

ORIGINAL ARTICLE

The evolution of trophic structure

G Bell

Redpath Museum and Biology Department, McGill University, Montreal, Quebec, Canada

The trophic relationships of an ecological community were represented by digital individuals consuming resources or prey within a simulated ecosystem and producing offspring that may differ from their parents. When individuals meet, a few simple rules are used to decide the outcome of their interaction. Trophically complex systems persist for long periods of time even in finite communities, provided that the strength of predator–prey interaction is sufficient to repay the cost of maintenance. The topology of the food web and important system-level attributes such as overall productivity follow from the rules of engagement: that is, the macroscopic properties of the ecosystem follow from the microscopic attributes of individuals, without the need to invoke the emergence of novel processes at the level of the whole system. Evolutionarily stable webs exist only

when the pool of available species is small. If the pool is large, or speciation is allowed, species composition changes continually, while overall community properties are maintained. Ecologically separate and topologically different source webs based on the same pool of resources usually coexist for long periods of time, through negative frequency-dependent selection at the level of the source web as a whole. Thus, the evolved food web of species-rich communities is a highly dynamic structure with continual species turnover. It both imposes selection on each species and itself responds to selection, but selection does not necessarily maximize stability, productivity or any other community property.

Heredity (2007) **99**, 494–505; doi:10.1038/sj.hdy.6801032; published online 8 August 2007

Keywords: food web; coevolution; predator–prey; complex systems; individual-based model

Introduction

The trophic structure of an ecological community is represented by the food web, showing the flow of material and energy among different kinds of organisms. It is a fundamental attribute of the community that relates its diversity to its productivity and stability, so it is important to understand how the topology of the food web is determined, and how it influences the properties of the ecosystem as a whole (Pimm *et al.*, 1991; Warren, 1994). One of the most severe limitations of most current community models is that the range of diversity and the pattern of interaction are specified in advance and cannot change, although there has been considerable interest recently in modelling communities whose attributes can evolve (Caldarelli *et al.*, 1998; Christensen *et al.*, 2002; Loeuille and Loreau, 2005; Stauffer *et al.*, 2005). In this article I shall describe an individual-based community model that operates through simple rules rather than through equations, and which is capable of effectively indefinite variation and evolution.

Evolvable systems, based on computer algorithms, which are able to replicate and mutate have proven to be useful in addressing issues where conventional analytical or numerical approaches fail. Conrad and Pattee (1970) (also see Conrad and Rizki, 1989) described a simple individual-based ecosystem model in which

organisms collected resources when their internal state matched external conditions, and more complex models of trophic interactions were developed as computing power increased (for example Smith, 1991). Holland (1992) introduced the notion of artificial chromosomes encoding the behaviour of digital organisms in models that were adapted for simulation of food webs by Schmitz and Booth (1997). Autonomous digital organisms were introduced by Ray (1991, 1994, 1998) in his artificial world ‘Tierra’, which follows the laws of population genetics (Yedid and Bell, 2001) and can be used to elucidate processes such as the repeatability of evolutionary processes (Yedid and Bell, 2002). The use of individual-based models in ecology and evolution has been reviewed by DeAngelis and Mooij (2005).

Trophic structure will depend on the degree of specialization of producers and predators, that is, on the degree to which substances, structures and behaviour have evolved through their effect on capturing prey or avoiding predators. Some devices that predators use to attack prey are highly specific. Snake venom, spider webs, the piercing and sucking mouthparts of mosquitoes and plant bugs, and the construction of sand pits by ant-lion larvae are all elaborate devices that have evolved solely for the purpose of capturing prey. In other cases, prey are captured by relatively little-modified structures or behaviours that continue to serve other functions. The ciliary tracts of ciliates and rotifers, for example, serve for locomotion as well as for particle-feeding; the inhalant current of tunicates and bivalves serves for respiration as well as for filter-feeding; and the active pursuit of individual prey, the grazing or browsing of plant tissue or the scraping of biofilms from surfaces all involve

Correspondence: Dr G Bell, Department of Biology, McGill University, 1205 Avenue Docteur Penfield, Montreal, Quebec, Canada H3A 1B1.

E-mail: graham.bell@mcgill.ca

Received 1 December 2006; revised 25 May 2007; accepted 18 June 2007; published online 8 August 2007

rather unspecialized and non-specific structures and behaviours. Predator attack is resisted by equally complex or simple adaptations of the prey, each of which may be highly modified so as to antagonize some particular feature of a predator's armoury, or may be a more general ability to resist several different features. Many of the toxins elaborated by plants, for example, are highly active towards some herbivores but not towards others; precisely cryptic or dramatically aposematic colouration and behaviour is effective against predators hunting by sight but not against those hunting by sound or by scent. The shell of a snail, on the other hand, resists any biting or stabbing predator; and a filamentous growth habit, adherence to a substrate, or simply running away are all based on rather general features of organisms that will provide protection from many kinds of predators. Thus, a rough classification of interactions might recognize four categories: those in which both predator and prey are specialized, those in which neither are specialized, and those in which one participant, either the predator or prey, is specialized whereas the other is not. The model described in this article is an attempt to capture these properties of ecological communities with a small set of simple rules.

Modelling methods

The model I shall describe is an artificial electronic ecosystem containing organisms that consists of a few lines of computer code. This code specifies the identity of an individual, and determines how it will acquire resources, reproduce and interact with other individuals. The behaviour of the individual is otherwise under the control of the CPU, which sets system parameters and manipulates individuals according to the instructions they bear. The ecosystem is called Uqbar, after the

imaginary parallel world whose fragmentary records were described by Jorge Luis Borges.

Ecosystem

The habitable part of the ecosystem (Figure 1) comprises an ordered list of 'slots'. Each slot is either vacant or is occupied by a single entity. There are two types of entities: organisms and substrates. A substrate consists of a number of different non-substitutable resources in fixed proportions; different substrates are made up according to different recipes. Each exists in the form of single packets containing a fixed quantity of each resource, composed of ingredients obtained from the Resource pool. This is a separate compartment, which communicates with the ecosystem but is not habitable. Resources enter the pool either from the ecosystem, through recycling, or as exogenous input from an undefined outside world. The current contents of the resource pool determine the quantities of different resources available to the ecosystem, so that new packets of any given substrate can be formed only if the necessary resources are available from the resource pool. A new substrate packet enters the ecosystem from the resource pool with given probability in each move. It is eventually removed from its slot when it decomposes or is consumed by an organism.

Organisms

Organisms exist as individuals that are able to acquire resources. Some species (collectively called Producers) grow exclusively by consuming substrate packets; others (called Predators) by consuming other organisms. Growth beyond a certain point is followed by reproduction, when the parent gives rise to a newborn individual that is allocated a separate ecosystem slot. Individuals eventually die when they starve to death, having been

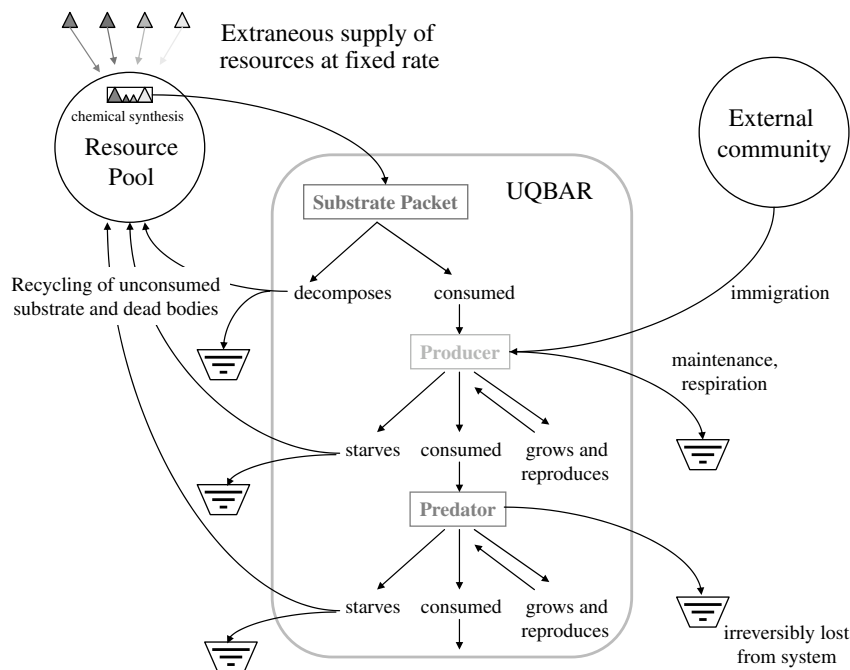


Figure 1 The main ecological processes in Uqbar. The ways in which resource supply, decomposition, recycling, maintenance, consumption, reproduction, death and immigration are implemented are described in the text.

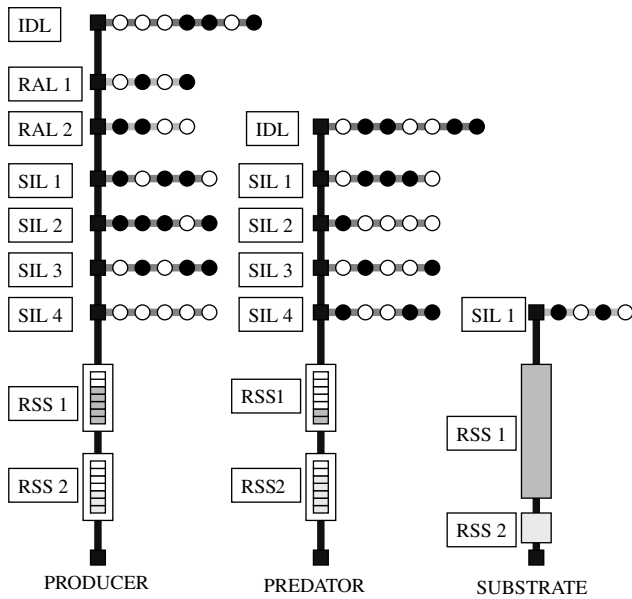


Figure 2 The structure of the three kinds of entity in Uqbar, showing the identity locus (IDL), resource acquisition loci (RAL), social interaction loci (SIL) and resource storage sites (RSS). The circles along IDL, RAL and SIL are examples of the bit string (open = 0 and filled = 1) that defines the genotype at a given locus. The shading of the RSS for Producer and Predator denotes the quantity of each resource possessed by an individual. The size of the RSS for the substrate indicates the quantity of that resource that the substrate contains.

unable to find sufficient or appropriate substrates, or when they are consumed by a predator. The resources contained in dead bodies may leave the system or may be recycled to the resource pool. Entities are constructed on a common plan, but regions present in some types of entity are lacking in others. The two principal types of entity are organisms and substrates, but an organism may be a Producer or a Predator (Figure 2). The most complicated entity is a Producer, which consists of four distinct regions.

- (1) A single identity locus (IDL). This bears a binary number of given length (that is, given number of bits) that labels the individual as belonging to a particular species. All individuals bearing the same identity number have identical genotypes, that is, they are identical in state at all other loci.
- (2) A set of social interaction loci (SIL). Each SIL bears a binary number of given length. The SIL determine the outcome of pairwise interactions between individual organisms.
- (3) A set of resource acquisition loci (RAL). This is a second set of loci each bearing a binary number of given length. The RAL determine the outcome of an interaction between an individual Producer and a resource package. They occur only in Producers.
- (4) The resource storage sites (RSS). A series of sites each of which holds the individual's supply of a certain resource. Resources are acquired and expended through interactions or attempted interactions with other organisms or with substrates. The current value and composition of RSS determine whether an individual will reproduce and whether it will die.

A Predator has precisely the same structure, except that it lacks RAL. A substrate package consists of two regions only: a single SIL and a set of RSS. The RSS of a substrate packet have fixed contents that correspond to a particular substrate SIL.

Ecosystem operation

The ecosystem list is used as an ordered queue, with execution proceeding from lower numbered to higher numbered slots. In each new move, the entity occupying the next highest numbered slot is selected as the 'Agent'. A second slot is then selected from the ecosystem list. In the present study, the second slot is selected at random, but the program provides the option of selecting it from a limited 'home range' centred on the slot occupied by the agent. The entity that occupies this slot is called the 'Patient'. Agent and Patient then interact according to the rules described below, and their status adjusted appropriately. When the top of the ecosystem list is reached, terminating a 'cycle' of moves, execution continues from the bottom; the list is thus essentially circular, with no definite limit to the number of cycles that may occur.

Each move consists of a series of external events that may affect the ecosystem, followed by an internal procedure. Each external event occurs with some defined probability in each move: exogenous supply to the resource pool (a given quantity of each resource enters the resource pool from outside), substrate supply to the ecosystem (a single random substrate packet is added to the ecosystem from the resource pool) and immigration (an individual chosen at random from the initial community enters the ecosystem from an external species pool). Uqbar is also subject to system-wide disturbance and species invasions, but these are not considered in this report. The internal procedure depends on the nature of the Agent. If the Agent is a vacant slot, execution proceeds directly to the next slot in sequence. If it is a substrate packet, then the packet may decompose spontaneously and its constituent resources may then be recycled to the resource pool. The probabilities of decomposition and recycling influence the availability of resources in the ecosystem. If the Agent is an organism, then it incurs a metabolic cost represented by the expenditure of fixed quantities of each resource during the move, even if the Agent encounters only a vacant slot. The cost of maintenance depends on the resource-gathering potential of an individual: the maintenance of a Producer is proportional to the number of RAL it bears, and that of a Predator on the number of SIL. In addition, SIL may be more (or less) expensive than RAL, so that heterotrophy may be intrinsically more (or less) expensive than autotrophy. These resources are recycled to the resource pool with a specified probability. If the Agent encounters an entity, then the consequences depend on the outcome of the interaction between Agent and Patient. Whether or not consumption occurs, however, the maintenance cost is exacted from any organism involved in a move, whether as Agent or as Patient.

Interaction

When two creatures encounter one another, one may consume the other if their SIL are complementary,

meaning that the sequence borne by one has the opposite bit at corresponding positions in the other (such as 00101 and 11010), that is, a bit-wise binary XOR operation on the two sequences will return a value greater than zero. If they bear more than one SIL, they may be tested for complementarity for all combinations of SIL (interaction in *trans*) or only by comparing SIL in corresponding positions (interaction in *cis*). Complementarity at several SIL increases the probability that the prey is captured. In principle, any kind of activity involving two individuals (for example, helping, trading or mating) can be governed by the complementarity or identity of the SIL, but in this report, only consumption of one by the other is considered. When a Producer encounters a substrate packet, the Producer consumes the substrate if any of its RAL are complementary to the single SIL of the substrate. The substrate packet is not consumed if the agent is a noncomplementary Producer or a Predator. When two organisms encounter one another, their interaction is antagonistic if one or more of their SIL are complementary.

The specificity of interactions between predator and prey is governed by the length of SIL and the rules for complementarity. Long SIL represent specialized prey, in the sense that the longer the sequence, the less likely is a predator to possess an exactly complementary sequence. Moreover, if the prey population is currently vulnerable, it is more likely to evolve resistance with long SIL because longer SIL are more likely to undergo mutation. Given a vulnerable prey population, the probability of a resistance mutation appearing in a given SIL of length L is Lu per genome per replication, where u is the mutation rate per bit, whereas the probability of a mutation restoring voracity is u . This reflects the fundamental asymmetry between prey and predators: any mutation in the relevant SIL of a vulnerable prey genotype will produce resistance, but only the precisely corresponding mutation in the predator will restore voracity. On the other hand, many SIL represent generalized predators, in the sense that having more SIL makes it more likely that a predator will find a complementary sequence in the prey it encounters. It also makes the evolution of voracity more likely, once a point mutation for resistance has become common in the prey, because with N SIL similar in state the probability of a mutation restoring voracity will be Nu . In practice, interaction in *cis* or *trans* is used to represent variation in the number of SIL, *cis* being equivalent to few and *trans* to many SIL, because this allows the physical genome size (total number of bits in SIL) to be kept constant. In short, long SIL and interaction in *cis* represent specialized prey, which readily evolve resistance, whereas many SIL and interaction in *trans* represent generalized predators that readily evolve voracity. In this way, although the SIL are clearly a highly formalized representation of the mechanisms involved in predator-prey interaction, they are nevertheless capable of being configured so as to imitate a broad range of ecological and evolutionary scenarios (Figure 3).

When two complementary individuals meet, a rule must decide which consumes the other. A Predator always consumes a Producer rather than *vice versa*. When two Predators meet, three kinds of rules might operate. The first is phenotypic and invokes a state that varies

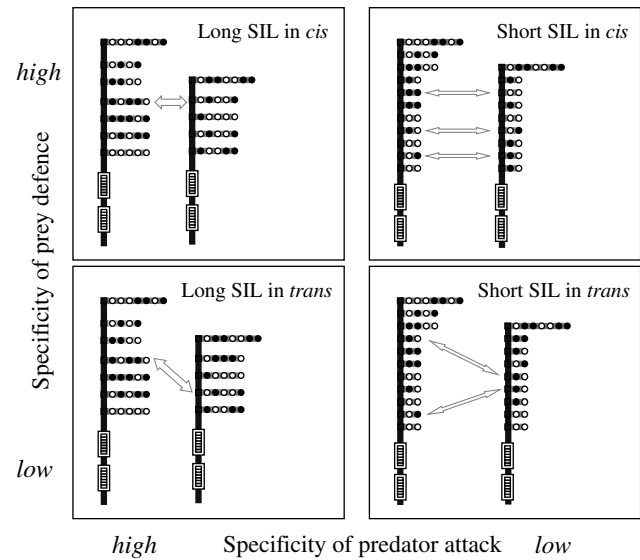


Figure 3 Rules of engagement in Uqbar, represented formally by short vs long SIL and *cis* vs *trans* interaction. The arrows indicate some of the complementary interactions between SIL of predator and Producer.

among individuals of the same species: the larger or stronger individual (that with greater RSS contents) consumes the smaller or weaker. RSS contents are not inherited, so the system cannot evolve. The second possible rule is a genetic criterion that is invariant within a species and consistent for pairwise combinations of species: the individual with the greater value at IDL consumes the other. This has two drawbacks: IDL is not heritable in any straightforward way (because it is a label that distinguishes a descendant from an ancestral taxon), and interactions are forced to be transitive (in the sense that if A eats B and B eats C then C cannot eat A). The final possibility is a genetic criterion that is invariant within a species but not necessarily consistent for pairwise combinations of species: the individual with the greater value (larger binary number) at complementary SIL consumes the other. This allows heritability and does not force transitivity, and for these reasons, it is the rule adopted here, although the other two can be invoked as options in Uqbar.

Life cycle

Individuals reproduce when they possess more than some minimum amount of each resource; the overall rate of reproduction thus depends on both the required quantity of each and on the amount of each available in substrate packets. When the RSS of an individual all exceed the threshold quantity, an ecosystem slot is allocated to the offspring. Only slots that are not occupied by organisms are available. If the slot is occupied by a substrate packet, this is removed, and its contents may be recycled to the resource pool. The genome of the parent (IDL, SIL and, in the case of a Producer, RAL) are then copied to the vacant slot and the appropriate header affixed. Finally, the newborn is provided with a basic ration of each resource, this amount being subtracted from the parent's stores. If no slot is available, no reproduction can occur.

Genetics

The contents of the SIL and RAL are copied with a certain probability of error, representing mutation. Three types of mutation may occur. The first is deletion, the omission of a complete SIL (from a Predator) or RAL (from a Producer). The second is duplication, the addition of a new SIL or RAL with the same structure as the parental locus. The third kind of mutation is point mutation, with each bit of each SIL and RAL being liable to be flipped from one state to the other. When mutation occurs, the IDL may be reset to a new value not possessed by any other organism, subject to the species concept in operation.

Two aspects of mutational variation need to be emphasized. The first is that unequivocally beneficial mutations are not allowed. Thus, increasing the number of RAL increases the range of substrates that a Producer could consume, but it also increases the cost of maintenance, and is allowed for this reason. Decreasing the number of SIL, on the other hand, would reduce the vulnerability of a Producer to predation with no countervailing disadvantage, and for this reason is not allowed. Deletion of SIL in Predators, on the other hand, is allowed because although it reduces both the cost of maintenance and the range of possible predators, it also reduces the range of possible prey. The second aspect is that macroevolution is not permitted. Producer and Predator are fixed categories, and organisms are not permitted to evolve from one into the other. Instead, a Producer that loses all its RAL or a Predator that loses all its SIL by deletion is unable to feed and eventually starves to death.

Uqbar is a general-purpose system that can be used to investigate a wide range of community processes. In this first report, I shall address five basic issues. First, in what circumstances do trophically complex systems persist? Second, how is the topology of the food web determined? Third, are food webs evolutionarily stable (ES)? Fourth, how do the specificity and the strength of interactions evolve through time? Finally, when topologically different food webs compete, which will persist?

Results

Randomly assembled communities

Uqbar can be inoculated with a specified number of predator taxa and a specified number of Producer taxa, each of which has a randomly specified genotype. Randomly connected food webs have been intensively studied since the classical work of May (1973). Randomly assembled communities in Uqbar, however, never have randomly connected food webs. Instead, overall connectedness and the pattern of connection are determined by the rules of engagement—the way in which predator and prey interact with one another (Figure 4).

In long-*cis* systems with high specificity, the complementarity of any two random types is likely to be zero, and connectance is correspondingly low. Many Producers have no predator, and most vulnerable Producers are consumed only by a single predator. Most Predators consume only one or a few types of prey, and few Predators are capable of consuming other Predators. The food web is thus dominated by very short, largely unconnected food chains, and resembles a lawn. In

long-*trans* and short-*cis* systems, complementarity and connectance are considerably greater. Isolated Predators that lack a predator are scarce, food chains lengthen, and trophic diagrams take on a bushy or 'webby' appearance. It is often difficult or impossible to assign a type to a single trophic level unambiguously; indeed, recursive loops involving non-transitive trophic relationships occur quite frequently. In short-*trans* systems, connectance is complete, although complementarity and thus the strength of interaction may vary widely among predator-prey couples. The food web resembles a tree, with a single long food chain. Among Predators, interactions are completely hierarchical, meaning that any Predator can consume all (or almost all) the Predators at lower levels.

The trophic structure of randomly assembled communities follows from the rules of engagement. When interactions are in *cis*, a given taxon may interact with others at different SIL; its relationship with any one is a poor predictor of its relationship with another. When interactions are in *trans*, the value (magnitude of binary number represented by bit sequence) of a given SIL will render it generally superior or inferior to complementary SIL in another individual, and a taxon will tend to behave consistently as predator or prey, depending on the values of all its SIL. The outcome of encounters will then be predominantly transitive, and this tendency will become stronger as the number of SIL increases. When there are many SIL they will generate many complementary pairwise interactions, and so their average value will be a reliable indicator of performance. When there are very few SIL, on the other hand, most encounters will involve only a single pair of SIL, and interaction in *trans* is little different from interaction in *cis*. At one extreme, short-*trans* systems are both consistently transitive and densely connected; the result is necessarily an almost perfectly hierarchical web with a single Predator food chain. This strict hierarchy breaks down in short-*cis* systems where interactions are less consistent and connections less dense, to give a branched, incompletely hierarchical web. It breaks down further in long-*trans* systems, which have short-branched chains where recursive loops often occur, along with isolated Predators which consume Producers and have no predators. Long-*cis* systems are at the other extreme, where the high degree of specificity dictates low connectance, and the outcome is a set of very short, isolated and unbranched food chains.

Evolution of trophic structure through sorting

When an isolated community is allowed to proliferate, poorly adapted species starve or are consumed, and become extinct. This process of sorting dissects out a simpler persistent community from the original inoculum. This community maintains trophic structure only if the rate of consumption by Predators is sufficient to pay for their cost of maintenance (Figure 5). The rate of consumption is governed by the relationship between the degree of complementarity and the probability of prey capture. Provided that the probability of capturing the prey when it is encountered is high enough, Predators paying some level of maintenance are able to persist. If the interaction is too weak, the Predators die of starvation. If the interaction is very strong, however,

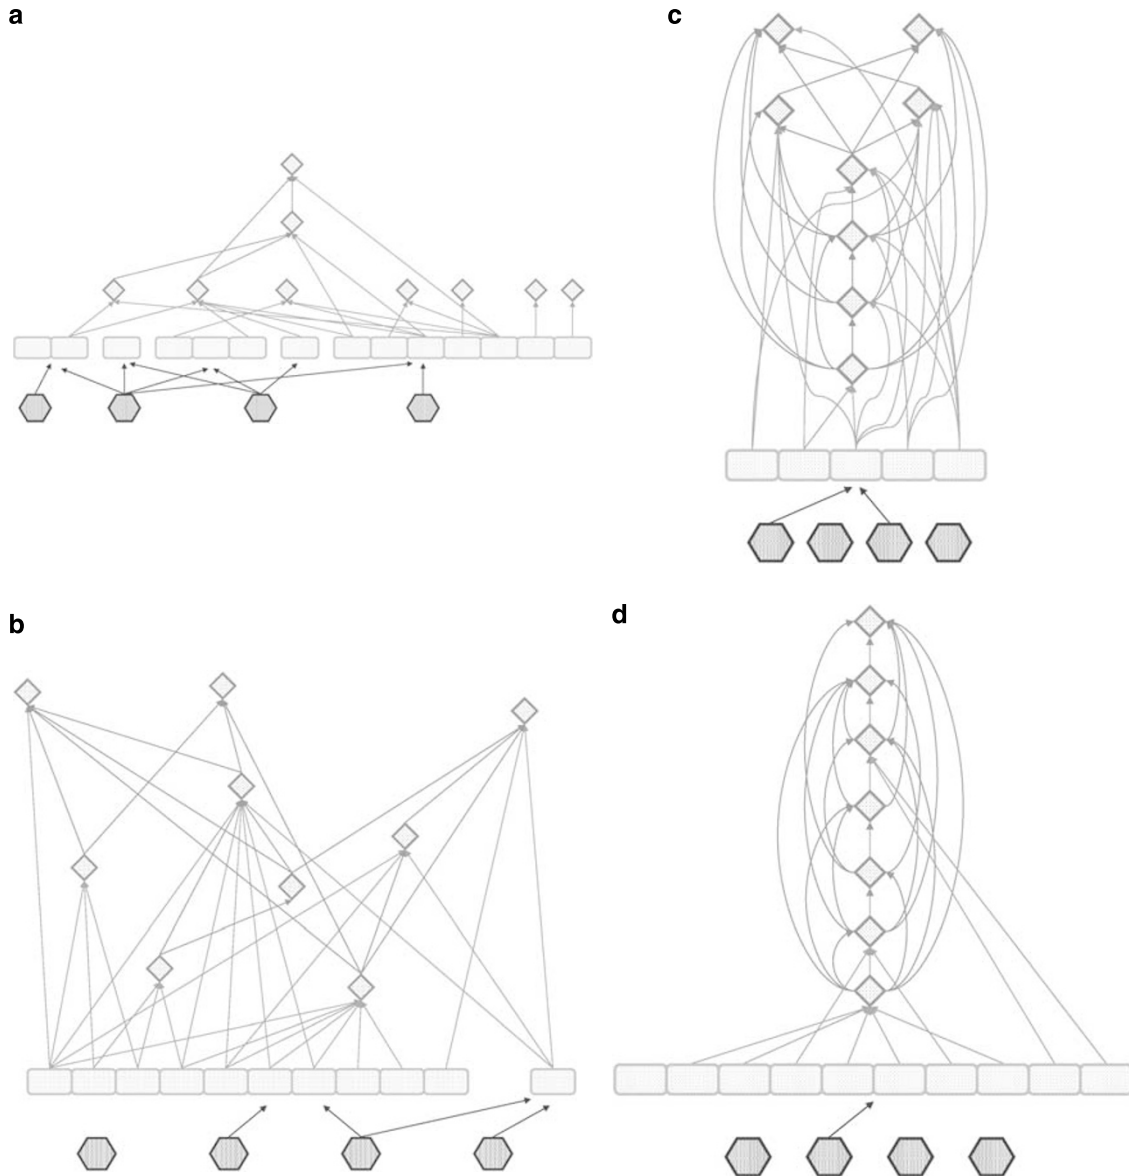


Figure 4 Schematic illustrations of food web topology for randomly assembled communities under different rules of engagement. Diamonds represent Predator species, rounded rectangles represent Producer species and hexagons represent substrates. Arrows indicate complementarity and the direction of flow of resources; they constitute a binary source web, although the output from Uqbar is quantitative, so that the abundance of species and flux along paths could be shown if desired. (a) Long-*cis* interactions; (b) long-*trans*; (c) short-*cis*; and (d) short-*trans*.

Predator abundance fluctuates with large amplitude and Predators often become extinct when rare through demographic stochasticity. A low level of immigration from an external pool of species rescues all taxa from permanent extinction and may support trophic complexity even when predator-prey interaction is weak. The general result is that the pattern of connection in webs under immigration, like that under pure sorting, resembles that of the randomly assembled communities from which they were derived, for instance highly specialized long-*cis* systems retain a lawn-like food web, whereas generalized short-*trans* systems remain hierarchical (see Figure 4).

Two simple statistics that can be used to describe the food web are the mean and variance of the number of links per species, standardized as the 'connectance' (fraction of all possible links realized) and 'dispersion'

(coefficient of variation of number of links; compare with Williams and Martinez, 2000). Sorting with immigration reduces connectedness by eliminating prey species that are vulnerable to many predators, and predators that have few or no prey. The effect is more pronounced when connectedness is initially very high (as in short systems), so sorted communities tend to have low or moderate values of connectedness. Dispersion tends to increase, largely because the variance is little changed, or somewhat less, whereas the mean is consistently much less. Thus, the overall topology of the web is little changed by sorting in open communities, except by removal of the most vulnerable and the least effective species. Food webs from real communities resemble those that develop in Uqbar (Figure 6), even though model parameters were chosen solely because they yield persistent communities, and not because they reflect values characteristic of

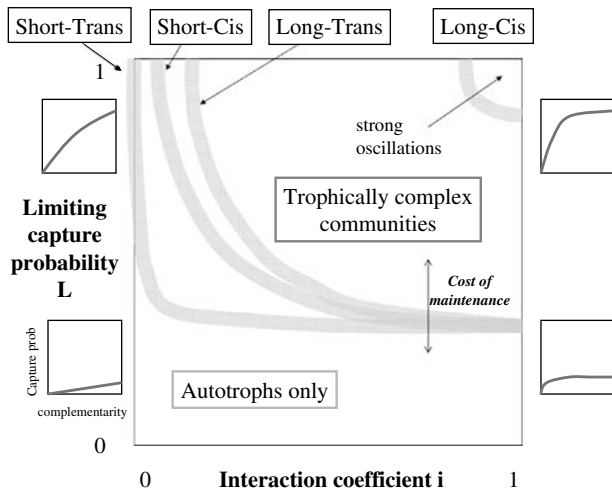


Figure 5 Conditions for the persistence of trophic structure in terms of affinity and the cost of maintenance. Affinity is defined as probability of capture = $L[1 - \exp(-i * \text{complementarity})]$ where L is the limiting (maximal) probability, i is an interaction coefficient and complementarity is the number of SIL at which predator and prey are complementary. The small boxes next to the axes illustrate capture probability as a function of complementarity for different combinations of L and i . Each thick band in the figure is the approximate limit of the parameter space above the band within which trophic complexity persists after the initial phase when poorly adapted species quickly become extinct through starvation. The double-headed arrow indicates that these bands can be tuned up or down by manipulating the cost of maintenance.




Figure 6 Linkage patterns in food webs after sorting in open communities. Top connectedness and top dispersion refer to Predators alone. Twenty replicate communities with different initial composition were followed for each of the four sets of rules. The system analysed in Figures 6–8 had a capacity of 10 000 individuals and was seeded with 64 Producer and 64 Predator species, with a subsequent probability of immigration by a single individual of random species of 0.0001 per move. This was run for 2000 cycles (that is 2×10^7 moves), with output values calculated as means over the final 100 cycles. The crosses denote estimates from real communities: Silwood Broom (Memmott *et al.*, 2000), Ythan Estuary (Hall and Raffaelli, 1991), Chesapeake Bay (Baird and Ulanowicz, 1989), Little Rock Lake (Martinez, 1991) and Skipwith Pond (Warren, 1989).

real communities. Highly specialized systems, such as those dominated by parasitoids (Silwood Broom), map in the region occupied by long-cis systems, as expected. Less specialized aquatic systems with vertebrate pre-

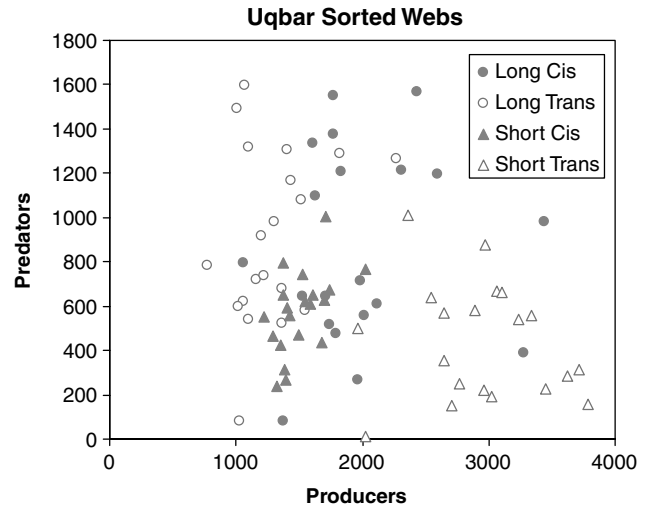


Figure 7 The abundance of Producers and Predators in relation to the rules of engagement, after sorting in open communities.

dators (Chesapeake Bay, Ythan Estuary, Little Rock Lake) map near the region typical of long-trans systems, whereas a pond lacking vertebrates (Skipwith Pond) may represent a highly generalized short-trans situation. A wide range of other metrics has been developed to describe the detailed structure of food webs (see Montoya *et al.*, 2006), but a description of how they are expressed in Uqbar is deferred to a subsequent publication. The fundamental issue of compartmentalization, however, is dealt with in the final section of the Results.

Community metabolism

The simplest feature of a trophically complex community is the balance between Producers and Predators, which in Uqbar is set primarily by the rate of resource input. At very low rates, Producer individuals are too rare to sustain a persistent community of Predator individuals. As rates increase Producer density increases, until it reaches a threshold where Predators can survive and reproduce. Above this threshold, the abundance of Predators increases with resource supply, whereas there is no trend in the abundance of Producers: that is, all the marginal resource supply is captured by Predators. Within this zone, where trophically complex communities can persist, the balance between Producers and Predators is strongly affected by the rules of engagement (Figure 7). Long-trans systems tend to have few Producers and many Predators, whereas short-trans systems have many Producers and few Predators. Short-cis systems have low numbers of both Producers and Predators; the most specialized long-cis systems are highly variable, but may have large numbers of both.

The output of the ecosystem, in terms of the quantity of resources transformed into offspring, is unrelated to the abundance of Producers, but increases with the abundance of Predators. At any given level of Predator abundance, production depends on the rules of engagement (Figure 8). In general, systems with unspecialized Predators are the more productive, because they do not include a large proportion of specialists whose prey is so rare that they seldom reproduce. The overall production of the ecosystem is thus affected in opposite senses by

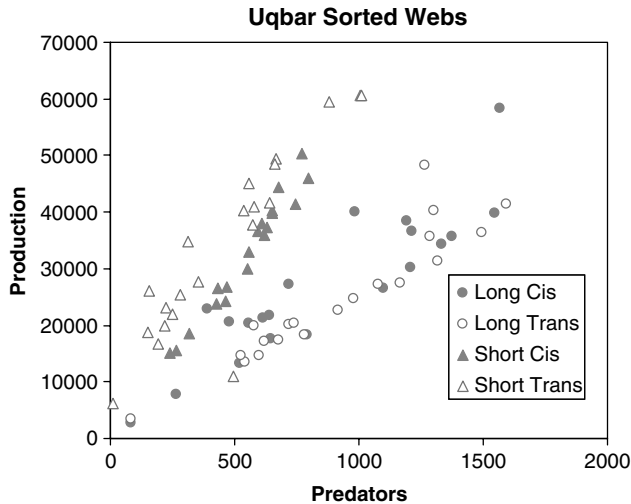


Figure 8 Productivity of ecosystems in relation to the abundance of predators and the rules of engagement. Productivity is expressed as the overall resource content of offspring.

the rules of engagement: generalized predators are less abundant but more active. Consequently, production does not differ much on average between systems with different rules, but the processes that determine production are nonetheless governed by these rules, given adequate resource supply.

Evolutionarily stable webs

The outcome of sorting may be a simpler persistent community, but it is not necessarily the only such community, or the most stable; some or all of the components of a more stable system may have been lost stochastically early in the sorting process. An ES web is a stable community that cannot be invaded by any other member of a defined initial pool of taxa. Whether or not the community acquires a more or less fixed composition depends on the diversity of the inoculum. With a low-diversity inoculum (several tens of taxa) an ES web usually emerges within a few thousand cycles. With a high-diversity inoculum (several hundreds of taxa), a stable community of roughly constant composition has never been observed, even though the topology of the food web remains the same. The turbulent dynamics of trophically complex systems is illustrated in Figure 9: a taxon often invades, flourishes for a few hundred cycles, then dwindles and disappears, only to invade successfully at some later time and enjoy another short-lived period of prosperity. Inspecting the genotypes segregating in the community shows that the underlying mechanism is coevolutionary, involving the spread of predators able to consume the most abundant types of prey, depressing the numbers of these prey and consequently favouring other prey species; this leads to the loss of these predators but provides an opportunity for other types to spread. The dynamics are fast, perhaps because the prey have no refuge. Community similarity declines rapidly over the first few hundred cycles (tens of generations) and then more slowly (Figure 10). Producers are more similar than Predators for any lag, expressing the slower turnover of Producers, perhaps because they are adapting in part to a fixed spectrum of resource packets. Rapid change and

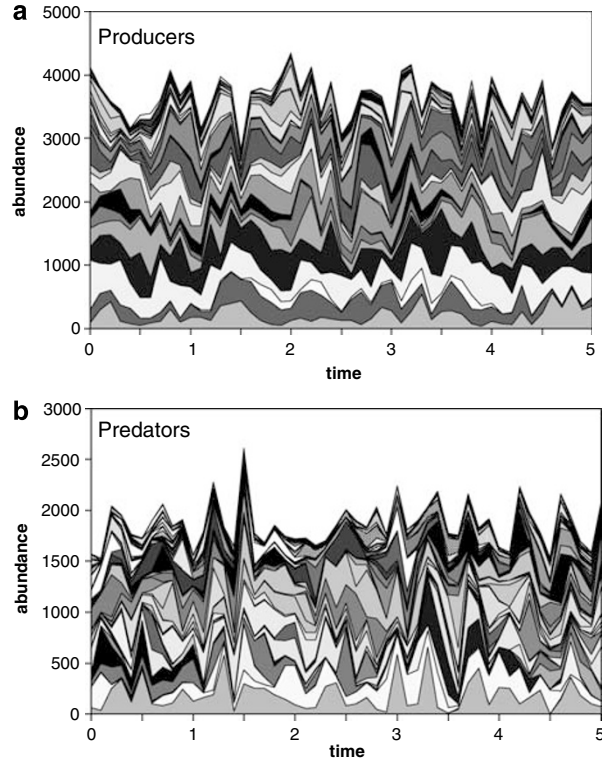


Figure 9 Fluctuation in community composition over time. This shows 5000 cycles of a long-trans system, beginning 10 000 cycles after inoculation and plotted at intervals of 100 cycles. Important system parameters are: extent 25 000 slots; inoculum 128 Producer taxa, 128 Predator taxa; organism specification four SIL each of length 5 bits, two RAL of length 2 bits; immigration 0.0001 per move. (a) Producer taxa and (b) Predator taxa.

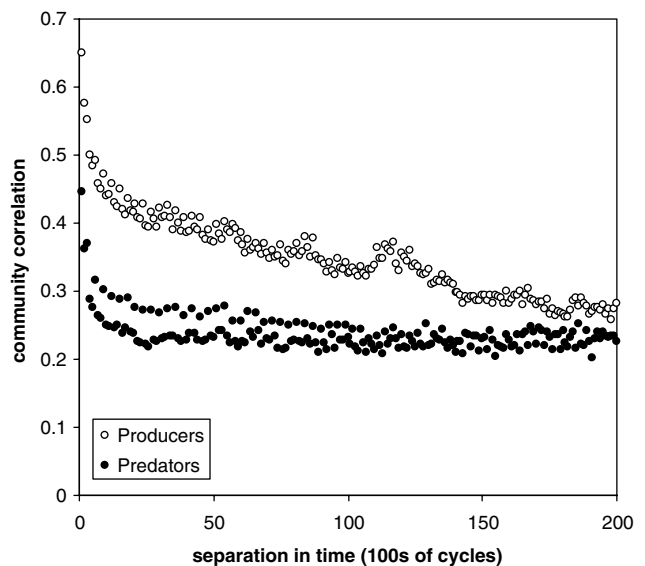


Figure 10 Community similarity through time. The plot shows the binary correlation coefficient expressing similarity in terms of presence or absence of species at 100-cycle intervals for lags of up to 20 000 cycles in the long-trans experiment illustrated in Figure 9.

the greater lability of Predators appear to hold for all systems and do not seem to require highly specific predator-prey interactions.

Evolved webs

Mutation prolongs the existence of trophically complex communities through continued coevolution. This is a very prominent effect in long-*cis* systems with very high specificity, where communities quickly become dominated by resistant Producers and so lose trophic complexity. With no mutation, there is an irreversible decline in Predator diversity in closed systems, culminating in complete extinction. When Predators are allowed to mutate they can recover voracity, with the result that trophically complex communities persist for much longer. The magnitude of the effect depends on the flux of mutants and thus on community size. In large systems, point mutation seems to immortalize Predator communities that would otherwise disappear within a few hundred cycles. The underlying dynamics can be investigated by tracking the abundance of particular genotypes, or categories of genotypes. The Producer community is relatively stable, changing slowly in composition over periods of thousands of cycles. The Predator community, by contrast, fluctuates violently in composition over periods of tens to hundreds of cycles. Most Predator taxa, most of the time, are rare specialists barely able to persist through occasional encounters with equally rare vulnerable prey. Many Producer taxa, on the other hand, are abundant, and are invulnerable to any existing Predator. From time to time, a mutation confers on the offspring of a Predator the ability to consume one of these abundant Producers. This novel Predator proliferates rapidly, wiping out its prey as it does so. It therefore crashes quite soon, but during its brief period of rapid expansion it will have given rise to a cloud of new mutants, some of which will persist and may in turn give rise to the next successful type. At the level of the community as a whole, this process cannot be perceived; the overall abundance of Predators and Producers does not fluctuate strongly through time. In systems with less highly specific interactions, the fluctuations of Predators are less extreme, but the predator community is always more dynamic than the Producers, as with immigration and sorting. The main reason for this is likewise that the fitness of Producers depends in part on the non-evolving composition of available substrates.

Evolution of affinity

Under selection for optimal number of SIL, mean affinity (the probability that a potential prey individual will be captured when encountered) may either increase or decrease. In long-*cis* systems, almost all captures involve a single complementary pair of SIL. This cannot be reduced, and any increase through duplication would be unprofitable, so mean affinity remains high, corresponding to a probability of capturing any prey encountered in excess of 0.9. In long-*trans* systems, mean affinity increases through time under point mutation. This is because individual Producers tend to bear only a single kind of SIL, so that Predators encounter vulnerable prey less often but, when they do, interact with them more strongly. Consequently, the probability of capture evolves towards high values of about 0.8. Duplication and deletion make no difference to this outcome. In short systems, on the other hand, selection for optimal numbers of SIL in Predators leads to much lower

capture probabilities of 0.25 in *trans* and 0.1 in *cis*. This is driven solely by deletion or duplication, and the same result is obtained whether or not point mutation is permitted.

Complex systems

In the systems described above, all individuals have SIL of the same length. Individuals having SIL of different length cannot interact, so when the length of SIL is allowed to vary the system initially consists of a number of subsystems, isolated food webs each with its own rules of engagement. All Producers in the system compete for the same pool of substrates, but Predators can consume only those individuals belonging to the same food web. Subsystems may differ in any way, including length and mode of operation of SIL, affinity and maintenance cost. When alternative subsystems all comprise Producers alone, then the most frugal subsystem (that capable of reducing substrate concentration to the lowest value) replaces all others. This is merely an extension of the well-known result for competing species (Tilman, 1982). The situation is more complex when the food web as a whole consists of several source webs with similar or dissimilar rules of engagement. These subsystems will compete because they draw on a common pool of resources, and selection among them will modify the properties of the ecosystem as a whole. It might be imagined that the more frugal or the more stable subsystems would be favoured by selection, and by replacing all others would enhance ecosystem properties such as average productivity, variation in productivity, trophic transfer efficiency, constancy of composition or robustness to perturbation. This does not seem to be generally true: the outcome of competition between trophically complex subsystems is usually coexistence rather than replacement. The reason for this is that as frugal Producers become very abundant their Predators increase, thus raising the death rate of Producers and reducing their abundance. Hence, selection among subsystems is negatively frequency-dependent, and the most frugal subsystem will replace all others only if the difference in maintenance costs is sufficiently high. Subsystems that differ only in affinity will coexist, provided that each is capable of persisting in pure culture in similar conditions. These persistent subsystems may become more or less abundant: when rules of engagement are similar, those with weaker interactions tend to become the more abundant. This is because the Predators of subsystems with strong interactions drive down Producer numbers and thereby restrict their own abundance. When the rules of engagement differ among subsystems, there seems to be no general rule specifying the replacement of one by another, or the prevalence of a particular kind of web. Systems that are able to persist as isolated systems in pure culture are likely to coexist as mixtures, given comparable conditions of growth. In any particular case, component webs with different rules of engagement may differ in frequency, but I have not been able to find any general relationship between the pattern of specialization and relative abundance. There is no tendency in Uqbar, therefore, for one web topology to replace all others, and the ecosystem will usually consist of component webs with a variety of different topologies.

Summary of results

The five questions posed earlier can now be answered, at least with respect to Uqbar.

- (1) Trophically complex systems persist even in finite communities if the strength of predator–prey interaction is sufficient to repay the cost of maintenance (Figure 5).
- (2) Many features of the topology of food webs are attributable to the rules of engagement governing the outcome of encounters between individuals (Figure 4).
- (3) ES webs exist only when the pool of available species is small. If the pool is large, or speciation is allowed, species composition changes continually while overall community properties are maintained.
- (4) Affinity is optimized through selection among individuals and shows no consistent tendency to increase or to decrease.
- (5) Topologically different food webs usually coexist, through negative frequency-dependent selection at the level of the source web as a whole.

Discussion

Rule-driven systems like Uqbar offer a strategic approach to understanding the behaviour of complex ecological systems that bridges the gap between equation-driven Lotka–Volterra models (for example see Pimm, 1980; Montoya and Sole, 2003) and detailed tactical models of particular communities based on systems ecology (see for example Baretta-Bekker *et al.*, 1997). The most thoroughly investigated evolvable model seems to be the equation-driven Webworld of Caldarelli *et al.* (1998), further elaborated by Drossel *et al.* (2001) and reviewed by McKane (2004). Each species possesses a set of characteristics chosen from a list, and may transmit a modified set to its offspring. Summing over these characteristics provides a rule of engagement assigning predator and prey roles to any pair of species. The physical environment is represented as an additional species that can be consumed by species which have appropriate attributes; these constitute primary Producers, and are sustained by exogenous resource input to the system. The resources distributed among primary Producers are captured by predators in a fashion governed by the match between the character states of predator and prey. Each predator species receives a quantity of resources proportional to the population size of each of its prey species and to its ability to consume it relative to other species of predator. The model is governed by difference equations that update the population size of each species at each time step. This generates a food web that rapidly converges on a stable and predictable configuration if there is no mutation. The community evolves through speciation, a new species being formed by altering the state of a single character of its ancestor. The appearance of a new species may drive the community to a different configuration, which is then used as the basis for the next speciation episode. This results in continually changing community composition and food web structure. Basal species turn over less frequently, however, because they are adapting to a fixed resource input, and, consequently, predators can become increasingly well adapted over time. The probability that

a new species will become established then declines over time, and Caldarelli *et al.* (1998) believed that the community might eventually reach a stable state. Subsequent work with more realistic population dynamics has shown that the community approaches a dynamic equilibrium at which bulk properties (such as total species richness) are conserved although there is a continual turnover of species (McKane, 2004; Quince *et al.*, 2005). The food webs emerging from Webworld resemble the bushy webs characteristic of long-*trans* systems in Uqbar, with largely transitive food chains of up to four or five links and frequent omnivory. Quince *et al.* (2005) explain the evolution of this structure in terms of the abundance of potential prey at the trophic level below. At high levels, prey abundance is restricted and predators tend to be nonoverlapping generalists, whereas at lower levels, there are more prey available and predators tend to be specialized, each prey species being consumed by several species of predator. As resource supply rate increases, the population at each level becomes larger and trophic chains become longer.

Christensen *et al.* (2002) have described a more abstract individual-based ‘Tangled Nature’ model in which individuals are represented by strings of binary digits. Each string interacts in an arbitrary fashion with any other given string, without any consistent rule governing the kind and strength of interaction between strings. The fitness of an individual is then determined by the sum of its interactions with all other extant individuals. If these enable it to reproduce, mutation may switch any digit in the string defining its offspring, which then has an arbitrarily different pattern of interaction with other individuals. Regardless of these interactions, each individual dies with constant probability in each time step. The community evolves through changes in the frequency of members of the set of possible strings, driven by coevolution, although without either trophic structure or ecological context. The outcome is an unexpected two-phase dynamic in which long-lasting communities are disrupted by turbulent periods of fluctuating composition, after which a new persistent community is established. These transitions become less frequent as time advances, but the community never attains a final stationary configuration (Hall *et al.*, 2002). Recknagel (2003) has also discussed individual-based models with evolvable trophic interactions that descend from economics theory (see Holland and Miller, 1991), but the trophic structure they produce has not yet been described in any detail.

Loeuille and Loreau (2005) describe a very simple equation-driven model in which species of similar size compete, whereas larger species are able to consume smaller species, at the cost of a lower specific metabolic rate. New species of different size appear occasionally through mutation. A single ancestral species consuming inorganic resources thereby diversifies to form a cascade-type food web. If consumption is strong relative to competition, trophic levels are well-defined and with few species, whereas strong interference competition leads to diffuse trophic structure, many species and a high level of omnivory (see Lässig *et al.*, 2001). The system converges to a stationary state that appears to be a dynamic equilibrium, as species composition continues to change.

A detailed comparison of the properties of evolvable model ecosystems is beyond the scope of this paper, except to note that ES webs have seldom been reported. In most conventional models, persistent trophically complex communities seldom comprise more than a handful of species. Law and Morton (1996) found ES webs of only four–six species from an initial pool of 50 species or more. This occurs in Uqbar when the range of variation is very restricted, either because the species pool is small or because there are very few possible configurations of SIL. In more diverse systems, sorting and continued selection fuelled by immigration or mutation lead to a dynamic equilibrium at which overall system properties are maintained, while the underlying composition of the community is in continual flux. The Tangled Nature, Webworld, Loeuille–Loreau and Uqbar models all seem to approach a dynamic equilibrium at which whole-system properties are conserved despite a continual turnover of species and trophic links. This reconciles the inherent instability of predator–prey dynamics with the persistence of trophically structured communities, and may be a fundamental property of complex ecological systems that will be captured by a wide range of models with different modes of implementation. There are also indications of more specific rules, such as the greater lability of predator taxa in both Webworld and Uqbar. Whether more detailed features of food web architecture can be robustly predicted remains to be seen.

There have been several attempts to interpret food web topology in terms of ecological processes (reviewed by Warren, 1994). For example, connectance and dispersion might affect ecosystem stability, with more stable web configurations being favoured (see Chen and Cohen, 2001). Within the limits set by resource supply rates and metabolic costs, however, the main features of the shape of food webs in Uqbar follow from the simple rules that govern how two individuals behave when they encounter one another (Figure 3). The length of food chains is likewise governed by the rules of engagement, being shortest in highly specialized (long-*cis*) systems. Food chains are longest in generalized systems where the length of the longest chain depends on the abundance of predators, which in turn depends on the rate of resource supply for a system of given extent. It was suggested long ago that the length of natural food chains is governed by the rate of resource supply (Hutchinson, 1959) but although this has been demonstrated in some cases it does not seem to be the general rule (see Townsend *et al.*, 1998 and Post *et al.*, 2000). Food chain length increases with resource supply in Uqbar, however, only when systems with the same rules of engagement are compared, and this may account for the variable reports from natural communities.

Interest in food webs has led to a large and rapidly growing body of work that seeks to infer the nature of ecological processes in trophically complex communities by identifying the leading statistical regularities of food web architecture. This statistical approach was pioneered with the use of graph theory by Cohen and Briand (reviewed in Cohen *et al.*, 1990), and has continued by characterizing properties such as link invariability, scale-dependence, small-worldness and compartmentalization (Sugihara *et al.*, 1989; Havens, 1992; Martinez, 1994; Murtaugh and Kollath, 1997; Watts and Strogatz, 1998;

Dunne *et al.*, 2002; Krause *et al.*, 2003; Montoya *et al.*, 2006). It has been speculated that properties such as these are attributes of complex systems that emerge at the level of the system as a whole, so an holistic approach is needed to uncover rules that apply to complex interconnected systems of all kinds, whether they be food webs, gene interactions, power grids or entire economies (for example Melian and Bascompte, 2002; Garlaschelli *et al.*, 2003; see also Winemiller *et al.*, 2001). Trophic structure in Uqbar, however, arises from a few simple rules of engagement operating microscopically at the level of individuals, and from selection among individuals for more effective prey capture or predator avoidance. It can be interpreted in a completely reductionist manner, without the need to invoke novel emergent properties of the system as a whole.

It has sometimes been held that complex ecological systems will tend to become more stable or productive through time, for example because weaker interactions are favoured (see Paine, 1992; McCann *et al.*, 1998; McCann, 2000) or because interactions tend to become compartmentalized (Krause *et al.*, 2003). Montoya and Sole (2003) found that random Lotka–Volterra systems resembled real food webs after sorting only if maximum affinity was low. It might be imagined that, at some level, selection will favour more stable over less stable sets of interacting species in any given part of the overall food web, thus tending to stabilize the system as a whole through the evolution of weak interactions. This does not appear to be the case in Uqbar. When the rules of engagement of Uqbar are fixed, selection among individuals will favour more effective predators and more resistant prey, and the overall strength of interactions may either increase or decrease as a consequence. For any particular predator–prey interaction, selection will always favour predator genotypes with stronger interactions and prey genotypes with weaker interactions. When deletion and duplication are allowed, for example, selection acts to optimize expenditure on SIL (or in Producers on RAL). In any move, the cost is a fixed metabolic expenditure, whereas the benefit depends on the probability of encountering prey, the probability of capturing an individual once it has been encountered, and the nutritive value of the prey. The optimal number of SIL can then be calculated for any given capture function (see Supplementary Information Part 7), and when deletion and duplication are permitted communities evolve towards this optimal value from above or below within a few thousand cycles. Thus, simple systems comprising a single food web do not naturally tend to become more stable, because selection among individuals will not generally produce this outcome. Where there are several isolated systems with different rules of engagement, however, selection may operate among these trophically isolated communities. In general, those with weaker interactions between predator and prey tend to become the more abundant, and in this case the overall strength of interaction may fall, relative to that of a simpler system.

References

- Baird D, Ulanowicz RE (1989). The seasonal dynamics of the Chesapeake Bay ecosystem. *Ecol Monogr* 59: 329–364.

- Baretta-Bekker JG, Baretta JW, Ebenhijh W (1997). Microbial dynamics in the marine ecosystem model ERSEM II with decoupled carbon assimilation and nutrient uptake. *J Sea Res* **38**: 195–211.
- Caldarelli G, Higgs PG, McKane AJ (1998). Modelling coevolution in multispecies communities. *J Theor Biol* **193**: 345–358.
- Chen X, Cohen JE (2001). Global stability, local stability and permanence in model food webs. *J Theor Biol* **212**: 223–235.
- Christensen K, di Collobiano SA, Hall M, Jensen HJ (2002). Tangled nature: a model of evolutionary ecology. *J Theoret Biol* **216**: 73–84.
- Cohen JE, Briand F, Newman CM (1990). *Community Food-Webs: Data and Theory*. Springer: Berlin.
- Conrad M, Pattee HH (1970). Evolution experiments with an artificial ecosystem. *J Theor Biol* **28**: 393–409.
- Conrad M, Rizki MM (1989). The artificial worlds approach to emergent evolution. *BioSystems* **23**: 247–260.
- DeAngelis DL, Mooij WM (2005). Individual-based modelling of ecological and evolutionary processes. *Annu Rev Ecol Evol Systemat* **36**: 147–168.
- Drossel B, Higgs PG, McKane AJ (2001). The influence of predator–prey dynamics on the long-term evolution of food web structure. *J Theor Biol* **208**: 91–107.
- Dunne JA, Williams RJ, Martinez ND (2002). Food-web structure and network theory: the role of connectance and size. *Proc Natl Acad Sci USA* **99**: 12917–12922.
- Garlaschelli D, Caldarelli G, Pietronero L (2003). Universal scaling relations in food webs. *Nature* **423**: 165–168.
- Hall M, Christensen K, di Collobiano SA, Jensen HJ (2002). Time-dependent extinction rate and species abundance in a tangled-nature model of biological evolution. *Phys Rev E* **66**: 1–10.
- Hall SJ, Raffaelli D (1991). Food-web patterns: lessons from a species-rich web. *J Anim Ecol* **60**: 823–842.
- Havens K (1992). Scale and structure in natural food webs. *Science* **257**: 1107–1109.
- Holland JH (1992). *Adaptation in Natural and Artificial Systems*. MIT Press: Cambridge, MA, USA.
- Holland JH, Miller JH (1991). Artificial adaptive agents in economic theory. *Am Econ Rev* **81**: 365–370.
- Hutchinson GE (1959). Homage to Santa Rosalia, or why are there so many kinds of animals? *Am Nat* **93**: 145–159.
- Krause AE, Frank KA, Mason DM, Ulanowicz RE, Taylor WW (2003). Compartments revealed in food-web structure. *Nature* **426**: 282–285.
- Lässig M, Bastolla U, Manrubia SC, Valleriani A (2001). Shape of ecological networks. *Phys Rev Lett* **86**: 4418–4421.
- Law R, Morton RD (1996). Permanence and the assembly of ecological communities. *Ecology* **77**: 762–775.
- Loeuille N, Loreau M (2005). Evolutionary emergence of size-structured food webs. *Proc Natl Acad Sci USA* **102**: 5761–5766.
- Martinez N (1994). Scale-dependent constraints on food web structure. *Am Nat* **155**: 935–953.
- Martinez ND (1991). Artifacts or attributes? Effects of resolution on the Little Rock Lake food web. *Ecol Monogr* **61**: 367–392.
- May RM (1973). *Stability and Complexity in Model Ecosystems*. Princeton University Press: Princeton, NJ.
- McCann K, Hastings A, Huxel G (1998). Weak trophic interactions and the balance of nature. *Nature* **395**: 794–798.
- McCann KS (2000). The diversity–stability debate. *Nature* **405**: 228–233.
- McKane MJ (2004). Evolving complex food webs. *Eur Phys J B* **38**: 287–295.
- Melian CJ, Bascompte J (2002). Complex networks: two ways to be robust? *Ecol Lett* **5**: 705–708.
- Memmott J, Martinez ND, Cohen JE (2000). Predators, parasitoids and pathogens: species richness, trophic generality and body sizes in a natural food web. *J Anim Ecol* **69**: 1–15.
- Montoya JM, Pimm SL, Dolé RV (2006). Ecological networks and their fragility. *Nature* **442**: 259–264.
- Montoya JM, Sole RV (2003). Topological properties of food webs: from real data to community assembly models. *Oikos* **102**: 614–622.
- Murtaugh PA, Kollath JP (1997). Variation of trophic fractions and connectance in food webs. *Ecology* **78**: 1382–1387.
- Paine RT (1992). Food-web analysis through field measurements of per-capita interaction strengths. *Nature* **355**: 73–75.
- Pimm SJ (1980). Food web design and the effect of species deletion. *Oikos* **35**: 139–149.
- Pimm SL, Lawton JH, Cohen JE (1991). Food web patterns and their consequences. *Nature* **350**: 669–674.
- Post DM, Pace ML, Hairston NG (2000). Ecosystem size determines food-chain length in lakes. *Nature* **405**: 1047–1049.
- Quince C, Higgs PC, McKane AJ (2005). Topological structure and interaction strengths in model food webs. *Ecol Modell* **187**: 389–412.
- Ray TS (1991). An approach to the synthesis of life. In: Farmer DJ, Langton C, Rassmussen S, Taylor C (eds). *Artificial Life II: Santa Fe Institute Studies in the Sciences of Complexity*, Vol 11. Addison-Wesley: Redwood City, CA, USA. pp 371–408.
- Ray TS (1994). An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artif Life* **1**: 195–226.
- Ray TS (1998). Tierra 5.0., with documentation. <ftp://alife.santafe.edu>.
- Recknagel F (2003). Simulation of aquatic food web and species interactions by adaptive agents embodied with evolutionary computation: a conceptual framework. *Ecol Modell* **170**: 291–302.
- Schmitz OJ, Booth G (1997). Modelling food-web complexity: the consequences of individual-based, spatially explicit behavioural ecology on trophic interactions. *Evol Ecol* **11**: 379–398.
- Smith M (1991). Using massively parallel supercomputers to model stochastic spatial predator–prey systems. *Ecol Modell* **58**: 347–367.
- Stauffer D, Kunwar A, Chowdhury D (2005). Evolutionary ecology *in silico*: evolving food webs, migrating population and speciation (sic). *Physica A* **352**: 202–215.
- Sugihara G, Schoenly K, Trombla A (1989). Scale invariance in food web properties. *Science* **245**: 48–52.
- Tilman D (1982). *Resource Competition and Community Structure*. Princeton University Press: Princeton, NJ.
- Townsend CR, Thompson RM, McIntosh AR, Kilroy C, Edwards E, Scarsbrook MR (1998). Disturbance, resource supply and food-web architecture in streams. *Ecol Lett* **1**: 200–209.
- Warren PH (1989). Spatial and temporal variation in the structure of a freshwater food web. *Oikos* **55**: 299–311.
- Warren PH (1994). Making connections in food webs. *Trends Ecol Evol* **4**: 136–140.
- Watts DJ, Strogatz SH (1998). Collective dynamics of ‘small-world’ networks. *Nature* **393**: 440–442.
- Williams RJ, Martinez ND (2000). Simple rules yield complex food webs. *Nature* **404**: 180–183.
- Winemiller KO, Pianka ER, Vitt LJ, Joern A (2001). Food web laws or niche theory? Six independent empirical tests. *Am Nat* **158**: 193–199.
- Yedid G, Bell G (2001). Microevolution in an electronic microcosm. *Am Nat* **157**: 465–487.
- Yedid G, Bell G (2002). Macroevolution simulated with autonomously replicating computer programmes. *Nature* **420**: 810–812 (News & Views 420: 756).

Supplementary Information accompanies the paper on Heredity website (<http://www.nature.com/hdy>)