

Phase transition in random catalytic networks

Rudolf Hanel,¹ Stuart A. Kauffman,^{2,3} and Stefan Thurner^{4,*}

¹*Institute of Physics, University of Antwerp, Groenenborgerlaan 171, 2020 Antwerp, Belgium*

²*Biocomplexity and Informatics, The University of Calgary, Calgary, AB T2N 1N4, Canada*

³*Santa Fe Institute, 1399 Hyde Park, Santa Fe, New Mexico 87501, USA*

⁴*Complex Systems Research Group, HNO, Medical University of Vienna, Währinger Gürtel 18-20, A-1090, Austria*

(Received 29 April 2005; published 16 September 2005)

The notion of (auto)catalytic networks has become a cornerstone in understanding the possibility of a sudden dramatic increase of diversity in biological evolution as well as in the evolution of social and economical systems. Here we study catalytic random networks with respect to the final outcome diversity of products. We show that an analytical treatment of this long-standing problem is possible by mapping the problem onto a set of nonlinear recurrence equations. The solution of these equations shows a crucial dependence of the final number of products on the initial number of products and the density of catalytic production rules. For a fixed density of rules we can demonstrate the existence of a phase transition from a practically unpopulated regime to a fully populated and diverse one. The order parameter is the number of final products. We are able to fully understand the origin of this phase transition as a crossover from one set of solutions from a quadratic equation to the other. We observe a remarkable similarity of the solution of the system and the PVT diagrams in standard thermodynamics.

DOI: [10.1103/PhysRevE.72.036117](https://doi.org/10.1103/PhysRevE.72.036117)

PACS number(s): 89.75.Hc, 87.10.+e, 02.10.Ox, 05.70.Ln

I. INTRODUCTION

Chemicals act on chemicals to produce new chemicals, goods act on goods to produce new goods, and ideas act on ideas to produce new ideas. The concept that elements of a set act on other elements of the same set to produce new elements which then become part of this set is ubiquitous not only in nature but also in social systems and processes. We might think of the development of modern chemistry, where the invention of a new compound leads to possibilities to use this compound (as a catalyst) to produce yet another compound. The same is true for economy, where one newly introduced good can be used as a tool to produce new goods and tools. Other famous examples are recent models of evolution; maybe even our whole concept of history in general can be seen as a process of this type.

Maybe the most fascinating question associated with these processes is under which conditions self-sustaining systems can emerge—i.e., that the newly produced chemicals, goods, or ideas find adequate other new or old chemicals, goods, and ideas such that they can act on each other to produce yet new products and so on. For any scientific approach to this subject it is clear that it is necessary to specify rules, which product can act on another product to produce a third one. For example there is a chemical rule that oxygen and hydrogen will produce water but there is no rule that gold and helium can form a compound. There is a rule that a hammer acting on a block of iron will produce sheet metal, but no rule that welding together two blocks of uranium 238 will lead to a big block of uranium. If one imagines that all existing and nonexisting—but possible—products are listed in a high-dimensional vector, then the rules how those ele-

ments can act on each other can be thought of elements in an interaction matrix, where a zero entry means no interaction and a nonzero element gives the interaction strength.

In history there have been plenty of instances where a system of the above type underwent a transition from a state with very few products to a state of vast abundance of products. These transitions happen over relatively short time scales. An example from biology is the Cambrian explosion [1], where an unprecedented number of new taxa emerged within very short time scales. An economical example is the industrial revolution, where the number of industrial goods exploded to previously unimaginable numbers, a social example is the explosion of culture with the advent of the Renaissance, or a more modern example is the explosion of the number chemical compounds in the last century.

All of these “explosive” processes share the same structure: There are possibly constructively interacting elements; the interaction is governed by a set of rules (natural laws, social consensus, religious restrictions) and a number of initially existing products. What does trigger the explosive event; why is this event sometimes missing for very long time periods? The only parameters at hand are the number of rules, the number of initially present elements, and a possible structure in the interaction matrix. It is likely that the cultural explosion was driven by an increase of rules which was possible by driving back religious restrictions. The explosion of chemical compounds was possible by discovering the rules of modern chemistry. Let us mention here that for well-studied systems the set of rules can in principle be known completely. However, for large systems, it might be wise as a first step to model them stochastically—i.e., let the interaction matrix be a random matrix. In this case there will of course be no structure in the interaction matrix, and the only parameters will be rule density r and initial number a_0 of existing elements.

*Electronic address: thurner@univie.ac.at

Since quite some time there has been conjectured the existence of a phase transition in the above systems—e.g., in [2]. By this we mean that in the r - a_0 plane there exist well-defined regions, which are practically unpopulated or almost fully populated, with a sharp transition between the regimes. This means that fixing the number of initial elements there exists a critical density of rules, r_{crit} . When the system is below r_{crit} the number of elements will remain relatively low compared to the total number of possible products. Above r_{crit} the system will become self-sustaining and drive towards a heavily populated state.

Even though this setting seems to be fundamental to a variety of disciplines and its importance has long been noticed [3,4], the progress of a systematic scientific treatment of these problems is limited. Relevant contributions to this field come from chemistry and biology. The adequate mathematical treatment of such processes is so-called catalytic equations, which are sets of coupled, quadratic, ordinary differential equations. Special cases of those equations are, for example, the class of Lotka-Volterra replicators, the hypercycle [5], or more recent ideas like the Turing gas [6]. Replicator dynamics, which is linear, is obtained from catalytic dynamics by a proper scaling of time. It is needless to mention that the nonlinear catalytic equations carry a tremendous potential of complex dynamics; however, in earlier studies some robustness in terms of fixed points seems to have been observed. For more details see [7].

The aim of the present work is to prove the existence and study details of the nature of the above-mentioned phase transition. This is possible by introducing some concepts of set theory and—making use of the random structure of the production rules—by mapping the size of a consecutive series of sets onto a set of update equations, which can be solved and analyzed. The analytic formulas give insights into what is happening at the transition. A practical aspect of this present work is that we cover analytically the large-system limit, which is beyond numerical reachability and has so far not been possible to study.

The intuition for this work is based on a bit-string formulation of the above problem, which is a generalization of models recently termed random grammars [4]. In the bit-string model there is a set of initial bit strings and a set of strings which act on these strings by either combining strings or substituting substrings. These latter strings can be seen as catalysts; their existence constitutes the presence of rules of what can be combined and/or substituted. The model which is presented below can be shown to be one-to-one compatible with the substitution and combination rules in the Kauffman bit-string model.

The paper is structured as follows: In Sec. II, the formal problem is stated and necessary notation and concepts are developed. Section III, the set of update equations is derived. Section IV presents the solution of these equations, and in Sec. V, the results are presented, which are further discussed in Sec. VI.

II. NETWORK EQUATION AND SOME NOTATION

Our basic concept is to view the abundances of all possible products as entries in a d -dimensional time-dependent

vector $x(t)$; i.e., $x_i(t)$ is the quantity of product i present at time t . We drop the time notation in the following. The population dynamics of a system, where product i under the influence of product j produces product k , is given by network equations of the type

$$\dot{x}_k = \sum_{i,j} \alpha_{ij}^k x_i x_j - x_k \Phi, \quad \Phi = \sum_{i,j,k} \alpha_{ij}^k x_i x_j, \quad (1)$$

with $k \in \Lambda$ and the component $\alpha_{ij}^k > 0$. Λ is the domain of nodes or—put in fancier terms—the index set of the dynamic process. x is the state vector with dimensionality $d = |\Lambda|$, the number of all possible products or nodes. The state vector is a vector of relative frequencies $0 \leq x_k < 1$, and $\sum_k x_k = 1$. Each α^k encodes the structural information about which binary combinations $(i, j) \in \Lambda^2$, the *substrates*, can interact in order to form *product* k . Each entry in α^k is a real number; however, for the sake of simplicity in the following we consider only 1 for interaction or 0 for no interaction. α^k can be thought of as the set of rules of how objects interact. In the following α^k is interpreted as a random matrix in the following way: We assume that for each product $k \in \Lambda$ there are

$M(k)$ pairs of substrates $(L(k), M(k)) \in \Lambda^2$ such that $L(k) \rightarrow \{k\}$. In words, this arrow means that k is produced by substrate $L(k)$ under the “influence” of substrate $M(k)$. Note that $L(k)$ and $M(k)$ do not have to be unique; there can be more than one pair $(L(k), M(k))$ producing a specific product k . Let us call the number of pairs leading to product k , $\mathcal{N}_{L,M}(k)$. We define the *production rule density* as the average number of pairs leading to one product:

$$r = \langle \mathcal{N}_{L,M}(k) \rangle_k. \quad (2)$$

Equations like Eq. (1) have long been known for their rather surprising robustness in terms of fixed points where the system converges to. This is not obvious, and one would rather expect a situation more dominated by more complicated orbits and limit cycles. For a more detailed discussion of this type of equations see [7]. Fixed points will therefore provide much information about the effective behavior of such systems, the fixed-point equation being

$$x_k^* = \frac{1}{\Phi^*} \sum_{i,j} \alpha_{ij}^k x_i^* x_j^*. \quad (3)$$

We are dealing with a nonlinear process, so that the solutions of the network equation will in general depend strongly on the initial conditions. Even assuming the process to be driven towards a stable fixed point does not imply the uniqueness of that fixed point.

Notation

We want to get some feeling as what to expect when investigating the population dynamics of randomly sampled fixed 0/1 networks. The dynamical properties of Eq. (1) are linked to specific topologies. Knowing the topological features will enable one to solve for the dynamics. Not knowing the interaction tensor α exactly and only given its statistical properties, following the classical concept of statistical phys-

ics, one still can understand the expected dynamics of the system and, in particular, its expected final outcome. We hence study the system in a probabilistic fashion—i.e., averaging over all possible topologies. The choices for $(L, M) \times (k)$ are equally distributed, and L and M can be seen as independent random variables.

The question we try to answer is, how many products do we expect to end up with when starting from a given number of randomly chosen initial substrate species? Before we proceed we need a number of definitions. We denote the number of elements contained in some set A by writing $|A|$ and define the *support* of a process at a given time:

$$S(x) := \{k | x_k > 0\}. \quad (4)$$

Suppose that the process is driven to a (stable) fixed point such that the final support is

$$S^*(x) := S(\lim_{t \rightarrow \infty} x(t)). \quad (5)$$

Inversely, we can ask which initial conditions end up in the same fixed point¹ and define the *body* of this fixed point to be the set

$$B(x^*) := \left\{ x \mid \sum_k x_k = 1 \wedge \lim_{t \rightarrow \infty} x(t) = x^* \right\}. \quad (6)$$

We further define the following operations on product sets: the *forward difference* of set A ,

$$\partial_+ A := \{k \in \Lambda \mid \{L(k), M(k)\} \subset A\} \setminus A, \quad (7)$$

and the *backward difference* of A ,

$$\partial_- A := \cup_{k \in A} \{L(k), M(k)\} \setminus A. \quad (8)$$

With these operations we can define the *forward closure* of set A ,

$$\bar{A}^+ := \cap \{B \mid A \subset B \wedge \partial_+ B \subset B\}, \quad (9)$$

and the *backward closure* of a set,

$$\bar{A}^- := \cap \{B \mid A \subset B \wedge \partial_- B \subset B\}. \quad (10)$$

These definitions can best be understood by viewing our dynamical system, Eq. (1), as a directed graph whose nodes are connected by arrows. A node looking at its edges can therefore distinguish between feather ends or arrowheads. A node holding a feather end of an arrow is a substrate; a node holding an arrowhead is a product with respect to this edge. Given some set of nodes we can identify for each node k all the nodes which will be formed due to the influence of k in the next time step. In this sense we say the nodes look “forward.” The forward difference of some set is therefore the set of all products the nodes of the set look forward to, excluding the products that are already present as nodes of the set we started out with. The backward difference follows the same idea only looking at the arrows from arrowhead to

wards the feathers. The corresponding closures are then simply obtained by adding the set differences iteratively to the initial set; i.e., we add the difference to the initial set and form a new difference on the first extension of the set, then add this difference to obtain the second extension of the initial set, and so on, until the difference (forward or backward) is empty and the iterative process comes to a halt. That this eventually may happen can be understood by looking at chains of arrows as an example. Take a node and add an arrow pointing at some other randomly chosen node in the set. In the beginning the chance to select a virginal node is high and the chain will grow, but as the chain grows the probability of sampling a member already in the chain increases, and even though the single-sample probability may still be small, the chance to sample into the chain eventually is not.

For clarity let us summarize our philosophy: All objects which can possibly act on each other are represented by the index set Λ of the processing system, Eq. (1), which contains all the “names” of the considered and “thinkable” objects. The index set provides the domain for all dynamical considerations. Properties of the system are implemented via the map $x: \Lambda \rightarrow \mathbf{R}_+$, which are the relative frequencies of the indexed species in the domain. We are not interested in particular weights but only in the directed network topology coded by the matrices α^k on the domain Λ . Not only does this topological approach provide us with the means to talk about randomly distributed productive substrate pairs, but also about their density of occurrence r . Even more important, if we are not interested in details of the dynamics, but decide to focus only on how large the expected final set will be (given some initial substrate set and density), we can drop dynamical considerations and pass to topological ones. We are not interested in how much of each object species we will effectively end up with, just if it got produced or not. The subset of the domain that is effectively populated is called the support $S \subset \Lambda$. To investigate the flow of the support under the network equation we may utilize set operations compatible with the topological structure of directed reaction graphs leading to the definition of the forward and the backward difference and their respective closures. It is intuitively clear that the forward closure of some set is the subset of the domain that is flooded by the initial set during the production process. It necessarily forms an upper limit for the size of possible self-sustainable subsets of the domain reachable from the initial condition. In a simplification considering an equally distributed random interaction topology we gain a notion of expected growth rates for the set differences based on expected sampling rates. This leads to equations of expected growth as demonstrated next.

III. GROWTH LAWS FOR CATALYTIC SETS

We now develop a method to compute the size of the forward closure $a_\infty = |\bar{A}^+|$ as a function of the production rule density and the size of the initial set A . a_∞ is the final number of products once the system has converged. The probability of choosing some random but specific substrate pair l and m among all possible pairs is obviously

¹If there is no stable fixed point, the limits in the definition of the final support are not converging and one has to specify convergent subsequences reflecting the structure of the limit cycle or the chaotic regime.

$$p = \frac{2}{d(d-1)}. \quad (11)$$

Imagine an initial random set A_0 of products with a number of $a_0=|A_0|$ elements. Given the dimensionality d of the system and a production rule density r , the expected number of produced elements in the next time step is the number of possible pairs in A_0 times the probability to find a productive pair—i.e.,

$$rd \binom{a}{2} p = \frac{ra(a-1)}{(d-1)} \sim \frac{ra^2}{d}. \quad (12)$$

Several of these newly created elements will already be $\in A_0$. The probability that one of these produced elements is not yet $\in A_0$ is $1-a_0/d$, leading to the actual size of the catalytic set at time 1, A_1 :

$$a_1 = a_0 + \Delta a_0 \quad \text{with} \quad \Delta a_0 = \frac{r}{d} \left(1 - \frac{a_0}{d}\right) a_0^2, \quad (13)$$

with $\Delta a_0 = |\partial_+ A_0|$ being the increment of elements. What will happen in the next time step? We now have $A_1 = A_0 \cup \partial_+ A_0$. The increment for the next time step will be made up of all the new products which are possible (and not already in A_1) by combining two elements from $\partial_+ A_0$ or by combining one element from $\partial_+ A_0$ with one from A_0 . The number of possible pairs for those combinations are $\Delta a_0(\Delta a_0 - 1)/2 \sim \Delta a_0^2/2$ and $a_0 \Delta a_0$, respectively. Note that the third possibility combining two elements from set A_0 will always lead to $\partial_+ A_0$, which is already $\in A_1$, and no new products can emerge from this. These possibilities multiplied by p , rd , and the probability that the new product lies in A_1 already, lead to the increment for the second time step:

$$\Delta a_1 = \frac{r}{d} \left(1 - \frac{a_1}{d}\right) (\Delta a_0^2 + 2a_0 \Delta a_0) = \frac{r}{d} \left(1 - \frac{a_1}{d}\right) (a_1^2 - a_0^2). \quad (14)$$

Now, continue the iterative scheme $A_{t+1} = A_t \cup \partial_+ A_t$. This means that the set at time $t+1$ is the old set plus the products newly generated in the time span $[t, t+1]$. We can finally write down a growth equation for catalytic sets:

$$a_{t+1} = a_t + \Delta a_t, \quad \Delta a_{t+1} = \frac{r}{d} \left(1 - \frac{a_{t+1}}{d}\right) (a_{t+1}^2 - a_t^2), \quad (15)$$

with initial conditions $a_0 = |A_0|$ and $a_{-1} = 0$. To sum up, we have to take care of all possible new pairs by looking at all new elements added in an iteration step and the new pairs they can build with themselves. We also have to take into account all the pairs they can build with elements produced earlier. We have to exclude all the pairs that already have been considered. This is all captured by Eq. (15), which should be noted to be scale invariant with respect to dimension d . To see this just scale $a \rightarrow a/d$, and the dimension drops out of the equation. It is therefore fully justified to drop d from Eq. (15), if wanted.

IV. SOLUTION OF THE UPDATE EQUATION: RESULT

The solution of the growth equations (15) with respect to the productive pair density r and the initial set size a_0 is given in Fig. 1(a). The immediate message is that there is a critical density r_{crit} above which a continuous increase of the initial set size a_0 does not correspond to a continuous increase of its forward closure size a_∞ , but displays a phase transition, a jump from small to very large forward closures at some $a_0^{crit}(r)$. a_0^{crit} vanishes for $r < r_{crit}$. This demonstrates unambiguously the existence of a phase transition in catalytic random systems. As a check of consistency of the above arguments we have also simulated Eq. (1) directly for $d = 50$ random networks. The corresponding results coincide very nicely with Fig. 1(a) and can be found in [8].

Analytical approximation of the forward-closure size

Let us define $c_t = \Delta a_{t+1} / \Delta a_t$. As an approximation imagine for a moment that c_t is a constant in time c . Then for the forward closure we have

$$a_\infty = \sum_t c^t a_0 = \frac{1}{1-c}. \quad (16)$$

Let us write c , starting from right part of Eq. (15) and dividing it by Δa_t to get

$$c = \frac{\Delta a_{t+1}}{\Delta a_t} = \frac{r}{d} \left(1 - \frac{a_{t+1}}{d}\right) (a_{t+1} + a_t). \quad (17)$$

For the $t \rightarrow \infty$ limit we naturally assume $\lim_{t \rightarrow \infty} a_{t+1} + a_t = 2a_\infty$, so that we get a candidate for c which produces the right asymptotic results:

$$c = 2r \left(1 - \frac{a_\infty}{d}\right) \frac{a_\infty}{d} \quad \text{with} \quad a_\infty = \frac{a}{1-c}. \quad (18)$$

We now use Eq. (18) as an Ansatz which is a polynomial of third degree, $a_\infty^3 - a_\infty^2 d + a_\infty (d^2/2r) - a (d^2/2r) = 0$, which—when written in terms of c —reads

$$c^3 - 2c^2 + c \left(1 + 2r \frac{a}{d}\right) + 2r \frac{a}{d} \left(\frac{a}{d} - 1\right) = 0. \quad (19)$$

Substituting $c = x + 2/3$ reduces to the equation $x^3 + px + q = 0$ with

$$p = -\frac{1}{3} \left(1 - 6r \frac{a}{d}\right), \quad q = \frac{2}{27} + \frac{2}{3} r \frac{a}{d} \left(3 \frac{a}{d} - 1\right), \quad (20)$$

so that Cardano's method can be used with $x = y + z$ to get

$$y = \left[-\frac{q}{2} + \sqrt{(q/2)^2 + (p/3)^3} \right]^{1/3}, \quad z = -\frac{p}{3y}, \quad (21)$$

providing us with three solutions $y+z$, $\rho^2 y + \rho z$, and $\rho y + \rho^2 z$ with $\rho = \exp(2\pi i/3)$. We are only interested in the real solution $y+z$ which yields the final result

$$a_\infty = \frac{a}{1-c} \quad \text{with} \quad c = 2/3 + y + z, \quad (22)$$

which is a good approximation of the true solution of the forward equation (15) and is plotted in Fig. 1(b) in the r - a_0

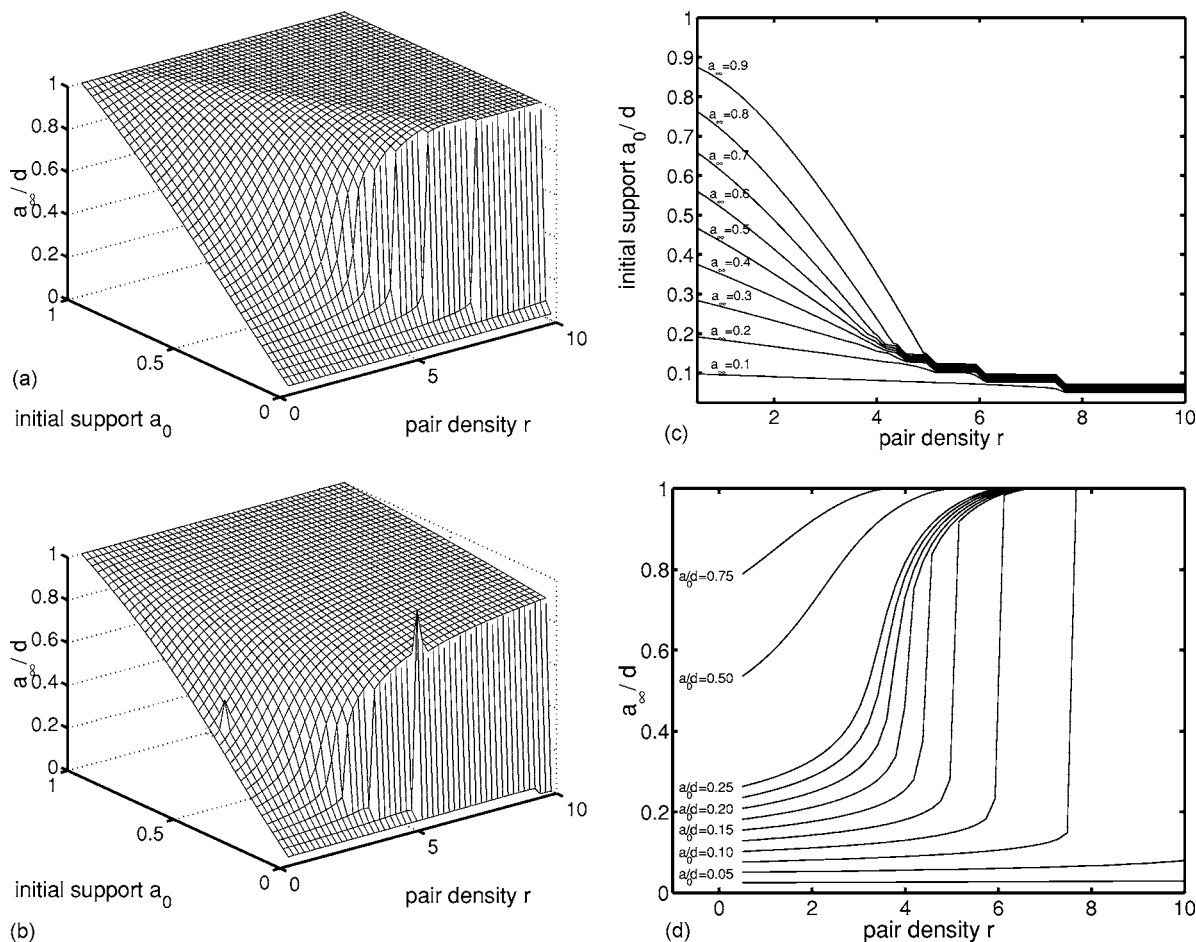


FIG. 1. Solution to the update equations (15) (a) and its analytical approximation (22) (b). There is clearly a critical line in the r - a_0 plane where a transition from an unpopulated to an almost fully populated regime occurs. The spikes in (b) are due to numerical problems in the plot with poles in the solution a_∞ . (c) Projection of the exact solution to the r - a_0 plane. The critical point is best seen in the “PV diagram” as the starting point of the phase plane at pair density $r=4$ and $a_0/d=0.65$. (d) Projection of the exact solution to the r - a_∞/d plane. The iso- a_0 lines correspond to Maxwell lines in the classical gas analogon.

plane. The comparison with Fig. 1(a) demonstrates the quality of the analytical solution.

The analytical form of the solution allows us to understand the origin of the phase transition. For this reason we observe that we can compute a_∞ by solving $c=2r(1-a_\infty)a_\infty$, providing us with two solutions

$$a_{\infty\pm} = \frac{d}{2} \left(1 \pm \sqrt{1 - \frac{c}{r}} \right). \quad (23)$$

These solutions are shown in Figs. 2(a) and 2(b). It becomes obvious that the phase transition actually always takes place by the system switching from the one solution to the other at $a_0^{crit}(r)$. Below a critical value of r_{crit} the one solution smoothly passes into the other. We observe a crossover of the probable and the *improbable* solution. Most likely this is due to a convexity argument, implying the monotonicity of the size of the forward closure with respect to the size of the initial condition. The size of the forward closure will not shrink with increasing the initial set size and therefore has to jump to the alternative solution at some critical line. As long as the third-order polynomial is strictly monotonic we have a

unique real solution for the zero of the polynomial. At r_{crit} the polynomial starts to have a local minimum and maximum and we have in fact two relevant real zeros (out of three), the large and the small solution. r_{crit} is determined by the triple zero of the polynomial. Note a famous analogy here: Just as in Vanderwaal’s gas we can draw Maxwell lines. At the phase transition, the small solution becomes instable and the large solution becomes stable. Here the productive pair density r takes the role of temperature in Vanderwaal’s gas. Below the critical value a system *freezes* into a small set of durable species, while above the critical production law density, supercritical but yet small sets of individuals *evaporate* into their forward closure.

V. DISCUSSION AND OUTLOOK

We have developed a way to map a conceptual system of (auto)catalytic agents into a quantitative framework. With this it is possible to show that the combination of the initial number of products a_0 and the density of mutual production rules r is crucially influencing the growth mode of sets of products. Our main result is that we are able to map the class

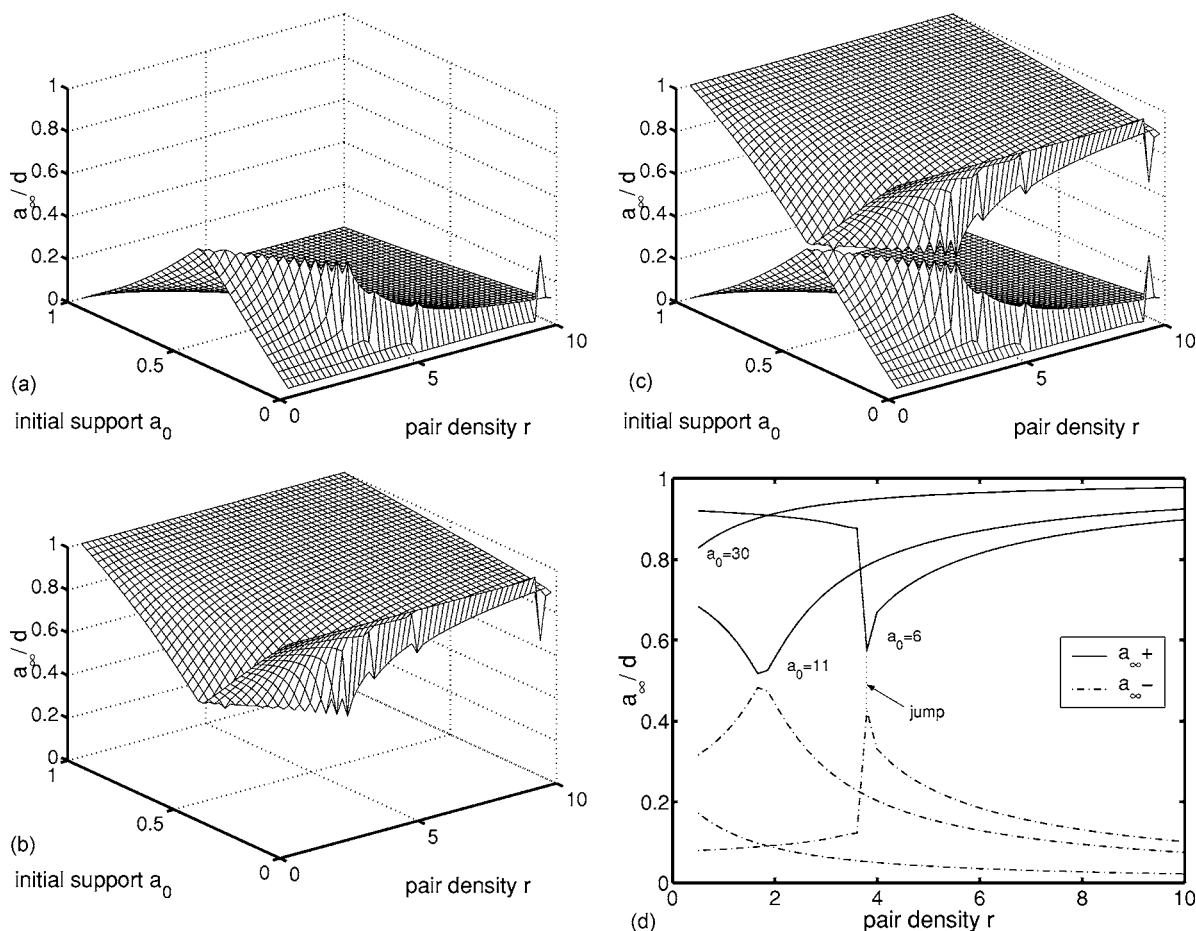


FIG. 2. Plot of the solutions $a_{\infty-}$ (a), $a_{\infty+}$ (b), and both of the above in one plot (c). In (d) a projection onto the r - a_{∞} plane is shown for several a_0 values. It is clear that the phase transition happens as the solutions switch from $a_{\infty+}$ to $a_{\infty-}$.

of random catalytic networks onto a set of growth equations, which allows us to study the expectation value of the final number of products. The resulting equation can be solved and shows a clear phase transition from practically unpopulated zones towards almost fully populated ones, in “rule density and initial number of products” space. The transition is a crossover between two sets of solutions to a quadratic equation.

We believe that this nicely resembles numerical-experimental findings of [7], where the authors find a decrease of species with an increase of p_{el} in their Fig. 1(d). The fact that the transition from small to full forward-closure sizes is gradual is either due to the limited system size (ten species) or that the initial support was high—i.e., above $a_0^{crit}(r)$ —so that no sharp increase could be found. In fact, systems of the type of Eq. (1) are straightforwardly solvable numerically up to system sizes of about $d=100$ within reasonable computing time. We note that the practical value of this present work is that it captures the large-system-size limit, which is out of numerical reach and was not tractable analytically before.

The present approach does not try to explain the detailed role of autocatalytic cycles or of key-node species in the context of understanding the beginning of an explosion of

species numbers as was very nicely done, for example, in [9]. The periods of fast extinction reported there are not incorporated into our consideration yet, since we have not taken any evolutionary hypotheses into account so far; nor did we incorporate evolutionary concepts in the sense that products are associated characteristics (some sort of fitness—e.g., the weight x_k) upon which they can get selected by some method. Taking these arguments into account seems a natural starting point for future research. We believe that it should be possible that a combination of backward- and forward-closure arguments can be used to estimate a critical density of autocatalytic cycles necessary for a system to become critical as studied here. We have not so far considered negative entries in the interaction matrices, which should also be present in natural or social systems. Finally, we mention that our arguments given here should—due to the absence of a characteristic scale in the update equations—not be limited to finite sets.

ACKNOWLEDGMENTS

S.T. would like to thank the SFI and in particular J. D. Farmer for their great hospitality and support in 2004. The project was funded in part by the Austrian Council and FWF Project No. P17621 G05.

- [1] S. J. Gould, *Wonderful Life* (Harvard University Press, Cambridge, MA, 1989).
- [2] J. D. Farmer, S. A. Kauffman, and N. H. Packard, *Physica D* **22**, 50 (1986).
- [3] J. A. Schumpeter, *Theorie der wirtschaftlichen Entwicklung* (Dunker and Humbolt, Muenchen und Leipzig, 1926).
- [4] S. A. Kauffman, *The Origins of Order* (Oxford University Press, London, 1993).
- [5] M. Eigen and P. Schuster, *The Hypercycle* (Springer-Verlag, Berlin, 1979).
- [6] W. Fontana, in *Artificial Life II*, edited by C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (Addison-Wesley, Redwood City, CA, 1992), pp. 159–210.
- [7] P. F. Stadler, W. Fontana, and J. H. Miller, *Physica D* **63**, 378 (1993).
- [8] <http://www.complex-systems.meduniwien.ac.at/supplementary2005B/>
- [9] S. Jain and S. Krishna, *Phys. Rev. Lett.* **81**, 5684 (1998); *Proc. Natl. Acad. Sci. U.S.A.* **99**, 2055 (2002).