateral triangles. *Proc. Camb. Phil. Soc.* **44**, 463–482. (doi:10.1017/S030500410002449X)
- Fiedler M, Sedláček J. 1958 O W-basích orientovaných grafů. *Časopis pro pěstování matematiky* **83**, 214–225. (doi:10.21136/cpm.1958.108301)
- Chaiken S. 1982 A combinatorial proof of the all-minors matrix-tree theorem. *SIAM J. Alg. Disc. Meth.* **3**, 319–329. (doi:10.1137/0603033)
- Michaelis L, Menten M. 1913 Die kinetik der invertinwirkung. *Biochem. Z.* **49**, 333–369.
- King EL, Altman C. 1956 A schematic method of deriving the rate laws for enzyme-catalyzed reactions. *J. Phys. Chem.* **60**, 1375–1378. (doi:10.1021/j150544a010)
- Monod J, Wyman J, Changeux JP. 1965 On the nature of allosteric transitions: a plausible model. *J. Mol. Biol.* **12**, 88–118. (doi:10.1016/S0022-2836(65)80285-6)
- Koshland DE, Némethy G, Filmer D. 1966 Comparison of experimental binding data and theoretical models in proteins containing subunits. *Biochemistry* **5**, 365–385. (doi:10.1021/bi00865a047)
- Ackers GK, Johnson AD, Shea MA. 1982 Quantitative model for gene regulation by lambda phage repressor. *Proc. Natl Acad. Sci. USA* **79**, 1129–1133. (doi:10.1073/pnas.79.4.1129)
- Prabakaran S, Lippens G, Steen H, Gunawardena J. 2012 Post-translational modification: nature's escape from genetic imprisonment and the basis for cellular information processing. *Wiley Interdiscip. Rev. Syst. Biol. Med.* **4**, 565–583. (doi:10.1002/wsbm.1185)
- Walsh CT. 2006 *Posttranslational modification of proteins*. Englewood, CO: Roberts and Company.

42. Johnson KA, Goody RG. 2011 The original Michaelis constant: translation of the 1913 Michaelis-Menten paper. *Biochemistry* **50**, 8264–8269. (doi:10.1021/bi201284u)
43. Xu Y, Gunawardena J. 2012 Realistic enzymology for post-translational modification: zero-order ultrasensitivity revisited. *J. Theor. Biol.* **311**, 139–152. (doi:10.1016/j.jtbi.2012.07.012)
44. Blüthgen N, Bruggeman FJ, Legewie S, Herzog H, Westerhoff HV, Kholodenko BN. 2006 Effects of sequestration on signal transduction cascades. *FEBS J.* **273**, 895–906. (doi:10.1111/j.1742-4658.2006.05105.x)
45. Ortega F, Acerenza L, Westerhoff HV, Mas F, Cascante M. 2002 Product dependence and bifunctionality compromise the ultrasensitivity of signal transduction cascades. *Proc. Natl Acad. Sci. USA* **99**, 1170–1175. (doi:10.1073/pnas.022267399)
46. Fersht A. 1985 *Enzyme structure and mechanism*. New York, NY: W. H. Freeman and Company.
47. Johnson KA (ed.). 2003 *Kinetic analysis of macromolecules: a practical approach*. Oxford, UK: Oxford University Press.
48. Srinivasan B. 2021 Explicit treatment of non-Michaelis-Menten and atypical kinetics in early drug discovery. *ChemMedChem* **16**, 899–918. (doi:10.1002/cmdc.202000791)
49. Segel IH. 1993 *Enzyme kinetics: behaviour and analysis of rapid equilibrium and steady-state enzyme systems*. Hoboken, NJ: Wiley Interscience.
50. Okar DA, Wu C, Lange AJ. 2004 Regulation of the regulatory enzyme, 6-phosphofructo-2-kinase/fructose-2, 6-bisphosphatase. *Adv. Enzyme Regul.* **44**, 123–154. (doi:10.1016/j.advenzreg.2003.11.006)
51. DeHart CJ, Chahal JS, Flint SJ, Perlman DH. 2014 Extensive post-translational modification of active and inactivated forms of endogenous p53. *Mol. Cell. Proteomics* **13**, 1–17. (doi:10.1074/mcp.M113.030254)
52. Cox D, Little J, O'Shea D. 1997 *Ideals, varieties and algorithms*, 2nd edn. Berlin, Germany: Springer.
53. Markevich NI, Hoek JB, Kholodenko BN. 2004 Signalling switches and bistability arising from multisite phosphorylation in protein kinase cascades. *J. Cell Biol.* **164**, 353–359. (doi:10.1083/jcb.200308060)
54. Bates DJ, Hauenstein JD, Sommese AJ, Wampler CW (eds). 2013 *Numerically solving polynomial systems with Bertini: software, environment and tools*. Philadelphia, PA: SIAM.
55. Dickenstein A. 2020 Algebraic geometry tools in systems biology. *Not. Am. Math. Soc.* **67**, 1706–1715.
56. Gross E, Harrington HA, Rosen Z, Sturmfels B. 2016 Algebraic systems biology: a case study for the Wnt pathway. *Bull. Math. Biol.* **78**, 21–51. (doi:10.1007/s11538-015-0125-1)
57. Manrai A, Gunawardena J. 2008 The geometry of multisite phosphorylation. *Biophys. J.* **95**, 5533–5543. (doi:10.1529/biophysj.108.140632)
58. Pérez Millán M, Dickenstein A. 2018 The structure of MESSI biological systems. *SIAM J. Appl. Dyn. Syst.* **17**, 1650–1682. (doi:10.1137/17M1113722)
59. Feliu E, Wiuf C. 2013 Variable elimination in post-translational modification reaction networks with mass-action kinetics. *J. Math. Biol.* **66**, 281–310. (doi:10.1007/s00285-012-0510-4)
60. Sáez M, Wiuf C, Feliu E. 2017 Graphical reduction of reaction networks by linear elimination of species. *J. Math. Biol.* **74**, 195–237. (doi:10.1007/s00285-016-1028-y)
61. Astumian RD, Pezzato C, Feng Y, Qiu Y, McGonigal PR, Cheng C, Stoddart JF. 2020 Non-equilibrium kinetics and trajectory thermodynamics of synthetic molecular pumps. *Mater. Chem. Front.* **4**, 1310–1314. (doi:10.1039/D0QM00022A)
62. Hill TL. 1966 Studies in irreversible thermodynamics. IV. Diagrammatic representation of steady state fluxes for unimolecular systems. *J. Theor. Biol.* **10**, 442–459. (doi:10.1016/0022-5193(66)90137-8)
63. Schnakenberg J. 1976 Network theory of microscopic and macroscopic behaviour of master equation systems. *Rev. Mod. Phys.* **48**, 571–586. (doi:10.1103/RevModPhys.48.571)
64. Bauer M, Cornu F. 2015 Local detailed balance: a microscopic derivation. *J. Phys. A Math. Theor.* **48**, 015008. (doi:10.1088/1751-8113/48/1/015008)
65. Seifert U. 2011 Stochastic thermodynamics of single enzymes and molecular motors. *Eur. Phys. J. E* **34**, 26. (doi:10.1140/epje/i2011-11026-7)
66. Lewis GN. 1925 A new principle of equilibrium. *Proc. Natl Acad. Sci. USA* **11**, 179–183. (doi:10.1073/pnas.11.3.179)
67. Dill KA, Bromberg S. 2003 *Molecular driving forces: statistical thermodynamics in chemistry and biology*. New York, NY: Garland Science.
68. Christopoulos A, Kenakin T. 2002 G protein-coupled receptor allostery and complexing. *Pharmacol. Rev.* **54**, 323–374. (doi:10.1124/pr.54.2.323)
69. Bernardi O. 2012 On the spanning trees of the hypercube and other products of graphs. *Electron. J. Comb.* **19**, P51. (doi:10.37236/2510)
70. Chebotarev P, Agaev R. 2002 Forest matrices around the Laplacian matrix. *Linear Algebr. Appl.* **356**, 253–274. (doi:10.1016/S0024-3795(02)00388-9)
71. Karsenti E. 2008 Self-organization in cell biology: a brief history. *Nat. Rev. Mol. Cell Biol.* **9**, 255–262. (doi:10.1038/nrm2357)
72. Kolomeisky AB, Fisher ME. 2007 Molecular motors: a theorist's perspective. *Annu. Rev. Phys. Chem.* **58**, 675–695. (doi:10.1146/annurev.physchem.58.032806.104532)
73. Hopfield JJ. 1974 Kinetic proofreading: a new mechanism for reducing errors in biosynthetic processes requiring high specificity. *Proc. Natl Acad. Sci. USA* **71**, 4135–4139. (doi:10.1073/pnas.71.10.4135)
74. Ptashne M. 2003 Regulated recruitment and cooperativity in the design of biological regulatory systems. *Phil. Trans. R. Soc. A* **361**, 1223–1234. (doi:10.1098/rsta.2003.1195)
75. Dexter JP, Biddle JW, Gunawardena J. 2018 Model discrimination for Ca²⁺-dependent regulation of myosin light chain kinase in smooth muscle contraction. *FEBS Lett.* **592**, 2811–2821. (doi:10.1002/1873-3468.13207)
76. Engel PC. 2013 A hundred years of the Hill equation. *Biochem. J.* **35**, 40–43. (doi:10.1042/BJ20131164)
77. Weiss JN. 1997 The Hill equation revisited: uses and misuses. *FASEB J.* **11**, 835–841. (doi:10.1096/fasebj.11.11.9285481)
78. Martínez-Corral R, Nam K-M, DePace AH, Gunawardena J. In preparation. The Hill function is the universal Hopfield barrier for sharpness.
79. Wodak SJ *et al.* 2019 Allostery in its many disguises: from theory to applications. *Structure* **27**, 566–578. (doi:10.1016/j.str.2019.01.003)
80. Wong F, Thomas I, Cetiner U, Martínez-Corral R, Nam K-M, Biddle JW, Gunawardena J. In preparation. Unlimited ultrasensitivity from a single binding site away from thermodynamic equilibrium.
81. Bott R, Mayberry JP. 1954 Matrices and trees. In *Economic activity analysis* (ed. O Morgenstern), pp. 391–400. New York, NY: Wiley.
82. Jaffe A. 1965 Divergence of perturbation theory for bosons. *Commun. Math. Phys.* **1**, 127–149. (doi:10.1007/BF01646496)
83. Freidlin MI, Wentzell AD. 2012 *Random perturbations of dynamical systems*, 3rd edn. Heidelberg, Germany: Springer.
84. Feinberg M. 1987 Chemical reaction network structure and the stability of complex isothermal reactors I. The deficiency zero and deficiency one theorems. *Chem. Eng. Sci.* **42**, 2229–2268. (doi:10.1016/0009-2509(87)80099-4)
85. Nam K-M, Gunawardena J. In preparation. Algebraic formulas for first-passage times of Markov processes in the linear framework.
86. Lucas T, Tran H, Perez Romero CA, Guillou A, Fradin C, Coppey M, Walczak AM, Dostatni N. 2018 3 min to precisely measure morphogen concentration. *PLoS Genet.* **14**, e1007676. (doi:10.1371/journal.pgen.1007676)
87. Nam K-M, Nasser J, Gunawardena J. In preparation. A graph-theoretic framework for integrating gene regulation and expression.