

Von Neumann's growth model: statistical mechanics and biological applications

Andrea De Martino¹, Enzo Marinari², and Andrea Romualdi²

¹ CNR-IPCF, Dipartimento di Fisica, Sapienza Università di Roma, Roma, Italy

² Dipartimento di Fisica, Sapienza Università di Roma, Roma, Italy

Abstract. We review recent work on the statistical mechanics of Von Neumann's growth model and discuss its application to cellular metabolic networks. In this context, we present a detailed analysis of the physiological scenario underlying optimality à la Von Neumann in the metabolism of the bacterium *E. coli*, showing that optimal solutions are characterized by a considerable microscopic flexibility accompanied by a robust emergent picture for the key physiological functions. This suggests that the ideas behind optimal economic growth in Von Neumann's model can be helpful in uncovering functional organization principles of cell energetics.

1 Introduction

In the late 1930's John Von Neumann published (originally in German) a simple linear model to describe optimal economic growth [1,2]. Quite generally, he conceived a system in which N technologies can operate by combining M commodities. Technologies are distinguished by the amounts of goods they use as inputs and return as outputs. Moreover, they are linear (or constant returns to scale in economic jargon), in the sense that doubling the amount of inputs allows to double that of outputs (which implies that, for instance, saturation effects are neglected). Crucially, the system is required to be self-sustained, so that all goods necessary for production at each period must have been produced at the previous period. Given the input/output specifics of production processes, one poses the question of establishing which processes would be operated and which would instead be kept inactive if the system has to maximize the net production rate of commodities. It has been noted that Von Neumann, who had been trained in chemistry and chemical engineering, might have been inspired by chemical reactors for such a construction [3]. In any case, his growth model has become a cornerstone of the theory of economic growth [4,5].

Von Neumann's general idea can easily be transported to other contexts outside economics, where input/output processes of various kinds are wired together to form a heterogeneous network [6,7,8]. A very notable example is found in cells. Metabolism, namely the complex web of chemical reactions underlying energy transduction in living organisms [9,10,11], becomes an ideal candidate to test Von Neumann's ideas once the cross-membrane flow of matter is included in the above scheme. Given the fundamental nature of Von Neumann's assumptions, it is natural to ask to which

degree the chemical activity of real cells can actually be described by an optimal growth (in the sense discussed above) framework.

On the other hand, the analysis of optimal growth in linear input/output systems is appealing for statistical physicists as well. As the system size increases, the optimal growth properties of ‘random’ input/output networks should hopefully become independent of the particular input-output relationships. Rather, it should be possible to highlight a set of structural characteristics or constraints that are sufficient to frame optimal growth factors and activity profiles. Random input-output systems then constitute a key benchmark against which one can evaluate the performances of single real instances, and the theory of spin glasses and neural networks provides the technical and conceptual tools for this type of analysis [12]. Indeed, their effectiveness to deal with linear problems close in spirit to Von Neumann’s setup has been demonstrated in various instances in the past (see e.g. the so-called knapsack problem [13,14] or even, more recently, the problem of general economic equilibrium [15]).

Our goal here is, on one hand, to look back on the recent work concerned with the statistical mechanics of Von Neumann’s growth model for random input-output systems. These studies have revealed a universal behaviour of the relevant quantities and allowed to highlight the key ingredients that can bring the basic framework closer to biologically important questions. Furthermore, we wish to discuss its application to cellular metabolism, which began more recently and is an ongoing and expanding challenge. To complement previous studies, we shall present a broad analysis of the physiological scenario underpinned by Von Neumann’s model applied to the cellular metabolism of the bacterium *Escherichia coli*. In our view, such studies demonstrate that the basic postulate of optimal growth à la Von Neumann might provide a useful key to understand the organization of cell metabolism and energetics.

2 Von Neumann’s model

The basic setup of Von Neumann’s growth model is as follows. One considers a system of N distinct processes (technologies, reactions) that are capable of operating on M compounds (commodities, chemical species). The system is assumed to be closed, i.e. there is no external supply of compounds to the system. In other terms, compounds can only be produced from each other. Each process i is specified by a vector $\mathbf{b}_i = \{b_i^\mu\}$, representing for each $\mu = 1, \dots, M$ the amount of compound μ process i uses as input (the substrate), and by a vector $\mathbf{a}_i = \{a_i^\mu\}$, representing for each $\mu = 1, \dots, M$ the amount of compound μ process i outputs (the product). In the simplest case, it is assumed that $a_i^\mu, b_i^\mu \geq 0$ for each i and μ . For sakes of definiteness, each process should have at least one input, whereas each compound should be the output of at least one process. In turn, the input and output vectors form the columns of the $M \times N$ input and output matrices, which we shall denote as $\mathbf{B}$ and $\mathbf{A}$, respectively. Processes are assumed to operate linearly, i.e. when run at scale $s_i \geq 0$ they use amounts $s_i \mathbf{b}_i$ of inputs to produce $s_i \mathbf{a}_i$ outputs. In this setup, the production side (its dual will be discussed later on) of Von Neumann’s growth problem (also known as the ‘technological expansion problem’) is formulated in the following terms:

$$\max_{\rho > 0} (\rho) \quad , \quad \text{subject to } (\mathbf{A} - \rho \mathbf{B})\mathbf{s} \geq \mathbf{0} \quad (1)$$

In other words, one wants to find the largest $\rho > 0$ (which we shall denote as ρ^*) such that the system of inequalities

$$c^\mu \equiv \sum_{i=1}^N (a_i^\mu - \rho b_i^\mu) s_i \geq 0 \quad \forall \mu \quad (2)$$

admits a solution $\mathbf{s} = \{s_i\}$ with $s_i \geq 0$ for all i (with the trivial solution $s_i = 0$ being ruled out). The physical meaning of the above conditions becomes clear by noting that the M -vector $\mathbf{B}\mathbf{s}$ (resp. $\mathbf{A}\mathbf{s}$) encodes the total amount of each compound μ used as input (resp. returned as output) by the various technologies. Hence for fixed ρ , a vector $\mathbf{s}$ satisfying (2) guarantees that for each compound the total output is at least ρ times the total input. Maximizing ρ then amounts to finding the largest growth factor sustainable by the system. Note that the growth factor is defined uniformly over compounds, and one does not try to find an optimal vector of growth factors, one for each compound. In such a setting (which would make for an interesting generalization), the ρ^* defined above would characterize the compounds forming the production bottlenecks, namely those with the smallest feasible growth factor.

In a discrete-time dynamical setup [16], one can imagine a scenario in which an operation scale is chosen for every technology at each time step t . Then the input vector $\mathbf{I}(t) \equiv \mathbf{B}\mathbf{s}(t)$ generates an output vector $\mathbf{O}(t) \equiv \mathbf{A}\mathbf{s}(t)$, part of which is recycled into production (i.e. as the input at the next time step). The difference, namely

$$\mathbf{C}(t) = \mathbf{O}(t) - \mathbf{I}(t+1) \quad (3)$$

is available e.g. for consumption or for other uses outside the production system strictly defined, including as waste. This condition ensures to be focusing on self-sustaining production patterns. A minimal constraint to be imposed for stability is that $\mathbf{C}(t) \geq 0$, otherwise an external source of compounds would be needed to overcome the input shortages at time step $t+1$. In economics, a central issue at this point is that of maximizing the discounted utility of the overall production over the possible time evolutions. A class of results known as ‘turnpike theorems’ [17] show that optimal paths essentially coincide with the paths of maximal expansion described by Von Neumann’s model. Therefore one may restrict to trajectories $\{\mathbf{s}(t)\}$ ensuring that $\mathbf{I}(t+1) = \rho\mathbf{I}(t)$ with $\rho > 0$ (i.e. such that $I^\mu(t) = I^\mu(0)\rho^t$). It is easily seen that such paths correspond, for operation scales, to $\mathbf{s}(t+1) = \rho\mathbf{s}(t)$, implying

$$\mathbf{C}(t) = \rho^t(\mathbf{A} - \rho\mathbf{B})\mathbf{s}(0) \quad (4)$$

The stability requirement $\mathbf{C}(t) \geq 0$ for all t immediately leads to consider the problem (2) and, in turn, (1). It is now clear that if the solution ρ^* of (1) is larger than one the system is expanding (i.e. trajectories are stationary states of constant growth at a growth factor at least equal to ρ^*); if $\rho^* < 1$ the system is instead contracting.

Note that a linear combination $\lambda\mathbf{s} + \lambda'\mathbf{s}'$ of two solutions $\mathbf{s}$ and $\mathbf{s}'$ with non negative coefficients $\lambda > 0$ and $\lambda' > 0$ always satisfies the constraints, i.e. the solution space of (2) is convex. As ρ increases (clearly, all vectors $\mathbf{s}$ are solutions for $\rho = 0$) one expects its volume to shrink. As we shall see, a key question concerns the size of the solution space that survives at ρ^* .

A classical and easy to prove result concerns the existence of ρ^* [4]. Surely, if ρ is sufficiently small the system (2) is bound to admit a solution, i.e. it is possible to find a $\rho > 0$ for which $\mathbf{A}\mathbf{s} \geq \rho\mathbf{B}\mathbf{s}$. On the other hand, it is also possible to choose $\rho > 0$ so large that the sum of the row entries of the matrix $\mathbf{A} - \rho\mathbf{B}$ is negative for each row, which would imply the absence of a solution of (2). ρ^* then coincides with the least upper bound of the set of ρ ’s for which (2) admits a solution. Establishing the value of ρ^* is certainly the first task to be carried out. An equally important challenge is however that of profiling the way in which processes are employed in states of optimal growth. Because not all processes need to be operated, one expects that some of them will actually be employed, whereas others will be shut down being unnecessary. Finally we shall focus on the characterization of the set of compounds that will become available for consumption (if any).

3 Random input-output networks

The rationale for studying random input-output networks is simply that the macroscopic behaviour of a system of many interacting entities (technologies in Von Neumann's case) can be understood at the statistical level when the number of degrees of freedom diverges. In this limit, in fact, the key macroscopic observables are expected to become substantially independent of the specific realization of the couplings [18]. Large systems with random couplings are therefore expected to be good approximations of non-random ones. In this sense, they also constitute a fundamental benchmark against which real systems should be compared. Luckily, a host of analytical and conceptual tools are available to deal with the statistical mechanics of systems with random interactions of the type relevant for Von Neumann's problem [12,14].

3.1 Fully connected networks

The fully connected Von Neumann model is defined by [16]

$$a_i^\mu = \bar{a}(1 + \alpha_i^\mu) \quad (5)$$

$$b_i^\mu = \bar{b}(1 + \beta_i^\mu) \quad (6)$$

where $\bar{a}$ and $\bar{b}$ are constants (both > 0) whereas α_i^μ and β_i^μ are independent, identically distributed quenched Gaussian random variables of mean zero and variances $\overline{(\alpha_i^\mu)^2}$ and $\overline{(\beta_i^\mu)^2}$, respectively. In brief, processes in such a system employ every compound in a finite amount to produce every compound in a finite amount. The topology underlying this problem is hence that of a fully connected bipartite graph. Inserting the above definitions into (2) one arrives at the conditions

$$(\bar{a} - \rho\bar{b}) \sum_{i=1}^N s_i + \sum_{i=1}^N s_i (\bar{a}\alpha_i^\mu - \rho\bar{b}\beta_i^\mu) \geq 0 \quad \forall \mu \quad (7)$$

Because $s_i > 0$, for $N \gg 1$ the first term dominates over the second term, which scales like $\sqrt{N}$ (assuming no correlations between s_i 's and the Gaussian disorder). Therefore, to leading order in N ,

$$\rho^* = \bar{a}/\bar{b} \quad (8)$$

This tells us that the leading part of ρ^* is independent of the details of the input/output vectors. Rather, it is determined only by the average input and output coefficients. For sakes of simplicity, we shall henceforth set $\bar{a} = \bar{b} = 1$. To account for subleading corrections, it is convenient to set

$$\rho = 1 + \frac{g}{\sqrt{N}} \quad (9)$$

so that, for $N \gg 1$, $\rho \simeq e^{g/\sqrt{N}}$ and g can be interpreted as a growth rate. In turn, (2) yields the conditions

$$c^\mu \equiv \sum_{i=1}^N s_i \left[\alpha_i^\mu - \left(1 + \frac{g}{\sqrt{N}} \right) \beta_i^\mu - \frac{g}{\sqrt{N}} \right] \geq 0 \quad \forall \mu \quad (10)$$

One is interested in studying the largest g for which (10) has solutions. The calculation, detailed in [16] follows standard steps of spin-glass-like systems (specifically, it is similar to a Gardner calculus [19]) and we limit ourselves to a brief sketch here.

The solution space volume at fixed disorder and fixed g can be expressed as

$$V(g) = \text{Tr}_{\mathbf{s}} \prod_{\mu} \theta(c^{\mu}) \delta \left(\sum_{i=1}^N s_i - N \right) \quad (11)$$

where a hard constraint $\sum_{i=1}^N s_i = N$ has been included so as to remove the linear degeneracy inherent in Von Neumann's problem (this is equivalent to setting a scale for the s_i 's). The typical solution space volume $V_{typ}(g) \simeq \exp(Nv(g))$ can be obtained from

$$v(g) = \lim_{N \rightarrow \infty} \frac{1}{N} \overline{\log V(g)} \quad (12)$$

where the over-bar denotes an average over the quenched disorder (5,6). The evaluation of the disorder average can be carried out resorting to the replica trick. The relevant thermodynamic limit in this case is that where both N (number of processes) and M (number of compounds) diverge with a fixed ratio $n = N/M$. Here, after averages are computed, one can obtain an exact expression for $v(g)$ via a saddle-point integration and a replica-symmetric Ansatz. An interesting outcome of the above calculation is that ρ^* (or more precisely g^*) depends on the disorder statistics only via the parameter

$$k = \overline{(\alpha_i^{\mu} - \beta_i^{\mu})^2} \quad (13)$$

Specifically, one obtains that $g^*/\sqrt{k}$ is a function of n only. This implies that, generically, larger growth rates require larger spread in the input/output coefficients, suggesting that optimization (and thus process selection) can be operated by choosing maximally diversified technologies. It turns out that for $g \rightarrow g^*$ $v(g)$ can be expressed in terms of the average Euclidean distance between solutions of (2) (denoted as $\chi \equiv \chi(g)$). In particular, one finds

$$v(g) \simeq \chi^{\gamma N} \quad (14)$$

where γ is a positive exponent. Optimal growth (i.e. $g \rightarrow g^*$) can be retrieved in the limit $\chi \rightarrow 0$, in which case $v(g)$ reduces to a single point (i.e. there is a single configuration of operation scales that corresponds to an optimally growing system). The resulting (n, g) phase diagram is shown in Fig. 1. In brief, $g^* < 0$ for $n < 1$, implying that all viable configurations of operation scales are contracting. When $n > 1$, instead, expanding states become viable. The growth rate generically increases with n . However the rate of increase changes drastically between the phase where optimal states are contracting to that where expansion is possible. Indeed one gets

$$\frac{\delta \rho}{\delta n} \simeq - \frac{g}{n^{3/2} \sqrt{M}} \quad (15)$$

which suggests that an increase in N (i.e. in the number of available technologies) can benefit growth for sufficiently small n . By contrast, when n is large technological innovation appears to bear a much smaller effect on the growth properties. The saddle point conditions also allow to derive a simple algebraic equation valid at g^* that relates the fraction ψ_0 of inactive processes (i.e. processes such that $s_i = 0$) to the fraction ϕ_0 of compounds μ such that $c^{\mu} = 0$ (i.e. for which the overall output exactly matches the overall input):

$$\psi_0 = 1 - \phi_0/n \quad (16)$$

In different terms, (16) simply states that for $g = g^*$ (when a single vector of operation scales survives satisfying (2)) the number of variables $s_i \neq 0$ (namely $N(1 - \psi_0)$)




Fig. 1. Phase structure of the random fully connected Von Neumann model. Values of n and g in the shaded region are viable, i.e. there exist configurations of s_i 's satisfying (2). The continuous line is $g^*/\sqrt{kn}$. For $n < 1$ expanding states are unfeasible. Only for $n > 1$ can one obtain a feasible expanding state [16].

should equal the number of equalities (namely $M\phi_0$). Both ψ_0 and ϕ_0 increase with n in a roughly sigmoidal way, with $\psi_0(n=1) = \phi_0(n=1) = 1/2$. In words, optimal states in the contracting regime ($n \ll 1$) are characterized by having almost all processes operated and almost all compounds available for consumption. In optimal expanding states, instead ($n \gg 1$), a small fraction of processes is profitable while almost all compounds are perfectly recycled with limited production of consumable goods or waste.

3.2 Finitely connected networks

A mathematically more compelling scenario is obtained when the topology underlying the production networks is characterized by finite connectivity, i.e. when each process uses a finite number of compounds as substrates to produce a finite number of other (different) compounds [20]. This more realistic case can be conveniently dealt with by the cavity method under the assumption that the bipartite graph that underlies the production network has a locally tree-like structure. In essence, cavity theories exploit the fact that on such graphs the removal of a node of one type (either a compound or a reaction) makes the nodes of the other types (reactions or compounds) that were linked to the node we removed roughly statistically independent, since short loops causing feedbacks and correlations are assumed to be absent (or anyway negligible). So if we denote by $s_{i \in \mu}$ the scale variables of processes connected to a compound μ and by $\mathbf{c}^{\mu \in i}$ the variables defined in (2) for compounds connected to a process i , then their probability distributions are such that

$$p(s_{i \in \mu}) \neq \prod_{i \in \mu} p_i(s_i) \quad \text{and} \quad q(\mathbf{c}^{\mu \in i}) \neq \prod_{\mu \in i} q^\mu(c^\mu) \quad (17)$$

unless compound μ or, respectively, process i are removed, in which case one has

$$p^{(\mu)}(s_{i \in \mu}) = \prod_{i \in \mu} p_i^{(\mu)}(s_i) \quad \text{and} \quad q_{(i)}(\mathbf{c}^{\mu \in i}) = \prod_{\mu \in i} q_{(i)}^\mu(c^\mu) \quad (18)$$

where the subscripts (i) and (μ) indicate the removal of process node i and of compound μ , respectively. Note however that $q_{(i)}^\mu(c^\mu)$ can be written explicitly as

$$q_{(i)}^\mu(c^\mu) = \text{Tr}_{\mathbf{s}_{j \in \mu \setminus i}} \delta \left(c^\mu - \sum_{j \in \mu \setminus i} s_j (a_j^\mu - \rho b_j^\mu) \right) \prod_{j \in \mu \setminus i} p_j^{(\mu)}(s_j) \quad (19)$$

where $j \in \mu \setminus i$ indicates that the operations extend over all processes connected to μ except i . Similarly, introducing the fictitious energy function

$$E = - \sum_{\mu=1}^M \theta(c^\mu) \quad (20)$$

one has

$$p_i^{(\mu)}(s_i) \propto \text{Tr}_{\mathbf{c}^{\nu \in i \setminus \mu}} \exp \left[-\beta \sum_{\nu \in i \setminus \mu} \theta[-c_{(i)}^\nu - s_i(a_i^\nu - \rho b_i^\nu)] \right] \prod_{\nu \in i \setminus \mu} q_{(i)}^\nu(c^\nu) \quad (21)$$

(ultimately, the limit of large β must be considered). In turn, knowledge of the so-called cavity distributions (19) and (21) allows to reconstruct the entire probability distribution since

$$p_i(s_i) \propto \text{Tr}_{\mathbf{c}^{\nu \in i} q_{(i)}(\mathbf{c}^{\nu \in i})} \exp \left[-\beta \sum_{\nu \in i} \theta[-c_{(i)}^\nu - s_i(a_i^\nu - \rho b_i^\nu)] \right] \quad (22)$$

$$q^\mu(c^\mu) = \text{Tr}_{\mathbf{s}_{j \in \mu}} p^{(\mu)}(\mathbf{s}_{j \in \mu}) \delta \left(c^\mu - \sum_{j \in \mu} s_j (a_j^\mu - \rho b_j^\mu) \right) \quad (23)$$

Equations (19) and (21) can be solved self-consistently by population dynamics. The numerical procedure to extract ρ^* from this setup is non-trivial and we refer the reader to [20] for details. The lesson one learns from this analysis is that the picture obtained in the fully connected case is rather robust to dilution. Finite connectivity does however bear a strong quantitative impact. In the contracting regime ($\rho^* < 1$) optimal growth rates generically decrease with increasing dilution while the opposite happens in expanding regimes with $\rho^* > 1$. Moreover, optimal growth rates turn out to be consistently larger when processes have heterogeneous (as opposed to regular) connectivities, independently of the degree distribution of compounds. Hence, generically dilution (in expanding regimes) and stochastic connectivities allow to achieve larger growth rates with respect to the case of fully connected input/output systems. The phase structure shown in Fig. 1 is however qualitatively preserved.

3.3 Reversible networks

An important generalization considers processes as reversible, i.e. such that the roles of input and output coefficients can be interchanged [21]. This simple extra ingredient generates a surprisingly rich and considerably more complex scenario. Let us first notice that by allowing processes to revert Von Neumann's problem becomes intrinsically non-linear. Indeed it can be written compactly upon defining a matrix $\mathbf{R}$ with entries

$$R_i^\mu = \begin{cases} a_i^\mu - \rho b_i^\mu & \text{for } s_i > 0 \\ b_i^\mu - \rho a_i^\mu & \text{for } s_i < 0 \end{cases} \quad (24)$$

whereby one sees that $\mathbf{R} \equiv \mathbf{R}(\rho, \mathbf{s})$. (The fact that process can be reverted is reflected also in the possibility that s_i becomes negative.) Under this condition, (1) now reads

$$\max_{\rho > 0} \rho \quad \text{subject to } \mathbf{R}(\rho, \mathbf{s})\mathbf{s} \geq \mathbf{0} \quad (25)$$

There are two basic properties of the reversible Von Neumann problem that are both easy to prove and rich of consequences. The first is that in general reversing all processes is not equivalent to time reversal, i.e. the naive expectation that if a scale vector $\mathbf{s}$ guarantees a growth factor ρ then the reversed vector $\mathbf{s}'$ should guarantee a growth factor $\rho' = 1/\rho$ turns out to be false. To see this, define

$$c_d^\mu(\mathbf{s}) = \sum_{i=1}^N s_i (a_i^\mu - \rho b_i^\mu) \quad (26)$$

$$c_r^\mu(\mathbf{s}) = \sum_{i=1}^N s_i (b_i^\mu - \rho' a_i^\mu) \quad (27)$$

and note that

$$c_d^\mu(\mathbf{s}) + \rho c_r^\mu(\mathbf{s}) = (1 - \rho\rho') \sum_{i=1}^N s_i a_i^\mu \quad (28)$$

Clearly, $c_d^\mu(\mathbf{s}) = c_r^\mu(\mathbf{s}) = 0$ implies $\rho' = 1/\rho$ but if $c_r^\mu(\mathbf{s}) \geq 0$ then

$$\rho\rho' \leq 1 - \frac{c_d^\mu(\mathbf{s})}{\sum_{i=1}^N s_i a_i^\mu} \quad (29)$$

which tells us that $c_d^\mu(\mathbf{s}) = 0$ for all μ is a necessary but not sufficient condition for ρ' to be equal to $1/\rho$. In general,

$$\rho\rho' \leq \min_{\mu} \left(1 - \frac{c_d^\mu(\mathbf{s})}{\sum_{i=1}^N s_i a_i^\mu} \right) \quad (30)$$

The second notable fact is that the solution space ceases to be convex, in particular for $\rho < 1$. Re-writing (2) in this case as

$$c^\mu(\mathbf{s}) = \sum_{i=1}^N s_i [a_i^\mu(\theta(s_i) + \rho\theta(-s_i)) - b_i^\mu(\theta(-s_i) + \rho\theta(s_i))] \geq 0 \quad \forall \mu \quad (31)$$

one easily sees that

$$c^\mu(\lambda\mathbf{s} + \lambda'\mathbf{s}') = \lambda c^\mu(\mathbf{s}) + \lambda' c^\mu(\mathbf{s}') + A^\mu(\mathbf{s}, \mathbf{s}') \quad (32)$$

with

$$A^\mu(\mathbf{s}, \mathbf{s}') = (1 - \rho) \sum_{i=1}^N (a_i^\mu + b_i^\mu) [\lambda s_i \theta(-s_i) + \lambda' s'_i \theta(-s'_i) - (\lambda s_i + \lambda' s'_i) \theta(-\lambda s_i - \lambda' s'_i)] \quad (33)$$

If $\mathbf{s}$ and $\mathbf{s}'$ are both solutions, then a linear combination will again be a solution if $A^\mu \geq 0$. $s_i s'_i > 0$ indeed implies $A^\mu = 0$ for each ρ . However for $s_i s'_i < 0$ one finds that $A^\mu \geq 0$ only for $\rho \geq 1$. In other words convexity is guaranteed only for $\rho \geq 1$ and may not occur even at ρ^* when $\rho^* < 1$. Note however that the solution space




Fig. 2. Phase structure of the random fully connected reversible Von Neumann model. The continuous (resp. dashed and dot-dashed) line is $g^*/\sqrt{kn(1+\phi)}$ for $\phi = 0$ (resp. $\phi = 1/2$ and $\phi = 1$). Markers correspond to optimal growth rates estimated numerically as explained in [21]. (For clarity, the shaded region is only shown for $\phi = 0$.) For $n(1+\phi) < 1$ expanding states are unfeasible. However one sees slight disagreements between the predicted (replica-symmetric) value of g^* and computed one.

is obviously convex for $\rho = 0$, showing that $\rho < 1$ is a necessary but not sufficient condition for the solution space to be non-convex (e.g. fragmented in disjoint parts). It is very important to remark that non-convexity poses a serious problem for the design of efficient algorithms to sample the solution space in the contracting regime [21].

The statistical mechanics analysis of this case in a fully connected topology can proceed along steps similar to those performed for the random fully connected network, with (11) replaced by

$$V(g) = \text{Tr}_{\mathbf{s}} \prod_{\mu} \theta(c^{\mu}) \delta \left(\sum_{i=1}^N |s_i| - N \right) \quad (34)$$

If ϕ denotes the fraction of reversible processes (i.e. the effective number of processes that can be operated in the system is $N+N\phi$), one finds an expanding phase with $g^* > 0$ for $n > n_c \equiv 1/(1+\phi)$ (see Figure 2). This suggests, as expected, that being able to run processes in both directions extends the range of the regime where expanding solutions exist. The price to be paid is that optimal growth rates in the expanding regime are comparatively smaller when processes are reversible. For $n(1+\phi) < 1$, instead, the system is confined to a contracting regime. Here, larger ϕ 's imply larger growth rates, meaning that increasing the availability of reversible processes benefits the system's growth properties. As seen above, however, the contracting phase is characterized by lack of convexity. This is hinted by the fact that the computed and predicted (replica-symmetric) values of g^* disagree (albeit just slightly) for $n < n_c$ when $\phi \neq 0$. To see this more explicitly one can measure directly the probability that the uniform linear combination of two solutions solves Von Neumann's problem (Figure 3). Inspecting results, one sees immediately that the solution space displays marks of convexity for sufficiently small $g < g^*$ (recall that the solution space is certainly convex for $\rho = 0$). In the intermediate range of values of g up to g^* , instead,




Fig. 3. Probability that the linear combination of two solutions with coefficients $1/2$ solves the fully reversible ($\phi = 1$) random fully connected von Neumann problem for $n = 1/2$. The solution space appears to lose convex values of $g < g^* = 0$. Different curves correspond to different values of N increasing in the direction of the arrow from 50 to 500.

convexity appears to break down. The replica-symmetric theory one develops for this case is thus expected to fail in a range of values of $g < 0$ (as seen in Fig 2).

As before, an increase in the optimal growth rate causes a decrease in the number of compounds with $c^\mu > 0$. This corresponds to a ‘waste’ reduction and, on the other hand, to an impoverishment in the global output profile. Quite strikingly, the reversible part of the process set displays a neat second-order transition as n exceeds n_c : all reversible processes are employed in the contracting phase, where a finite fraction of them is inactive when $g^* > 0$ [21].

4 Towards metabolic networks

Von Neumann’s problem finds a natural ground of application in chemical reaction networks, in which the input-output relations linking technologies to goods in production systems are replaced by the wiring between reactions and chemical species governed by stoichiometry. Perhaps most importantly, in this context living cells present us the opportunity to test Von Neumann’s ideas directly in a real system, namely their energy metabolism.

In cells, an enormous number of genetic, regulatory and biochemical mechanisms interact on different spatial and temporal scales to accomplish a wide range of tasks, from growth and reproduction to motility, homeostasis and exterior sensing. The rapid development of high-throughput genome sequencing techniques together with the availability of refined gene annotations has allowed, over the past few years, to set up consistent, controlled protocols for reconstructing the cellular metabolic networks of different organisms to an unprecedented degree of detail [9,22]. Such networks underlie the basic energy harvesting processes in cells, by which the energy derived from nutrients is transduced into usable forms of mechanical or chemical energy. The output of metabolism is constituted essentially by the so-called building blocks for functionally relevant biological macromolecules (amino-acids, fatty acids, nucleotides) as well as by waste products like e.g. carbon dioxide and acetate (mostly for bacteria). Building blocks are then used in a series of processes downstream of metabolism to form functional proteins (including reaction-catalyzing enzymes), membrane structures and nucleic acids. To a large degree, the physiological outcome of this can be read off

from the cell's proteome, i.e. from the composition of the repertoire of proteins that the organism was able to produce, which in turn effects the regulation of metabolism.

It is immediately clear that, once the structure of the network (encoded in the reactions' stoichiometry) is known, Von Neumann's stability criterion can serve as a minimal model for the organization of metabolism. However, before discussing this application more in detail (Sec. 5), it is worth trying to bring the theory a few steps forward (in a direction relevant to biology) by addressing a couple of peculiar aspects of biochemical reaction networks that, on one hand, have a direct influence on the solutions of Von Neumann's problem and, on the other, generalize it. These are, respectively, (a) the emergence of the so-called 'conserved pools of reagents', which translate into the need to consider constraints on the quenched disorder (the input-output coefficients), and (b) the highly important issue of thermodynamic stability.

4.1 Stoichiometric quenched disorder

In first place, addressing optimality in reaction networks requires a more careful evaluation of the role of stoichiometry, since real networks are of course not random: stoichiometric coefficients enforce chemical balance at each reaction node in the network. This set of local relations has been shown to give rise, at a larger scale, to conservation laws for the aggregate concentration of pools of reagents [23]. Such laws are expressed by the relation

$$\sum_{\mu \in P} (a_i^\mu - b_i^\mu) = 0 \quad \forall i \quad (35)$$

where $P \subseteq \{1, \dots, M\}$ is the set containing the chemical species that form the conserved pool. In concrete, the total number of molecules in a pool is conserved over time, so that the aggregate concentration does not change in time (while the concentrations of individual metabolites can change relative to each other). These dynamical invariants have a purely topological origin in the structure of stoichiometric coefficients. As a byproduct, one easily understands that metabolites belonging to a conserved pool cannot be global network outputs (a detailed theory of this fact with applications is developed in [24,25]). In other terms, in presence of a conserved metabolite pool one expects that $\rho^* \leq 1$.

A first step in driving Von Neumann's model towards real cellular networks consists therefore in studying how the existence of conserved metabolite pools affects the picture derived in the case of random networks. The simplest way to implement this ingredient is to impose a constraint on the quenched disorder. Specifically, one can introduce a 'random' conserved pool formed by a finite fraction ϕ of reagents by explicitly adding to the volume defined in (11) a hard constraint of the form

$$\sum_{\mu=1}^M z^\mu (a_i^\mu - b_i^\mu) = 0 \quad (36)$$

(where z^μ is a quenched random variable equal to one with probability ϵ and zero otherwise). In order to check explicitly that the constraint (36) implies $\rho^* \leq 1$, i.e. that at best optimal growth states are characterized by constant operation scales for reactions (fluxes), it suffices to multiply each term in (2) by z^μ and sum over μ (we ignore reversibility). One obtains:

$$(1 - \rho) \sum_{i=1}^N s_i \sum_{\mu=1}^M z^\mu b_i^\mu + \sum_{i=1}^N s_i \sum_{\mu=1}^M z^\mu (a_i^\mu - b_i^\mu) \geq 0 \quad (37)$$




Fig. 4. Phase structure of the random fully connected constrained Von Neumann model. The continuous line marks the separation between the regions with $g^* = 0$ (stationary regime) and $g^* < 0$ (contracting regime). The expanding regime lies on the $\epsilon = 0$ axis ($n > 1$) and is represented by a green strip [26].

If the vector $\mathbf{z}$ identifies a conserved pool then the second term vanishes. The first term allows for $\rho > 1$ only if all reactions connected to metabolites in the conserved pool are inactive (i.e. if the pool is effectively removed from the network). In particular this leads to the trivial solution $s_i = 0$ for each i in a fully connected network. Hence necessarily $\rho^* \leq 1$. So the states of expansion recovered in the previous sections are a ‘pathology’ of random systems. When stoichiometric quenched disorder is used they cannot be present. Notice that this does not mean that the network cannot have a positive net production of compounds. The introduction of (36) slightly modifies the disorder averaging but the replica theory developed for random input/output systems can be conveniently extended to this constrained case [26]. The (ϵ, n) phase structure one obtains, in agreement with the previous observations, is shown in Figure 4.

A noteworthy consequence of the fact that ρ^* cannot exceed 1 is that for large enough n , when an expanding phase would be expected in absence of the conserved pool, multiple optima survive, i.e. the average distance between solutions does not vanish as $g \rightarrow g^*$. So stoichiometric disorder appears to increase robustness in reaction networks, since many (microscopic) arrangements of reaction operation scales turn out to be compatible with optimal growth. It is interesting to note that real cellular metabolic networks typically have values of n larger than one. This suggest that the ability to respond to perturbations (e.g. by re-arranging scales) without losing the desirable property of maximum productive capacity may be a key ingredient for biological stability and resilience.

4.2 Von Neumann’s dual problem and thermodynamics

Thus far, we have limited ourselves to considering Von Neumann’s problem in a production perspective. Our focus has been set on finding ways to operate the available processes that ensure self-sustainability and optimality. In his original work however J. Von Neumann has given equal attention to the dual problem (in the sense of linear programming duality) [2]. In elementary economic terms, it can be introduced upon defining an M -vector $\mathbf{p}$ of prices ($p^\mu \geq 0$ for each μ). The cost associated to running

a certain process i at unit scale is simply $\sum_{\mu=1}^M p^\mu b_i^\mu$. This generates a corresponding revenue given by $\sum_{\mu=1}^M p^\mu a_i^\mu$. The problem dual to (1) consists in solving

$$\min_{\sigma > 0} \sigma \quad \text{subject to } \mathbf{p}(\mathbf{A} - \sigma \mathbf{B}) \leq \mathbf{0} \quad (38)$$

(compare this with (1)). The interpretation of the parameter σ is not straightforward. The simplest is in terms of an interest factor. Imagining that the operation of processes is financed by borrowing cash at a rate r , at the end of each period $\sigma = 1 + r$ units of cash must be returned to the financiers. In this setting, the conditions

$$\sum_{\mu=1}^M p^\mu (a_i^\mu - \sigma b_i^\mu) \leq 0 \quad \forall i \quad (39)$$

simply mean that the net extra profit generated by each process can at most be zero. (Zero profit is a standard scenario in models of economic equilibrium, see e.g. [27] or, in a statistical physics perspective, [28].)

It is simple to show that, in the above settings of (1) and (38) and under broad assumptions, $\rho^* \geq \sigma^*$ (with σ^* the solution of (38)) [4]. Furthermore, restricting slightly the definition of the model it can be proven that indeed $\rho^* = \sigma^*$ (i.e. that the growth factor equals the interest factor, as would perhaps be intuitively expectable).

It is however tempting to search for a physical interpretation of Von Neumann's dual problem. The simplest connection resides in the thermodynamic scenario underlying biochemical activity. It is well known that in non equilibrium steady states with non-zero fluxes, reactions must proceed in directions of decreasing free energies (which in turn implies that thermodynamically feasible configurations of reaction fluxes may not contain directed cycles) [29]. The condition for thermodynamic feasibility reads

$$v_i \Delta G_i \leq 0 \quad \forall i \quad (40)$$

where ΔG_i is the (Gibbs) free energy change associated to reaction i and v_i is the net flux (the difference between the forward and reverse rates) of reaction i . ΔG_i can be written as the sum of the chemical potentials of the compounds weighted by their respective stoichiometric coefficients. It then follows that thermodynamic feasibility is related to the existence of a chemical potential vector $\mathbf{g}$ such that

$$v_i \sum_{\mu=1}^M g^\mu (a_i^\mu - b_i^\mu) \leq 0 \quad \forall i \quad (41)$$

Connection to (39) is now straightforward. In this respect, prices play the role of chemical potentials and (39) have the form of a thermodynamic feasibility constraint.

It would be interesting to construct a more stringent physical analogy for (38). Quite recently an attempt in this direction has been performed in conjunction with Flux-Balance-Analysis (FBA) [9,30]. It was shown in specific that the dual to the linear programming problem that arises in FBA can be re-formulated as a free energy consumption minimization problem conditioned on a given free energy drain (representing, in the case of cellular systems, growth) [31].

5 Application to cellular metabolism

Modeling a cell's metabolic activity is a central problem in systems biology. In principle, knowledge of the network structure and of basic thermodynamic parameters

like reaction constants allows to formulate non-linear differential equations for the enzyme-mediated joint evolution of reactant concentrations and reaction rates (or fluxes) that fully account for the kinetics of a biochemical reaction network [29]. Unfortunately at genome scale this route is most often barred by uncertainties about kinetic parameters and reaction or transport mechanisms. Constraint-based models of metabolism (like Flux-Balance Analysis [9,29,30,32]) provide a feasible alternative that has marked a considerable breakthrough in the quest to reconstruct and simulate global cellular functions. In a nutshell, they rely on steady state assumptions to impose sets of physically motivated linear constraints (enforcing mass balance) on the network nodes corresponding to reagents, and retrieve the viable flux states as the configurations of reaction operation scales that ensure that all constraints are satisfied. Regulatory and thermodynamic restrictions are included either in the allowed ranges of variability of fluxes or in the a priori reversibility assignments of reactions. Such a constrained linear system suffices to define a space of feasible flux configurations (whose characterization is in itself a challenging problem [33,34]). Quite crucially, physiologically relevant states are usually assumed to maximize *ad hoc* objective functions that provide a further functional constraint on the cell’s metabolic activity (see e.g. [35,36,37,38] for some examples).

Applying Von Neumann’s idea to cellular metabolism amounts to taking a rather different (though not unrelated) viewpoint. Recalling that in real stoichiometric systems $\rho^* = 1$, finding optimal solutions in the Von Neumann sense means finding flux vectors $\mathbf{s}$ such that

$$c^\mu \equiv \sum_{i=1}^N (a_i^\mu - b_i^\mu) s_i \geq 0 \quad \forall \mu \quad (42)$$

where the index i runs over reactions, μ over chemical species and a_i^μ and b_i^μ denote the output and input (respectively) stoichiometric coefficients of metabolite μ in reaction i . As explained above, (42) encodes in essence a stability constraint. If fluxes that transport matter (e.g. nutrients) into the cell are included in the stoichiometric matrices to make the system self-sustainable (in such a way that the nutrient consumption equals its influx), then solving (42) allows to retrieve information about what the cell is in principle capable of net-producing *in a given environment*, as well as on the corresponding flux configurations. Metabolites for which $c^\mu = 0$ indeed turn out to be mass-balanced, so that their production and consumption fluxes exactly match. On the contrary, if $c^\mu > 0$ then metabolite μ is globally producible, meaning that it can become available either for macromolecular processes or as waste. The availability of a neural network learning-based algorithm for the efficient generation of solutions of (42) (see [20], based on [39]) then permits a detailed statistical analysis of the cell’s metabolic capabilities.

Work performed along these lines for the bacterium *Escherichia coli* [40,41] has shown that environment selection restricts the feasible states according to (42) to flux configurations that both reproduce well the available empirical evidence on E.coli’s metabolic fluxes and guarantee that the correct physiological task (in that case, growth or biomass production) is carried out by the cell without the need of additional *ad hoc* functional constraints. We want to add here further evidence supporting the relevance of Von Neumann’s observations in the context of cell metabolism.

We have computed and analyzed optimal flux configurations (in Von Neumann’s sense) for the *Escherichia coli* metabolic network reconstructed in [43], preparing it in a minimal growth medium with glucose as its main carbon source. At odds with the work presented in [40,41], we have modified the model to be able to measure the levels of some of the physiological objective functions used in the computational biology literature (see e.g. [38]). Our goal is to characterize the solution space defined




Fig. 5. Von Neumann-optimal states of the *Escherichia coli* metabolism. (Left) Distribution of the flux overlap (black) and production profile overlaps (red) over distinct solution pairs obtained from a set of 250 solutions of (42). (Right) Flux distributions obtained from (42). Black markers denotes the distribution of average fluxes, the orange curve corresponds to the power-law fit with exponent -1 and red markers correspond to a convenient binning of the data. Inset: flux distributions in various solutions (different colors). Flux units are arbitrary.

by (42) from a view point of how disperse solutions are in terms of flux arrangements and biological functionality.

The similarity between two solutions α and β with flux vectors $\mathbf{s}_\alpha$ and $\mathbf{s}_\beta$ respectively is conveniently quantified by their overlap. For each reaction i , we define

$$q_{\alpha\beta}(i) = \begin{cases} 1 & \text{for } s_{i\alpha} < \epsilon \text{ and } s_{i\beta} < \epsilon \\ 1 - s_{i\alpha} & \text{if } \epsilon < s_{i\alpha} < 1 \text{ and } s_{i\beta} < \epsilon \\ 0 & \text{if } s_{i\alpha} \geq 1 \text{ and } s_{i\beta} < \epsilon \\ \frac{2s_{i\alpha}s_{i\beta}}{s_{i\alpha}^2 + s_{i\beta}^2} & \text{otherwise} \end{cases} \quad (43)$$

where ϵ is a cut-off value below which we treat fluxes as zero (results do not depend on the choice of the cut-off as long as ϵ is 10^{-5} or less). This parameter measures in essence the variation of flux i in solutions α and β . For a complete flux configuration, we define

$$q_{\alpha\beta} = \frac{1}{N} \sum_i q_{\alpha\beta}(i) \quad (44)$$

as a characterization of the “distance” between solutions α and β : the closer q is to 1 (resp. 0) the more similar (resp. dissimilar) the two configurations are. Likewise, we define the overlap between the vectors $\mathbf{c}$ that correspond to the production profiles in each solution of (42). The distributions of these overlaps are displayed in Figure 5. One sees that the distance between solutions has a broader distribution than that between production profiles, indicating that a degree of flexibility in the flux organization is allowed that nevertheless ensure a remarkable robustness in the emerging production profile. The finding of [41] that indeed the most probable output of the network is formed by the biomass constituents [44] is hereby fully confirmed (data now shown), as is the fact that the distribution of fluxes (see Figure 5) retains the empirically observed [45] power-law shape with exponent -1 . Single solutions however display a certain degree of variability, as can be seen again in Figure 5.

Let us now analyze how this scenario reflects in the values of some of the recurrent physiological objective functions used in the literature (see Figure 6). We




Fig. 6. Von Neumann-optimal states of the *Escherichia coli* metabolism. Probability distributions of the values of commonly employed physiological objective functions obtained from the solutions of (42): (a) bare ATP yield; (b) bare NADH consumption; (c) total net flux; (d) bare glucose consumption; (e) bare NADH consumption per unit of GLC consumed; (f) bare ATP yield per flux unit; (g) bare ATP yield per unit of GLC consumed. Plots (1) and (2) display the flux pertaining to a biomass production reaction computed in a network reconstruction where such a reaction was explicitly included. This flux represents an addition over the biomass that the network spontaneously yields: (1) bare extra biomass flux; (2) bare extra biomass flux per flux unit. Flux units are arbitrary.

have monitored in particular: the total production of ATP (adenosine-triphosphate, the key molecular energy carrier in cells), the total consumption of NADH (nicotinamide adenine dinucleotide, the key reducing agent in electron transfer processes in metabolism, with several other important functional roles at different levels in cells), the total net flux (which is assumed to be a good proxy for enzymatic efficiency) and the glucose (GLC) uptake (measuring nutrient consumption). ATP production, NADH consumption and total net flux display a roughly unimodal shape, pointing to the fact that under Von Neumann's hypothesis the network is capable of tuning these functions within a relatively small range of values. Notice that the total flux is usually assumed to be minimized by cells, to account for optimal use of the available pool of enzymes. On the other hand, it is reasonable to think that cells under certain conditions maximize the ATP yield for optimal energetic efficiency. A similar picture holds for the ATP yield per flux unit, which is usually assumed to be maximized so as to guarantee optimal energetic efficiency at minimum enzyme usage. NADH consumption is instead thought to be minimized, to reduce the activity of redox processes. GLC consumptions displays instead a somewhat broader profile. This is consistent with previous results [41] and points to the fact that the emerging physiologic scenario (which appears to be robust from the energetic and enzymatic efficiency viewpoints) is




Fig. 7. Von Neumann-optimal states of the *Escherichia coli* metabolism. ATP yield (inset) and ATP yield per flux unit (main) versus extra biomass yield in states where an extra biomass flux is forced in (42).

indeed recoverable under widely varying GLC consumptions. In other words, glucose availability can fluctuate within a broad range without affecting optimality in the Von Neumann sense. This however does have an effect on other physiological observables like the ATP production or the consumption of NADH per unit of GLC consumed. These functions turn out to be broadly distributed as well, although with well defined peaks at low values. Finally, we have considered the effect of adding explicitly the biomass production reaction to the pool of processes encoded in the network reconstruction. Recalling that the system globally outputs biomass constituents even in absence of this additional reaction, Figure 6 shows that the extra biomass flux due to the inclusion of the additional process is essentially null. In other words, the unforced biomass output of (42) by itself fully describes the cell's production capacity in the environment we selected. In summary, the physiological scenario emerging from Von Neumann's model applied to *Escherichia coli*'s metabolism is one in which the variability of optimal flux configurations underlies considerable robustness at the level of production profiles as well as at the level of key metabolic functions, the main exception being glucose consumption, which instead is allowed to fluctuate over a large relative range. The latter aspect contributes to robustness as well.

To conclude, we have investigated the problem of if and how forcing an extra biomass flux from the cell (i.e. forcing a non-zero value for the flux of the biomass reaction) changes the physiological scenario with respect to the reference case in which the biomass output is obtained self-consistently from the solutions of (42). Figure 7 shows how the bare ATP yield and the ATP yield per flux unit (two elementary parameters to quantify the physiological state of the cell) change upon increasing the extra biomass flux. In first place, as should be expected, the bare ATP yield slowly increases initially as a non-zero additional biomass flux is forced, but its value tends to saturate as the latter increases, indicating that the system has reached a maximal ATP output capacity. The relative gain in ATP yield with respect to the unforced reference state is close to 30%. At the same time, however, the ATP yield

per flux unit, which as said above quantifies energetic efficiency against enzyme usage, is actually reduced as the biomass flux increases, meaning that the cell can achieve larger biomass outputs and a gain in bare ATP yield only at the cost of a dramatic decrease in the overall metabolic efficiency (the relative loss being close to 50%). Quite significantly, then, optimal states from the physiological viewpoint appear to be those defined by (42) *without* any additional constraint. Similarly (data not shown) NADH consumption increases upon increasing the forced biomass flux while, by contrast, GLC consumption roughly stays constant until the extra biomass flux reaches the value 0.5, when it starts increasing dramatically.

These results suggest that Von Neumann’s optimal states are on one hand able to provide a description of the metabolic activity of bacterial cells that is in good agreement with the empirical knowledge; on the other, they frame key physiological observables in contained ranges (implying robustness) so that forcing an improvement in the growth capacity leads to a decrease in metabolic efficiency, thereby pointing to a direct biological counterpart of Von Neumann’s economic optimality. Clearly however, further work is needed in order to clarify the extent to which Von Neumann’s model might be applicable in the context of metabolic networks, and above all for more complex cell types (like yeast or human cells), where regulation and kinetics may play more important roles and steady state models might turn out to be insufficient to identify the emerging physiological characteristics with a good degree of confidence.

6 Conclusions

Von Neumann’s linear growth model addresses the very general problem of characterizing a self-sustained network of interacting input-output processes in terms of (a) the maximum achievable growth rate, (b) the pattern of activity of processes and (c) the emergent net production profile. Despite its remarkably simple rationale, this setup makes for a challenging and easily generalizable statistical mechanics problem that gives rise to a rich phenomenology which can be studied both analytically (for ensembles of random networks) and numerically (for single instances). The theory described here covers in our view only the simpler and more straightforward modifications of the original model, all guided by a chemical rather than economical intuition. From a strictly theoretical viewpoint (but possibly also with direct relevance for many applications) it would be interesting to bring them farther beyond Von Neumann’s original definition, e.g. by relaxing or abandoning the assumption of constant returns to scale. Perhaps the most notable application of these ideas outside of economics (where they were originally conceived and where they form a basis for the theory of economic growth) can be located in cellular metabolism. Applying Von Neumann’s growth model to cell metabolism, in our view, corresponds to appraising to what extent and in what manner do environmental, stoichiometric and thermodynamic factors limit and determine the output of a cell’s biochemical machinery within a minimal physical assumption of stability. The quantitative analysis of intracellular chemical reaction networks performed along these lines has provided a physiologically sensible characterization of the metabolic capabilities of the bacterium *Escherichia coli* (as well as of other, simpler cell types [42]) in agreement with empirical evidence. This suggests that the function of the biochemical core of the energetics of cellular systems mirrors at least to some extent the basic notions of stability and optimality that are encoded in Von Neumann’s model and possibly underlie optimal economic growth. Understanding how this picture is affected by the cross-talk between metabolism and regulation, which could be the major source of non-linearity, is perhaps the most obvious next step.

Acknowledgments. We are indebted with M. Figliuzzi, D. Granata, F.A. Masucci and V. Van Kerrebroeck and, above all, M. Marsili, C. Martelli, R. Monasson, I. Perez Castillo for long-standing, pleasant and fruitful collaborations on the issues discussed here. Our views on these matters have been crucially shaped by the insight we derived from discussion with many colleagues, including M. Avalle, W. Bialek, A.C.C. Coolen, D. De Martino, F. Di Patti, M. Di Nuzzo, S. Franz, F. Giove, T. Hwa, O. Martin, A. Pagnani, G. Parisi, F. Ricci Tersenghi, A. Seganti, D. Segrè, M.A. Virasoro, M. Weigt and R. Zecchina, whom we gratefully acknowledge. This work is supported by a Seed Project (DREAM) of the Italian Institute of Technology (IIT) and by the joint IIT/Sapienza Lab “Nanomedicine”. The IIT Platform “Computation” is gratefully acknowledged.

References

1. Von Neumann, J. (1937) *Ergebnisse eines Mathematisches Kolloquium* **8** 73–83
2. Von Neumann, J. (1945) *Review of Economic Studies* **13** 1–9
3. Brody, A. (1989) Economics and thermodynamics. In *John Von Neumann and modern economics* (Dore, M. and Chakravarty, S. and Goodwin, R., eds.), Clarendon (Oxford)
4. Gale, D. (1960) *The theory of linear economic models* (The University of Chicago Press)
5. Dore, M. and Chakravarty, S. and Goodwin, R. (eds.) (1989) *John Von Neumann and modern economics* (Clarendon, Oxford)
6. Cohen J.E. and F. Briand and C. M. Newman (1990) *Community Food Webs: Data and Theory* (Springer, Berlin)
7. Helbing D. and S. Lämmer and T. Seidel and P. Seba and T. Platkowski (2004) *Physical Review E* **70** 066116
8. Rodriguez-Iturbe I. and A. Rinaldo (1996) *Fractal River Basins: Chance and Self-Organization* (Cambridge University Press, Cambridge)
9. Palsson B. (2006) *Systems Biology: Properties of Reconstructed Networks* (Cambridge University Press)
10. Heinrich R. and S. Schuster (1996) *The Regulation Of Cellular Systems* (Springer)
11. Fell D. (1996) *Understanding the Control of Metabolism* (Portland Press)
12. Dotsenko V. (1995) *An Introduction to the Theory of Spin Glasses and Neural Networks* (World Scientific Publishing Company)
13. Korutcheva E. and M Oppen and B Lopez (1994) *Journal of Physics A: Mathematical and General* **27** L645
14. Nishimori H. (2004) *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Oxford University Press)
15. De Martino A. and Marsili, M and Perez Castillo, I (2004) *Journal of Statistical Mechanics: Theory and Experiment* P04002
16. De Martino A. and M Marsili (2005) *Journal of Statistical Mechanics: Theory and Experiment* L09003
17. McKenzie L. W. (1986) Optimal economic growth, Turnpike theorems and comparative dynamics. In *Handbook of Mathematical Economics (vol 3)* (K J Arrow and M D Intriligator, eds.), North-Holland, Amsterdam
18. Wigner E.P. (1958) *Annals of Mathematics* **67** 325
19. Gardner E. (1988) *Journal of Physics A: Mathematical and General* **21** 257–270
20. De Martino A. and Martelli, C and Monasson, R and Perez Castillo, I (2007) *Journal of Statistical Mechanics: Theory and Experiment* P05012
21. De Martino A. and Figliuzzi, M and Marsili, M (2010) *Journal of Statistical Mechanics: Theory and Experiment* P07032
22. Thiele I and Palsson, B (2010) *Nature Protocols* **5** 93–121
23. Famili I and Palsson, B (2003) *Biophysical Journal* **85** 16–26
24. Imielinski M and Belta, C and Halasz, A and Rubin, H (2005) *Bioinformatics* **21** 2008–2016

25. Imielinski M and C. Belta and H. Rubin and A. Halasz (2006) *Biophysical Journal* **90** 2659
26. De Martino A. and Martelli, C and Massucci, FA (2009) *Europhysics Letters* **85** 38007
27. Debreu G. (1972) *Theory of Value: An Axiomatic Analysis of Economic Equilibrium* (Yale University Press)
28. De Martino, A and Marsili, M and Perez Castillo, I (2007) *Macroeconomic Dynamics* 1–28
29. Beard D and Qian, H (2008) *Chemical Biophysics: Quantitative Analysis of Cellular Systems* (Cambridge University Press)
30. Orth J and I Thiele and B Palsson (2010) *Nature Biotechnology* **28** 245–248
31. Warren PB and Jones, JL (2007) *Physical Review Letters* **99** 108101
32. Kauffman K and Prakash P and Edwards J (2003) *Current Opinion in Biotechnology* **14** 491–496
33. Braunstein A and Mulet, R and Pagnani, A (2008) *BMC Bioinformatics* **9** 240
34. Bianconi G (2008) *Physical Review E* **78** 035101(R)
35. Segrè D and Vitkup D and Church G (2002) *Proceedings of the National Academy of Sciences of the USA* **99** 15112–15117
36. Holzhütter HG (2004) *European Journal of Biochemistry* **271** 2905–2922
37. Shlomi T and Berkman O and Rupp E (2005) *Proceedings of the National Academy of Sciences of the USA* **102** 7695–7700
38. Schuetz R and Kuepfer L and Sauer U (2007) *Molecular Systems Biology* **3** 119
39. Krauth W and Mezard M (1987) *Journal of Physics A: Mathematical and General* **20** L745
40. Martelli C and De Martino A and Marinari E and Marsili M and Perez Castillo I (2009) *Proceedings of the National Academy of Sciences of the USA* **106** 2607–2611
41. De Martino A and Marinari E (2010) *Journal of Physics Conference Series* **233** 012019
42. De Martino A and D. Granata and E. Marinari and C. Martelli and V. Van Kerrebroeck (2010) *Journal of Biomedicine and Biotechnology* **2010** 415148
43. Reed JL and Vo, TD and Schilling, CH and Palsson, BO (2003) *Genome biology* **4** R54
44. Feist A and Palsson, Bernhard (2010) *Current Opinion in Biotechnology* **13** 344–349
45. Almaas E and Kovacs, B and Vicsek, T and Oltvai, ZN and Barabasi, AL (2004) *Nature* **427** 839–843