

CHAPTER 1

MODELS IN SYSTEMS BIOLOGY: THE PARAMETER PROBLEM AND THE MEANINGS OF ROBUSTNESS

With four parameters I can fit an elephant and with five I can make him wiggle his trunk.
—told by Enrico Fermi to Freeman Dyson and attributed to John von Neumann, [25]

1.1 INTRODUCTION

I co-teach a graduate course at Harvard called “*An Introduction to Systems Biology*”. It covers some of the mathematical methods used to build mechanistic models of molecular and cellular systems. Beginning students tend to ask two kinds of questions. Those with a biological background say “*Why do I need to use mathematical models? What can they tell me that conventional biological methods cannot?*”, while those from the physical sciences (mathematics, physics, engineering) or computer science say “*I know how to model. Why is biology any different from physics or engineering?*”. Broadly speaking, everyone wants to know, from very different perspectives, “*How do I do systems biology?*”. Students are usually under the misapprehension that the person standing in front of them knows the answers to such questions. In my case, I was only marginally less ignorant than the students themselves. It was their curiosity

and skepticism, along with a realisation that the field lacks a shared foundation for discussing such questions, that forced me to think more deeply about the issues.

This paper is the first of at least two in which I review some tentative conclusions. It sets out a framework for thinking about models in which I try to rise above the partisan assertions that are sometimes made—“*my kind of model is better than yours*”—and point to some of the broader themes and open problems. It should be obvious that this can be no more than a report of work in progress and is neither complete nor definitive. The next paper will discuss why models are being used in systems biology and what we should expect from them [37]. Ideally, this should not be treated separately but I found it difficult to do justice to everything in the bounds of a single paper.

For our purposes, systems biology may be defined as the emerging discipline that asks how physiology and phenotype emerge from molecular interactions [4, 50]. Mathematical models are being used in support of this, continuing a long tradition inherited from genetics [60, 65], physiology [39, 44], biochemistry [41, 48, 78], evolutionary biology [33, 68] and ecology [63]. Models, however, mean different things to physicists, mathematicians, engineers and computer scientists, not to mention to biologists of varying persuasions. These different perspectives need to be unravelled and their advantages distilled if model building is to fulfill its potential as an explanatory tool for studying biological systems. I begin in §1.2 by pointing out that most mechanistic models (as opposed to those arising from “omics”) can be thought of as some form of dynamical system. This provides a unified framework in which to compare different kinds of models. Mechanistic models are often complex, in the sense of having many undetermined parameters, and the parameter problem emerges as one of the central difficulties in the field. Different disciplines provide sharply contrasting approaches to this, as I discuss in §1.3, and this has tended to obscure the problem in the literature. Attempts are sometimes made to resolve the parameter problem by making assertions of “robustness”. This is generally regarded as a desirable feature—who could doubt that biology is robust? However, its wide usage is often uncorrelated with precise definition. I identify in §1.5 four kinds of robustness which arise in the dynamical systems framework and review some previous studies in terms of this classification. §1.4 outlines the qualitative view of dynamical systems that forms the basis for this discussion.

Parameters and robustness are concepts that have been widely studied in mathematics, engineering and statistics. My intention here is not to review this material, for which there are many standard texts—see, for instance, [97, 101]—but rather to show how these concepts are being used, and sometimes abused, in systems biology and to draw attention to some of the scientific issues that arise from that.

1.2 MODELS AS DYNAMICAL SYSTEMS

Two broad directions have emerged in systems biology. The first, “omics”, initiated by new technologies such as the microarray [79], relies on inferring causality from correlation in large datasets (see, for instance, [82]). To the extent that models are used, they are statistical in character. The second direction, which might be called

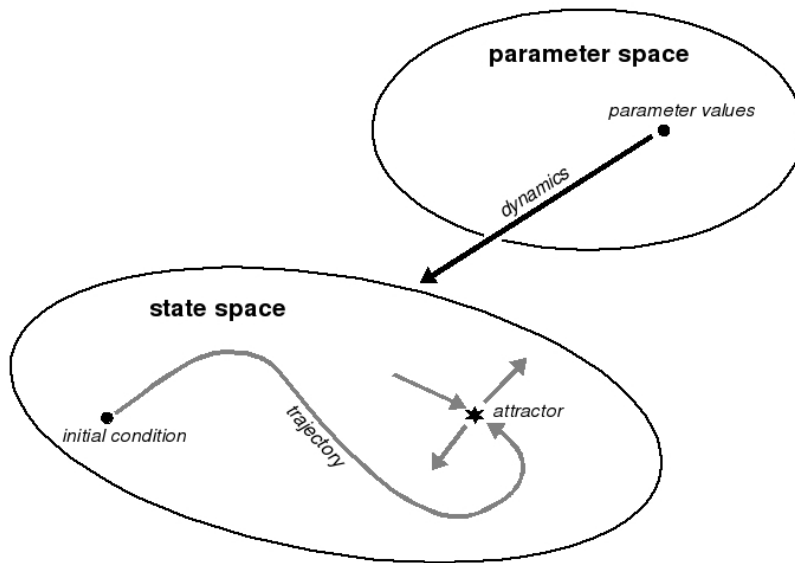


Figure 1.1 Dynamical system. A point in parameter space, given by a set of parameter values, defines a dynamics on the state space. If the system is prepared in an initial condition, then the dynamics typically lead to an attractor, pictured here as a star. Common attractors are steady states or periodic orbits but they can be much more complex [87]. Note that some trajectories leave the attractor, indicating that it is unstable, as discussed in §1.4.1. The parameter and state spaces are pictured as abstract sets. For ODE models, they usually correspond to Euclidean spaces, \mathbb{R}^k , of some dimension k but for other kinds of models the state space can be infinite dimensional (PDEs or stochastic models) or not have any linear structure (discrete models).

“mechanistic” systems biology, has been less visible but has deeper historical roots [39, 41, 44, 48, 78]. The resulting models specify molecules, cells and tissues and their interactions based on what is known or believed to be true. It is with the latter type of model that we will be concerned here. The subtleties of causal analysis are well discussed elsewhere [72].

Most mechanistic models in systems biology can be regarded as some form of *dynamical system*. A dynamical system describes the *states* of a biological system and how these states change in time. It can be abstractly visualised as in Figure 1.1 as a *state space*, upon which is imposed a temporal dynamics: given a particular state as an *initial condition*, the dynamics define the *trajectory* taken over time from that starting point. Not all models take this form. For instance, constraint-based models represent systems at steady state and have no explicit representation of time [70]. We focus here on models that do.

Dynamical systems usually depend on parameters. In abstractly visualizing a dynamical system, therefore, one should always keep in mind the *parameter space* that accompanies the state space, as in Figure 1.1. The dynamics on the state space

cannot be defined without first specifying the parameter values, thereby fixing a point in parameter space. As this point varies, so do the dynamics.

1.2.1 Continuous models

A type of model that is frequently used is one in which the state of a molecular component, x , is its concentration in some cellular compartment (cytoplasm, plasma membrane, etc), which we will also denote by x and treat as a function of time, $x(t)$. The temporal dynamics are then described by an ordinary differential equation (ODE) for the net rate of production of x . This is how the biochemistry of enzymes has been modelled [18], which provides a foundation for models of molecular networks [3, 94]. As an example, if x is produced at a (zero order) rate of a molar per second and consumed at a (first order) rate of b per second, then

$$\frac{dx}{dt} = a - bx. \quad (1.1)$$

In this case, the dynamical system has a 1 dimensional state space, consisting of the single state variable x , and a 2 dimensional parameter space, consisting of the two parameters a and b . Since (1.1) is linear, it can be readily solved [42]:

$$x(t) = \frac{a}{b} - \left(\frac{a}{b} - x_0\right) \exp(-bt), \quad (1.2)$$

where x_0 is the initial condition from which the system starts at time $t = 0$: $x(0) = x_0$. We see that no matter where the system is started from, it relaxes exponentially to the unique steady state, $x = a/b$, at which production and consumption are exactly balanced. As the values of the parameters a and b change, the steady state also changes but the dynamics remain “qualitatively” the same. Much of the difficulty in comprehending nonlinear, higher-dimensional systems lies in understanding how this very simple picture has to be refined; see §1.4.

Example (1.1) is unusual in that it is explicitly solvable in a closed form in which the parameters appear as symbols. Most dynamical systems arising in systems biology are nonlinear and cannot be solved in this way. (Except possibly at steady state; see §1.4.2.) They have to be studied by simulation, for which parameter values must be specified. The difficulties with this—the parameter problem—are discussed in §1.3.

Several kinds of differential equation models have proved useful in systems biology, reflecting the emergence of new experimental techniques. Fluorescent sensors have revolutionised cell biology, making it possible to image specific proteins in individual living cells in real time and revealing extraordinary dynamical complexity. Ionic calcium, Ca^{2+} , for instance, exhibits sparks, puffs, oscillations and travelling waves in certain cell types [24], reflecting its role as a second messenger linking external signals (first messengers) to a spectrum of cellular responses. To model this, spatial compartments need to be represented as two- or three-dimensional geometries, rather than as unstructured entities like “cytoplasm” and “membrane”, and the dynamics need to be described by partial differential equations, often of reaction-diffusion type, with the compartment geometry entering into the boundary conditions [84].

The same fluorescent technology has more recently made it possible to measure noise in individual cells, revealing the impact of both molecular stochasticity (“intrinsic noise”) and cell-to-cell variability (“extrinsic noise”) [29, 71]. Extrinsic noise can sometimes be modelled as a probability distribution on the initial conditions of a deterministic model or by adding external noise terms, as in the Langevin approach [96]. Molecular stochasticity, however, requires some form of stochastic master equation in which the state of a component is described by the probability distribution of the number of molecules of component x , as a function of time, and the dynamics are described by stochastic differential equations [96].

1.2.2 Discrete models

Differential equation models of the kinds discussed above are familiar in the physical sciences, biochemistry and physiology. Biologists, however, often find it convenient to describe gene expression in terms of discrete states—on/off or low/high—and the development of microarray technology allows mRNA levels to be quantified into multiple discrete levels, as in the familiar heat-maps. Genetic manipulations also lead naturally to causal inferences expressed in Boolean logic: “*in the absence of X, Y becoming low leads to high Z*”. These kinds of data and reasoning can be modelled by dynamical systems with discrete states, where the temporal dynamics are given by discrete transitions between states, rather than being parameterised by a global clock, t , that marks the passage of time. When states are composed of many discrete variables (for example, many genes), state transitions may take place synchronously, with each variable being updated simultaneously, or asynchronously, with variables being updated independently of each other.

Discrete models often permit abstraction from the mechanistic details [59]. Such abstraction may lead to an absence of visible parameters, which is sometimes touted as an advantage of discrete models over continuous models. Such assertions should be treated skeptically. Parameters are usually insidiously hiding in the unstated assumptions that accompany a discrete model. For instance, for states composed of many discrete variables, the assumption of asynchronous timing gives equal opportunity to each interleaved sequence of updates. In reality, each variable may have its own rate of change and a model that took these rates into account as parameters would select some interleaved sequences in preference to others. Such distinctions are beyond the scope of unparameterised discrete models but may sometimes have serious biological implications.

Discrete models have a long history in biology [49, 90], prior to the recent resurgence of interest in them via computer science [32]. Theoretical computer scientists view discrete models as computing machines [45]. A Turing machine, for instance, is a discrete state/transition system coupled to a read/write memory. Computer scientists are concerned, among other issues, with methods for constructing such machines; for instance, for building complex machines out of simpler ones. This introduces a syntactic capability that is lacking from the physical science perspective but which becomes important in building models [43, 61].

One way to construct a complex model is to regard it as emerging from the collective interactions of independent agents, each of which has its own internal state and can undertake computations based on rules about its state and the state of other agents in the system [34, 43, 76]. For instance, each individual molecule could be an agent and the computations undertaken by agents could represent chemical reactions between molecules. Such agent-based systems capture molecular fluctuations and can reproduce stochastic models but their syntactic structure permits additional forms of analysis such as model-checking or abstract interpretation, which have been important in computer science [22, 32].

This last example illustrates the limitations of Figure 1.1. In an agent-based system, the state space may unfold with the dynamics and is then no longer a static entity. More generally, cells produce new cells, organisms produce new organisms; one of the characteristic features of biology is its capability for self-reproduction. There exists no general mathematical framework for dynamical systems in which the dynamics reconstruct the state space as they progress. Hybrid models, which combine discrete and continuous dynamics, provide only a partial kludge [7].

Thinking in terms of dynamical systems draws attention to the state of the system. Deciding how the state should be represented, whether coarsely as Boolean levels or at fine grain in an agent-based description or somewhere in between as concentrations, and how time and space should be modelled, should depend not on the disciplinary prejudices of the modeller but on the nature of the experimental data and the kinds of biological questions that are being asked. No one type of model is best for all purposes.

1.3 THE PARAMETER PROBLEM

Biological systems have many “moving parts”, whose collective interactions produce the physiology or phenotype of interest. Two general strategies have emerged to model this complexity. One seeks to bring the model’s assumptions close to reality by embracing the details of components and interactions. The resulting models are *thick*, with many states and more parameters. The other strategy moves in the opposite direction and seeks to abstract the essentials from the details, giving rise to *thin* models with fewer parameters. Despite parochial assertions to the contrary, both strategies have provided biological insight; their pros and cons are discussed in the companion paper to this [37]. In both cases, but most especially with thicker models, the problem arises of determining parameter values in a way that maintains credibility in a model’s conclusions. The importance of this problem has tended to be obscured in the literature for several reasons. On the one hand, it is easier to assert (particularly to an experimental audience) “*This model accounts for the data*” than “*This model, with these parameter values, accounts for the data*”. The latter formulation invites awkward questions as to why those parameter values were chosen and not others. (One might have included “initial conditions” along with parameter values but since the initial conditions are values of state variables, they share the same level of measurability and are, therefore, usually easier to determine than parameter values.) Even if

editors and reviewers are aware of the problem—and it seems they are mostly not—they are generally disinclined to ferret about in the Supplementary Information, to which graveyard such technical details are usually consigned. Finally, such a variety of approaches have something to say about the problem that it is hardly surprising to find confusion as to best practice. Here, we emphasize the significance and centrality of the parameter problem by contrasting different disciplinary perspectives of it.

1.3.1 Parameterphobia

Parameters are anathema to physicists, who take the view expressed in the quotation from von Neumann that, with enough parameters, any behaviour can be modelled. Of course, von Neumann was joking: a weighted sum of increasing functions with positive weights (parameters) can never fit a decreasing function, no matter how many parameters are used. (See §1.4.2 for a more relevant example.) However, the truth behind the joke distills a long tradition of modelling the inanimate world on the basis of the fundamental laws of physics. Biology, while founded entirely upon these laws, is not modelled in terms of them. Molecular or cellular behaviour is not deduced from Schrödinger's equation. At best, a model may be based on chemical principles, such as the law of mass action. At worst, it may rely on some ad-hoc guess that is only tenuously related to specific biological knowledge, let alone an underlying molecular mechanism. We have, in such cases, no systematic methodology for avoiding parameters.

While physicists are familiar with parameters and keep them firmly in their place, computer scientists (at least those of a theoretical disposition) are less acquainted with them. The discrete models used in theoretical computer science, like finite automata or Turing machines, have no parameters [45]. (They may have labels but these are passive adornments that do not effect the rate of state transitions.) When discrete models are parameterised they transmogrify into Markov chains, whose properties are more commonly studied elsewhere than in computer science. In consequence, computer science has had little to say about the parameter problem.

1.3.2 Measuring and calculating

Ideally, parameter values should be independently measured. In practice, our limited ability to make quantitative measurements of molecular states makes this difficult if not impossible for many parameters. Even when parameters have been measured, the conditions may have been sufficiently different as to raise doubts as to the relevance of the measurements. In-vitro values, for instance, may differ substantially from those in vivo, while in-vivo measurements themselves may require very careful interpretation [85]. Nevertheless, such measurements as do exist are often useful for initial analysis. Molecular dynamics (MD) calculations—arising from atomic-scale models—can now provide illuminating explanations of intra-molecular behaviour [88]. Certain kinds of parameters, such as binding constants, might be calculated from such MD models. Since these calculations are limited largely by computational power, it would be unwise to bet against them in the long run, but it seems unlikely

that they will yield a systematic approach anytime soon. They will, in any case, be limited to only certain kinds of parameters and to molecules whose atomic structures are well understood.

1.3.3 Counter fitting

Engineers are accustomed to building thick models with many parameters—of chemical reactors or combustion chambers, for instance—and determining parameter values by fitting to quantitative data [101]. This is the strategy most widely adopted in systems biology when sufficient data of the right kind is available. The development of nonlinear optimisation algorithms has made parameter fitting easy to undertake but has also concealed its dangers. These take several forms. The structure of a model may render it non-identifiable *a priori*: it may not be possible, even in principle prior to any data fitting, to determine certain parameter values. Even if a model is identifiable, the fitting process itself may need to be carefully examined. The reported optimum may be only local. Even if a global optimum is found, there may be several parameter sets which yield roughly similar optimal values. In other words, the energy landscape underlying the optimization may be undulating with many optimal valleys rather than a broad funnel leading to a single optimum. A classic example is that of fitting a sum of two exponentials; see, for instance, Figure 4.6 of [57].

The second and more serious danger in model fitting brings us back to the broader significance of von Neumann's quip. How is a model to be rejected? The answer "*when there are no parameter values that fit the data*" would not have satisfied von Neumann because, in his view, a model that is complex enough may fit all manner of data. In other words, the rejection criterion is inadequate. As we will see in §1.4.2, the behaviour of biochemical models is more subtle than this: models with arbitrary many parameters may sometimes have the simple qualitative behaviour shown by equation (1.2). The core issue may be restated in terms of explanatory power. A model does not explain the data to which it is fitted; the process of fitting already incorporates the data into the model.

Of course, parameter fitting is widely used in other areas of science. An X-ray crystal structure, for instance, is obtained by fitting an atomic model to diffraction data, with many free parameters (bond angles, bond lengths, etc). In such cases, independent cross-validation is used [14]. The data are partitioned into two sets: "test" data and "working" data. Parameters are determined by fitting on the working data. Having been fitted, they are used to account for the test data. If they do, the model is accepted; if not, it is rejected. Hodgkin and Huxley used a similar strategy for their famous model of the action potential in the squid giant axon [44]. The parameters were fitted in independent experiments on each of the three ion channels. Once fitted, the model, with those parameter values, was shown to numerically reproduce the time course of the action potential. Another strategy is to use wild-type data as working data and mutant data to test it by computationally mimicking the effect of the mutation [2]. As these examples make clear, a model's explanatory power comes from being able to account for data to which it has not been fitted.

Merely showing that quantitative data can be accounted for with some choice of parameter values can be such an effort, particularly with thick models, that it is often regarded as sufficient in itself. While this is easy to get away with, at least at present, it is not a good foundation for a new discipline.

1.3.4 Beyond fitting

Determining a specific set of parameter values and accounting for novel data is only part of the parameter problem. We have a general suspicion of models that are fine-tuned, for which some parameters require precise values. They are not “robust”. (Much the same argument is made about unstable steady states; see §1.4.1.) Robustness is a good feature, so the argument goes, because there are always errors, often substantial errors, in measuring and fitting data. Related systems might also be expected to show qualitatively similar behaviour but not have quite the same parameter values. If a model can be shown to be robust to changes in parameter values, then one can be more confident in drawing conclusions from it despite such uncertainties. There may also be properties of a model which are robust to variation in certain parameter values, like temperature compensation in circadian oscillators. Identifying such properties may yield biological insight; see §1.5.3. Aside from such robustness, which we will discuss further in §1.5, there may not always be sufficient quantitative data, or data of the right type, to fit all parameter values. The available data may, for instance, not be numerical but qualitative, as in developmental patterns. Finally, models can also be used in an exploratory way to understand how to think about a system in the first place, prior to any determination of parameter values. In all these cases, it becomes important to know how the model’s behaviour varies as a function of parameter values. This is the broader aspect of the parameter problem. To address it, a more qualitative view of dynamical systems becomes necessary.

1.4 THE LANDSCAPES OF DYNAMICS

1.4.1 Qualitative dynamics

Although the general ideas outlined in this section apply to most forms of dynamical system, they are best understood for ODE models [87, 94]. Figure 1.2 illustrates, in a simple case, the kind of behaviour to be expected of a model similar to example (1.1), in which

$$\frac{dx}{dt} = f(x; a), \quad (1.3)$$

where $x \in \mathbb{R}^n$ is a vector of state variables, $a \in \mathbb{R}^m$ is a vector of parameters and $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is the vector rate function expressing the balance between production and consumption of each x_i . Biological state variables are frequently non-negative (concentrations, for instance) and the state space may then be taken to be the non-negative orthant of \mathbb{R}^n . For any given set of parameter values, the trajectory starting from a given initial condition will typically converge upon an *attractor*: a limited

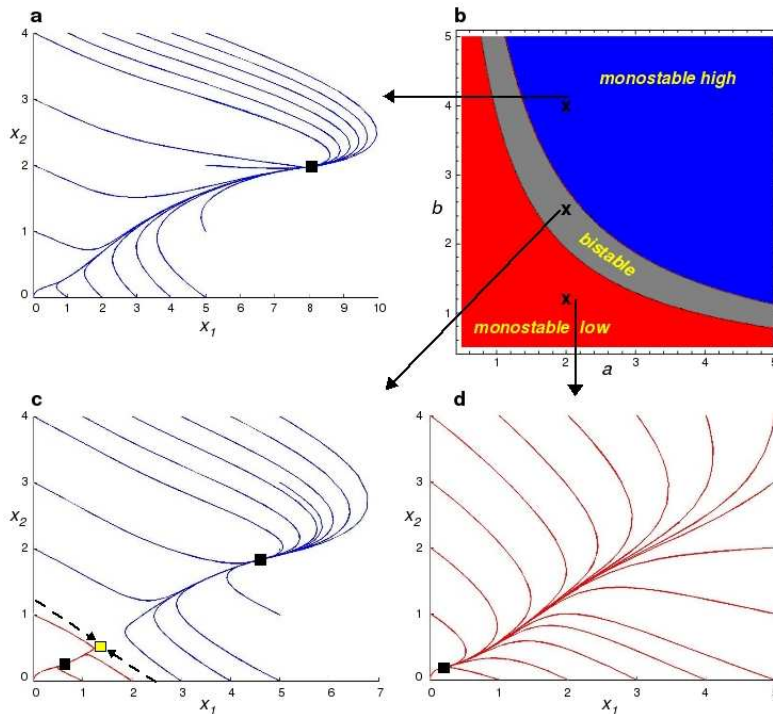


Figure 1.2 Qualitative dynamics. **a**, **c**, **d** show different patterns of trajectories on the state space—the nonnegative quadrant of \mathbb{R}^2 —of the ODE model $dx_1/dt = bx_2 - x_1$, $dx_2/dt = a(1 + x_1^3)/(10 + x_1^3) - x_2$, adapted from [69]. Each figure shows the trajectories starting from the initial conditions with integer coordinates on the boundary of the box defined by the origin and $(5, 4)$. Note that the vertical axis is the same in each figure but the horizontal axis varies. Black square denotes a stable steady state, yellow square an unstable saddle point. Blue trajectories go to the state with high x_1 value, red trajectories to the state with low x_1 value. Figures **a**, **d** have only a single basin of attraction leading to a stable state. Figure **c** has three basins of attraction corresponding to bistability. The dashed line marks the approximate location of the (1 dimensional) basin of attraction of the saddle point, which provides the boundary between the two larger (2 dimensional) basins leading to the stable states. **b** shows the parameter space for the two parameters a and b divided into regions corresponding to parameter values with qualitatively similar dynamics. Note the bifurcations—creation or destruction of steady states—that arise as the boundaries of regions are crossed, a behaviour that was absent in example (1.1). Both basins of attraction and parameter regions can be much more complex than in this simple example, particularly in higher dimensions.

region of the state space within which trajectories become confined. For instance, the trajectory may reach a steady state, as in example (1.1), or a periodic orbit, as in models of the cell cycle [94], circadian rhythms [64] or developmental clocks [58]. Chemical systems can also have more complex attractors and exhibit behaviours like bursting and chaos [15], which may have some biological role in the excitable tissues found in cardiac and neural systems [54]. A dynamical system may have several different attractors for a given set of parameter values. A familiar instance in systems biology is bistability [31, 69, 94], in which a dynamical system has three attractors, consisting of two stable steady states and one unstable steady state (Figure 1.2c). In this case, different initial conditions may reach different attractors and each attractor will have its own *basin of attraction* consisting of those initial conditions which lead to it. The state space breaks up into multiple disjoint basins of attraction, each leading to a unique attractor.

The geometry of a basin of attraction reveals something of the dynamics leading to the corresponding attractor. For instance, a steady state is stable if its basin of attraction has the same dimension as that of the ambient state space (dimension 2 for the two stable states in Figure 1.2c). If its dimension is lower, then moving away from the attractor along one of the missing dimensions leads outside the basin of attraction and towards some other attractor. This is the case for the saddle point in Figure 1.2c for which the basin of attraction has dimension 1. The argument is made that an unstable steady state is never found experimentally because random perturbations (“noise”) would destabilise it. Stable states are “robust” to such perturbation. Consequently, a steady state of a model that is claimed to represent some observed behaviour should always be checked to be stable. However, if only a few dimensions among hundreds are missing from a basin of attraction, then it may be possible for the system to linger in the corresponding steady state for an appreciable time, relative to the noise time scales in the system, before becoming destabilised. Our experience of high-dimensional systems is still too limited to know how significant this might be.

The dynamics may also satisfy constraints, which complicate the above picture. We will return to this in §1.5.2.

The dimension of a basin of attraction can often be estimated in the local vicinity of an attractor. For instance, the Hartman-Grobman theorem [23] tells us that for a reasonable (“nondegenerate”) steady state, $x = x^*$, the local dynamics are qualitatively the same as those of the linearised system, in which the full dynamics represented by $f(x)$ is replaced by the linearised dynamics

$$\frac{dx}{dt} = J(x^*)x \quad (1.4)$$

where $J(x)$ is the $n \times n$ matrix of first partial derivatives (the Jacobian), $J(x) = \partial f_i / \partial x_j$. Since linear equations are solvable, (1.4) gives considerable information about the local vicinity of steady states, including the (local) dimension of the basin of attraction [42]. In contrast, rather little is known, in general, about the global geometry of basins of attraction. Are they large or small and is their shape long and thin or short and squat? A characteristic difficulty in dynamical systems derived from differential equations is that local behaviour may be accessible (a derivative is a




Figure 1.3 Waddington's epigenetic landscapes. As he put it, "A multidimensional phase space is not very easy for the simple-minded biologist to imagine or to think about." [99, Page 27]. (He refers to the state space by its alternative name of "phase" space.) **a** Waddington abstracted dynamics on the high-dimensional state space of a developing embryo into a picture of a ball rolling down an inclined landscape into a branching fan of valleys and coming to rest at the end of one of them [99, Figure 4]. The end points represent different attractors, corresponding to different differentiation states of the organism. **b** Waddington's analogy for the action of genes on development shows the underside of the landscape being maintained by guy ropes in tension [99, Figure 5]. Each peg represents a gene which can have multiple effects and each attachment point can have multiple genes influencing it. Changes to a single gene may sometimes have little effect on the dynamics, depending on the background of the other genes.

measure of local slope) but global behaviour can be very challenging to analyse. For new developments in this direction see, in particular, [38, 75].

Systems biology forces us to confront the subtleties of global dynamics in high-dimensional spaces. This was already apparent to Conrad Waddington over fifty years ago [99]. His "epigenetic landscape" (Figure 1.3a) was an attempt to create a visualisable analogy for the complex dynamics through which an egg gives rise to an adult organism. (Sewall Wright's earlier "adaptive landscape" had a similar heuristic intent for the dynamics of genotypes during evolution but lacked the moving parts [104].) The epigenetic landscape continues to provide a conceptual basis for thinking about biological dynamics in high dimensions [47]. While many biologists are now familiar with the ball rolling down the valleys, fewer are aware of the mathematical models that Waddington used to arrive at this analogy [99, Chapter 2].

The picture of trajectories in state space holds for a given set of parameter values. If we now imagine moving through the parameter space, the pattern of trajectories will change (Figure 1.2b). In general, the parameter space itself also falls into disjoint regions. Within each region the pattern of trajectories remains qualitatively ("topologically") similar. It is as if the trajectories were inscribed on rubber and the rubber is stretched: while distances change, the connectivities remain the same. Different parameter regions, however, exhibit qualitatively different patterns (Figure 1.2b). In moving between regions, attractors may appear or disappear or change their dynamical characteristics, for instance from stable to unstable, and the trajectories may

reorganise themselves accordingly. Such *bifurcations* are usually key features of the overall behaviour [87].

Waddington was well aware of the role of parameters and illustrated them through the use of *guy ropes*, representing genes with pleiotropic effects (Figure 1.3b). While this gives a vivid illustration of systems behaviour, it is less satisfactory in giving a sense of the landscape of parameter space. Waddington was a remarkable scientist, who, more than any other, anticipated modern systems biology [66] and dismantled some of the barriers between biology and mathematics—see §1.5.4. He was marginalised in his own time partly because he was so far ahead of it in thinking about development, genetics and evolution as an integrated system. It is good to see his reputation restored for a modern audience [83].

For a given dynamical system, it would be useful to know, at least, the number of parameter regions and, for each region, the number of attractors and their types. No general methods are known for eliciting such details but some partial insights have come from different mathematical approaches.

1.4.2 Steady state attractors of ODE models

Chemical Reaction Network Theory Example (1.1) has only a single parameter region and only a single attractor—a stable steady state—for all parameter values in that region. Remarkably, more complex models may still exhibit similar behaviour. This emerges from Feinberg’s Chemical Reaction Network Theory (CRNT) [30]; see [36] for an overview and other references. CRNT applies to the ODE model coming from a network of chemical reactions by applying the principle of mass action. It associates to such a network a nonnegative integer called the “deficiency”, which does not depend on the values of the parameters but only on the underlying network of reactions. The deficiency is the dimension of a certain linear subspace, reflecting one of the key insights of CRNT: behind the nonlinearity of mass-action kinetics, there exists a remarkable degree of hidden linearity [36]. Under reasonable conditions, deficiency zero networks behave like example (1.1): provided constraints are respected (see §1.5.2 for an explanation of constraints), there is a single parameter region and only a single stable steady state for all parameter values in that region [30, 36]. This theorem is important because it shows that thick models, with many parameters, may nevertheless have simple qualitative dynamics. One cannot always fit an elephant! Having said that, the “deficiency zero theorem” is too restricted to be widely useful in systems biology, where parameter values have typically been found to influence the qualitative dynamics. Recent developments in CRNT may be more relevant [20] and the full implications of CRNT for systems biology remain to be worked out.

Monotone systems The dependence of the qualitative dynamics on the parameters can often be calculated for ODE models with only two state variables. The method of nullclines provides a geometric guide to the existence of steady states and there are mathematical theorems, like that of Poincaré-Bendixson, that help identify more complex attractors like periodic orbits [87]. Such methods are strictly limited

to two dimensional systems. Sontag and others have shown, nevertheless, that the steady states of certain high-dimensional ODE systems, with many state variables and parameters, correspond to those of an associated two-dimensional system [8]. There are several requirements for this method to work; among the most crucial is that the high-dimensional system is *monotone*, meaning, roughly speaking, that its dynamics preserve an underlying order on the state space (for full details, see [8]). Powerful mathematical results are known for such monotone systems, upon which is based the reduction from many dimensions to two. For a model that satisfies the requirements, monotone theory shows that the steady state behaviour and its parameter dependence is no more complex than would be expected for the associated two-dimensional model. This can be an useful tool when it can be applied.

If an enzymatic reaction is modelled in the standard biochemical manner [18], with an enzyme-substrate complex and mass-action kinetics, then it is not monotone. It becomes monotone in the quasi steady-state approximation, which leads to the familiar Michaelis-Menten rate function. While continuing to be widely used in complex models, the Michaelis-Menten function is suspect for at least two reasons. First, in the context of a single enzyme acting on a single substrate, it emerges through a singular perturbation based on a separation of time scales, which is only known to be accurate under certain conditions on the enzyme and substrate [80, 95]. Second, because the enzyme-substrate complex is removed from the dynamics (which is what makes the perturbation singular), the approximation cannot capture enzyme sequestration when there are many substrates present. This can readily lead to errors. The “total quasi steady state” approximation appears safer in both respects [16]. It would be interesting if a separation of time scales argument could be found that was both broadly accurate and also resulted in monotonicity.

Algebraic geometric methods As the previous discussion suggests, there is much to be said for constructing a model directly from a network of chemical reactions using the principle of mass action. This is a systematic procedure which allows the biochemistry to be modelled in a realistic form. (Of course, the sheer complexity of biology makes this infeasible in general.) Mass action has one other consequence, which has, until recently, been largely overlooked. If the rate function $f(x; a)$ in equation (1.3) comes from some network of chemical reactions by mass action, then it is always a polynomial function of the state variables, x_1, \dots, x_n . Accordingly, the steady states of the system, at which $dx/dt = 0$, correspond to an *algebraic variety* [19]. One of the interesting features of algebraic geometry, which it shares with linear algebra, is that it can be undertaken over an arbitrary coefficient field. In particular, the set of steady state solutions, $\{f_i(x; a) = 0\}$, can be regarded as an algebraic variety over the field $\mathbb{R}(a)$ of real rational functions in the parameters, a_1, \dots, a_m . In other words, the parameters can be treated as uninterpreted symbols, rather than as actual numbers, to which can be applied, nevertheless, all the usual arithmetic operations of addition, subtraction, multiplication and division.

While this possibility is evident, it has not previously been exploited because there appeared to be nothing one could say about the geometric structure of the steady state variety. Recently, we have shown that for multisite phosphorylation systems, the

steady state variety forms a *rational algebraic curve* over $\mathbb{R}(a)$ [62, 92]. Rationality provides an explicit description of the steady states, which, together with the ability to roam algebraically over the parameter space, leads to unexpected insights. We show that such systems can have a parameter region with multiple stable steady states, whose maximum number increases with the number of sites, suggesting that multisite phosphorylation, which plays a key regulatory role in most cellular processes, can implement complex information processing [91]. The method also yields stringent quantitative predictions which, nevertheless, do not require parameter values to be known or estimated [62]. While these results are currently limited to multisite phosphorylation, they suggest that algebraic geometric methods and symbolic parameter analysis may have wider application to the parameter problem in systems biology.

The freedom to treat parameters as algebraic symbols applies only to the steady state; the dynamics, which depend upon derivatives and infinitesimal procedures, are fundamentally non-algebraic. It remains an interesting question, however, to what extent other attractors, such as periodic orbits, can also be analysed symbolically.

1.5 THE MEANINGS OF ROBUSTNESS

Robustness is one of the themes to have emerged in systems biology [6, 11, 52, 86] and it is particularly relevant to the parameter problem. Unfortunately, it is also one of those concepts whose wide usage has not been matched by precise definition. Robustness means, broadly, that some property of the system remains the same under perturbation. To make this precise, it is necessary to say what the property is, in what sense it remains the same and what kinds of perturbations are being considered. The property might be the overall qualitative dynamics of a system, in which case “remaining the same” could mean that the number and type of attractors and the connectivity and shape of the trajectories remain the same under perturbation. Alternatively, the property could be a quantitative function evaluated on an attractor, like the period of a periodic orbit. In this case, “remaining the same” could mean that the property remains quantitatively unchanged under perturbation (“exact robustness”) or that it only changes by a limited amount (“approximate robustness”). As for perturbations, at least three different kinds can be distinguished: changes to parameter values, changes to initial conditions and changes to the functional form that describes the dynamics (ie: the f in equation (1.3) for an ODE model). These perturbations have distinct mathematical and biological implications. We will discuss the first two as preparation for reviewing some influential studies of robustness and then return to the third.

1.5.1 Parameter biology

Consider an ODE model derived by the principle of mass action from a network of biochemical reactions. In this case, the parameters are rate constants of various kinds: association rates, disassociation rates, catalytic rates, etc. Such rates are, hopefully (see the next paragraph), intrinsic features of the corresponding proteins

and would not be expected to change except through alterations to their amino acid sequences. This could happen on an evolutionary time scale, so that different species may have different parameter values, but this would not be expected to happen in different cells of the same organism or tissue or clonal population of cells in cell culture. The situation could be different in a polyclonal population, such as a tumour or a natural population of outbred organisms, in which there could be substantial genetic polymorphism. Depending on which loci exhibit polymorphism and how it affects protein function, this genetic variation could give rise to rate constant variation between different cells or different organisms.

(A caveat is essential here. Rate constants are not solely determined by intrinsic features of a protein. They also depend on the ambient conditions in the cell—temperature, pH, other ionic strengths—as well as, potentially, post-translational modifications such as disulphide bridges or glycosylations, or the presence of accessory molecules such as chaperones or scaffolds, none of which might have been included in a model. The reductionist approach commonly used in systems biology, in which the properties of a system are deduced from its components, is always at risk of the system biting back: the properties of the components may depend on that of the system [37]. To put it another way, the boundary of a system has to be drawn somewhere, with the implicit assumption that what is outside the boundary is irrelevant to the behaviour inside. Such assumptions tend to be taken for granted until they fail.)

Models are not always deduced from mass action. For instance, separation of time scales is often convenient, if not essential, in reducing complexity. Whether this is achieved through the suspect “quasi steady state”, or the safer “total quasi steady state”, approximations discussed in §1.4.2, it necessarily leads to parameters which are no longer rate constants. Similarly, models of allosteric enzymes [56, 67] or rate functions for gene expression in terms of transcription factor binding [1] are also based on separation of time scales and lead to rational algebraic rate functions resembling the ubiquitous Hill functions. (Despite their very wide usage, Hill functions are not derived from any approximation and have no well-founded mechanistic interpretation [18].) The basic issues can be discussed for the Michaelis-Menten formula

$$\frac{rx}{k+x}, \quad (1.5)$$

in which r , the maximal rate, and k , the Michaelis-Menten constant, are the two parameters. Of these, k is derived from rate constants [18] and may hence be assumed to vary only under the same conditions. Notice, however, that this depends on the underlying mechanistic derivation of (1.5) and on the assumptions behind it. As for r , it is, in terms of the usual derivation [18], a product of a catalytic rate and an enzyme concentration. The enzyme is not formally part of the dynamics but its concentration can change on multiple time scales. On a physiological time scale the concentration is set by the balance between synthesis and degradation and could readily vary from cell to cell within a single organism, tissue or clonal population through differences in cell volumes, intrinsic noise in transcription/translation and stochastic partitioning of molecules during cell division. In polyclonal populations, genetic variation or gene

copy number variation could introduce additional variation in concentration levels. These factors would also play a role on a longer evolutionary time scale. As we see, the biological interpretation of changes to parameter values depends both on the model and the nature of its parameters as well as on the biological context that is being modelled.

1.5.2 Robustness to initial conditions

If the property thought to be robust is associated to an attractor, such as a steady state, then its robustness to initial conditions would seem to follow from the stability of the attractor, in the sense discussed in §1.4.1. However, it is often the case that the dynamics satisfy additional constraints. For instance, an enzyme suffers no net change in concentration in any reaction that it catalyses. If it is not being otherwise synthesized or degraded then its total concentration remains constant at all times. Similarly, if a substrate exists in many states of modification—multisite phosphorylation, for instance—and is also not synthesized or degraded, then its total concentration remains constant. (Note that these constraints are linear in the state variables; non-linear constraints may also be possible.) If there are k independent constraints, they confine the dynamics to lie within a subspace of dimension $d = n - k$, where n is the dimension of the ambient space. The state space thereby becomes divided into “slices” of dimension d , each corresponding to a set of constraint values (Figure 1.4). Within each slice, the dynamics behave as they did in Figure 1.1, with attractors, basins of attraction and stability, as appropriate to an ambient space of dimension d (not n). However, its qualitative character can change with the constraint values. Hence, the constraint space also becomes divided into regions, within each of which the dynamics in the corresponding slices remain qualitatively similar (Figure 1.4).

Unlike variation of parameters, rather little seems to be known about variation of constraints. Parameters and constraints are mathematically distinct. Parameters can be chosen independently of initial conditions, while constraints cannot. Parameters define the dynamics; constraints confine the dynamics. The biological implications of the two forms of robustness can be quite distinct; see §1.5.3.

In summary, for properties associated to an attractor, robustness to initial conditions may take two forms. If initial conditions are varied within the same set of constraint values, then it corresponds to stability of the attractor and the dimension and shape of the basin of attraction in state space (with respect to the effective ambient space of dimension d) provide measures of it. If constraint values are varied, then robustness goes beyond stability and the dimension and shape of the appropriate region in constraint space become relevant.

1.5.3 Robustness in reality

With this background, let us review some particularly interesting and influential demonstrations of robustness in different biological systems.

Signalling in bacteria is typically implemented by two-component systems consisting of a sensor kinase coupled to a response regulator protein [102]. The sensor




Figure 1.4 Dynamical system with constraints. The state space becomes divided into “slices”, represented by the blue lines, each slice corresponding to a set of constraint values, represented by a point in the space of constraints. Note that if the invariants are nonlinear, then the slices may be curved spaces. The dynamics are confined within the slices. If an initial condition is chosen within a slice, then the trajectory remains within that slice for all time; trajectories never cross between slices. The dynamics within a slice can have attractors, represented by stars, and other features as described in Figure 1.1 but their qualitative character can change as the constraints vary, as illustrated by the appearance and disappearance of attractors.

autophosphorylates in response to a signal, using ATP as the phosphate donor. It then transfers the phosphate to the response regulator, which initiates the signalling response, by, for instance, stimulating gene transcription. In some two component systems involved in homeostasis—such as the EnvZ/OmpR system which regulates osmolarity in *Escherichia coli*—the sensor also catalyses the dephosphorylation of the response regulator. This unusual bifunctional mechanism has been studied in several models [12, 77, 81], whose general conclusion is that the mechanism enables the amount of phosphorylated response regulator at steady state to be constraint robust with respect to changes in the total amounts of sensor and response regulator. The initial analysis by Russo and Silhavy using Michaelis-Menten kinetics [77], which provided the first indication of this robustness, was subsequently refined using mass-action kinetics by Batchelor and Goulian [12]. Their analysis showed approximate constraint robustness when the amount of sensor kinase is much less than the amount of response regulator, which, indeed, corresponds well to *E. coli*'s normal operating regime. In their accompanying experimental analysis they varied the total amounts of EnvZ and OmpR and found good agreement with their model. Shinar *et al* incorporated a further element into the mechanism by noting that in certain bifunctional two-component systems [81, Table 1], including the EnvZ/OmpR system in *E. coli*, ATP acts as a cofactor in the dephosphorylation of the response regulator. Their model for this shows exact constraint robustness of the amount of phosphorylated response regulator, with respect to changes to the total amounts of sensor, response regulator

and ATP, provided the amount of response regulator remains above a threshold. These predictions were also borne out by experiment.

E. coli has also been a model bacterium for the study of chemotaxis. It moves by rotating its multiple flagella. Rotation in one direction brings the flagella into alignment, allowing the bacterium to “run” in a straight line. Rotation in the other direction drives the flagella apart, causing the bacterium to “tumble” and randomly reorient its direction. By regulating its tumbling frequency, the bacterium can efficiently seek out nutrients and escape poisons (chemotaxis) in environments that lie outside its control. Because *E. coli* is so small, it has to sense changes in ligand concentration over time, not space. It has been found to adapt its sensitivity to such changes across a remarkably broad range of background concentrations. Unravelling the mechanism behind this has been a triumph of systems biology [13].

Barkai and Leibler [10] studied robustness of the precision of adaptation in *E. coli* by simulation of an ODE model. For a given chemotactic ligand concentration the system appears to reach a steady state, presumably stable, irrespective of the other parameter values. To measure the precision of adaptation, the activated state of the receptor was evaluated at steady state for zero ligand and for a fixed saturating concentration of ligand (1mM) and the ratio of the latter to the former, denoted p , was taken as a measure of the precision of adaptation. If $p = 1$, the adaptation is “perfect”. There are three constraints in the model, corresponding to the total amounts of the receptor and the two chemotactic enzymes CheR and CheB, which implement adaptation by methylating and demethylating the receptor. The constraints and the parameters were randomly sampled and it was shown that the precision of adaptation remains close to 1 despite substantial perturbation around a reference model with physiologically realistic parameters and constraints.

This analysis reveals both constraint robustness and parametric robustness. In his commentary on [10], Hartwell invoked them both (implicitly) by suggesting that the robustness explains why chemotactic behaviours are buffered against the extensive polymorphism seen in outbred natural populations [40]. This is an attractive argument but it would be bolstered by knowing how much the genes in the chemotactic network are specifically affected by this polymorphism. How much of this variation contributes to variation in rate constants and how much to variation in concentration levels? Is the robustness in the model consistent with the actual level of variation seen in natural populations? Because of the implications for human physiology and disease, there are increasing data on polymorphisms in human populations [17] but few studies on how this affects the function of specific molecular systems or even individual proteins (see, for instance, [93]). In a subsequent paper with Alon and Surette [5], only constraint robustness was experimentally verified. Concentrations of chemotactic proteins were varied and the precision of adaptation was measured in individual bacteria. In these circumstances, each bacterium in a clone would be expected to have different concentrations of proteins through intrinsic noise in the transcriptional machinery and stochastic partitioning between daughter cells. The precision of adaptation was found to be very close to 1 [5, Table 1]. The experimental data are well explained by constraint robustness. However, since the parameters are

all rate constants, the data are not at all explained by parametric robustness. We see that these two types of robustness are distinct both mathematically and experimentally.

Morphogens are spatial signals that direct patterning in embryonic development [103]. They have been found to exhibit remarkable levels of robustness between different embryos [46, 26]. In a series of penetrating studies in *Drosophila*, Barkai and others used robustness as a design principle to identify molecular mechanisms that implemented it [26, 27]; for overviews, see [11, 28] and [4, Chapter 9]. They assumed a network of molecular interactions based on what was known in the literature and that the concentrations of the network components could vary between embryos because of polymorphisms in the population. By sampling points in parameter space (“numerical screening”), they identified regions in which the spatial profile of the morphogen exhibited robustness to changes in initial concentration levels of the network components. While these parametric regions were tiny ($< 1\%$ of the sampled points) they could be interpreted as particular kinds of mechanisms. The underlying models in these spatial studies are PDEs rather than ODEs but the qualitative framework of §1.4 can be used in much the same way and we see that the robustness here is to changes in the state space rather than the parameter space.

It would be interesting to know whether robustness to changes in the state space on a physiological time scale arises from the same mechanisms as robustness to changes in the parameter space on an evolutionary time scale, or whether different aspects of the molecular circuitry are responsible. As Waddington recognised [99], physiological robustness may lay the foundation for evolutionary adaptation (“genetic assimilation” as he put it); see also [51]. The different types of robustness discussed here may provide a framework for studying such questions.

1.5.4 Structural stability

Robustness with respect to functional variation—perturbing the f in equation (1.3)—has not been as widely utilised as the kinds of robustness described above. However, it was the basis for a remarkable historical episode which still has resonance for us today. Waddington’s distillation of biological dynamics inspired the distinguished French pure mathematician René Thom to develop a mathematical framework for describing it [100]. Thom made two general assumptions. First, that the dynamics arose from descending down a gradient, so that $f(x; a) = -\nabla g(x; a)$, where $\nabla = \sum_{i=1}^n \partial/\partial x_i$ is the gradient operator. Waddington’s epigenetic landscape has just such a gradient dynamics but for Thom the assumption arose from technical necessity rather than analogy and, in his case, the parameters play a key role. In gradient dynamics, steady states correspond to minima of the gradient function, g , which provides a crucial simplification. Second, Thom assumed that, in the absence of detailed knowledge about the underlying molecular mechanisms that gives rise to g , it was reasonable to focus on *structurally stable* behaviours. That is, those behaviours that remained qualitatively the same if the function g was perturbed, $g \rightarrow g + h$, where h is “small”. Under these assumptions, Thom proved that, for small numbers of parameters ($m \leq 5$), there were only finitely many—in fact, just eleven—different types of structurally stable bifurcations [74, Chapter 7]. Note that the state space can be of any dimension.

Furthermore, most bifurcations that have been studied tend to depend on only a few parameters, with the others playing only a background role. Hence, in practice, the restriction to $m \leq 5$ is not limiting.

The subtitle to Thom's book, "*An outline of a general theory of models*" [89], reflects the broad view he took of the scope of these results. The theory remains deep and difficult. Thom was himself a Fields medallist but could only guess parts of the argument and had to enlist the help of other mathematicians to complete the details. Later work filled in some gaps and clarified its place within mathematics, where it is now largely absorbed into bifurcation theory and singularity theory [9, 23]. Poston and Stewart remains the most accessible account [74]. Thom's own book [89], "*transcends the world of numbers*", as the back cover puts it.

What is important about Thom's theorem is that it gives the first hint that even very complex dynamics may still be composed of only a small finite number of key "motifs". At the same time, from our vantage point, the difficulties with Thom's assumptions become much clearer. First, the dynamics arising from molecular networks are rarely of gradient type. Second, it is not reasonable to perturb f in an arbitrary way, since the resulting perturbed function may not have arisen from any molecular network. What is required, instead, is a restricted notion of structural stability in which perturbations are confined to a biochemically realistic sub-class.

We unwittingly undertook a computational study of this in the context of developmental patterning in the *Drosophila* embryo [61]. In an influential paper, von Dassow *et al* had found that the segment polarity gene regulation network was parametrically robust [98] (using the language developed here) and the evolutionary implications of this were widely cited [103]. However, their model was based on a regular hexagonal lattice of cells, which is far from the normal structure of an epithelium [35]. Moreover, because cellularisation in *Drosophila* takes place late in embryonic development, the segment polarity network has to operate without knowing in advance what lattice of cells has emerged. If the cellular lattice is changed, the effect on the model is to change f in a biochemically realistic manner, through alterations in cell-to-cell communication. Hence, robustness to lattice variation is a form of restricted structural stability. Our paper was concerned with computational infrastructure for building models rather than with structural stability (the relevance of which was unclear at the time) but our limited analysis suggested that the segment polarity network was structurally unstable despite being parametrically robust. We speculated that small changes to the underlying molecular network might render it robust to lattice variation but were unable to pursue this further.

In our analysis, the robustness operates on the physiological and not the evolutionary time scale. However, molecular networks can be reorganised during evolution, which can change both nodes and links, as well as parameter values and expression levels. Restricted structural stability might be the appropriate type of robustness with which to study this.

We see from this that Thom's ideas remain relevant, despite being largely forgotten. His approach came to be called "catastrophe theory" and garnered great celebrity, being compared to Newton's theory of gravitation and mechanics. The resulting fall from grace was predictably brutal [55, 105]; for a more balanced perspective, see

[9, 106]. While most of those who know of it think it dead and buried, I think, in contrast, that it has merely been dormant, waiting for systems biology to provide a more fertile landscape for Thom's ideas to germinate again.

1.5.5 Classifying robustness

One reason why robustness has attracted such attention is that it may be a biological design principle, [52]. This is an appealing idea but to make sense of it, robustness needs to be precisely defined and grounded in the kind of careful experiments discussed in §1.5.3. As we have shown, there are different types of robustness, which may be classified according to which aspect of the dynamical system is changed.

- Type I: dynamical stability. Robustness to change of initial conditions within a fixed set of constraint values.
- Type II: constraint robustness. Robustness to change of constraint values.
- Type III: parametric robustness. Robustness to change of parameter values.
- Type IV: structural stability. Robustness to change of the dynamical function.

No doubt there are others. As noted in §1.5.1, the interpretation of these mathematical properties depends crucially on the biological context that is being modelled. Robustness could be quantified if we could estimate the size and shape of various regions in high-dimensional spaces: basins of attraction, constraint regions, parameter regions. Many studies can be seen as attempts to do this by random sampling [10, 98]. Lack of space precludes a discussion of robustness trade-offs [21, 52] and new methods of global sensitivity analysis [38, 75]. Kitano has remarked on the need for a theory of biological robustness [53]. The dynamical systems framework outlined here may provide a basis for this.

1.6 CONCLUSION

One of the difficulties for students of systems biology is to make sense of the many different concepts and techniques that are coming into the subject from the physical sciences and computer science. Those of us who have been trained in these other disciplines necessarily take a particular perspective (as will be evident to readers of this paper) and it is ultimately our students who bear the burden of harmonising this cacophony. Steven Pinker tells the story [73] of indentured labourers from different language groups being brought together on some remote island under colonial occupation. The first generation cobbles together a form of communication—a “pidgin”—which suffices for getting along on an everyday basis. It is the second generation who, spontaneously and magically, creates a fully-fledged natural language, a “creole”. It should be evident that this paper is written in systems biology pidgin. Let us hope that, in time, our students will teach us how to write systems biology creole.

I thank the students of MCB195, SB101 and SB200 for their questioning enthusiasm, which prompted this paper; Uri Alon for many helpful and perceptive comments; an anonymous reviewer for pointing out some issues of scope that needed clarification; Rebecca Ward for her unerring editorial eye; and the editors for their patience. They bear no responsibility for any of the paper's remaining mistakes, obscurities or omissions, for which I alone must apologise.

REFERENCES

1. G. K. Ackers, A. D. Johnson, and M. A. Shea. Quantitative model for gene regulation by lambda phage repressor. *Proc. Natl. Acad. Sci. USA*, 79:1129–33, 1982.
2. J. G. Albeck, J. M. Burke, B. B. Aldridge, M. Zhang, D. A. Lauffenburger, and P. K. Sorger. Quantitative analysis of pathways controlling extrinsic apoptosis in single cells. *Mol. Cell*, 30:11–25, 2008.
3. B. B. Aldridge, J. M. Burke, D. A. Lauffenburger, and P. K. Sorger. Physicochemical modelling of cell signalling pathways. *Nat. Cell Biol.*, 8:1195–203, 2006.
4. U. Alon. *An Introduction to Systems Biology: Design Principles of Biological Circuits*. Chapman and Hall/CRC, 2006.
5. U. Alon, M. G. Surette, N. Barkai, and S. Leibler. Robustness in bacterial chemotaxis. *Nature*, 397:168–71, 1999.
6. I. Amit, R. Wides, and Y. Yarden. Evolvable signalling networks of receptor tyrosine kinases: relevance of robustness to malignancy and cancer therapy. *Mol. Syst. Biol.*, 3:151, 2007.
7. A. R. A. Anderson, A. M. Weaver, P. T. Cummings, and V. Quaranta. Tumor morphology and phenotypic evolution driven by selective pressure from the microenvironment. *Cell*, 127:905–15, 2006.
8. D. Angeli, J. E. Ferrell, and E. D. Sontag. Detection of multistability, bifurcations, and hysteresis in a large class of biological positive-feedback systems. *Proc. Natl. Acad. Sci. USA*, 101:1822–7, 2004.
9. V. I. Arnold, V. F. Afrajmovich, Yu. S. Il'yashenko, and L. P. Shil'nikov. *Bifurcation Theory and Catastrophe Theory*. Encyclopedia of Mathematical Sciences. Springer, 1999.
10. N. Barkai and S. Leibler. Robustness in simple biochemical networks. *Nature*, 387:913–7, 1997.
11. N. Barkai and B.-Z. Shilo. Variability and robustness in biomolecular systems. *Mol. Cell*, 28:755–60, 2007.
12. E. Batchelor and M. Goulian. Robustness and the cycle of phosphorylation and dephosphorylation in a two-component regulatory system. *Proc. Natl. Acad. Sci. USA*, 100:691–6, 2003.
13. H. C. Berg. *E. coli in Motion*. Springer-Verlag, New York, NY, USA, 2004.
14. A. Brünger. Free R value: a novel statistical quantity for assessing the accuracy of crystal structures. *Nature*, 355:472–5, 1992.

15. E. Di Cera, P. E. Phillipson, and J. Wyman. Limit cycle oscillations and chaos in reaction networks subject to conservation of mass. *Proc. Natl. Acad. Sci. USA*, 86:142–6, 1989.
16. A. Ciliberto, F. Capuani, and J. J. Tyson. Modeling networks of coupled enzymatic reactions using the total quasi-steady state approximation. *PLoS Comp. Biol.*, 3:e45, 2007.
17. The International Hap Map Consortium. The international HapMap project. *Nature*, 426:789–96, 2003.
18. A. Cornish-Bowden. *Fundamentals of Enzyme Kinetics*. Portland Press, London, UK, 2nd edition, 1995.
19. D. Cox, J. Little, and D. O’Shea. *Ideals, Varieties and Algorithms*. Springer, 2nd edition, 1997.
20. G. Craciun, Y. Tang, and M. Feinberg. Understanding bistability in complex enzyme-driven reaction networks. *Proc. Natl. Acad. Sci. USA*, 103:8697–02, 2006.
21. M. E. Csete and J. C. Doyle. Reverse engineering of biological complexity. *Science*, 295:1664–9, 2002.
22. V. Danos, J. Feret, W. Fontana, and J. Krivine. Abstract interpretation of cellular signalling networks. In *Proceedings VMCAI 2008*, volume 4905 of *Lecture Notes in Computer Science*. Springer-Verlag, 2008.
23. M. Demazure. *Bifurcations and Catastrophes*. Universitext. Springer, 2000.
24. G. Dupont, S. Swillens, C. Clair, T. Tordjmann, and L. Combettes. Hierarchical organization of calcium signals in hepatocytes: from experiments to models. *Biochim. Biophys. Acta*, 1498:134–52, 2000.
25. F. Dyson. A meeting with Enrico Fermi. *Nature*, 427:297, 2004.
26. A. Eldar, R. Dorfman, D. Weiss, H. Ashe, B.-Z. Shilo, and N. Barkai. Robustness of the BMP morphogen gradient in *Drosophila* embryonic patterning. *Nature*, 419:304–8, 2002.
27. A. Eldar, D. Rosin, B.-Z. Shilo, and N. Barkai. Self-enhanced ligand degradation underlies robustness of morphogen gradients. *Dev. Cell*, 5:635–46, 2003.
28. A. Eldar, B.-Z. Shilo, and N. Barkai. Elucidating mechanisms underlying robustness of morphogen gradients. *Curr. Opin. Genet. Dev.*, 14:435–9, 2004.
29. M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain. Stochastic gene expression in a single cell. *Science*, 297:1183–6, 2002.
30. M. Feinberg. Chemical reaction network structure and the stability of complex isothermal reactors I. The deficiency zero and deficiency one theorems. *Chem. Eng. Sci.*, 42(10):2229–68, 1987.
31. J. E. Ferrell. Self-perpetuating states in signal transduction: positive feedback, double-negative feedback and bistability. *Current Opinion in Chemical Biology*, 6:140–8, 2002.
32. J. Fisher and T. Henzinger. Executable cell biology. *Nat. Biotech.*, 25:1239–49, 2007.
33. R. A. Fisher. *The Genetical Theory of Natural Selection*. Clarendon Press, Oxford, 1930.
34. W. Fontana. Algorithmic chemistry. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*. Addison-Wesley, Redwood City, CA, USA, 1992.
35. M. C. Gibson, A. B. Patel, R. Nagpal, and N. Perrimon. The emergence of geometric order in proliferating metazoan epithelia. *Nature*, 442:1038–41, 2006.

36. J. Gunawardena. Chemical Reaction Network Theory for *in-silico* biologists. Lecture notes, Harvard Univ, 2003. <http://vcp.med.harvard.edu/papers/crnt.pdf>.
37. J. Gunawardena. Models in systems biology: thick and thin models and the dilemmas of reductionism. In preparation, 2008.
38. R. N. Gutenkunst, J. J. Waterfall, F. P. Casey, K. S. Brown, C. R. Myers, and J. P. Sethna. Universally sloppy parameter sensitivities in systems biology models. *PLoS Comput. Biol.*, 3:1871–8, 2007.
39. A. C. Guyton, T. G. Coleman, and H. J. Granger. Circulation: overall regulation. *Annu. Rev. Physiol.*, 34:13–44, 1972.
40. L. Hartwell. A robust view of biochemical pathways. *Nature*, 387:855–7, 1997.
41. R. Heinrich and T. A. Rapoport. A linear steady-state treatment of enzymatic chains. General properties, control and effector strengths. *Eur. J. Biochem.*, 42:89–95, 1974.
42. M. W. Hirsch and S. Smale. *Differential Equations, Dynamical Systems and Linear Algebra*. Pure and Applied Mathematics. Academic Press, San Diego, USA, 1974.
43. W. S. Hlavacek, J. R. Faeder, M. L. Blinov, R. G. Posner, M. Hucka, and W. Fontana. Rules for modeling signal-transduction systems. *Sci. STKE*, 344:re6, 2006.
44. A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.*, 117:500–44, 1942.
45. J. E. Hopcroft, R. Motwani, and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison Wesley, Boston, MA, USA, 2006.
46. B. Houchmandzadeh, E. Wieschaus, and S. Leibler. Establishment of developmental precision and proportions in the early *Drosophila* embryo. *Nature*, 415:798–802, 2002.
47. S. Huang and D. E. Ingber. A non-genetic basis for cancer progression and metastasis: self-organizing attractors in cell regulatory networks. *Breast Dis.*, 26:27–54, 2006-7.
48. H. Kacser and J. A. Burns. The control of flux. *Biochem. Soc. Trans.*, 23:341–66, 1995. Reprint of 1973 paper in Symp. Soc. Exp. Biol.
49. S. A. Kauffman. Metabolic stability and epigenesis in randomly constructed genetic nets. *J. Theor. Biol.*, 22:437–67, 1969.
50. M. Kirschner. The meaning of systems biology. *Cell*, 121:503–4, 2005.
51. M. W. Kirschner and J. C. Gerhart. *The Plausibility of Life*. Yale University Press, 2005.
52. H. Kitano. Biological robustness. *Nat. Rev. Genet.*, 5:826–37, 2004.
53. H. Kitano. Towards a theory of biological robustness. *Mol. Syst. Biol.*, 3:137, 2007.
54. C. Koch. *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press, 1999.
55. G. Kolata. Catastrophe theory: the emperor has no clothes. *Science*, 196:287–351, 1977.
56. D. E. Koshland, G. Nemethy, and D. Filmer. Comparison of experimental binding data and theoretical models in proteins containing subunits. *Biochemistry*, 5:365–85, 1966.
57. J. A. Lakowicz. *Principles of Fluorescence Spectroscopy*. Kluwer Academic, New York, 2nd edition, 1999.
58. J. Lewis. Autoinhibition with transcriptional delay: a simple mechanism for the Zebrafish somitogenesis oscillator. *Curr. Biol.*, 13:1398–1408, 2003.

59. S. Li, S. M. Assmann, and Réka Albert. Predicting essential components of signal transduction networks: a dynamic model of guard cell abscisic acid signaling. *PLoS Biol.*, 4:e312, 2006.
60. S. E. Luria and M. Delbrück. Mutations of bacteria from virus sensitivity to virus resistance. *Genetics*, 28:491–511, 1943.
61. A. Mallavarapu, M. Thomson, B. Ullian, and J. Gunawardena. Programming with models: modularity and abstraction provide powerful capabilities for systems biology. *J. R. Soc. Interface*, doi:10.1098/rsif.2008.2005, 2008.
62. A. Manrai and J. Gunawardena. The geometry of multisite phosphorylation. *Biophys. J.*, doi:10.1529/biophysj.108.140632, 2008.
63. R. May. *Stability and Complexity in Model Ecosystems*. Princeton University Press, 1973.
64. A. Mehra, C. I. Hong, M. Shi, J. J. Loros, J. C. Dunlap, and P. Ruoff. Circadian rhythmicity by autocatalysis. *PLoS Comput. Biol.*, 2:0816–23, 2006.
65. G. Mendel. Versuche über Pflanzen-Hybriden. *Verhandlungen des naturforschenden Vereines, Abhandlungen, Brünn*, 4:3–47, 1866.
66. G. Mitchison. Theory in biology. Happy days here again? *Curr. Biol.*, 14:R97–8, 2004.
67. J. Monod, J. Wyman, and J. P. Changeux. On the nature of allosteric transitions: a plausible model. *J. Mol. Biol.*, 12:88–118, 1965.
68. M. A. Nowak. *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap Press, 2006.
69. E. M. Ozbudak, M. Thattai, H. N. Lim, B. I. Shraiman, and A. van Oudenaarden. Multistability in the lactose utilization network of *Escherichia coli*. *Nature*, 427:737–40, 2004.
70. B. Ø. Palsson. *Systems Biology: Properties of Reconstructed Networks*. Cambridge University Press, New York, NY, USA, 2006.
71. J. Paulsson. Summing up the noise in gene networks. *Nature*, 427:415–8, 2004.
72. J. Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, Cambridge, UK, 2000.
73. S. Pinker. *The Language Instinct: How the Mind Creates Language*. William Morrow and Co., 1994.
74. T. Poston and I. Stewart. *Catastrophe Theory and its Applications*. Dover Publications, 1996.
75. D. A. Rand. Mapping the global sensitivity of cellular network dynamics: sensitivity heat maps and a global summation law. *J. Roy. Soc. Interface*, 5 Suppl 1:S59–69, 2008.
76. A. Regev, W. Silverman, and E. Shapiro. Representation and simulation of biochemical processes using the pi-calculus process algebra. *Pac. Symp. Biocomput.*, pages 459–70, 2001.
77. F. D. Russo and T. J. Silhavy. The essential tension: opposed reactions in bacterial two-component regulatory systems. *Trends Microbiol.*, 1:306–10, 1993.
78. M. A. Savageau. *Biochemical Systems Analysis: A Study of Function and Design in Molecular Biology*. Addison-Wesley, 1976.
79. M. Schena, D. Shalon, R. W. Davis, and P. O. Brown. Quantitative monitoring of gene expression patterns with a complementary DNA microarray. *Science*, 270:467–70, 1995.

80. L. Segel. On the validity of the steady-state assumption of enzyme kinetics. *Bull. Math. Biol.*, 50:579–93, 1988.
81. G. Shinar, R. Milo, M. R. Martínez, and U. Alon. Input-output robustness in simple bacterial signaling systems. *Proc. Natl. Acad. Sci. USA*, 104:19931–5, 2007.
82. S. K. Sieberts and E. E. Schadt. Moving towards a system genetics view of disease. *Mamm. Genome*, 18:389–401, 2007.
83. J. M. W. Slack. Conrad Hal Waddington: the last Renaissance biologist. *Nat. Rev. Genet.*, 3:889–95, 2002.
84. B. M. Slepchenko, J. C. Schaff, J. H. Carson, and L. M. Loew. Computational cell biology: spatiotemporal simulation of cellular events. *Annu. Rev. Biophys. Biomol. Struct.*, 31:423–41, 2002.
85. B. L. Sprague, R. L. Pego, D. A. Stavreva, and J. G. McNally. Analysis of binding reactions by fluorescence recovery after photobleaching. *Biophys. J.*, 86:3473–95, 2004.
86. J. Stelling, U. Sauer, Z. Szallasi, F. J. Doyle III, and J. Doyle. Robustness of cellular functions. *Cell*, 118:675–85, 2004.
87. S. H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry and Engineering*. Perseus Books, 2001.
88. E. Tajkhorshid, P. Nollert, M. O. Jensen, L. J. Miercke, J. O’Connell, R. M. Stroud, and K. Schulten. Control of the selectivity of the aquaporin water channel family by global orientational tuning. *Science*, 296:525–30, 2002.
89. R. Thom. *Structural Stability and Morphogenesis*. W. A. Benjamin, Inc., Reading, MA, USA, 1975. English translation of the 1972 French edition.
90. R. Thomas and R. D’Ari. *Biological Feedback*. CRC Press, Boca Raton, Florida, USA, 1990.
91. M. Thomson and J. Gunawardena. A new method of steady state analysis reveals unlimited multistability in multisite phosphorylation systems. Submitted, 2008.
92. M. Thomson and J. Gunawardena. The steady states of a multisite kinase, phosphatase, substrate system form a rational projective algebraic curve. In preparation, 2008.
93. M. Tiseo, M. Capelletti, G. De Palma, V. Franciosi, A. Cavazzoni, P. Mozzoni, R. R. Alfieri, M. Goldoni, M. Galetti, B. Bortesi, C. Bozzetti, M. Loprevite, L. Boni, R. Camisa, G. Rindi, P. G. Petronini, and A. Ardizzoni. Epidermal growth factor receptor intron-1 polymorphism predicts gefitinib outcome in advanced non-small cell lung cancer. *J. Thorac. Oncol.*, 3:1104–11, 2008.
94. J. J. Tyson, Kathy Chen, and Bela Novak. Network dynamics and cell physiology. *Nature Rev. Mol. Cell Bio.*, 2:908–916, 2001.
95. R. A. Tzafiriri. Michaelis-Menten kinetics at high enzyme concentration. *Bull. Math. Biol.*, 65:1111–29, 2003.
96. N. G. van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, Amsterdam, The Netherlands, 1992.
97. A. Varma, M. Morbidelli, and H. Wu. *Parametric Sensitivity in Chemical Systems*. Cambridge University Press, 2005.
98. G. von Dassow, E. Meir, E. M. Munro, and G. M. Odell. The segment polarity network is a robust developmental module. *Nature*, 406:188–92, 2000.

99. C. H. Waddington. *The Strategy of the Genes: A Discussion of Some Aspects of Theoretical Biology*. George Allen & Unwin Ltd., London, 1957.
 100. C. H. Waddington, editor. *Towards a Theoretical Biology. 1. Prolegomena*. Edinburgh University Press, 1968.
 101. E. Walter and L. Pronzato. *Identification of Parametric Models from Experimental Data*. Springer, 1997.
 102. A. H. West and A. M. Stock. Histidine kinases and response regulator proteins in two-component signaling systems. *Trends Biochem. Sci.*, 26:369–76, 2001.
 103. L. Wolpert. *Principles of Development*. Oxford University Press, 2001.
 104. S. Wright. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In *Proceedings of the 6th International Conference on Genetics, Vol. 1*, pages 356–66. 1932.
 105. R. S. Zahler and H. J. Sussmann. Claims and accomplishments of applied catastrophe theory. *Nature*, 269:759–63, 1977.
 106. E. C. Zeeman, R. Bellairs, B. Goodwin, M. R. Mackley, I. Stewart, M. Berry, J. Guckenheimer, and A. E. R. Woodcock. In support of catastrophe theory. *Nature*, 270:381–4, 2005.
- ...