1
Can we recognize an innovation?: Perspective from an evolving network model

Sanjay Jain and Sandeep Krishna

'Innovations' are central to the evolution of societies and the evolution of life. But what constitutes an innovation? We can often agree after the event when its consequences and impact over a long term are known whether something was an innovation or not, and whether it was a 'big' innovation or a 'minor' one. But can we recognize an innovation 'on the fly' as it appears? Successful entrepreneurs often can. Is it possible to formalize that intuition? We discuss this question in the setting of a mathematical model of evolving networks. The model exhibits self-organization, growth, stasis, and collapse of a complex system with many interacting components, reminiscent of real world phenomena. A notion of 'innovation' is formulated in terms of graph-theoretic constructs and other dynamical variables of the model. A new node in the graph gives rise to an innovation provided it links up 'appropriately' with existing nodes; in this view innovation necessarily depends upon the existing context. We show that innovations, as defined by us, play a major role in the birth, growth and destruction of organizational structures. Furthermore, innovations can be categorized in terms of their graph-theoretic structure as they appear. Different structural classes of innovation have potentially different qualitative consequences for the future evolution of the system, some minor and some major. Possible general lessons from this specific model are briefly discussed.

1.1
Introduction

In everyday language, the noun innovation stands for something new that brings about a change; it has a positive connotation. Innovations occur in all branches of human activity – in the world of ideas, in social organization, in technology. Innovations may arise by conscious and purposeful activity, or serendipitously; in either case, innovations by humans are a consequence of cognitive processes. However, the word innovation does not always re-

Titel. Author(s)

ISBN: 3-527-XXXXX-X

fer to a product of cognitive activity. In biology, we say, for example, that photosynthesis, multicellularity, and the eye were evolutionary innovations. These were products not of any cognitive activity, but of biological evolution. It nevertheless seems fair to regard them as innovations; these novelties certainly transformed the way organisms made a living. The notion of innovation seems to presuppose a context provided by a complex evolutionary dynamics; for example, in everyday language the formation of the earth, or even the first star, is not normally referred to as an innovation.

Innovations are a crucial driving force in chemical, biological and social systems, and it is useful to have an analytical framework to describe them. This subject has a long history in the social sciences (see, e.g., [1, 2]). Here we adopt a somewhat different approach. We give a mathematical example of a complex system that seems to be rich enough to exhibit what one might intuitively call innovation, and yet simple enough for the notion of innovation to be mathematically defined and its consequences analytically studied. The virtue of such a stylized example is that it might stimulate further discussion about innovation, and possibly help clarify the notion in more realistic situations.

Innovations can have 'constructive' and 'destructive' consequences at the same time. The advent of the automobile (widely regarded as a positive development) was certainly traumatic for the horse drawn carriage industry and several other industries that depended upon it. When aerobic organisms appeared on the earth, their more efficient energy metabolism similarly caused a large extinction of several anaerobic species [3]. The latter example has a double irony. Over the first two billion years of life on earth, there was not much oxygen in the earth's environment. Anaerobic creatures (that did not use free oxygen for their metabolism) survived, adapted, innovated new mechanisms (e.g., photosynthesis) in this environment, and spread all over the earth. Oxygen in the earth's environment was largely a by-product of photosynthetic anaerobic life, a consequence of anaerobic life's 'success'. However, once oxygen was present in the environment in a substantial quantity, it set the stage for another innovation, the emergence of aerobic organisms which used this oxygen. Because of their greater metabolic efficiency the aerobic organisms out-competed and decimated the anaerobic ones. In a very real sense, therefore, anaerobic organisms were victims of their own success. Innovation has this dynamic relationship with 'context': what constitutes 'successful' innovation depends upon the context, and successful innovation then alters the context. Our mathematical example exhibits this dynamic and explicitly illustrates the two-faced nature of innovation. We show that the ups and downs of our evolutionary system as a whole are also crucially related to innovation.

1.2
A framework for modelling innovation: Graph theory and dynamical systems

Systems characterized by complex networks are often represented in terms of a graph consisting of nodes and links. The nodes represent the basic components of the system, and links between them their mutual interactions. A graph representation is quite flexible and can represent a large variety of situations [4]. For a society, nodes can represent various agents such as individuals, firms and institutions, as well as goods and processes. Links between nodes can represent various kinds of interactions, such as kinship or communication links between individuals, inclusion links (e.g., a directed link from a node representing an individual to a node representing a firm implying that the individual is a member of the firm), production links (from a firm to a good that it produces), links that specify the technological web (for every process node, incoming links from all the goods it needs as input and outgoing links to every good it produces), etc. In an ecological setting, nodes can represent biological species, and links their predator-prey or other interactions. In cellular biology, nodes might represent molecules such as metabolites and proteins as well as genes, and links their biochemical interactions.

A graph representation is useful for describing several kinds of innovation. Often, an innovation is a new good, process, firm, or institution. This is easily represented by inserting a new node in the graph, together with its links to existing nodes. Of course, not every such insertion can be called an innovation; other conditions have to be imposed. The existing structure of the graph provides one aspect of the 'context' in which a prospective innovation is to be judged, reflecting its 'location' or relationship with other entities. In this formulation it is clear that innovations such as the ones mentioned above are necessarily a change in the graph structure. Thus a useful modelling framework for innovations is one where graphs are not static but change with time. In real systems graphs are always evolving: new nodes and links constantly appear, and old ones often disappear as individuals, firms and institutions die, goods lose their utility, species become extinct or any of these nodes lose some of their former interactions. It is in such a scenario that certain kinds of structures and events appear that earn the nomenclature 'innovation'. We will be interested in a model where a graph evolves by the deletion of nodes and links as well as the insertion of new ones. Insertions will occasionally give rise to innovations. We will show that innovations fall in different categories that can be distinguished from each other by analysing the instantaneous change in the graph structure caused by the insertion, locally as well as globally. We will argue that these different 'structural' categories have different 'dynamical' consequences for the 'well-being' of other nodes and the evolution of the system as a whole in the short as well as long run.

In addition to an evolving graph, another ingredient seems to be required for modelling innovation in the present approach: a graph dependent dynamics of some variables associated with nodes or links. In a society, for example, there are flows of information, goods and money between individuals that depend upon their mutual linkages, which affect node attributes such as individual wealth, power, etc. The structure of an ecological food web affects the populations of its species. Thus, a change in the underlying graph structure has a direct impact on its 'node variables'. Deciding whether a particular graph change constitutes an innovation must necessarily involve an evaluation of how variables such as individual wealth, populations, etc., are affected by it. Changes in these variables in turn trigger further changes in the graph itself, sometimes leading to a cascade of changes in the graph and other variables. E.g., the decline in wealth of a firm (node variable) may cause it to collapse; the removal of the corresponding node from the market (graph change) may cause a cascade of collapses. The invention of a new product (a new node in the graph) which causes the wealth of the firm inventing it to rise (change in a node variable) may be emulated by other firms causing new linkages and further new products.

In order to 'recognize an innovation on the fly' it thus seems reasonable to have a framework which has (a) a graph or graphs representing the network of interactions of the components of the system, (b) the possibility of graph evolution (the appearance and disappearance of nodes and links) and (c) a graph dependent dynamics of node or link variables that in turn has a feedback upon the graph evolution. The example discussed below has these features. They are implemented in simple framework that has only one type of node, one type of link and only one type of node variable.

1.3
Definition of the model system

The example is a mathematical model [5] motivated by the origin of life problem, in particular, the question of how complex molecular organizations could have emerged through prebiotic chemical evolution [6, 7, 8, 9, 10]. There are s interacting molecular species in a 'prebiotic pond', labelled by $i \in S \equiv \{1, 2, \ldots, s\}$. Their interactions are represented by the links of a directed graph, of which these species are nodes. The graph is defined by its $s \times s$ adjacency matrix $C = (c_{ij})$, with $c_{ij} = 1$ if there exists a link from node j to node i (chemically that means that species j is a catalyst for the production of species i), and $c_{ij} = 0$ otherwise. c_{ii} is assumed zero for all i: no species in the pond is self-replicating. Initially the graph is chosen randomly, each c_{ij} for $i \neq j$ is chosen to be unity with a small probability p and zero with probability $1 - p$. p represents the ‘catalytic probability’ that a given molecular species will catalyse the production of another randomly chosen one [11].

The pond sits by the side of a large water body like a sea or river, and periodically experiences tides or floods which can flush out molecular species from the pond and bring in new ones, changing the network. We use a simple graph update rule in which exactly one node is removed from the graph (along with all its links) and one new node is added whose links with the remaining $s-1$ nodes are chosen randomly with the same probability p. We adopt the rule that the species with the least relative population (or, if several species share the least relative population, one of them chosen randomly) is removed. This is where selection enters the model: species with smaller populations are punished. This is an example of ‘extremal’ selection [10] in that the least populated species is removed; the results of the model are robust to relaxing the extremality assumption [12, 13].

In order to determine which node will have the least population, we specify a population dynamics which depends upon the network. The dynamics of the relative populations, x_i ($0 \leq x_i \leq 1$, $\sum_{i=1}^{s} x_i = 1$) is given by

$$\dot{x}_i = \sum_{j=1}^{s} c_{ij} x_j - x_i \sum_{k=1}^{s} \sum_{j=1}^{s} c_{kj} x_j. \tag{1.1}$$

This is a set of rate equations for catalysed chemical reactions in a well stirred chemical reactor[1]. They implement the approximate fact that under certain simplifying assumptions a catalyst causes the population of whatever it catalyses to grow at a rate proportional to its own (i.e., the catalyst’s) population [5, 14]. Between successive tides or floods the set of species and hence the graph remains unchanged, and the model assumes that each x_i reaches its attractor configuration X_i under (1.1) before the next graph update. The species with the least X_i is removed at the next graph update.

Starting from the initial random graph and random initial populations, the relative populations are evolved according to (1.1) until they reach the attractor $\mathbf{X}$, and then the graph is updated according to the above rules. The new incoming species is given a fixed relative population x_0, all x_i are perturbed about their existing values (and rescaled to restore normalization). This process is iterated several times. Note that the model has two inbuilt time scales, the population dynamics relaxes on a fast time scale, and graph evolution on a slow time scale. The above model may be regarded as an evolutionary model in non-equilibrium statistical mechanics.

1) See *Derivation of equation 1.1 in Appendix A.*

1.4
Time evolution of the system

A sample run is depicted in Figs. 1 and 2. For concreteness, we will discuss this run in detail, describing the important events, processes, and the graph structures that arise, with an emphasis on the role of innovation. The same qualitative behaviour is observed in hundreds of runs with the various parameter values. Quantitative estimates of average time scales, etc., as a function of the parameters s and p are discussed in [5, 14] and Appendix A. The robustness of the behaviour to various changes of the model is discussed in [12, 13].

Broadly, Figs. 1 and 2 exhibit the following features: Initially, the graph is sparse and random (see Figs 2a-d), and remains so until an autocatalytic set (ACS), defined below, arises by pure chance. On average the ACS arrives on a time scale $1/(p^2 s)$ in units of graph update time[2]; in the exhibited run it arrives at $n = 2854$ (Fig. 2e). In this initial regime, called the 'random phase', the number of populated species, s_1, remains small. The appearance of the ACS transforms the population and network dynamics. The network self-organizes, its density of links increases (Fig. 1a), and the ACS expands (Figs. 2e-n) until it spans the entire network (as evidenced by s_1 becoming equal to s, at $n = 3880$, Figs. 1b,2n). The ACS grows across the graph exponentially fast, on a time scale $1/p$ [5]. This growth is punctuated by occasional drops (e.g., Fig. 1b at $n = 3387$, see also Figs. 2h,i). The period between the appearance of a small ACS and its eventual spanning of the entire graph is called the 'growth phase'. After spanning, a new effective dynamics arises, which can cause the previously robust ACS to become fragile, resulting in crashes (the first major one is at $n = 5041$) in which s_1 as well as the number of links drops drastically. The system experiences repeated rounds of crashes and recoveries (Figs. 2o-u, see [15, 12] for a longer time scale.) The period after a growth phase and upto a major crash (more precisely, a major crash that is a 'core-shift', defined below) is called the 'organized phase'. After a crash, the system ends up in the growth phase if an ACS still exists in the graph (as at $n = 5042$, Fig. 2r) or the random phase if it does not (as at $n = 8233$, Fig. 2t). Below we argue that most of the crucial events in the evolution of the system, including its self-organization and collapse, are caused by 'innovation'.

2) See *Timescale for appearance and growth of the dominant ACS* in Appendix A.

1.5 Innovation

The word 'innovation' certainly connotes something new. In the present model at each graph update a new structure enters the graph: the new node and its links with existing nodes. However not every new thing qualifies as an innovation. In order for a novelty to bring about some change, it should confer some measure of at least temporary 'success' to the new node. (A mutation must endow the organism in which it appears some extra fitness, and a new product must have some sale, in order to qualify as an innovation.) In the present model, after a new node appears, the population dynamics takes the system to a new attractor of (1.1), which depends upon the mutual interactions of all the nodes. In the new attractor this node (denoted k) may go extinct, $X_k = 0$, or may be populated, $X_k > 0$. The only possible criterion of individual 'success' in the present model is population. Thus we require the following minimal 'performance criterion' for a new node k to give rise to an innovation: X_k should be greater than zero in the attractor that follows after that particular graph update. That is, the node should 'survive' at least till the next graph update.

This is obviously a 'minimal' requirement, a necessary condition, and one can argue that we should require of an innovation more than just this 'minimal performance'. A new node that brings about an innovation ought to transform the system or its future evolution in a more dramatic way than merely surviving till the next graph update. Note, however, that this minimal performance criterion nevertheless eliminates from consideration a large amount of novelty that is even less consequential. Out of the 9999 new nodes that arise in the run of Fig. 1, as many as 8929 have $X_k = 0$ in the next population attractor; only 1070 have $X_k > 0$. Furthermore, the set of events with $X_k > 0$ can be systematically classified in the present model using a graph theoretic description. Below we describe an exhaustive list of 6 categories of such events, each with a different level of impact on the system (see Fig. 3, discussed in detail below). One of these categories consists of nodes that disappear after a few graph updates leaving no subsequent trace on the system. Another category consists of nodes that have only an incremental impact. The remaining four categories cause (or can potentially cause) more drastic changes in the structure of the system, its population dynamics and its future evolution.

In view of this classification it is possible to exclude one or more of these categories from the definition of innovation and keep only the more 'consequential' ones. However, we have chosen to be more inclusive and will regard all the above categories of events as innovations. In other words we will regard the above 'minimal' performance criterion as a 'sufficient' one for innovation. Thus we will call the introduction of a new node k and the graph structure so formed an *innovation* if $X_k > 0$ in the population attractor that immediately follows the event. i.e., if the node 'survives' at least until the next graph update. This definition then includes both 'small' and 'big' innovations that can be recognized based on their graph theoretic structure upon appearance.

As will be seen below, it turns out that a new node generates an innovation only if it links 'appropriately' to 'suitable' structures in the existing graph. Thus the above definition makes the notion of innovation context dependent. It also captures the idea that an innovation rests on new linkages between structures.

1.6 Six categories of innovation

1.6.1 A shortlived innovation: Uncaring and unviable winners

There are situations where a node, say an agent in society or a species in an ecosystem, acquires the ability to parasite off another, without giving the system anything substantive in return. The parasite gains as long as the host survives, but often this situation doesn't last very long. The host dies, and eventually so does the parasite that is dependent on it. It is debatable whether the acquiring of such a parasitic ability should be termed an innovation, but from the local vantage point of the parasite, while the going is still good, it might seem like one.

Figs. 2b,c show an example of an innovation of such a type that appears in the random phase of the model. Node 25 is the node which is replaced at $n = 78$ (see Fig. 2b, where node 25 is coloured white, implying that $X_{25} = 0$ at $n = 78$.) The new node that replaces it (also numbered 25 in Fig. 2c) receives a link from node 23 thus putting it at the end of a chain of length 2 at $n = 79$. This is an innovation according to the above definition, for, in the attractor configuration corresponding to the graph of Fig. 2c, node 25 has a non-zero relative population; $X_{25} > 0$. This is because for any graph that does not contain a closed cycle, one can show that the attractor $\mathbf{X}$ of (1.1), for generic initial conditions, has the property that only those X_i are non-zero whose nodes i are the endpoints of the longest chains in the graph[3]. The X_i for all other nodes is zero [16]. Since the longest chains in Fig. 2c are of length 2, node 25 is populated. (This explains why a node is grey or white in Figs. 2a-d.) Note that node 25 in Fig. 2c has become a parasite of node 23 in the sense that there is a link from 23 to 25 but none from 25 to any other node. This means

3) See *The attractor of equation 1.1* in Appendix A.

that node 25 receives catalytic support for its own production from node 23, but does not give support to any other node in the system.

However this innovation doesn't last long. Nodes 20 and 23, on whom the well being of node 25 depends, are unprotected. Since they have the least possible value of X_i, namely, zero, they can be eliminated at subsquent graph updates, and their replacements in general do not feed into node 25. Sooner or later selection picks 23 for replacement, and then 25 also gets depopulated. By $n = 2853$ (Fig. 2d) node 25 and all others that were populated at $n = 79$, have joined the ranks of the unpopulated. Node 25 (and others of its ilk) are doomed because they are 'uncaring winners': they do not feed into (i.e., do not catalyse) the nodes upon whom their well being depends. That is why when there are no closed cycles, all structures are transitory; the graph remains random. Of the 1070 innovations, 115 were of this type.

1.6.2
Birth of an organization: Cooperation begets stability

At $n = 2853$ node 90 is an unpopulated node (Fig. 2d). It is eliminated at $n = 2854$ and the new node 90 forms a 2-cycle with node 26 (Fig 2e). This is the first time (and the only time in this run) an innovation forms a closed cycle in a graph that previously had no cycles. A closed cycle between 2 nodes is the simplest *cooperative* graph theoretical structure possible. Nodes 26 and 90 help each other's population grow; together they form a self-replicating system. Their populations grow much faster than other nodes in the graph; it turns out that in the attractor for this graph only nodes 26 and 90 are populated, with all other nodes having $X_i = 0$ ([16, 17]; see also Appendix A). Because node 90 is populated in the new attractor this constitutes an innovation. However, unlike the previous innovations, this one has a greater staying power, because nodes 26 and 90 do well *collectively*. At the next graph update *both* nodes 26 and 90 will be immune to removal since one of the other nodes with zero X_i will be removed. Notice that nodes 26 and 90 do not depend on nodes which are part of the least fit set (those with the least value of X_i). The cycle has all the catalysts it needs for for the survival of each of its constituents. This property is true not just for cyclic subgraphs but for a more general cooperative structure, the *autocatalytic set* (ACS).

An ACS is a set of species which contains a catalyst for each species in the set [18, 11, 19]. In the context of the present model we can define an ACS to be a subgraph each of whose nodes has at least one incoming link from a node of the same subgraph. While ACSs need not be cycles, they must contain at least one cycle. If the graph has (one or more) ACSs, one can show that the set of populated nodes ($X_i > 0$) must be an ACS, which we call the dominant ACS[4] [5, 16]. (In Figs. 2e-s, the subgraph of the grey and black nodes is the dominant ACS.) Therefore none of the nodes of the dominant ACS can be hit in the next graph update as long as there is any node outside it. In other words, the collective well being of all the constituents of the dominant ACS, ensured by cooperation inherent within its structure, is responsible for the ACS's relative robustness and hence longevity.

In societies, this kind of event is akin to the birth of an organization wherein two or more agents improve upon their performance by entering into an explicit cooperation. A booming new township or industrial district perhaps can be analysed in terms of a closure of certain feedback loops. In prehistory, the appearance of tools that could be used to improve other tools may be regarded as events of this kind which probably unleashed a lot of artifact building. On the prebiotic earth one can speculate that the appearance of a small ACS might have triggered processes that eventually led to the emergence of life [14].

If there is no ACS in the graph then the largest eigenvalue of the adjacency matrix of the graph, λ_1 is zero. If there is an ACS then $\lambda_1 \geq 1$[5] [5, 16]. In Fig 1b, λ_1 jumped from zero to one when the first ACS was created at $n = 2854$.

1.6.3
Expansion of the organization at its periphery: Incremental innovations

Consider Figs. 2f and g. Node 3, which is unpopulated at $n = 3021$, gets an incoming link from node 90 and an outgoing link to node 25 at $n = 3022$ which results in three nodes adding onto the dominant ACS. Node 3 is populated in the new attractor and hence this is an innovation. This innovation has expanded the 'periphery' of the organization, defined below.

Every dominant ACS is a union of one or more 'simple ACSs' each of which have a substructure consisting of a 'core' and 'periphery'. E.g., the dominant ACS in Fig. 1g has one simple ACS and in Fig. 1k it has two. For every simple ACS there exists a maximal subgraph, called the *core* of that ACS, from each of whose nodes there is a directed path to every node of that ACS. The rest of that ACS is termed its *periphery*. In Figs. 2e-s, the core is coloured black, and the periphery grey. Thus in Fig. 2g, the 2-cycle of nodes 26 and 90 is the core of the dominant ACS, and the chain of nodes 3, 25 and 18 along with the incoming link to node 3 from 26 constitutes its periphery. The core of a simple ACS is necessarily an *irreducible subgraph*. An irreducible subgraph is one that contains a directed path from each of its nodes to every other of its nodes [20]. When the dominant ACS consists of more than one simple ACSs, its core is the union of their cores, and its periphery the union of their peripheries.

4) See *Dominant ACS of a graph* in Appendix A.
5) See *Graph-theoretic properties of ACSs* in Appendix A.

Note that the periphery nodes by definition do not feed back into the core; in this sense they are parasites that draw sustenance from the core. The core, by virtue of its irreducible property (positive feedback loops within its structure, or cooperativity), is self-sustaining, and also supports the periphery. The λ_1 of the ACS is determined solely by the structure of its core [15, 12].

The innovation at $n = 3022$, one of 907 such innovations in this run, is an 'incremental' one in the sense that it does not change the core (and hence does not change λ_1). However such incremental innovations set the stage for major transformations later on. The ability of a core to tolerate parasites can be a boon or a bane, as we will see below.

1.6.4 Growth of the core of the organization: Parasites become symbionts

Another kind of innovation that occurs in the growth phase is illustrated in Figs 2l and m. In Fig. 2l, the dominant ACS has two disjoint components. One component, consisting of nodes 41 and 98, is just a 2-cycle without any periphery. The other component has a 2-cycle (nodes 26 and 90) as its core that supports a large periphery. Node 39 in Fig. 2l is eliminated at $n = 3489$. The new node 39 (Fig. 2m) gets an incoming link from the periphery of the larger component of the dominant ACS and an outgoing link to the core of the same ACS. This results in expansion of the core, with several nodes getting added to it at once and λ_1 increasing. This example illustrates two distinct processes:

(i) This innovation co-opts a portion of the parasitic periphery into the core. This strengthens cooperation: 26 contributes to the well being of 90 (and hence to its own well being) along two paths in Fig. 2m instead of only one in Fig. 2l. This is reflected in the increase of λ_1; $\lambda_1 = 1.15$ and 1 for Figs. 2m and 2l respectively. The larger the periphery, the greater is the probability of such core-enhancing innovations. This innovation is an example of how tolerance and support of a parasitic periphery pays off for the ACS. Part of the parasitic periphery turns symbiont. Note that this innovation builds upon the structure generated by previous incremental innovations. In Fig 1b each rise in λ_1 indicates an enlargement of the core [15, 12]. There are 40 such events in this run. As a result of a series of such innovations which add to the core and periphery, the dominant ACS eventually grows to span the entire graph at $n = 3880$, Fig 2n, and the system enters the 'organized phase'.

(ii) This example also highlights the competition between different ACSs. The 2-cycle of nodes 41 and 98 was populated in Fig. 2l, but is unpopulated in Fig. 2m. Since the core of the other ACS becomes stronger than this two cycle, the latter is driven out of business.

1.6.5
Core-shift 1: Takeover by a new competitor

Interestingly, the same cycle of nodes 41 and 98 that is driven out of business at $n = 3489$, had earlier (when it first arose at $n = 3387$) driven the 2-cycle of nodes 26 and 90 out of business. Upto $n = 3386$ (Fig. 2h), the latter 2-cycle was the only cycle in the graph. At $n = 3387$ node 41 was replaced and formed a new 2-cycle with node 98 (Fig. 2i). Note that at $n = 3387$ only the new 2-cycle is populated; all the nodes of the ACS that was dominant at the previous time step (including its core) are unpopulated. We call such an event, where there is no overlap between the old and the new cores, a *core shift* (a precise definition is given in ref. [15]). This innovation is an example of how a new competitor takes over.

Why does the new 2-cycle drive the old one to extinction? The reason is that the new 2-cycle is downstream of the old one (node 41 has also acquired an incoming link from node 39; thus there exists a directed path from the old cycle to the new one, but none from the new to the old). Both 2-cycles have the same intrinsic strength, but the new 2-cycle does better than the old because it draws sustenance from the latter without feeding back. In general if the graph contains two non-overlapping irreducible subgraphs A and B, let $\lambda_1(A)$ and $\lambda_1(B)$ be the largest eigenvalues of the submatrices corresponding to A and B. If $\lambda_1(A) > \lambda_1(B)$, then A wins (i.e., in the attractor of (1.1), nodes of A and all nodes downstream of A are populated), and nodes of B are populated if B is downstream of A and unpopulated otherwise. When $\lambda_1(A) = \lambda_1(B)$, then if A and B are disconnected, both are populated, and if one of them is downstream of the other, it wins and the other is unpopulated [15, 12]. At $n = 3387$ the latter situation applies (the λ_1 of both cycles is 1, but one is downstream of the other; the downstream cycle wins at the expense of the upstream one). Examples of new competitors taking over because their λ_1 is higher than that of the existing ACS are also seen in the model.

In the displayed run, two core shift of this kind occurred. The first was at $n = 3387$ which has been discussed above. One more occurred at $n = 6062$ which was of an identical type with a new downstream 2-cycle driving the old 2-cycle to extinction. Both these events resulted in a sharp drop in s_1 (Fig. 1b). A core-shifting innovation is a traumatic event for the old core and its periphery. This is reminiscent of the demise of the horse-drawn carriage industry upon the appearance of the automobile, or the decimation of anaerobic species upon the advent of aerobic ones.

At $n = 3403$ (Fig 2k) an interesting event (that is not an innovation) happens. Node 38 is hit and the new node 38 has no incoming link. This cuts the connection that existed earlier (see Fig. 2j) between the cycle 98-41 and the cycle 26-90. The graph now has two disjoint ACSs with the same λ_1 (see Fig 2k). As mentioned above, in such a situation both ACSs coexist; the cycle

26-90 and all nodes dependent on it once again become populated. Thus the old core has staged a 'come-back' at $n = 3402$, levelling with its competitor. As we saw in the previous subsection, at $n = 3489$ the descendant of this organization strengthens its core and in fact drives its competitor out of business (this time permanently).

It is interesting that node 38, though unpopulated still plays an important role in deciding the structure of the dominant ACS. It is purely a matter of chance that the core of the old ACS, the cycle 26-90, did not get hit before node 38. (All nodes with $X_i = 0$ have an equal probability of being replaced in the model.) If it had been destroyed between $n = 3387$ and 3402, then nothing interesting would have happened when node 38 was removed at $n = 3403$. In that case the new competitor would have won. Examples of that are also seen in the runs. In either case an ACS survives and expands until it spans the entire graph. It is worth noting that while overall behaviour like the growth of ACSs (including their average time scale of growth) is predictable, the details are shaped by historical accidents.

1.6.6 Core-shift 2: Takeover by a dormant innovation

A different kind of innovation occurs at $n = 4696$. At the previous time step, node 36 is the least populated (Fig. 2o). The new node 36 forms a two cycle with node 74 (Figs. 2p). This 2-cycle remains part of the periphery since it does not feed back into the core; this is an incremental innovation at this time since it does not enhance λ_1. However, because it generates a structure that is intrinsically self-sustaining (a 2-cycle) this innocuous innovation is capable of having a dramatic impact in the future.

At $n = 5041$, Fig 2q, the core has shrunk to 5 nodes (the reasons for this decline are briefly discussed later). The 36-74 cycle survives in the periphery of the ACS. Now it happens that node 85 is one of those with the least X_i and gets picked for removal at $n = 5042$. Thus of the old core only the 2-cycle 26-90 is left. But this is now upstream from another 2-cycle 74-36 (see Fig 2r). This is the same kind of structure as discussed above, with one cycle downstream from another. The downstream cycle and its periphery wins; the upstream cycle and all other nodes downstream from it except nodes 36, 74 and 11 are driven to extinction. This event is also a core shift and is accompanied by a huge crash in the s_1 value (see Fig. 1b). This kind of an event is what we call a 'takeover by a dormant innovation' [12]. The innovation 36-74 occurred at $n = 4696$. It lay dormant until $n = 5042$ when the old core had become sufficiently weakened so that this dormant innovation could take over as the new core.

In this run 5 of the 1070 innovations were dormant innovations. Of them only the one at $n = 4696$ later caused a core shift of the type discussed above. The others remained as incremental innovations.

At $n = 8233$ a 'complete crash' occurs. The core is a simple 3-cycle (Fig. 2s) at $n = 8232$ and node 50 is hit, completely destroying the ACS. λ_1 drops to zero accompanied by a large crash in s_1. Within $O(s)$ time steps most nodes are hit and replaced and the graph has become random like the initial graph. The resemblance between the initial graph at $n = 1$ (Fig 2a) and the graph at $n = 10000$ (Fig 2u) is evident. This event is not an innovation but rather the elimination of a 'keystone species' [12].

1.7 Recognizing innovations: a structural classification

The 6 categories of innovations discussed above occur in all the runs of the model and their qualitative effects are the same as described above. The above description was broadly 'chronological'. We now describe these innovations structurally. Such a description allows each type of innovation to be recognized the moment it appears; one does not have to wait for its consequences to be played out. The structural recognition in fact allows us to predict qualitatively the kinds of impact it can have on the system. A mathematical classification of innovations is given in Appendix B; the description here is a plain English account of that (with some loss of precision).

As is evident from the discussion above, positive feedback loops or cooperative structures in the graph crucially affect the dynamics. The character of an innovation will also depend upon its relationship with previously existing feedback loops and the new feedback loops it creates, if any. Structurally an 'irreducible subgraph' captures the notion of feedback in a directed graph. By definition, since there exists a directed path (in both directions) between every pair of nodes belonging to an irreducible subgraph, each node 'exerts an influence' on the other (albeit possibly through other intermediaries).

Thus the first major classification depends on whether the new node creates a new cycle and hence a new irreducible subgraph, or not. One way of determining whether it does so is to identify the nodes 'downstream' of the new node (namely those to which there is a directed path from this node) and those that are 'upstream' (from which there is a directed path to this node). If the intersection of these two sets is empty the new node has not created any new irreducible subgraph, otherwise it has.

A. *Innovations in which the new node does not create any new cycles and hence no new irreducible subgraph is created.* These innovations will have a relatively minor impact on the system. There are two subclasses here which depend upon the context: whether an irreducible subgraph already exists somewhere else in the graph or not.

A1. *Before the innovation, the graph does not contain an irreducible subgraph.* Then the innovation is a shortlived one discussed in section 6.1 (Figs. 2b,c). There is no ACS before or after the innovation. The largest eigenvalue λ_1 of the adjacency matrix of the graph being zero both before and after such an innovation is a necessary and sufficient condition for it to be in this class. Such an innovation is doomed to die when the first ACS arises in the graph for the reasons discussed in the previous section.

A2. *Before the innovation an irreducible subgraph already exists in the graph.* One can show that such an innovation simply adds to the periphery of the existing dominant ACS, leaving the core unchanged. Here the new node gets a non-zero X_k because it receives an incoming link from one of the nodes of the existing dominant ACS; it has effectively latched on to the latter like a parasite. This is an incremental innovation (section 6.3, Figs. 2f,g). It has a relatively minor impact on the system at the time it appears. Since it does not modify the core, the ratios of the X_i values of the core nodes remain unchanged. However, it does eat up some resources (since $X_k > 0$) and causes an overall decline in the X_i values of the core nodes. λ_1 is nonzero and does not change in such an innovation.

B. *Innovations that do create some new cycle.* Thus a new irreducible subgraph gets generated. Because these innovations create new feedback loops, they have a potentially greater impact. Their classification depends upon whether or not they modify the core and the extent of the modification caused; this is directly correlated with their immediate impact.

B1. *The new cycles do not modify the existing core.* If the new irreducible subgraph is disjoint from the existing core and its intrinsic λ_1 less than that of the core, then the new irreducible subgraph will not modify the existing core but will become part of the periphery. Like incremental innovations, such innovations cause an overall decline in the X_i values of the core nodes but do not disturb their ratios and the value of λ_1. However, they differ from incremental innovations in that the new irreducible subgraph has self-sustaining capabilities. Thus in the event of a later weakening of the core (through elimination of some core nodes), these innovations have the potential of causing a core-shift wherein the irreducible graph generated in the innovation becomes the new core. At that point it would typically cause a major crash in the number of populated species, as the old core and all its periphery that is not supported by the new core would become depopulated. Such innovations are the dormant innovations (section 6.6, Figs. 2o,p). Note that not all dormant innovations cause core-shifts. Most in fact play the same role as incremental innovations.

B2. *Innovations that modify the existing core.* If the new node is part of the new core, the core has been modified. The classification of such innovations depends on the kind of core that exists before and the nature of the modification.

B21. *The existing core is non-empty,* i.e., an ACS already exists before the innovation in question arrives.

B211. *The innovation strengthens the existing core.* In this case the new node receives an incoming link from the existing dominant ACS and has an outgoing link to the existing core. The existing core nodes get additional positive feedback, and λ_1 increases. Such an event can cause some members of the parasitic periphery to be co-opted into the core. These are the core-enhancing innovations discussed in section 6.4 (Figs. 2l,m).

B212. *The new irreducible subgraph is disjoint from the existing core and 'stronger' than it.* 'Stronger' means that the intrinsic λ_1 of the new irreducible graph is greater than or equal to the λ_1 of the existing core, and in the case of equality it is downstream from the existing core. Then it will destabilize the existing core and become the new core itself, causing a core-shift. The takeovers by new competitors, discussed in section 6.5 (Figs. 2h,i) belong to this class.

B22. *The existing core is empty*, i.e., no ACS exists before the arrival of this innovation. Then the new irreducible graph is the core of the new ACS that is created at this time. This event is the beginning of a self-organizing phase of the system. This is the birth of an organization discussed in section 6.2 (Figs. 2d,e). This is easily recognized graph theoretically as λ_1 jumps from zero to a positive value.

Note that the 'recognition' of the class of an innovation is contingent upon knowing graph theoretic features like the core, periphery, λ_1, and being able to determine the irreducible graph created by the innovation.

The above rules are an analytic classification of all innovations in the model, irrespective of values of the parameters p and s. Note, however, that their relative frequencies depend upon the parameters. In particular, innovations of class A require the new node to have at least one link (an incoming one) and class B require at least two links (an incoming and an outgoing one). Thus as the connection probability p declines, for fixed s, the latter innovations (the more consequential ones) become less likely.

1.8 Some possible general lessons

In this model, due to the simplicity of the population dynamics, it is possible to make an analytic connection between the graph structure produced by various innovations and their subsequent effect on the short and long term dynamics of the system. In addition, we are able to completely enumerate the different types of innovations and classify them purely on the basis of their graph structure. Identifying innovations and understanding their effects is much more complicated in real world processes in both biological and social systems. Nevertheless, the close parallel between the qualititive categories of innovation we find in our model and real world examples means that there may be some lessons to be learnt from this simple mathematical model.

One broad conclusion is that in order to guess what might be an innovation, we need an understanding of how the patterns of connectivity influence system dynamics and vice versa. The inventor of a new product or a venture capitalist asks: what inputs will be needed, and whose needs will the product connect to? Given these potential linkages in the context of other existing nodes and links, what flows will actually be generated along the new links? How will these new flows impact the generation other new nodes and links and the death of existing ones and how that will feed back into the flows again? The detailed rules of this dynamics are system dependent, but presumably successful entrepreneurs have an intuitive understanding of this very dynamics.

In our model, as in real processes, there are innovations which have an immediate impact on the dynamics of the system (e.g., the creation of the first ACS and core-shifting innovations) and ones which have little or no immediate impact. Innovation in real processes analogous to the former are probably easier to identify because they cause the dynamics of the system to immediately change dramatically (in this model, triggering a new round of self-organized growth around a new ACS). Of the latter, the most interesting innovations are the ones which eventually do have a large impact on the dynamics: the dormant innovations. In this model dormant innovations sometimes lead to a dramatic change in the dynamics of the system at a later time. This suggests that in real world processes too it might be important, when observing a sudden change in the dynamics, to examine innovations which occurred much before the change. Of course, in the model and in real processes, there are innovations which have nothing to do with any later change in the dynamics. In real processes it would be very difficult to distinguish such innovations from dormant innovations which do cause a significant impact on the dynamics. The key feature distinguishing a dormant innovation from incremental innovations in this model is that a dormant innovation creates an irreducible structure which can later become the core of the graph.

This suggests that in real world processes it might be useful to find an analogy of the core and periphery of the system and then focus on innovations or processes which alter the core or create structures which could become the core. In the present model, it is possible to define the core in a purely graph theoretic manner. In real systems it might be necessary to define the core in terms of the dynamics. One possible generalization is based on the observation that removal of a core node causes the population growth rate to reduce (due to the reduction of λ_1) while the removal of a periphery node leaves λ_1 unchanged. This could be used as an algorithmic way of identifying core nodes or species in more complex mathematical models, or in real systems where such testing is possible.

1.9 Discussion

As in real systems, the model involves an interplay between the force of selection that weeds out underperforming nodes, the influx of novelty that brings in new nodes and links, and an internal (population) dynamics that depends upon the mutual interactions. In an environment of non-autocatalytic structures, a small ACS is very successful and drives the other nodes to the status of 'have-nots' ($X_i = 0$). The latter are eliminated one by one, and if their replacements 'latch on' to the ACS, they survive, else they suffer the same fate. The ACS 'succeeds' spectacularly: eventually all the nodes join it. But this sets the stage for enhanced internal competition between the members of the ACS. Before the ACS spanned the graph, only have-nots, nodes outside the dominant ACS, were eliminated. After spanning the eliminated node must be one of the 'haves', a member of the ACS (whichever has the least X_i). This internal competition weakens the core and enhances the probability of collapse due to core-transforming innovations or elimination of keystone species. Thus the ACS's very success creates the circumstances that bring about its destruction[6]. Both its success, and a good part of its destruction, is due to innovations (see also [12]).

6) For related discussion of discontinuous transitions in other systems, see [21, 22, 23]

It is of course true that we can describe the behaviour of the system in terms of attractors of the dynamics as a function of the graph without recourse to the word 'innovation'. The advantage in introducing the notion of innovation as defined above is that it captures a useful property of the dynamics in terms of which many features can be readily described. Further, we hope that the examples discussed above make out a reasonable case that this notion of innovation is sufficiently close (as close as is possible in an idealized model such as this) to the real thing to help in discussions of the latter.

In the present model, the links of the new node are chosen randomly from a fixed probability distribution. This might be appropriate for the prebiotic chemical scenario for which the model was constructed, but is less appropriate for biological systems and even less for social systems. While there is always some stochasticity, in these systems the generation of novelty is conditioned by the existing context, and in social systems also by the intentionality of the actors. Thus the ensemble of choices from which the novelty is drawn also evolves with the system. This feedback from the recent history of system states to the ensemble of system perturbations though not implemented in the present version of the model, certainly deserves future investigation.

Appendix A: Definitions and Proofs

In this Appendix we collect some useful facts about the model. These and other properties can be found in [5, 16, 17, 13]

Derivation of equation 1.1

Let $i \in \{1, \ldots, s\}$ denote a chemical (or molecular) species in a well-stirred chemical reactor. Molecules can react with one another in various ways; we focus on only one aspect of their interactions: catalysis. The catalytic interactions can be described by a directed graph with s nodes. The nodes represent the s species and the existence of a link from node j to node i means that species j is a catalyst for the production of species i. In terms of the adjacency matrix, $C = (c_{ij})$ of this graph, c_{ij} is set to unity if j is a catalyst of i and is set to zero otherwise. The operational meaning of catalysis is as follows:

Each species i will have an associated non-negative population y_i in the reactor that changes with time. Let species j catalyze the ligation of reactants A and B to form the species i, $A + B \xrightarrow{j} i$. Assuming that the rate of this catalyzed reaction is given by the Michaelis-Menten theory of enzyme catalysis, $\dot{y}_i = V_{max} ab \frac{y_j}{K_M + y_j}$ [24], where a, b are the reactant concentrations, and V_{max} and K_M are constants that characterize the reaction. If the Michaelis constant K_M

is very large this can be approximated as $\dot{y}_i \propto y_j ab$. Combining the rates of the spontaneous and catalyzed reactions and also putting in a dilution flux ϕ, the rate of growth of species i is given by $\dot{y}_i = k(1 + \nu y_j)ab - \phi y_i$, where k is the rate constant for the spontaneous reaction, and ν is the catalytic efficiency. Assuming the catalyzed reaction is much faster than the spontaneous reaction, and that the concentrations of the reactants are non-zero and fixed, the rate equation becomes $\dot{y}_i = K y_j - \phi y_i$, where K is a constant. In general because species i can have multiple catalysts, $\dot{y}_i = \sum_{j=1}^{s} K_{ij} y_j - \phi y_i$, with $K_{ij} \sim c_{ij}$. We make the further idealization $K_{ij} = c_{ij}$ giving:

$$\dot{y}_i = \sum_{j=1}^{s} c_{ij} y_j - \phi y_i. \tag{1.2}$$

The relative population of species i is by definition $x_i \equiv y_i / \sum_{j=1}^{s} y_j$. As $0 \leq x_i \leq 1, \sum_{i=1}^{s} x_i = 1$, $\mathbf{x} \equiv (x_1, \ldots, x_s)^T \in J$. Taking the time derivative of x_i and using (1.2) it is easy to see that $\dot{x}_i$ is given by (1.1). Note that the ϕ term, present in (1.2), cancels out and is absent in (1.1).

The attractor of equation 1.1

A graph described by an adjacency matrix, C, has an eigenvalue $\lambda_1(C)$ which is a real, positive number that is greater than or equal to the modulus of all other eigenvalues. This follows from the Perron-Frobenius theorem[20] and this eigenvalue is called the Perron-Frobenius eigenvalue of C.

The attractor $\mathbf{X}$ *of equation 1.1 is an eigenvector of* C *with eigenvalue* $\lambda_1(C)$.
Since (1.1) does not depend on ϕ, we can set $\phi = 0$ in (1.2) without loss of generality for studying the attractors of (1.1). For fixed C the general solution of (1.2) is $\mathbf{y}(t) = e^{Ct}\mathbf{y}(0)$, where $\mathbf{y}$ denotes the s dimensional column vector of populations. It is evident that if $\mathbf{y}^\lambda \equiv (y_1^\lambda, \ldots, y_s^\lambda)$ viewed as a column vector is a right eigenvector of C with eigenvalue λ, then $\mathbf{x}^\lambda \equiv \mathbf{y}^\lambda / \sum_i^s y_i^\lambda$ is a fixed point of (1.1). Let λ_1 denote the eigenvalue of C which has the largest real part; it is clear that $\mathbf{x}^{\lambda_1}$ is an attractor of (1.1). By the theorem of Perron-Frobenius for non-negative matrices [20] λ_1 is real and ≥ 0 and there exists an eigenvector $\mathbf{x}^{\lambda_1}$ with $x_i \geq 0$. If λ_1 is nondegenerate, $\mathbf{x}^{\lambda_1}$ is the unique asymptotically stable attractor of (1.1), $\mathbf{x}^{\lambda_1} = (X_1, \ldots, X_s)$.

The attractor of equation 1.1 when there are no cycles

For any graph with no cycles, in the attractor only the nodes at the ends of the longest paths have non-zero X_i. *All other nodes have zero* X_i.
Consider a graph consisting only of a linear chain of $r + 1$ nodes, with r links, pointing from node 1 to node 2, node 2 to 3, etc. The node 1 (to which there

is no incoming link) has a constant population y_1 because the r.h.s of (1.2) vanishes for $i = 1$ (taking $\phi = 0$). For node 2, we get $\dot{y}_2 = y_1$, hence $y_2(t) = y_2(0) + y_1 t \sim t$ for large t. Similarly, it can be seen that y_k grows as t^{k-1}. In general, it is clear that for a graph with no cycles, $y_i \sim t^r$ for large t (when $\phi = 0$), where r is the length of the longest path terminating at node i. Thus, nodes with the largest r dominate for sufficiently large t. Because the dynamics (1.1) does not depend upon the choice of ϕ, $X_i = 0$ for all i except the nodes at which the longest paths in the graph terminate.

Graph-theoretic properties of ACSs

(i) An ACS must contain a closed walk.
(ii) If a graph, C, has no closed walk then $\lambda_1(C) = 0$.
(iii) If a graph, C, has a closed walk then $\lambda_1(C) \geq 1$. Consequently,
(iv) If a graph C has no ACS then $\lambda_1(C) = 0$.
(v) If a graph C has an ACS then $\lambda_1(C) \geq 1$.

(i) Let A be the adjacency matrix of a graph that is an ACS. Then by definition, every row of A has at least one non-zero entry. Construct A' by removing, from each row of A, all non-zero entries except one that can be chosen arbitrarily. Thus A' has exactly one non-zero entry in each row. Clearly the column vector $\mathbf{x} = (1, 1, \ldots, 1)^T$ is an eigenvector of A' with eigenvalue 1 and hence $\lambda_1(A') \geq 1$. Proposition 2.1 therefore implies that A' contains a closed walk. Because the construction of A' from A involved only removal of some links, it follows that A must also contain a closed walk.

(ii) If a graph has no closed walk then all walks are of finite length. Let the length of the longest walk of the graph be denoted r. If C is the adjacency matrix of a graph then $(C^k)_{ij}$ equals the number of distinct walks of length k from node j to node i. Clearly $C^m = 0$ for $m > r$. Therefore all eigenvalues of C^m are zero. If λ_i are the eigenvalues of C then λ_i^k are the eigenvalues of C^k. Hence, all eigenvalues of C are zero, which implies $\lambda_1 = 0$. This proof was supplied by V. S. Borkar.

(iii) If a graph has a closed walk then there is some node i that has at least one closed walk to itself, i.e. $(C^k)_{ii} \geq 1$, for infinitely many values of k. Because the trace of a matrix equals the sum of the eigenvalues of the matrix, $\sum_{i=1}^{s}(C^k)_{ii} = \sum_{i=1}^{s} \lambda_i^k$, where λ_i are the eigenvalues of C. Thus, $\sum_{i=1}^{s} \lambda_i^k \geq 1$, for infinitely many values of k. This is only possible if one of the eigenvalues λ_i has a modulus ≥ 1. By the Perron-Frobenius theorem, λ_1 is the eigenvalue with the largest modulus, hence $\lambda_1 \geq 1$. This proof was supplied by R. Hariharan.

(iv) and (v) follow from the above.

Dominant ACS of a graph

If a graph has (one or more) ACSs, i.e., $\lambda_1 \geq 1$, then the subgraph corresponding to the set of nodes i for which $X_i > 0$ is an ACS. Renumber the nodes of the graph so that $x_i > 0$ only for $i = 1, \ldots, k$. Let C be the adjacency matrix of this graph. Since **X** is an eigenvector of the matrix C, with eigenvalue λ_1, we have $\sum_{j=1}^{s} c_{ij} X_j = \lambda_1 X_i \Rightarrow \sum_{j=1}^{k} c_{ij} X_j = \lambda_1 X_i$. Since $X_i > 0$ only for $i = 1, \ldots, k$ it follows that for each $i \in \{1, \ldots, k\}$ there exists a j such that $c_{ij} > 0$. Hence the $k \times k$ submatrix $C' \equiv (c_{ij})$, $i, j = 1, \ldots, k$ has at least one non-zero entry in each row. Thus each node of the subgraph corresponding to this submatrix has an incoming link from one of the other nodes in the subgraph. Hence the subgraph is an ACS. We call this subgraph the 'dominant ACS' of the graph.

Timescales for appearance and growth of the dominant ACS.

The probability for an ACS to be formed at some graph update in a graph which has no cycles, can be closely approximated by the probability of a 2-cycle (the simplest ACS with 1-cycles being disallowed) forming by chance, which is $p^2 s$ (= the probability that in the row and column corresponding to the replaced node in C, any matrix element and its transpose are both assigned unity). Thus, the 'average time of appearance' of an ACS is $\tau_a = 1/p^2 s$, and the distribution of times of appearance is $P(n_a) = p^2 s (1 - p^2 s)^{n_a - 1}$. This approximation is better for small p.

Assuming that possibility of a new node forming a second ACS is rare enough to neglect, and that the dominant ACS grows by adding a single node at a time, one can estimate the time required for it to span the entire graph. Let the dominant ACS consist of $s_1(n)$ nodes at time n. The probability that the new node gets an incoming link from the dominant ACS and hence joins it is ps_1. Thus in Δn graph updates, the dominant ACS will grow, on average, by $\Delta s_1 = ps_1 \Delta n$ nodes. Therefore $s_1(n) = s_1(n_a)\exp((n - n_a)/\tau_g)$, where $\tau_g = 1/p$, n_a is the time of appearance of the first ACS and $s_1(n_a)$ is the size of the first ACS. Thus s_1 is expected to grow exponentially with a characteristic timescale $\tau_g = 1/p$. The time taken from the appearance of the ACS to its spanning is $\tau_g \ln(s/s_1(n_a))$.

Appendix B: Graph Theoretic Classification of Innovations

In the main text we defined an innovation to be the new structure created by the addition of a new node, when the new node has a non-zero population in the new attractor. Here, we present a graph-theoretic hierarchical classification of innovations (see Fig. 3). At the bottom of this hierarchy we recover the six categories of innovations described in the main text.

Some notation: We need to distinguish between two graphs, one just before the new node is inserted, and one just after. We denote them by C_i and C_f respectively, and their cores by Q_i and Q_f. Note that a graph update event consists of two parts – the deletion of a node and the addition of one. C_i is the graph after the node is deleted and before the new node is inserted. The graph before the deletion will be denoted C_0; Q_0 will denote its core[7]. If a graph has no ACS, its core is the null set.

The links of the new node may be such that new cycles arise in the graph (that were absent in C_i but are present in C_f). In this case the new node is part of a new irreducible subgraph that has arisen in the graph. N will denote the maximal irreducible subgraph which includes the new node. If the new node does not create new cycles, $N = \Phi$. If $N \neq \Phi$, then N will either be disjoint from Q_f or will include Q_f (it cannot partially overlap with Q_f because of its maximal character). The structure of N and its relationship with the core before and after the addition determines the nature of the innovation. With this notation all innovations can be grouped into two classes:

A. Innovations that do not create new cycles, $N = \Phi$. This implies $Q_f = Q_i$ because no new irreducible structure has appeared and therefore the core of the graph, if it exists, is unchanged.

B. Innovations that do create new cycles, $N \neq \Phi$. This implies $Q_f \neq \Phi$ because if a new irreducible structure is created then the new graph has at least one ACS and therefore a non-empty core.

Class A can be further decomposed into two classes:

A1. $Q_i = Q_f = \Phi$. In other words, the graph has no cycles both before and after the innovation. This corresponds to shortlived innovations discussed in section 6.1 (Figs. 2b,c).

7) Most of the time the deleted node (being the one with the least relative population) is outside the dominant ACS of C_0 or in its periphery. Thus, in most cases the core is unchanged by the deletion: $Q_i = Q_0$. However sometimes the deleted node belongs to Q_0. In that case $Q_i \neq Q_0$. In most such cases, Q_i is a proper subset of Q_0. In very few (but important) cases, $Q_i \cap Q_0 = \Phi$ (the null set). In these latter cases, the deleted node is a 'keystone node' [15]; its removal results in a 'core shift'.

A2. $Q_i = Q_f \neq \Phi$. In other words, the graph had an ACS before the innovation, and its core was not modified by the innovation. This corresponds to incremental innovations discussed in section 6.3 (Figs. 2f,g).

Class B of innovations can also be divided into two subclasses:

B1. $N \neq Q_f$. If the new irreducible structure is not the core of the new graph, then N must be disjoint from Q_f. This can only be the case if the old core has not been modified by the innovation. Therefore $N \neq Q_f$ necessarily implies that $Q_f = Q_i$. This corresponds to dormant innovations discussed in section 6.6 (Figs. 2o,p).

B2. $N = Q_f$, i.e., the innovation becomes the new core after the graph update. This is the situation where the core is transformed due to the innovation.

The 'core-transforming theorem' [12, 17, 13] states that an innovation of type B2 occurs whenever either of the following conditions are true:

(a) $\lambda_1(N) > \lambda_1(Q_i)$ or,

(b) $\lambda_1(N) = \lambda_1(Q_i)$ and N is downstream of Q_i.

Class B2 can be subdivided as follows:

B21. $Q_i \neq \Phi$, i.e., the graph contained an ACS before the innovation. In this case an existing core is modified by the innovation.

B22. $Q_i = \Phi$, i.e., the graph had no ACS before the innovation. Thus, this kind of innovation creates an ACS in the graph. It corresponds to the birth of a organization discussed in section 6.2 (Figs. 2d,e).

Finally, class B21 can be subdivided:

B211. $Q_i \subset Q_f$. When the new core contains the old core as a subset we get an innovation that causes the growth of the core, discussed in section 6.4 (Figs. 2l,m).

B212. Q_i and Q_f are disjoint (Note that it is not possible for Q_i and Q_f to partially overlap, else they would form one big irreducible set which would then be the core of the new graph and Q_i would be a subset of Q_f). This is an innovation where a core-shift is caused due to a takeover by a new competitor, discussed in section 6.5 (Figs. 2h,i).

Note that each branching above is into mutually exclusive and exhaustive classes. This classification is completely general and applicable to all runs of the system. Fig. 3 shows the hierarchy obtained using this classification.

Acknowledgement
S. J. would like to thank John Padgett for discussions.

References

1 Schumpeter, J. A. (1934) *The Theory of Economic Development* (Harvard University Press, Cambridge, MA, USA).

2 Rogers, E. M. (1995) *The Diffusion of Innovations, 4th edition* (The Free Press, New York).

3 Falkowski, P. G. (2006) *Science* **311**, 1724-1725.

4 Bornholdt, S & Schuster, H. G., eds. (2003) *Handbook of Graphs and Networks: From the Genome to the Internet*, (Wiley-VCH, Weinheim).

5 Jain, S. & Krishna, S. (1998) *Phys. Rev. Lett.* **81**, 5684-5687.

6 Dyson, F. (1985) *Origins of Life* (Cambridge Univ. Press).

7 Kauffman, S. A. (1993) *The Origins of Order* (Oxford Univ. Press).

8 Bagley, R. J., Farmer, J. D. & Fontana, W. (1991) in *Artificial Life II*, eds. Langton, C. G., Taylor, C., Farmer, J. D. & Rasmussen, S. (Addison Wesley, Redwood City), pp. 141-158.

9 Fontana, W. & Buss, L. (1994) *Bull. Math. Biol.* **56**, 1-64.

10 Bak, P. & Sneppen, K. (1993) *Phys. Rev. Lett.* **71**, 4083-4086.

11 Kauffman, S.A. (1971) *J. Cybernetics* **1**, 71-96.

12 Jain, S. and Krishna, S. (2002) *Proc. Natl. Acad. Sci. (USA)* **99**, 2055-2060.

13 Krishna, S. (2003) Ph. D. Thesis, *http://www.arXiv.org/abs/nlin.AO/0403050*.

14 Jain, S. & Krishna, S. (2001) *Proc. Natl. Acad. Sci. (USA)* **98**, 543-547.

15 Jain, S. & Krishna, S. (2002) *Phys. Rev. E* **65**, 026103.

16 Jain, S. & Krishna, S. (1999) *Computer Physics Comm.* **121-122**, 116-121.

17 Jain, S.& Krishna, S. (2003) in *Handbook of Graphs and Networks*, eds. Bornholdt, S & Schuster, H. G. (Wiley-VCH, Weinheim), pp. 355-395.

18 Eigen, M. (1971) *Naturwissenschaften* **58**, 465-523.

19 Rossler, O. E. (1971) *Z. Naturforschung* **26b**, 741-746.

20 Seneta, E. (1973) *Non-Negative Matrices* (George Allen and Unwin, London).

21 Padgett, J. (2001) in *Networks and Markets*, eds. Rauch, J. E. & Casella A. (Russel Sage, New York), pp. 211-257.

22 Cohen, M. D., Riolo, R. L. & Axelrod, R. (2001) *Rationality and Society* **13**, 5-32.

23 Carlson, J. M. & Doyle, J., *Phys. Rev. E* **60**, 1412 (1999).

24 Gutfreund, H. (1995) *Kinetics for the Life Sciences* (Cambridge Univ. Press, Cambridge).

Figure legends

Figure 1. A run with parameter values $s = 100$ and $p = 0.0025$. The x-axis shows time, n (= number of graph updates). Fig. 1a shows the number of links in the graph as a function of time. In Fig. 1b, the continuous line shows s_1, the number of populated species in the attractor (= the number of non-zero components of X_i) as a function of time. The dotted line shows λ_1, the largest eigenvalue of C as a function of time. (The λ_1 values shown are 100 times the actual λ_1 value.)

Figure 2. The structure of the evolving graph at various time instants for the run depicted in Fig. 1. Examples of several kinds of innovation and their consequences for the evolution of the system are shown (see text for details). Nodes with $X_i = 0$ are shown in white; according to the evolution rules all

white nodes in a graph are equally likely to be picked for replacement at the next graph update. Black and grey nodes have $X_i > 0$. Thus the number of black and grey nodes in an graph equals s_1, plotted in Fig. 1b. Black nodes correspond to the core of the dominant ACS and grey nodes to its periphery. Only mutual links among the nodes are of significance, not their spatial location which is arranged for visual convenience. The graphs are drawn using LEDA.

Figure 3. A hierarchy of innovations. Each node in this binary tree represents a class of node addition events. Each class has a name; the small box contains the mathematical definition of the class. All classes of events except the leaves of the tree are subdivided into two exhaustive and mutually exclusive subclasses represented by the two branches emanating downwards from the class. The number of events in each class pertain to the run of Figure 1 with a total of 9999 graph updates, between $n = 1$ (the initial graph) and $n = 10000$. In that run, out of 9999 node addition events, most (8929 events) are not innovations. The rest (1070 events), which are innovations, are classified according to their graph theoretic structure. The classification is general; it is valid for all runs. X_k is the relative population of the new node in the attractor of (1.1). N stands for the new irreducible subgraph, if any, created by the new node. If the new node causes a new irreducible subgraph to be created, N is the *maximal* irreducible subgraph that includes the new node. If not, $N = \Phi$ (where Φ stands for the empty set). Q_i is the core of the graph just before the addition of the node and Q_f the core just after the addition of the node. The six leaves of the innovation subtree are numbered (below the corresponding box) according to the subsection in which they discussed in the main text. The graph theoretic classes A, B, A1, B1, etc., are described in section 7 and Appendix B.

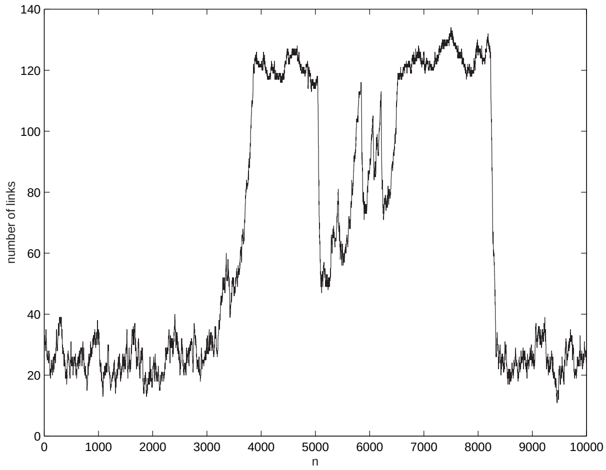

Figure 1a.

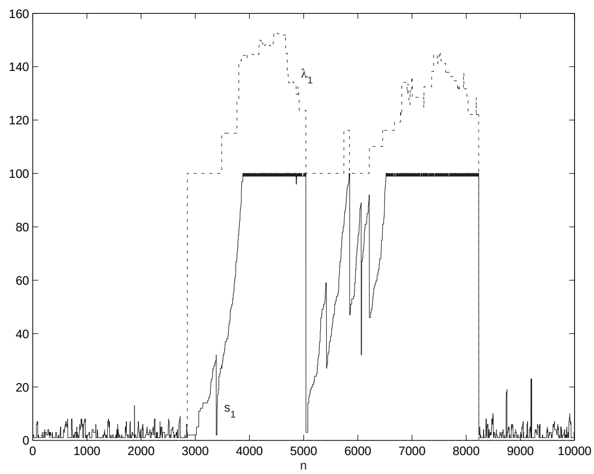

Figure 1b.

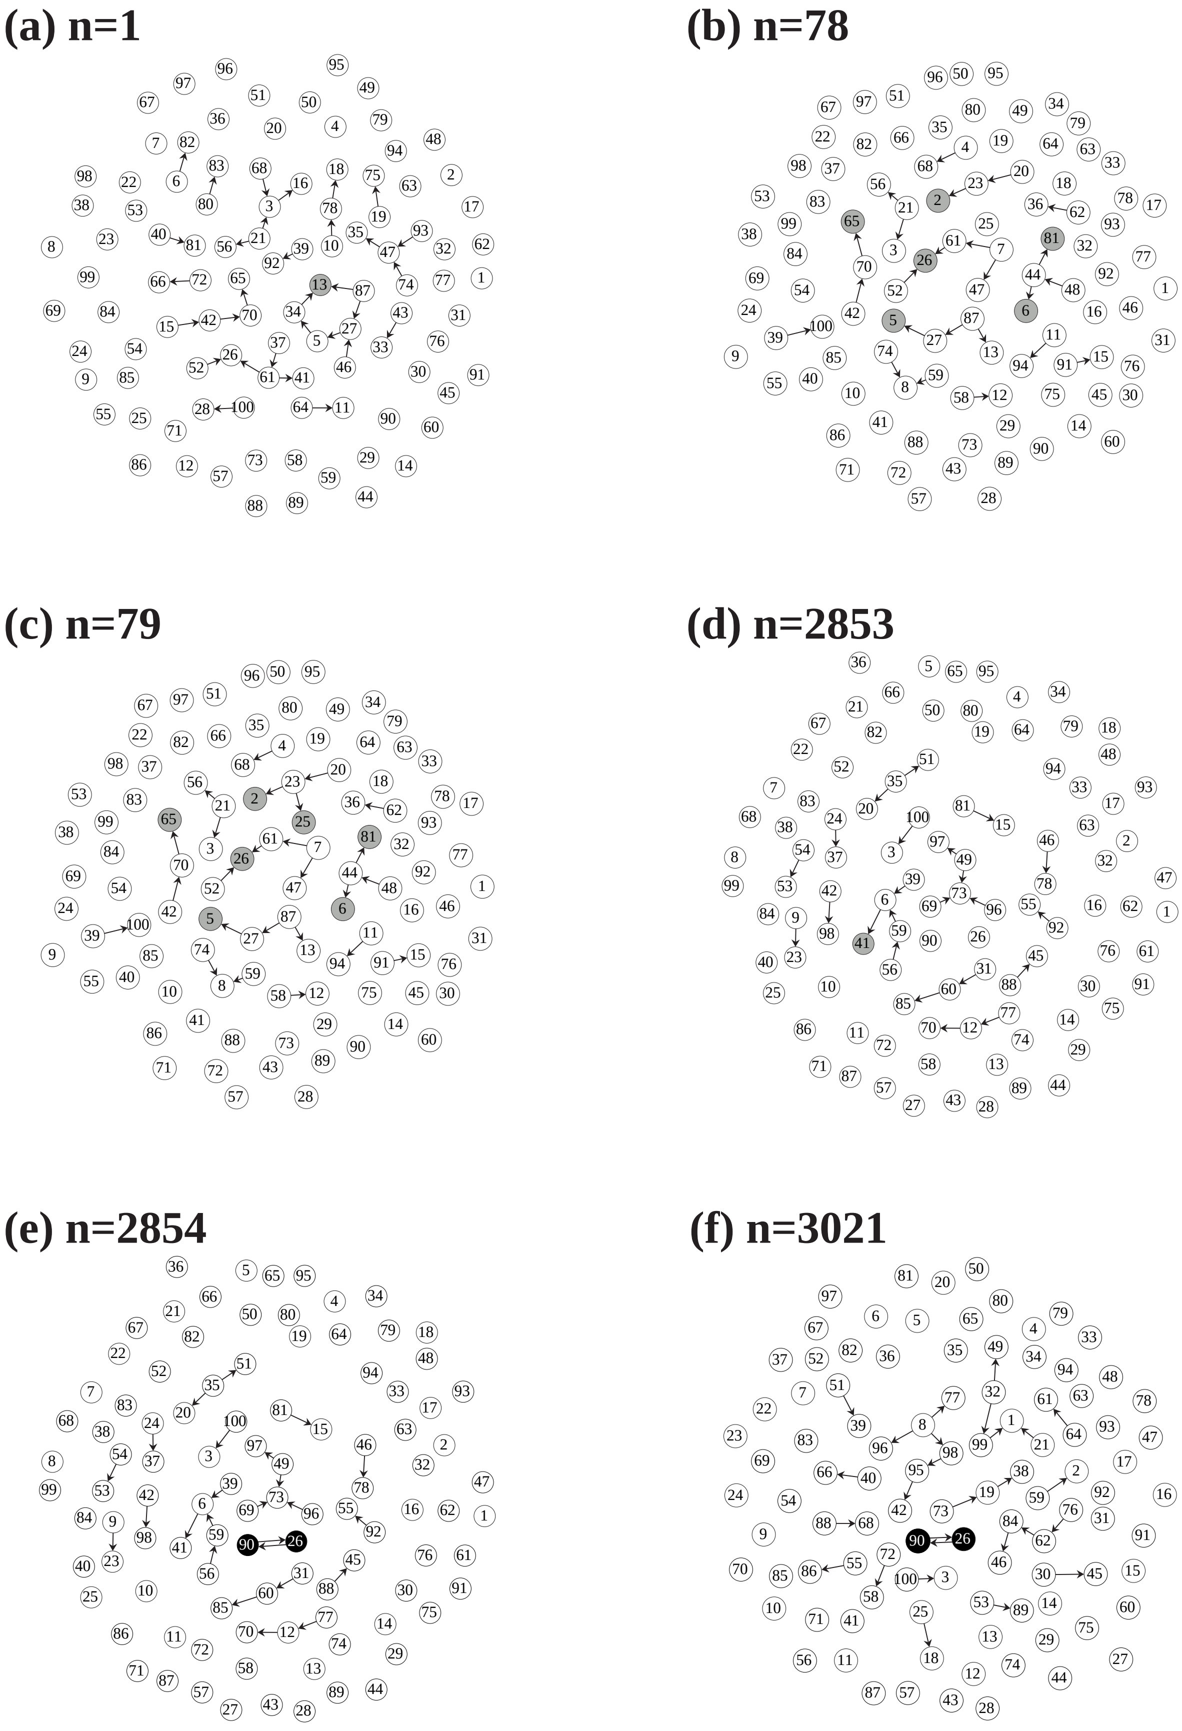

Figure 2.

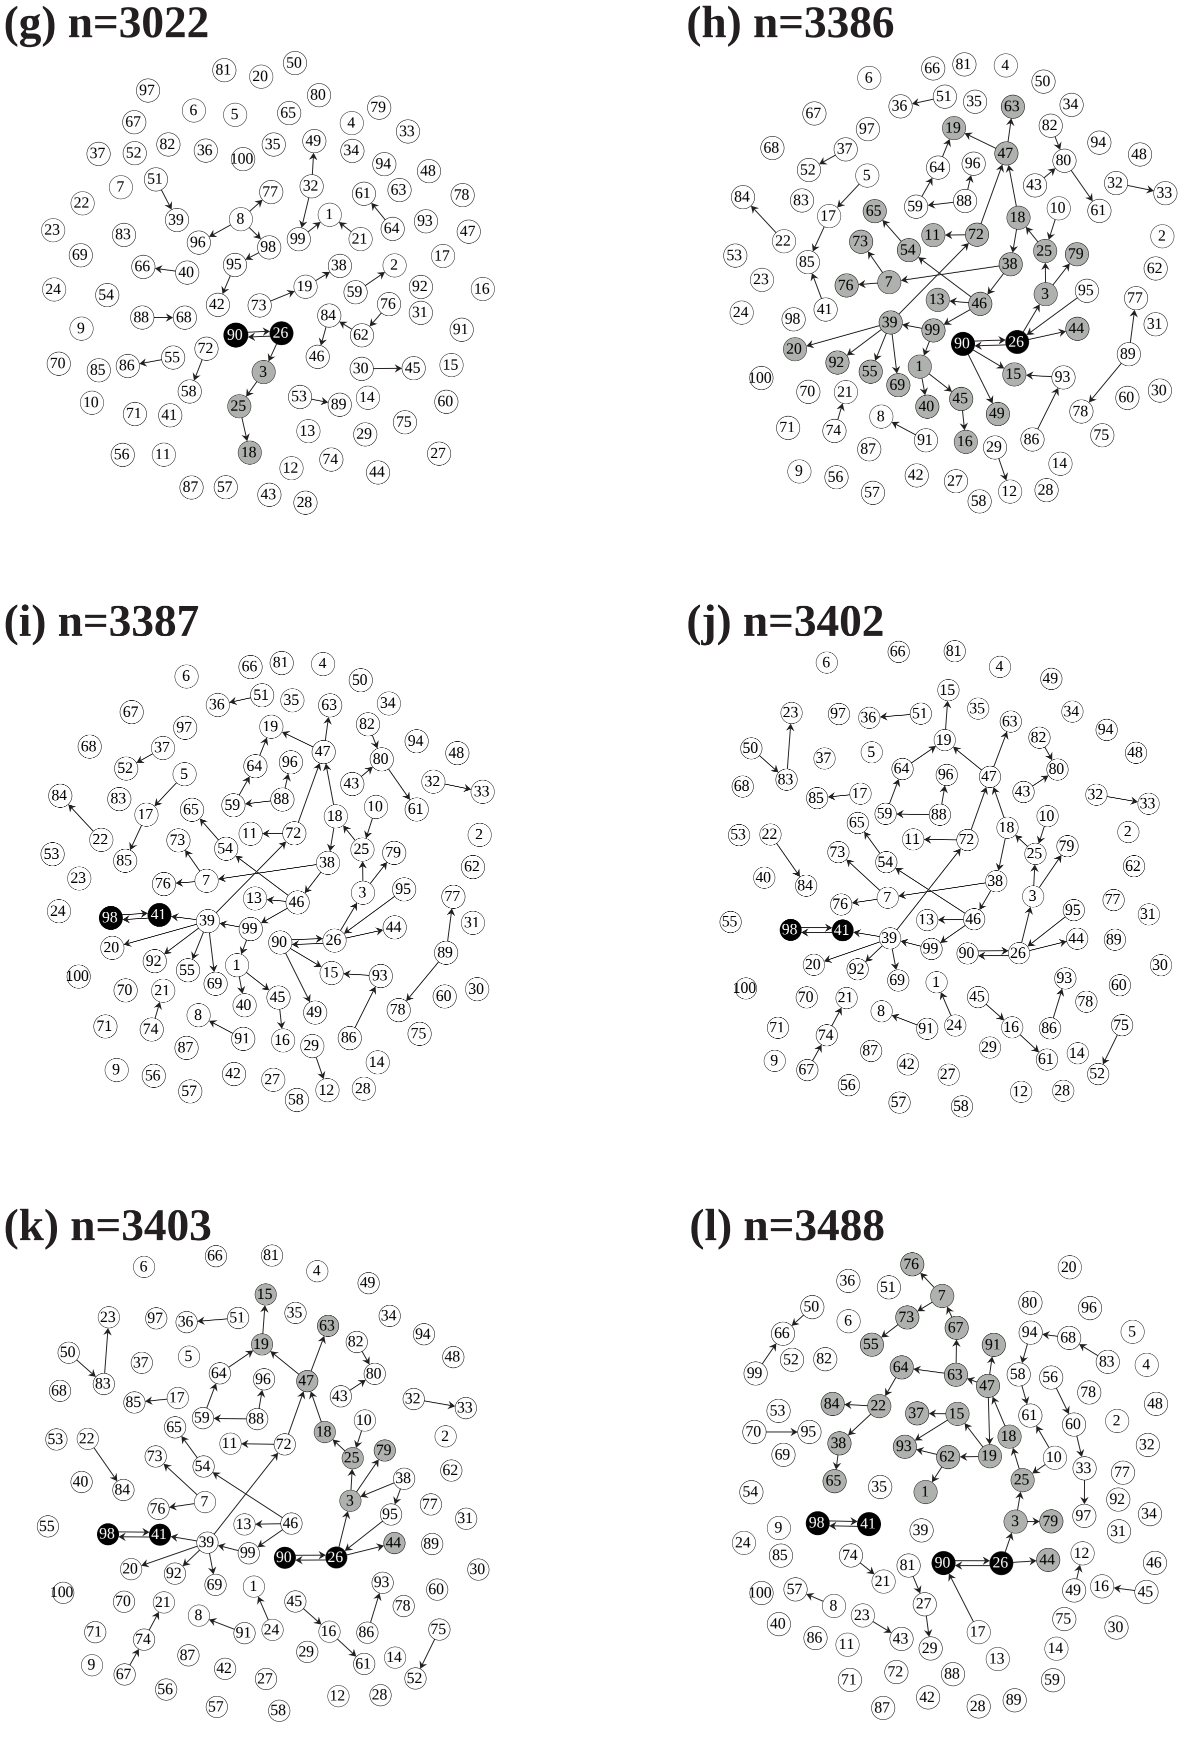

Figure 2. contd.

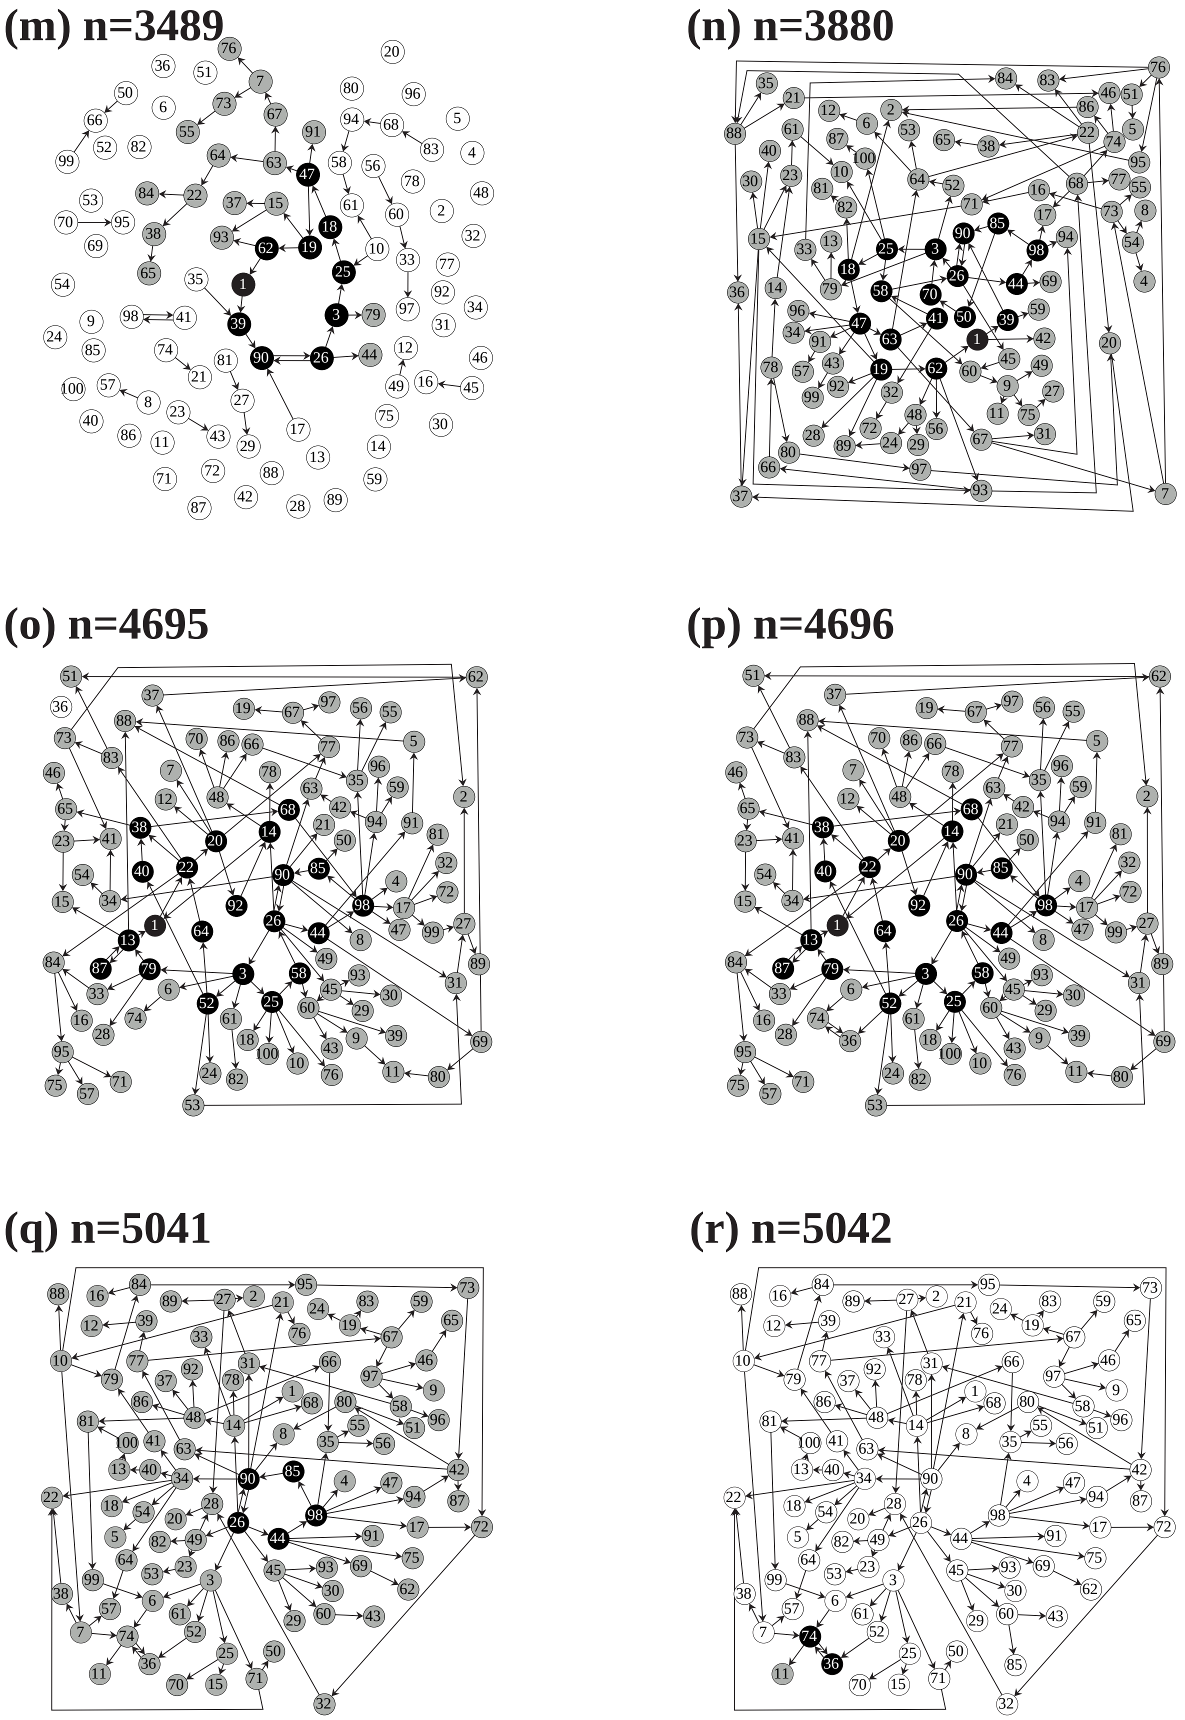

Figure 2. contd.

Figure 2. contd.

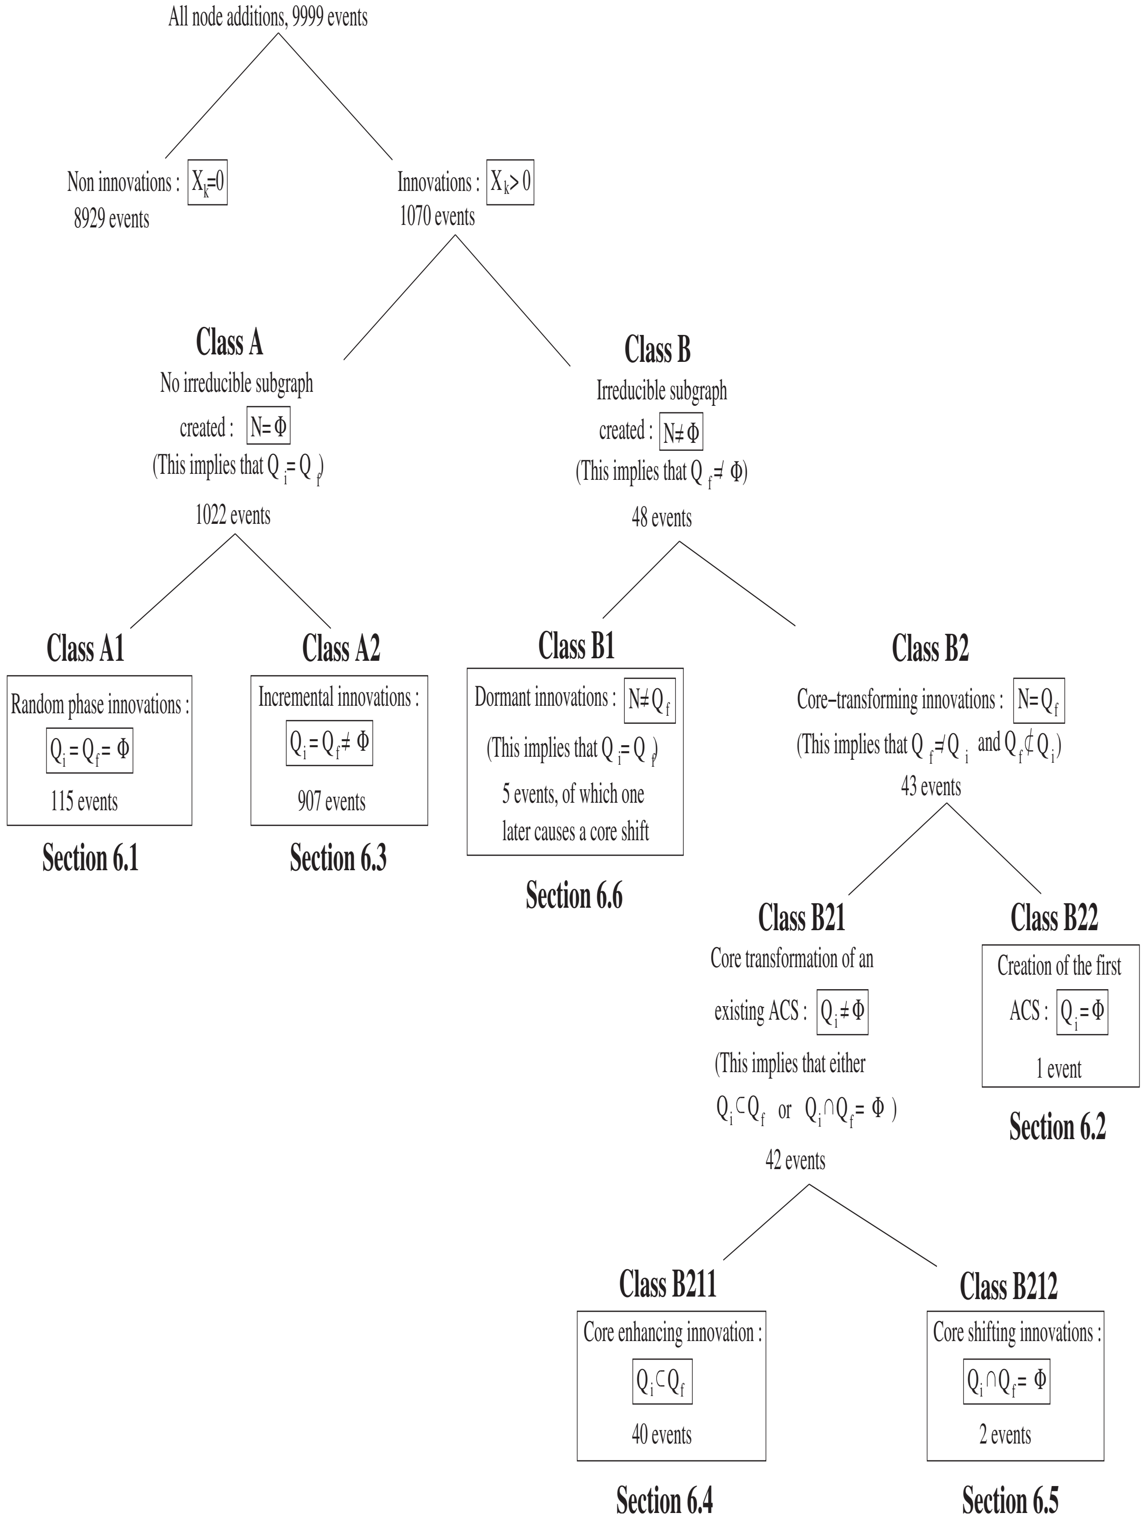
All node additions, 9999 events
Non innovations : $X_k=0$
8929 events
Innovations : $X_k>0$
1070 events
Class A
No irreducible subgraph
created : $N=\Phi$
(This implies that $Q_i=Q_f$)
1022 events
Class B
Irreducible subgraph
created : $N\neq\Phi$
(This implies that $Q_f\neq\Phi$)
48 events
Class A1
Random phase innovations :
$Q_i=Q_f=\Phi$
115 events
Section 6.1
Class A2
Incremental innovations :
$Q_i=Q_f\neq\Phi$
907 events
Section 6.3
Class B1
Dormant innovations : $N\neq Q_f$
(This implies that $Q_i=Q_f$)
5 events, of which one
later causes a core shift
Section 6.6
Class B2
Core−transforming innovations : $N=Q_f$
(This implies that $Q_f\neq Q_i$ and $Q_f\not\subset Q_i$)
43 events
Class B21
Core transformation of an
existing ACS : $Q_i\neq\Phi$
(This implies that either
$Q_i\subset Q_f$ or $Q_i\cap Q_f=\Phi$)
42 events
Class B22
Creation of the first
ACS : $Q_i=\Phi$
1 event
Section 6.2
Class B211
Core enhancing innovation :
$Q_i\subset Q_f$
40 events
Section 6.4
Class B212
Core shifting innovations :
$Q_i\cap Q_f=\Phi$
2 events
Section 6.5

Figure 3.