REVIEW AND SYNTHESIS

Economic game theory for mutualism and cooperation

Marco Archetti,[1,2]* István Scheuring,[3] Moshe Hoffman,[4] Megan E. Frederickson,[5] Naomi E. Pierce[1]* and Douglas W. Yu[6,7]*

Abstract

We review recent work at the interface of economic game theory and evolutionary biology that provides new insights into the evolution of partner choice, host sanctions, partner fidelity feedback and public goods. (1) The theory of games with asymmetrical information shows that the right incentives allow hosts to screen-out parasites and screen-in mutualists, explaining successful partner choice in the absence of signalling. Applications range from ant-plants to microbiomes. (2) Contract theory distinguishes two longstanding but weakly differentiated explanations of host response to defectors: host sanctions and partner fidelity feedback. Host traits that selectively punish misbehaving symbionts are parsimoniously interpreted as pre-adaptations. Yucca-moth and legume-rhizobia mutualisms are argued to be examples of partner fidelity feedback. (3) The theory of public goods shows that cooperation in multi-player interactions can evolve in the absence of assortment, in one-shot social dilemmas among non-kin. Applications include alarm calls in vertebrates and exoenzymes in microbes.

Keywords

Coevolution, common-pool resource, cooperation, game theory, host sanctions, mutualism, N-person prisoner's dilemma, public goods, symbiosis, volunteer's dilemma.

Ecology Letters (2011) **14**: 1300–1312

"I am an economist, but I am also what we might call an evolution groupie." Paul Krugman (1996).

THREE PROBLEMS OF COOPERATION

It is often stated that explaining the existence of cooperation in nature is one of the greatest challenges for evolutionary theory. While this is possibly true, it is useful to break the challenge into three, separate components.

Consider, for example, the interaction between bioluminescent bacteria, *Vibrio fischeri*, and the bobtail squid, *Euprymna scolopes*: the bacteria provide light (at a cost) for the squid, which in turn houses and feeds the bacteria (Ruby & McFall-Ngai 1999; Nyholm & McFall-Ngai 2004). At least three problems arise. First, before the interaction is established, there is a *hidden characteristics* problem: the qualities of potential bacterial symbionts are unknown to the squid, so how can the squid recruit the right ones? Second, after the bacteria have established, there is a *hidden actions* problem: how does the squid prevent the bacteria from misbehaving? Third, since bacteria colonise the squid in groups, there is a *collective action* problem: why do individual bacteria luminesce, when each could gain fitness by not lighting up?

Some solutions to these problems have long been known in evolutionary biology: signalling for the hidden characteristics problem;

host sanctions or partner fidelity feedback for the hidden actions problem; kin selection, spatial structure, punishment and/or repeated games for the collective action problem. However, this set of solutions is incomplete. Signalling is not always possible in partner choice, it has not been clear how to define and recognise host sanctions and partner fidelity feedback, and cooperation without iterations or spatial structure has been observed among non-kin.

We review recent work at the interface with microeconomics that provides some missing parts: *screening theory* solves the hidden characteristics problem when signalling is not possible; *contract theory* defines and differentiates host sanctions and partner fidelity feedback; and the theory of *public goods* shows that cooperation is possible even with one-shot interactions in well-mixed populations of non-kin. Deeper engagement with game theory will help evolutionary biology better explain cooperation and mutualism in nature.

PARTNER CHOICE VIA SCREENING

The problem of hidden characteristics

The study of partner choice, in which individuals 'differentially reward cooperative vs. uncooperative partners *in advance of any possible exploitation*' (Bull & Rice 1991, p. 63; italics added; Noë &

[1]*Department of Organismic and Evolutionary Biology and Museum of Comparative Zoology, 26 Oxford St., Harvard University, Cambridge, MA 02138, USA*

[2]*Department of Business and Economics, University of Basel, Peter Merian-Weg 6, Basel 4002, Switzerland*

[3]*Research Group of Theoretical Biology and Ecology, Department of Plant Taxonomy and Ecology, Eötvös University and HAS, Pázmány Péter sétány 1C, Budapest H-1117, Hungary*

[4]*Rady School of Management, University of California, San Diego, Otterson Hall, 9500 Gilman Drive 0553, La Jolla, CA 92093, USA*

[5]*Department of Ecology & Evolutionary Biology, University of Toronto, 25 Harbord Street, Toronto, ON M5S3G5, Canada*

[6]*Ecology, Conservation, and Environment Center (ECEC), State Key Laboratory of Genetic Resources and Evolution, 32 Jiaochang East Rd., Kunming Institute of Zoology, Chinese Academy of Sciences, Kunming, Yunnan 650223, China*

[7]*School of Biological Sciences, University of East Anglia, Norwich Research Park, Norwich, Norfolk NR47TJ, UK*

Correspondence: E-mail: archetti.marco@googlemail.com, npierce@oeb.harvard.edu, dougwyu@gmail.com

Hammerstein 1994), has a long tradition in evolutionary biology. The classic example is when females choose a partner from among different males. Information about male quality is asymmetrical: each male 'knows' his own quality, in the sense that he can behave accordingly, but the female does not know each male's quality. Because all potential partners would benefit from being chosen, the low-quality ones have no interest in revealing their true quality. How does a female choose the highest-quality partner? This problem of *hidden characteristics* is clearly not restricted to mate choice; it arises in any case of partner choice with asymmetrical information: for example, how does a host choose the right symbiont among many candidates when it is not possible to assess their qualities in advance?

A well-known solution is signalling (Grafen 1990; Maynard Smith & Harper 2003). Partner quality is revealed through the display of costly phenotypes whose production is profitable only for individuals of sufficiently high quality. Examples include males signalling to females in birds (Hamilton & Zuk 1982; Folstad & Karter 1992) and trees signalling to parasites through red autumn leaves (Archetti 2000, 2009a; Doring *et al.* 2009). But what happens when signalling seems impossible, such as between a bacterium and a host or an insect and a plant? (Nyholm & McFall-Ngai 2004; Heath & Tiffin 2009; Heil *et al.* 2009; Barke *et al.* 2011).

Screening theory

An alternative to signalling has been proposed recently (Archetti *et al.* 2011; Archetti 2011), based on the theory of games with asymmetric information, which has similarities with the concept of 'screening' in microeconomics (Rothschild & Stiglitz 1976). The partner choice problem is a sequential game. In signalling, the player whose characteristics are hidden (e.g. the male in mate choice) sends a signal about his own quality (implemented in the game by allowing him to move first), which the second mover (e.g. the female) evaluates to decide whether or not to accept the first mover.

In screening, instead, the first mover is the player that wants to discover the hidden characteristics (the 'principal', the host); the informed player (the 'agent', the symbiont) moves second (thereby ruling out any possibility of signalling). It might seem that efficient partner choice is impossible when the informed player cannot signal. Instead, one of the most remarkable features of games with asymmetric information is that a solution is still possible: the principal can set up rewards and costs for the agents to enter the interaction (in a job market, this would be a contract), which the agent then decides whether or not to accept (and hence whether or not to enter the interaction). If the principal designs the rewards and the costs appropriately, so that only high-quality agents find it profitable to enter the interaction, potential symbionts will screen themselves in or out according to their own quality, even though the principal observes neither quality nor signal (Fig. 1). In fact, the principal can achieve separation of high and low-quality types without ever observing their choice.

It is useful to note that screening is similar to signalling. In signalling, agents send costly information to principals; in screening, principals impose a strategically costly task on agents. In both games, effectiveness depends on the strength of the correlation between the cost and the quality desired by the principal: the cost (for the agent) must be higher for the low-quality (for the principal) partners. Which mechanism is used to solve the hidden characteristics problem depends on which is more feasible. For instance, in human societies,

signalling seems more feasible for, say, job searches, but screening is commonly used in insurance markets. Auto-rescue companies desire to separate low-quality from high-quality cars, and make the former pay more. Because car qualities are partly hidden, one screening mechanism is to charge a lower monthly premium if the customer is willing to forgo rescue at home (the cost of entering the interaction). Customers with balky cars that might fail to start in the morning reveal themselves by paying extra for a 'homestart' option, but customers with high-quality cars find that taking on the (for them, low) extra risk is worth the benefit of a lower premium. In contrast, a signalling approach would require expensive inspections.

Biological screening

The partner choice problem applies to a variety of mutualisms. One of the clearest examples would appear to be the marine bobtail squid, *Euprymna scolopes*, and its bioluminescent bacterial symbiont, *Vibrio fischeri* (reviews in Ruby & McFall-Ngai 1999; Nyholm & McFall-Ngai 2004; Archetti *et al.* 2011). What makes this system challenging is that the bacteria are not vertically transmitted but taken up anew from
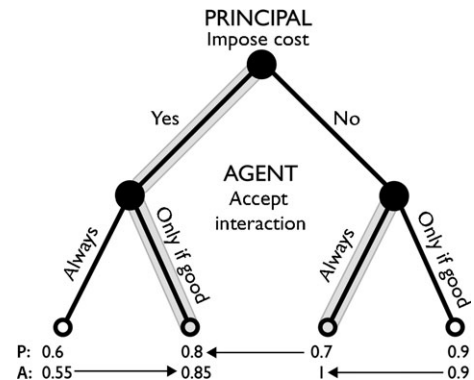

Figure 1 A screening game example. Two types of agents exist: good and bad quality. Quality is not observed by the principal. The principal can impose a cost for establishing the interaction, and the agents can 'always accept' the interaction (irrespective of being bad or good) or 'accept only if good quality.' For this example, let the frequency of good quality agents be ½. Black circles show a player's move; the principal moves first. Open circles show the payoffs to principal (P, top) and agent (A, bottom). Arrows show increases in payoff. The payoff for both agents and principal in the absence of any interaction is 0.8; a principal pays a cost 0.6 for interacting with a bad agent and no cost for interacting with a good agent; if the principal imposes a cost, a good agent pays 0.1 for the interaction, while a bad agent pays 0.8. The principal pays 0.1 for imposing a cost. The payoff for an agent that always accepts the interaction is 1 when the principal does not impose a cost and $[\frac{1}{2}(1 - 0.8) + \frac{1}{2}(1 - 0.1)] = 0.55$ when the principal imposes a cost. The payoff for an agent that accepts the interaction only when he is good is $[\frac{1}{2}(1 - 0.1) + \frac{1}{2}(0.8)] = 0.85$ when the principal imposes a cost and is $[\frac{1}{2}(1) + \frac{1}{2}(0.8)] = 0.9$ when the principal does not impose any cost. Assume that the principal encounters as many agents as he needs. When all agents accept only if good, the principal's payoff is $[\frac{1}{2}(0.8) + \frac{1}{2}(1 - 0) - 0.1] = 0.8$ if the principal imposes a cost and $[\frac{1}{2}(0.8) + \frac{1}{2}(1 - 0)] = 0.9$ if the principal does not. When all agents always accept the interaction, the principal's payoff is $[\frac{1}{2}(1 - 0.6) + \frac{1}{2}(1 - 0) - 0.1] = 0.6$ if the principal imposes a cost and $[\frac{1}{2}(1 - 0.6) + \frac{1}{2}(1 - 0)] = 0.7$ if the principal does not. Use backward induction (start from the agent's possible moves) to find the equilibria. The grey branches show the dominant strategies for each subgame. The best strategy for the agent is to 'accept only if good' if the principal imposes a cost (because 0.85 > 0.55) and to 'always accept' if the principal does not impose any cost (because 1 > 0.9). Therefore, the best strategy for the principal is to impose a cost (because 0.8 > 0.7). (Modified from Archetti *et al.* 2011).

seawater by each squid hatchling. *V. fischeri* is the only bacterial species that actively colonises the squid's light organ, even though the light organ provides enough food and housing to raise seawater densities of *V. fischeri* orders of magnitude higher than free-living bacteria. Note that the squid must screen out not only non-luminescent *V. fischeri* but also closely related pathogens.

Ruby & McFall-Ngai (1999) largely explained how the squid solves the hidden characteristics problem by pointing out that bacterial bioluminescence is a defence against the host's threat of reactive oxygen species (ROS), which are lethal to bacteria. The bacterial enzyme, luciferase, consumes O_2 in the squid's light organ and lowers concentrations to levels that prevent the squid's enzymes from converting O_2 to H_2O_2, hypochlorous acid and other ROSs. A by-product of luciferase action is light. Archetti *et al.* (2011) interpret the threat of ROS as a cost of interaction imposed by the squid host, which can be supported by high-quality (luciferase-producing) *V. fischeri*, given that they receive high levels of food. The cost of interaction cannot be supported by low-quality (non-luciferase-bearing) bacteria, which cannot prevent lethal ROS build-up. The constitutive presence of ROS is, in short, the mechanism evolved by the squid for screening good and bad bacteria. As a result, only luminescent *V. fischeri* should evolve to colonise a squid host that produces ROS.

ROS-based screening could have wide applicability in microbiomes. Selva *et al.* (2009) showed that sublethal H_2O_2 exposure can still trigger the lysis of a pathogenic bacterium, *Staphylococcus aureus*, because the H_2O_2 causes DNA damage that induces the SOS DNA repair mechanism. The SOS response inadvertently activates phage DNA that is often contained in *S. aureus* genomes, and the phage produces new viral particles and lyses the host cell. The connection to screening arises because virulence genes are preferentially located in phage DNA ('pathogenicity islands' Sui *et al.* 2009), which means that enhanced susceptibility to ROS damage is correlated with an increased probability of being pathogenic. Interestingly, many animal tissues that house bacterial symbionts contain high concentrations of the ROS-producing enzyme halide peroxidase (Ruby & McFall-Ngai 1999), and the roots of leguminous plants release ROS bursts during initiation of symbiosis with nitrogen-fixing bacteria (Soto *et al.* 2009). These might be mechanisms to screen out pathogens.

Alternatively, the host could indirectly impose an interaction cost by fomenting competition between symbiont types in ways that favour high-quality types. We call this 'competition-based screening.' The concept seems new for economics, but it has a number of possible examples in biology. One is given by several Mexican ant-acacia species, which are colonised by a variety of *Pseudomyrmex* ant species, some of which are mutualistic and engage in costly defence of their host and some of which are parasitic and do not defend (Heil *et al.* 2009). The identity of the ant symbiont is established when an ant queen colonises an ant-plant sapling, well before the plant can judge the quality of the ant, even if it could. Screening is suggested by the striking observation that acacia species that produce high levels of food and housing are $\sim$ 90% inhabited by mutualistic ant species, whereas acacia species that produce low levels are only $\sim$ 50% inhabited by mutualists, with the remainder colonised by parasites. Because ant-plant growth is dependent on successful protection by the ant and because competition between co-colonising ant symbionts is 'massive' (Heil *et al.* 2009), when ant queens colonise different parts of the same sapling, the colony that does a better job of protecting grows its portion of the ant-plant more quickly and evicts the other colony. This competition should be stronger when there is more food

(Heil *et al.* 2009), thereby increasing the probability of establishment by mutualistic ant species on high-reward hosts.

Competition-based screening can only occur if both symbiont types co-occur in a host, which injects stochasticity into the outcome, but symbionts can engage each other in ways that hosts cannot do, especially in mixtures. For instance, Eswarappa *et al.* (2009) have shown that the *lac* operon repressor gene, *lacI*, pleiotropically represses virulence in the pathogen *Salmonella enterica*. The fact that lactose is the major component of mammal milk (e.g. Zivkovic *et al.* 2010) might screen out pathogens from infant guts by conferring a growth advantage to microbes bearing the *lac* operon.

Testing screening models

It is a challenge to test screening models. A place to start is 'strategic waste,' which are resources expended solely for sending or inferring hidden information. In signalling, bright feathers constitute strategic waste because the carotenoid compounds responsible for the bright colouration are not available for immunostimulation. Likewise, if purely for screening, ROS-producing enzymes in the light organ of the bobtail squid should have no function other than to threaten ROS production, and in the *Acacia* species (Heil *et al.* 2009), high resource levels are wasted to the extent that they exist to exacerbate competition; e.g. more extra-floral nectar fuels worker aggression.

Another partial test is to map empirical systems onto detailed models. In screening, the design challenge for the host is to adjust the rewards and costs so that symbiont characteristics are revealed efficiently. It is possible to calculate optimal values with constrained optimisation models from economics and to compare parameter values with biological features (see Archetti *et al.* 2011 for a qualitative example).

A better test would be to contrast (microbial) evolution in experimentally created demanding and non-demanding host environments: symbionts would be expected to evolve screening (enter the interaction based on their quality) only when coevolving with demanding hosts. To our knowledge, no such experiments exist, but insect and coral symbioses with Actinobacteria, which produce antimicrobials to protect their hosts from pathogens (Ritchie 2006; Kaltenpoth 2009; Barke *et al.* 2011), might prove tractable. Antibiotics are thought to serve as interference competition mechanisms among bacteria, so we hypothesise that hosts selectively recruit Actinobacteria by fomenting competition for host resources in ways that favour interference competitors, despite the added cost.

A potentially important effect of screening is that it limits interactions of the host to screened-in symbionts, which could set the stage for pair wise co-adaptation, as suggested by the plethora of mutual recognition mechanisms in endosymbioses (e.g. Nyholm *et al.* 2009; Chaston & Goodrich-Blair 2010). In addition, screening is more likely to evolve if the host's demanding environment is derived from defence traits against parasites, because such defences can be pre-adaptations for screening (Archetti *et al.* 2011; Archetti 2011).

HOST SANCTIONS OR PARTNER FIDELITY FEEDBACK?

The problem of hidden actions

The problem of asymmetric information is not limited to partner choice. After a symbiont has colonised a host, what then prevents the symbiont from misbehaving? The problem is that, in most cases,

symbiont actions cannot be observed: this is a problem of *hidden actions*. But if the host is not able to observe symbiont behaviour and condition rewards and punishments on that behaviour, why do ants patrol, microbes light up, or yucca moths abstain from ovipositing? It might seem that mutants that do not do these things would have an advantage.

The hidden action problem, under different names, has attracted much attention in the mutualism literature (Connor 1986; Bull & Rice 1991; Pellmyr & Huth 1994; West *et al.* 2002a,b; Kiers *et al.* 2003; Sachs *et al.* 2004; Edwards *et al.* 2006; Foster & Kokko 2006; Foster & Wenseleers 2006; Kiers *et al.* 2006; Szilágyi *et al.* 2009; Weyl *et al.* 2010), and two broad concepts, host sanctions (HS) and partner fidelity feedback (PFF), have been developed to solve the problem.

To distinguish these two concepts, consider the mutualism between yucca moths and yucca flowers. Female moths carrying pollen visit yucca flowers, where they deposit pollen on stigmas and lay eggs inside ovaries. The developing moth larvae consume some of the plant's seeds, so there is a conflict of interest between yucca moths, whose fitness depends on the number of larvae that develop into adults, and yucca plants, whose fitness depends on the number of seeds produced. In a classic article, Pellmyr & Huth (1994) showed that yucca plants selectively abort flowers into which moths have laid too many eggs, thereby setting a limit to the number of larvae that can develop (larvae in aborted flowers die). The authors foreshadowed the mutualism literature when they wrote: 'It is not known whether the mechanisms [of abortion] represent co-evolved traits or pre-adaptations within the yuccas, but this does not affect their significance in contributing to maintaining a mutualistic equilibrium. It will be important to explore, however, for understanding the role of coevolution in highly specialized interactions.'

Neither the term host sanctions or partner fidelity feedback was yet in circulation, but in effect, Pellmyr and Huth were asking, 'is selective abortion (1) a response to cheating behaviour by moths or (2) a generalised response to floral damage, which is sometimes caused by oviposition?' We call the first HS, and the second PFF.

The distinction between PFF and HS host response also appears in the (within-species) social evolution literature. HS corresponds to 'punishment' or 'negative reciprocity' and is defined as a behaviour that benefits the punisher *as a result of* more cooperative behaviour of the punished individual (Clutton-Brock & Parker 1995). By contrast, PFF is 'negative pseudoreciprocity' (Bshary & Bergmüller 2008) and constitutes so-called 'self-serving punishment,' which is selected for *irrespective of* its effect on the punished individual. In other words, the cost of an HS response must be compensated for by its effects on symbiont behaviour, but PFF response need not be.

An example from ant-plants helps to clarify PFF further. Experimental removal of ants from new shoots results in higher levels of herbivory and eventual shoot death (e.g. Edwards *et al.* 2006 and included references), which is detrimental for the ants. The PFF interpretation is that once herbivory renders a shoot worthless to the plant, the plant's best (most self-serving) response is to abscise the remainder to save on maintenance costs. Since shoot loss happens to reduce ant housing, and thus fitness (Edwards *et al.* 2006), ants are selected to protect plant shoots. We can reverse signs; if the symbiont benefits the host, the most self-serving response could be growth or reproduction. In short, PFF hosts do not look back in anger at symbiont misbehaviour, but if the host's fitness and/or reallocation of resources feeds back sufficiently strongly to the fitness of the symbiont(s) that committed the action, PFF selects for mutualism.

The concept of credible commitment

Even this brief survey indicates the unwieldiness of the terms describing host response to misbehaviour and the difficulty of separating the two concepts. To further complicate matters, both terms, HS and PFF, have been used to describe the *concept* of PFF (e.g. West *et al.* 2002a use HS to describe the efficient reallocation of host resources). Therefore, to provide clarity and facilitate empirical work, Weyl *et al.* (2010) exploited economic contract theory to define and differentiate PFF and HS.

To describe this work, we first introduce the concept of credible commitment (Schelling 1960). A famous example is that of two drivers playing the game of Chicken (Rapoport & Chammah 1966): if neither veers, they crash into each other, and both lose, but each prefers to have the other veer. The winner is the first to commit credibly not to veer. For instance, a driver might toss her steering wheel out of the window in full view of the other. A credible commitment to limit one's own options can induce the opponent to adopt a certain behaviour.

In general, commitments are difficult to make credible and visible and, moreover, are understudied in biology, especially in mutualisms. However, examples of commitment do appear to exist in nature. For instance, threat displays are more likely to be honest when given within 'striking distance' from the receiver (Számadó 2008). Proximity acts as commitment because the option of escape is eliminated, and examples reviewed by Számadó (2008) suggest that only such close-up threats reliably elicit responses from receivers. Számadó (2011) has also argued that low-cost 'badges of status,' such as melanic patches, can be honest signals of quality when the badges commit the wearers to engage in multiple fights before winning a contested resource, like a territory or high rank. Since weak fighters are unlikely to win regularly against strong fighters, it is unprofitable to advertise a higher status than one is able to defend.

The presence or absence of commitment matters greatly to HS *in one-shot interactions*, because HS derives its benefits from increased cooperative behaviour by the symbiont. In repeated interactions, symbionts can change behaviour after the host's first move, but in one-shot interactions, the host's response to the symbiont's behaviour must be already in place before the symbiont moves (and the response must be evident and credible to the symbiont). If it were possible that the best response of the host to cheating by the symbiont were *not* to punish (e.g. to mature any remaining seeds in the case of yucca-moth mutualism), then the symbiont would have a temptation to cheat.

Contract theory differentiates host sanctions from partner fidelity feedback

Using an economics framework, an agent (symbiont) carries out an action, which results in one of many possible signals (outcomes) for the principal (host). The probability of any particular signal depends on the action (e.g. the signal of leaf loss is more probable if the symbiont ant shirks). (A signal in this context does not refer to communication). For consistency, we now use biological terms, but stick with the original variables from Weyl *et al.* (2010). HS says that the host responds to symbiont action, which has been inferred from the outcome (as the action itself is hidden, by definition). PFF says that the host just responds to the outcome.

Formally, let the symbiont A take an action $a \in \{a^*, a_1, a_2\}$, where a^* is the cooperative action, and a_1 and a_2 are two cheating actions. The host P observes an outcome $s \in S$ (S = all possible outcomes),

which is correlated with the action taken by the symbiont. Let $p(s|a)$ be the conditional probability of outcome s, given action a. The host then has a response (policy) π that can depend on s, $\pi(s) \in [0, \infty]$. Let the host's fitness (utility) for a given outcome and response be $U^P(s, \pi(s))$. Let the symbiont A's fitness U^A following its action a be the sum of the short-term fitness contribution from the action $u^s(a)$, plus any feedback from the effect of the action on the host P's fitness, plus any direct effect of the host response π on the symbiont. An exogenous parameter f scales the strength of the feedback. A pollinator might feel no host-fitness-feedback from its actions, whereas the fitness of a vertically transmitted endosymbiont is strongly coupled with its host's fitness.

$$U^A(a, s, \pi(s)) = u^s(a) + f\, U^P(s, \pi(s)) - \pi(s),$$

yielding the expected fitness of the agent for an action, given a host response,

$$U^A(a|\pi(\cdot)) = u^s(a) + \sum_{s \in S} \left[f\, U^P(s, \pi(s)) - \pi(s) \right] p(s|a).$$

PFF assumes that (*i*) the host cannot commit to punish, yet (*ii*) cooperation is nonetheless maintained. Formally, (*i*) implies that the host evolves the response $\pi^{PFF}(s)$ that is optimal for *each* s, and (*ii*) implies that the symbiont receives higher expected utility from cooperating than from cheating, given π^{PFF}, i.e.

$$U^A\big(a^*|\pi^{PFF}(\cdot)\big) \geq U^A\big(a_1|\pi^{PFF}(\cdot)\big), U^A\big(a_2|\pi^{PFF}(\cdot)\big),$$

where, for all s, $\pi^{PFF}(s) = \underset{\pi(s)}{argmax}\ U^P(s, \pi(s))$,

meaning that it is in the host's interest (self-serving) to respond in ways that happen to reduce symbiont fitness sufficiently when the symbiont cheats, so that the symbiont finds it in its interest not to cheat. Thus, natural feedback is sufficient to select for a^*.

In contrast, HS says that for some $a_i \in \{a_1, a_2\}$,

$$U^A\big(a^*|\pi^{PFF}(\cdot)\big) < U^A\big(a_i|\pi^{PFF}(\cdot)\big),$$

meaning that it is in the symbiont's short-term interest to cheat, unless the host responds to cheating with an action that is against the host's interest once the cheating has occurred. A similar problem exists in economics (known as *moral hazard*; Holmström 1979).

HS assumes that the host can commit, and commits to the optimal response policy that maintains cooperation, i.e.

$$\pi^{HS}(s) = \underset{\pi(s)}{argmax} \sum_{s \in S} U^P(s, \pi(s)) p(s|a^*)$$

subject to the constraint:

$$U^A\big(a^*|\pi(\cdot)\big) \geq U^A\big(a_1|\pi(\cdot)\big), U^A\big(a_2|\pi(\cdot)\big),$$

meaning that it is in the host's interest, before the symbiont has chosen an action, to commit to punish sufficiently in the case that the symbiont chooses to cheat, so as to deter cheating by the symbiont, given that the symbiont anticipates such punishment.

Testing for PFF and HS

Weyl *et al.* used this formalism to design tests for PFF and HS. In the first test, responding to outcomes (PFF) means that the host should respond equally to different experimental manipulations that result in a

symbiont taking the same value for the host, even if the natures of the two manipulations imply different symbiont actions. For instance, in the legume-rhizobia mutualism, supplementing the soil with a saturating level of nitrate fertiliser (manipulation 1) drives the value (to the host) of rhizobial bacteria to zero, *even if the rhizobia continue fixing*, because the nitrates are a free source of fixed nitrogen to the legume host, whereas rhizobia-provided ammonium is always costly. Removing N_2 from the atmosphere (manipulation 2) equally destroys the value of rhizobia. Manipulation 2 also implies that rhizobia are cheating (because plants cannot directly sense N_2), and, thus, an HS host should respond more strongly to manipulation 2. Kiers *et al.* (2006) actually conducted these experiments and reported the same final rhizobial nodule densities after both manipulations, which supports PFF.

The second test can be used on systems for which one can identify two cheating actions, a_1 and a_2. If a_1 and a_2 provide the *same benefits to symbionts*, then an HS host should respond equally to both (because HS hosts respond to symbiont action). But if, additionally, a_1 results in *more damage to the host*, a PFF host should respond more to a_1 (because PFF hosts respond to the effects of symbiont action like to any environmental setback). The yucca-moth mutualism provides an example. There are two defector moth types, deep (a_1) and shallow (a_2) ovipositor. 'Deeps' and 'shallows' receive approximately the same benefits, so HS hosts should respond similarly to both. However, only deep oviposition results in direct damage to seeds (damage from a_1 > damage from a_2). Consistent with PFF, *Yucca kanabensis* senesces flowers after 'deep' oviposition, but senesces no flowers after similar amounts of 'shallow' oviposition (Marr & Pellmyr 2003).

Evolutionary implications of PFF and HS

Why is it useful to differentiate response to action from response to outcome? The distinction seems subtle, but the two host responses imply different evolutionary origins.

PFF implies pre-adaptation. That is, it is parsimonious to infer that host response to outcomes evolved before the symbiosis. Floral abortion, for example, is widespread among angiosperms and is thought to be 'a pre-adaptation within the yuccas' (Pellmyr & Huth 1995, p. 558). Furthermore, yuccas abort flowers not only when they have too many moth eggs, but also in response to self-pollen, pollen deficit and feeding damage by florivores (Richter & Weis 1995, 1998; Huth & Pellmyr 2000). In plants generally, shoot, root and flower abscission are responses to damage or resource deficits, and this response policy $\pi^{PFF}(s)$ reduces symbiont fitness to the extent that abscission also happens to affect the symbiont that caused the damage or deficit. Because plants respond modularly to environmental shocks, PFF response tends to be directed at affected plant parts, and thus at the responsible symbionts. In vertically transmitted symbioses, the feedback arises because vertical transmission results in an exogenously high value of f.

In contrast, HS implies adaptation within the context of the mutualism, since the definition of HS requires the symbiont to be present. But theory has found that host sanctions do not invade when the symbiont population is uniformly composed of defectors or cooperators (West *et al.* 2002a; Foster & Kokko 2006; Weyl *et al.* 2010), so the origin of HS, if it exists, is currently unexplained. HS also implies commitment if host response is costly and if the interaction is one-shot, as in many mutualisms. As far as we know, there has been no attempt yet to find commitment in mutualisms, and any example will be of great interest. An intriguing candidate for HS

are the photoreceptors in the bobtail squid's light organ, which can monitor the luminescence (actions) of symbiotic bacteria (Tong *et al.* 2009).

All this said, it is also possible, in principle, for PFF response to be an adaptation within the context of the mutualism, as long as the host response represents a general and optimal response to any damage caused by the symbiont, such that if the same damage were caused by the environment, the host would react in the same way. In practise, pre-adaptation is the parsimonious explanation for PFF response.

This last point reveals some further advantages of having used game theory to model host response. Firstly, HS and PFF are now formally and clearly differentiated (response to action vs. response to outcome), which helps avoid the error of differentiating the two models by their implications (e.g. adaptive vs. pre-adaptive). Secondly, the possible role of commitment in host response is revealed. Thirdly, the formalism allows the design of experimental tests to differentiate the two models, whereas current practice typically argues in favour of one or the other on the grounds of plausibility. Importantly, HS and PFF both predict host response to cheating, so merely demonstrating host response cannot tell the two apart. One day, we might know the coevolutionary histories of the relevant molecular pathways and identify mechanisms directly, but until then, we must interpret the data we can get with theory designed to facilitate that interpretation.

PUBLIC GOODS: COEXISTENCE OF COOPERATORS AND DEFECTORS

The problem of collective action

When interactions occur among more than two players, the problem of cooperation is usually described as follows. Cooperators pay a cost for making a contribution that defectors do not pay; the contributions are gathered, transformed in some way, and then redistributed to everybody (including to any of the defectors). A *collective action problem* arises: defectors can free ride on the efforts of the cooperators, so how is cooperation stable (Olson 1965)? Social dilemmas of this kind occur at all levels of biological organisation, from self-promoting genetical elements (e.g. Burt & Trivers 2006), to microbes (e.g. Crespi 2001), to vertebrates (e.g. Creel 1997; Bednarz 1988), to human societies (e.g. Hardin 1968; Kollock 1998).

The solution to the problem of cooperation is usually believed to be some form of positive assortment among cooperators. Assortment can be due to kin selection (Hamilton 1964; reviewed by Frank 1998; Grafen 2009) brought about, for example, by kin discrimination or spatial structure (Grafen & Archetti 2008; Lehmann *et al.* 2007) or by repeated interactions, which allow reciprocation, reputation effects, or punishment (Axelrod & Hamilton 1981). This section reviews recent articles showing that cooperation in *N*-player social dilemmas can be maintained even in the absence of all assortment mechanisms.

To understand cooperation in *N*-player interactions, we must introduce the concept of the *public good* (Samuelson 1954), which is any good that is non-excludable (no one can be prevented from using it) and non-rivalrous (consumption does not reduce availability to other users). In practise, public goods are often partially excludable (*club* goods) or rivalrous (*common* goods), but club and common goods can be treated approximately as public goods.

Linear public goods

The simplest model of the production of a public good in *N*-player interactions is the *N*-player version of the prisoner's dilemma (NPD; Hamburger 1973): individuals can be cooperators or defectors; cooperators pay a contribution c that defectors do not pay. Contributions are summed, the sum is multiplied by a reward factor r (> 1), and then redistributed to all individuals (both cooperators and defectors). If groups of N individuals are formed at random, i.e. without assortment, the probability of having j cooperators among the other N-1 individuals is $f_j(x)$. In an infinitely large population, $f_j(x)$ would be the probability mass function of a binomially distributed random variable with parameters ($N-1$, x), where x is the frequency of cooperators in the population. The fitness of a defector and of a cooperator are, respectively:

$$W_D = \sum_{j=0}^{N-1} f_j(x) \frac{rcj}{N},$$

$$W_C = \sum_{j=0}^{N-1} f_j(x) \left[\frac{rc(j+1)}{N} - c \right] = W_D + c\left(\frac{r}{N} - 1 \right)$$

If $r/N < 1$, which is reasonable in any sizeable group, at equilibrium, everybody defects. Hence, in the NPD, cooperation requires additional forces like spatial structure, kin discrimination, or repeated interactions to allow cooperators to interact preferentially with cooperators (if $r/N > 1$, there is no dilemma, and everybody cooperates). The NPD, however, is only a special case of the public goods game. It relies on the critical assumption that the public good is a *linear* function of the individual contributions. This apparently innocuous assumption has fundamental consequences, which have been largely ignored until recently.

Is linearity at least justified from an empirical point of view? The answer is a firm no: nonlinear public goods are the rule in biology. For instance, in microbes producing extracellular enzymes (Crespi 2001; Lee *et al.* 2010; Gore *et al.* 2009; Chuang *et al.* 2010), the benefit of enzyme production is generally a saturating function (Hemker & Hemker 1969), a sigmoid function (Ricard & Noat 1986) or a step function (Mendes 1997; Eungdamrong & Iyengar 2004) of its concentration. Similarly, in vertebrate social behaviour, benefits from alarm calls, group defence and cooperative breeding are also saturating functions of the number of cooperators (Rabenold 1984; Bednarz 1988; Packer *et al.* 1990; Stander 1991; Creel 1997; Yip *et al.* 2008). As far as we know, no evidence of linear public goods exists except for human systems, and the NPD is therefore not a biologically realistic model of *N*-player interactions except in human experiments. Nonetheless, it has become so customary to use the NPD in the study of cooperation, that the two terms (NPD and public goods game) have become virtually synonymous in evolutionary biology. This is unfortunate, as nonlinear public goods games have completely different results. To see why, let us keep the assumption of random group formation (no assortment) and relax, instead, the assumption that the public good is a linear function of the individual contributions.

Nonlinear public goods

We define an *N*-player public goods game by a benefit $\beta(j)$ monotonically increasing in j (the number of cooperators) and a cost $\gamma(j)$ monotonically decreasing in j. Note that the benefits need not be

equal for cooperators and defectors, since any difference in benefit can be incorporated in the cost γ, without loss of generality (thus, we can include scenarios in which cooperators benefit more from their own contributions than do defectors). The average fitness of a defector and of a cooperator are, respectively:

$$W_D = \sum_{j=0}^{N-1} f_j(x)\beta(j)$$

$$W_C = \sum_{j=0}^{N-1} f_j(x)[\beta(j+1) - \gamma(j+1)]$$

As we are assuming no assortment, in an infinitely large population

$$f_j(x) = \binom{N-1}{j} x^j (1-x)^{N-1-j}.$$

The simplest nonlinear benefit function is the Heaviside step function:

$$\beta(i) = \begin{cases} 0 & \text{if } i < k \\ b & \text{otherwise} \end{cases},$$

and we can assume that $\gamma(i) = c$. In economics, this game is sometimes called the Volunteer's Dilemma if $k = 1$ (Diekmann 1985) or the Teamwork Dilemma if $k > 1$ (Myatt & Wallace 2009; see also Palfrey & Rosenthal 1984). For brevity, we use Volunteers' Dilemma (VD) to cover any N-player game in which the public good is produced only if at least k ($\leq N$) individuals pay a cost c. The dilemma is that each individual would rather avoid the cost c of volunteering and exploit the public good produced by others, but if the public good is not produced, everybody pays a cost higher than that of volunteering. The equilibrium frequency of cooperators x_{eq} can be found exactly (numerically) from (Archetti 2009b,c)

$$\binom{N-1}{k-1} x_{eq}^{k-1} (1-x_{eq})^{N-k} = \frac{c}{b}$$

In large groups, the approximate analytical solution (Archetti & Scheuring 2011) is

$$x_{eq}^{\pm} = \frac{k-1}{N-1} \pm \delta(k, N, c/b),$$

where $\delta(k, N, c/b) \ll \frac{k-1}{N-1}$, and x_{eq}^{-} is an unstable, while x_{eq}^{+} is a stable equilibrium.

At equilibrium, therefore, cooperators and defectors coexist (Fig. 2), and, in large groups, the frequency of cooperators is approximately the frequency of volunteers (k/N) necessary to produce the public good. This mixed equilibrium (or polymorphic equilibrium of defectors and cooperators), however, exists only if the c/b ratio is below a critical threshold (Fig. 2) (Archetti & Scheuring 2011). The equilibrium is still Pareto-inefficient (an improvement is possible by which at least one player could receive a higher payoff without reducing any other player's payoff: hence the inefficiency; no one, however, has an incentive to change their behaviour: hence the equilibrium), like in the NPD. But unlike in the NPD, there is no dominant strategy. Cooperation is maintained at intermediate levels (cooperators and defectors coexist) in the absence of any additional force.

The frequency of cooperators at the stable equilibrium x_{eq}^{+} can be very close to the frequency x_M required to produce maximum fitness, and fitness at x_{eq}^{+} can be close to fitness at x_M (Fig. 2). Although the existence of the mixed equilibrium depends on the c/b ratio, as long as this equilibrium exists, the amount of public good produced is affected only slightly by the value of c, especially in large groups (Archetti & Scheuring 2011). Variations of the game that allow the cost of cooperation to be inversely related to the number of cooperators (Souza *et al.* 2009), or in which the public good increases linearly only above a threshold (Pacheco *et al.* 2009) behave similarly. So far, however, we have replaced a biologically unlikely assumption on the shape of the public good (linearity: the NPD) with only a slightly more likely assumption (a step function). What happens if we use a more realistic benefit function?

The intuition that nonlinear benefits are important has been understood for a long time in biology (Cohen & Eshel 1976; Matessi & Jayakar 1976) and in the social sciences (Taylor & Ward 1982). Motro (1991) introduced explicit discounting and synergistic effects to N-player social dilemmas, showing the existence of mixed equilibria. Hauert *et al.* (2006) showed essentially the same, using a more tractable model. All these cases (linear benefits, step functions, synergistic or discounting effects) are special cases of the following general public goods game (Archetti & Scheuring 2011), which only assumes that the benefit is a monotonously increasing and saturating function of the individual contributions (Fig. 3). By changing the parameters k (the position of the inflection point) and s (the steepness of the function at the inflection point, which is the synergistic effect of volunteering), this model can generate all the cases described above,
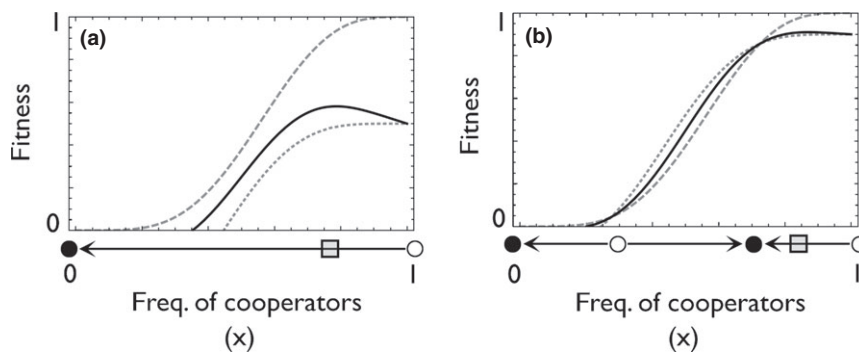

Figure 2 The mixed equilibrium in the volunteers' dilemma. The grey lines show the fitness for a cooperator (W_C, dotted line), for a defector (W_D, dashed line), and for the mixed strategy ($W_M = x W_C + (1-x) W_D$, black line), as a function of the frequency of cooperators (x). $N = 10$, $b = 1$, $k = 5$. Circles show the equilibria, stable (filled) or unstable (empty); arrows show the change in frequency of cooperators. The grey square shows the value of x for which the mixed strategy has the highest fitness ($W_{\max}$). (a) $c = 0.5$; there is no mixed equilibrium. (b) $c = 0.1$; at the mixed equilibrium, W_M is close to $W_{\max}$, and the frequency of cooperators is close to the frequency required for $W_{\max}$.
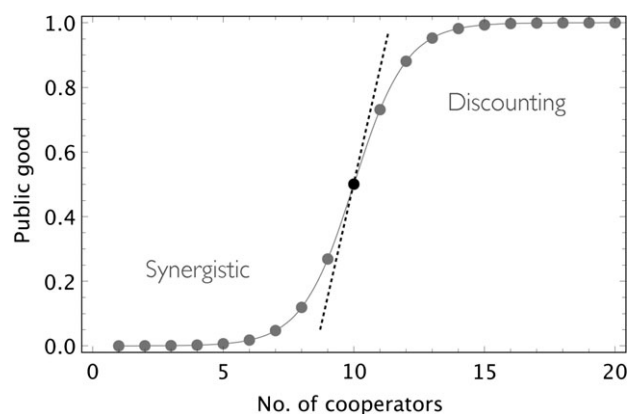
Figure 3 A general public goods game. The public good is a nonlinear function of the individual contributions: the benefit due to the public good increases monotonically with the number of cooperators; too few or too many contributions add little or nothing to the public good; at intermediate levels of contribution the increase is first synergistic (accelerating) and then discounting (decelerating). This type of public goods is very common in biology; in microbes that secrete extracellular enzymes, for example, the effect of enzyme production, and thus the level of benefit, is generally a saturating function of its concentration. The public-goods function has two characteristics, the inflection point k (black dot), where it shifts from accelerating to discounting, and the steepness s at the inflection point (dashed line). Group size is N. The public-goods function is $\beta(i) = b(\alpha(i) - \alpha(0))/(\alpha(N) - \alpha(0))$, where $\alpha(i) = (1 + \exp[s(k - i)])^{-1}$ and i is the number of cooperators. In this figure, $k = 10$, $s = 1$, and $N = 20$. When $s \rightarrow \infty$, the public-goods function becomes a step function of the individual contributions, with a threshold k (volunteers' dilemma). When $s \rightarrow 0$, the public good becomes a linear function of the individual contributions (N-player prisoner's dilemma) (see Fig. 4).

including increasing returns ($k \rightarrow N$) and diminishing returns ($k \rightarrow 0$), and it can be used to model the VD ($s \rightarrow \infty$) and the NPD ($s \rightarrow 0$).

It can be shown (Fig. 4; Archetti & Scheuring 2011) that changing the shape of the public good (s) changes the position of the mixed equilibrium only slightly, unless the function is almost linear ($s \rightarrow 0$), returning us to an NPD. Therefore, many social dilemmas can be approximated by the VD, and the results we have discussed for the VD hold for most public goods games. Seen in this general formulation, the NPD is just a degenerate case in which the mixed equilibrium does not exist and for which relatedness or other forms of assortment are needed to maintain cooperation. The existence of a mixed equilibrium in which cooperators and defectors coexist is the main result that emerges from the analysis of nonlinear public goods.

The reason why cooperators and defectors coexist in nonlinear public goods games can be explained intuitively. Consider first the simplest nonlinear public goods game: a public good is produced if and only if at least one individual cooperates. Why do cooperators and defectors coexist? Because if nobody else cooperates, the best strategy is to cooperate (not getting one's share of the public good is assumed to be worse than paying the cost of volunteering); but if someone else cooperates, the best strategy is to defect (the public good is produced anyway, and the cost of volunteering is spared). The result is that the best strategy is to cooperate with some probability (which depends on the cost/benefit ratio and group size). Now, imagine instead that the public good is produced if and only if at least k (>1) individuals volunteer: if k-1 other individuals volunteer, the best strategy is to volunteer, but if k or more other individuals already volunteer, the cost of volunteering is wasted. Again, the best strategy is to volunteer

with a probability that is a function of the cost/benefit ratio, group size, and k. The next step is to allow the public good to be a smooth sigmoid function rather than a step function. When the steepness (s) of this function is high, the public good is almost a step function, as in the previous case (Fig. 4) and the result is similar; volunteering with a probability. The coexistence of cooperators and defectors is possible as long as the cost of volunteering c is not too high. As the public good approaches a linear function, this critical cost c decreases, and when the public good is linear (NPD), no coexistence is possible for $c > 0$.

In summary, cooperation in N-player games should be modelled using the theory of general public goods. The NPD predicts that public goods can be produced only in the presence of some kind of assortment (due to kin selection or iterations). However, unlike the NPD, most public goods are nonlinear functions of the individual contributions. The VD or more general public goods games are the appropriate analytical tools in these cases, and these games have a mixed equilibrium in which cooperators and non-cooperators coexist in the absence of any form of assortment.

This does not mean that assortment is irrelevant. One reason is that, because the mixed equilibrium is generally not efficient, some assortment (but not too much – see Fig. 2) can improve fitness. Another reason is that assortment (for example genetic relatedness) can reduce the critical value of the cost/benefit ratio for which the mixed equilibrium exists; if a mixed equilibrium exists, however, increasing relatedness does not much improve the production of the public good (Archetti 2009a,b). Finally, assortment can help spread cooperation in a population of all defectors. It must be clear, however, that assortment is not necessary for the stability of public goods and the coexistence of cooperators and volunteers.

Implications for the evolution of cooperation

The theory of public goods is useful for understanding some recent developments in the study of cooperation in microbial systems.

While there is no doubt that assortment, and relatedness in particular, is an important determinant of cooperative behaviour, there is a tendency to attribute *a priori* any instance of cooperation to assortment. For example, Lee *et al.* (2010) recently reported that general antibiotic resistance in a non-structured population of *Escherichia coli* is due to a tiny minority of antibiotic-resistant cells that, at personal cost, export the chemical indole, which induces drug efflux and oxidative stress protection in the non-resistant cells. Lee *et al.* (2010) tentatively attributed their results to kin selection. However, their observations of the stable coexistence of cooperators and defectors and of higher group benefit (colony growth rate) occurring at intermediate frequencies of cooperators could be easily explained by the theory of nonlinear public goods without resorting to relatedness. Similar reasoning could be productively applied to virtually all articles that discuss cooperation in microbes. A review of these studies is not appropriate here, but it is worth relating one example in detail.

Consider invertase production by yeast, one of the most studied examples of public-goods cooperation in microbes; cells that secrete invertase are cooperators. This system was initially described as a prisoner's dilemma (Greig & Travisano 2004), but the prisoner's dilemma framework could not explain the coexistence of cooperators and defectors. The system was later described as a snowdrift game (Gore *et al.* 2009), which does explain coexistence. However,
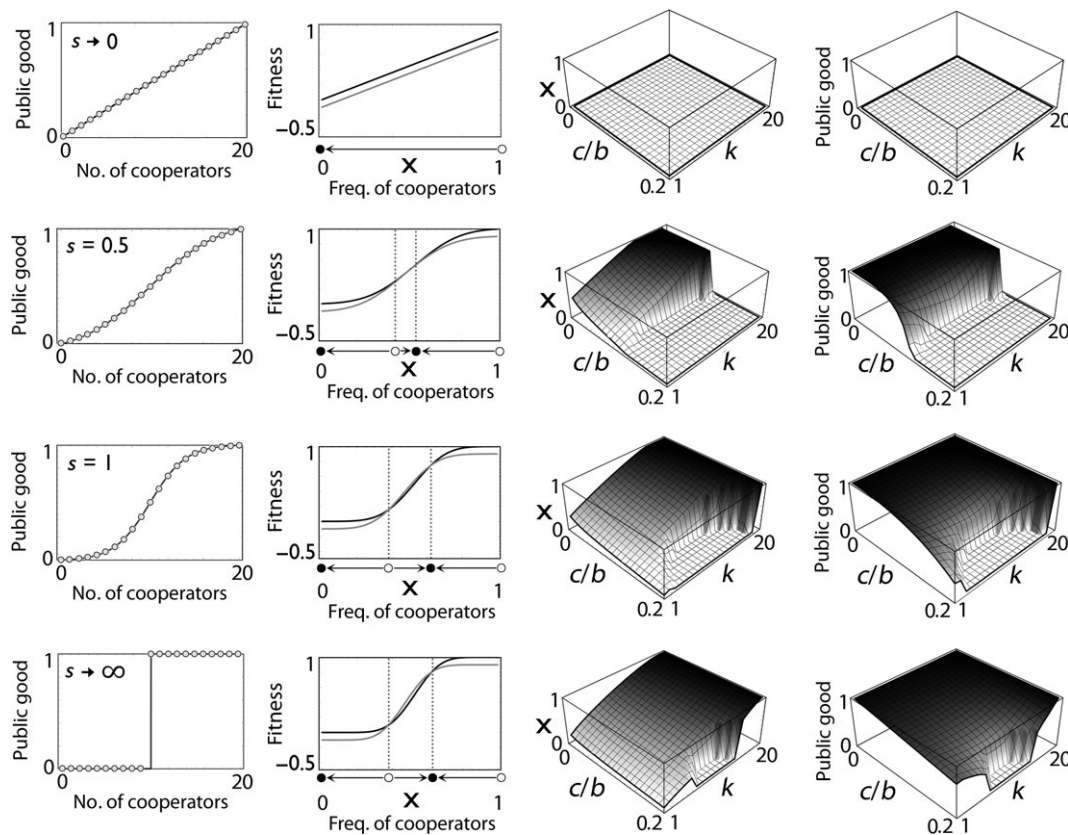
Figure 4 Cooperation in public goods games, from the N-player prisoner's dilemma to the volunteers' dilemma. For different values of *s*, the 2-dimensional plots show the amount of public good produced as a function of the number of cooperators (the continuous lines are only for guidance; the public good exists only for integer values of N – the grey dots) and the fitness of cooperators (W_C, grey) and defectors (W_D, black) as a function of the frequency of cooperators (x). $N = 20$, $c/b = 1/10$, $k = 10$. In the second column, the circles indicate equilibria, stable (black) or unstable (white), and the arrows show the change in the frequency of cooperators. The N-player prisoner's dilemma corresponds to $s \to 0$: every contribution increases the public good by a fixed amount proportional to $1/N$. The Volunteers' Dilemma corresponds to $s \to \infty$: the public good is produced only if at least *k* individuals volunteer to pay a contribution; an individual contribution adds nothing to the public good unless the number of other cooperators is *k*-1. In the intermediate cases, too few or too many contributions add little or nothing to the public good, and the increase is first synergistic and then discounting (see Fig. 3). If $s \to 0$ (N-player prisoner's dilemma), no mixed equilibrium exists. As *s* increases, two internal equilibria can exist (one stable, one unstable) because it is beneficial to be a defector only if there are too many or too few volunteers. Note that the position of the equilibria is almost the same for $s = 1$ and $s \to \infty$, so nonlinear public goods can be approximated by the volunteers' dilemma. In contrast, for $s \to 0$, the results are very different, and therefore the N-player prisoner's dilemma is not a good representation of public goods games in general. The 3-dimensional plots show the frequency of cooperators (x) at equilibrium and the amount of public good (PG) produced at equilibrium, as a function of the cost-benefit ratio of cooperation c/b and the parameter *k* (the position of the inflection point in the public goods function; the number of volunteers required for the production of the public good in the volunteers' dilemma). Note that, when the c/b ratio is such that the mixed equilibrium exists, the public good produced is substantial and almost independent of c/b (Modified from Archetti & Scheuring 2011).

describing invertase production as a 2-player game like the snowdrift game is clearly mistaken. MacLean *et al.* (2010) recently elaborated on the invertase system and found that, empirically, maximum group benefit occurs at an intermediate frequency of defectors, which, as they correctly point out, is not predicted by the snowdrift game. While MacLean *et al.* (2010) attribute the mismatch between theory and data to the failure of game theory to deal with complex systems, it is clear that the mismatch is simply due to the fact that invertase production is not a 2-player game but an *N*-player public goods game, for which, as we have shown, maximum group benefit is indeed expected at an intermediate frequency of cooperation, consistent with observation. Game theory can easily incorporate nonlinearities, differential costs and benefits and other elaborations.

CONCLUSION

The similarity between fitness maximisation in evolutionary biology and utility maximisation in microeconomics has been long recognised.

Modern biologists and economists use the same methods to analyse theoretical problems: optimisation (for non-strategic decisions) and game theory (for strategic interactions). It is surprising, therefore, that biology and economics have sometimes proceeded more or less unaware of each other to reach the same results (Hammerstein & Hagen 2005). Signalling theory, which is now standard in both microeconomics and evolution textbooks, is a case in point: the problem of honest signals was solved formally in evolutionary biology in the 1990's (Grafen 1990), when it was already well known in microeconomics, the problem first being solved by Spence (1973); economists have applied signalling theory in a variety of settings and also clarified signalling theory in biology (Nöldeke & Samuelson 1999, 2003). The theory of repeated games that started in biology with the repeated version of the prisoner's dilemma (Axelrod & Hamilton 1981) has a long tradition in economics (Aumann 1959; Mailath & Samuelson 2006). Examples of evolutionary methods adopted by economists are no less important. The concept of the Evolutionarily Stable Strategy (ESS; Maynard Smith & Price 1973) is an equilibrium

refinement of the Nash Equilibrium used in economics (an ESS has stricter conditions and can rule out some multiple equilibria). Evolutionary dynamics and ideas related to mutations and long-term stability are also used in economics (Samuelson 1997), and evolutionary game theory is now a standard approach in microeconomics.

We have discussed how recent work at the interface with microeconomics provides new solutions to three cooperation problems: partner choice without signalling, distinguishing host sanctions and partner fidelity feedback, and the production of public goods in one-shot interactions between non-kin (Fig. 5). While we treat each problem
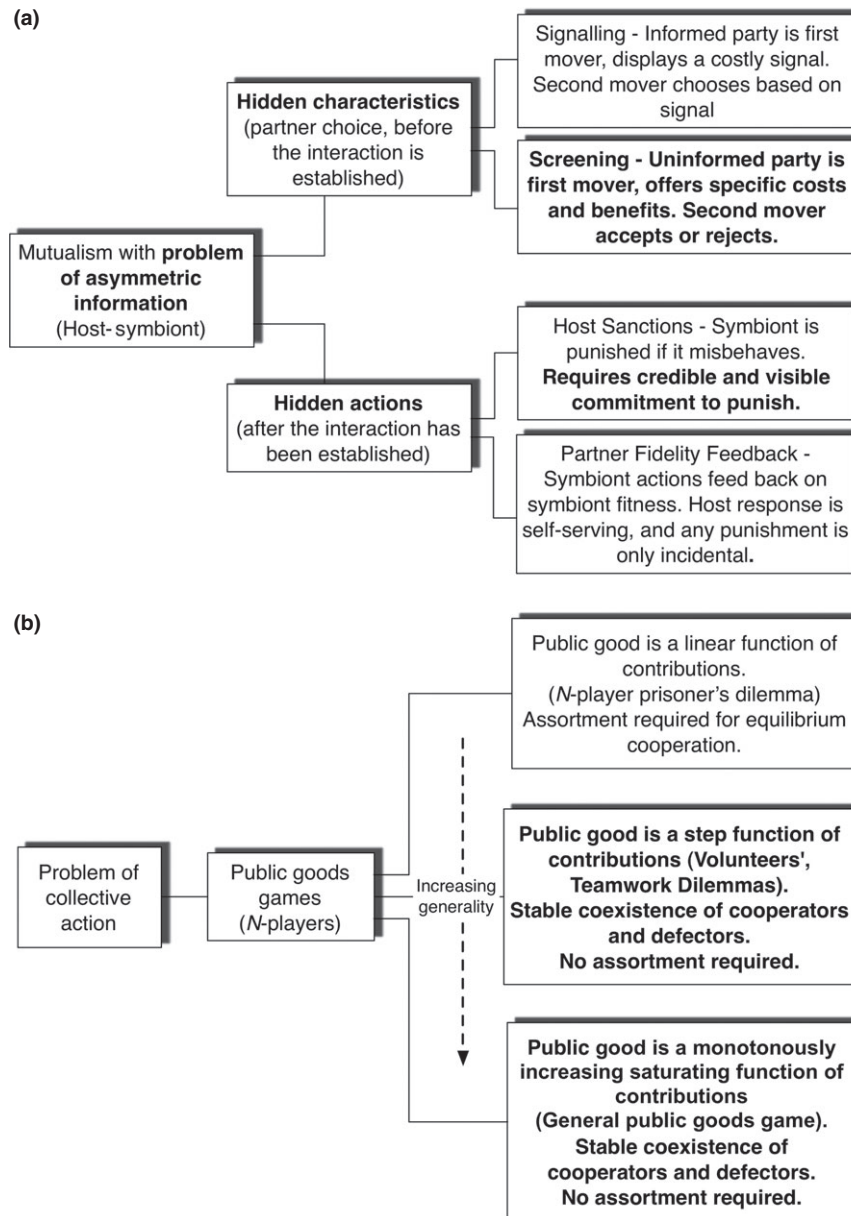
(a)

Mutualism with **problem of asymmetric information** (Host-symbiont)

Hidden characteristics (partner choice, before the interaction is established)

Signalling - Informed party is first mover, displays a costly signal. Second mover chooses based on signal

Screening - Uninformed party is first mover, offers specific costs and benefits. Second mover accepts or rejects.

Hidden actions (after the interaction has been established)

Host Sanctions - Symbiont is punished if it misbehaves. **Requires credible and visible commitment to punish.**

Partner Fidelity Feedback - Symbiont actions feed back on symbiont fitness. Host response is self-serving, and any punishment is only incidental.

(b)

Problem of collective action

Public goods games (*N*-players)

Increasing generality

Public good is a linear function of contributions. (*N*-player prisoner's dilemma) Assortment required for equilibrium cooperation.

Public good is a step function of contributions (Volunteers', Teamwork Dilemmas). Stable coexistence of cooperators and defectors. No assortment required.

Public good is a monotonously increasing saturating function of contributions (General public goods game). Stable coexistence of cooperators and defectors. No assortment required.

Figure 5 Outline of the concepts discussed in this paper. Terms and concepts in bold are mostly new to biology. (a) Games with asymmetric information. Before an interaction is established, there is a problem of *hidden characteristics*: how to choose the right partner? In signalling games, the non-informed individual observes a costly signal and chooses based upon the signal. In screening games, the non-informed individual imposes a strategic cost for entry, and neither observes any signal nor chooses. Both mechanisms allow hosts to separate symbionts of different qualities. After the interaction has been established, there is a problem of *hidden action*: how to prevent the symbiont from misbehaving? Host response is either committed to punish misbehaviour (host sanctions: provides future benefits to the host due to more cooperative behaviour of the punished symbiont) or simply takes the best self-serving action (partner fidelity feedback: provides benefits to the host irrespective of the effects on the punished individual). If host fitness after the best self-serving action automatically feeds back sufficiently strongly to the symbiont, the symbiont is selected to be mutualistic. Most mechanisms proposed to stabilise hidden action problems in mutualism theory are forms of self-serving actions with feedback. (b) If more individuals can contribute to a public good a *collective action problem* arises: why volunteer to contribute when one could free ride on the contributions of the others? In these social dilemmas, if the public good is a nonlinear function of individual contributions, cooperators and defectors can coexist at equilibrium in the absence of any assortment (kin selection, iterations). In the special case where the public good is a linear function of the individual contributions (the much-abused *N*-player prisoner's dilemma), defection is the only equilibrium, in which case assortment mechanisms are essential to achieve production of the public good.

separately for clarity, the problems can co-occur in nature. For example, partner choice does not exclude hidden action problems after hosts recruit symbionts. Moreover, we do not discard well-established concepts (like signalling and assortment); we simply provide explanations that can be invoked to bolster or replace established concepts when necessary. For example, nonlinear public goods game can be repeated, played between relatives, or in viscous populations.

Game theory provides a clear conceptual framework to study diverse strategic interactions and helps guide the design of empirical tests. Economists and evolutionary biologists thus still have much to learn from each other, but articles in ecology and evolution journals almost never cite current articles in economics journals, and *vice versa* (Rosvall & Bergstrom 2008). This is unfortunate because each field does not take advantage of the other's progress. There is still much to explore, and a larger set of conceptual tools and experimental models will open our eyes further to the stunningly diverse ways by which symbioses and societies evolve.

AUTHORSHIP

DY and NP conceived of and wrote a grant under the auspices of The Evolution of Cooperation and Trading (TECT) programme of the European Science Foundation, and the United States National Science Foundation supported much of the research synthesised herein. DY and MA wrote the section on hidden characteristics; DY, MH, MF, and NP wrote the section on hidden action; MA and IS wrote the public goods section; all authors contributed revisions.

REFERENCES

Archetti, M. (2000). The origin of autumn colours by coevolution. *J. Theor. Biol.*, 205, 625–630.

Archetti, M. (2009a). Evidence from the domestication of apple for the maintenance of autumn colours by coevolution. *Proc. R. Soc. B*, 276, 2575–2580.

Archetti, M. (2009b). The volunteer's dilemma and the optimal size of a social group. *J. Theor. Biol.*, 261, 475–480.

Archetti, M. (2009c). Cooperation as a volunteer's dilemma and the strategy of conflict in public goods games. *J. Evol. Biol.*, 22, 2192–2200.

Archetti, M. (2011). Contract theory for the evolution of cooperation: the right incentives attract the right partners. *J. Theor. Biol.*, 269, 201–207.

Archetti, M. & Scheuring, I. (2011). Coexistence of cooperation and defection in public goods games. *Evolution*, 65, 1140–1148.

Archetti, M., Úbeda, F., Fudenberg, D., Green, J., Pierce, N.E. & Yu, D.W. (2011). Let the right one in: a microeconomic approach to partner choice in mutualisms. *Am. Nat.*, 177, 75–85.

Aumann, R.J. (1959) Acceptable points in general cooperative N-person games. In: *Contributions to the Theory of Games IV* (eds Luce, R.D. & Tucker, A.W.) Annals of Mathematics Study 40, Princeton University Press, Princeton NJ, pp. 287–324.

Axelrod, R. & Hamilton, W.D. (1981). The evolution of cooperation. *Science*, 211, 1390–1396.

Barke, J., Seipke, R.F., Yu, D.W. & Hutchings, M.I. (2011). A mutualistic microbiome: how do fungus-growing ants select their antibiotic-producing bacteria? *Comm. Integ. Biol.*, 4, 1–3.

Bednarz, J.C. (1988). Cooperative hunting in Harris' hawks (*Parabuteo unicinctus*). *Science*, 239, 1525–1527.

Bshary, R. & Bergmüller, R. (2008). Distinguishing four fundamental approaches to the evolution of helping. *J. Evol. Biol.*, 21, 405–420.

Bull, J.J. & Rice, W.R. (1991). Distinguishing mechanisms for the evolution of co-operation. *J. Theor. Biol.*, 149, 63–74.

Burt, A. & Trivers, R. (2006). *Genes in Conflict*. Harvard University Press, Cambridge MA.

Chaston, J. & Goodrich-Blair, H. (2010). Common trends in mutualism revealed by model associations between invertebrates and bacteria. *FEMS Microbiol. Rev.*, 34, 41–58.

Chuang, J.S., Rivoirel, O. & Leibler, S. (2010). Cooperation and Hamilton's rule in a simple synthetic microbial system. *Mol. Syst. Biol.*, 6, 398–404.

Clutton-Brock, T.H. & Parker, G.A. (1995). Punishment in animal societies. *Nature*, 373, 209–216.

Cohen, D. & Eshel, I. (1976). On the founder effect and the evolution of altruistic traits. *Theor. Popul. Biol.*, 10, 276–302.

Connor, R.C. (1986). Pseudoreciprocity: investing in mutualism. *Anim. Behav.*, 34, 1652–1654.

Creel, S. (1997). Cooperative hunting and group size, assumptions and currencies. *Anim. Behav.*, 54, 1319–1324.

Crespi, B.J. (2001). The evolution of social behavior in microorganisms. *Trends Ecol. Evol.*, 16, 178–183.

Diekmann, A. (1985). Volunteer's dilemma. *J. Conflict Resolut.*, 29, 605–610.

Doring, T., Archetti, M. & Hardie, J. (2009). Autumn leaves seen through herbivore eyes. *Proc. R. Soc. Lond. B*, 276, 121–127.

Edwards, D.P., Hassall, M., Sutherland, W.J. & Yu, D.W. (2006). Selection for protection in an ant-plant mutualism: host sanctions, host modularity, and the principal-agent game. *Proc. R. Soc. Lond. B*, 273, 595–602.

Eswarappa, S., Karnam, G., Nagarajan, A., Chakraborty, S. & Chakravortty, D. (2009). *lac* repressor is an antivirulence factor of *Salmonella enterica*: its role in the evolution of virulence in *Salmonella*. *PLoS ONE*, 4, e5789.

Eungdamrong, N.J. & Iyengar, R. (2004). Modeling cell signalling networks. *Biol. Cell*, 96, 355–362.

Folstad, I. & Karter, A.J. (1992). Parasites, bright males, and the immunocompetence handicap. *Am. Nat.*, 139, 603–622.

Foster, K.R. & Kokko, H. (2006). Cheating can stabilize cooperation in mutualisms. *Proc. R. Soc. Lond. B*, 273, 2233–2239.

Foster, K.R. & Wenseleers, T. (2006). A general model for the evolution of mutualisms. *J. Evol. Biol.*, 19, 1283–1293.

Frank, S.A. (1998). *Foundations of Social Evolution*. Princeton University Press, Princeton, N.J.

Gore, J., Youk, H. & van Oudenaarden, A. (2009). Snowdrift game dynamics and facultative cheating in yeast. *Nature*, 459, 253–256.

Grafen, A. (1990). Biological signals as handicaps. *J. Theor. Biol.*, 144, 517–546.

Grafen, A. (2009). Formalizing Darwinism and inclusive fitness theory. *Philos. Trans. R. Soc. B*, 364, 3135–3141.

Grafen, A. & Archetti, M. (2008). Natural selection of altruism in inelastic homogeneous viscous populations. *J. Theor. Biol.*, 252, 694–710.

Greig, D. & Travisano, M. (2004). The Prisoner's Dilemma and polymorphism in yeast SUC genes. *Proc. R. Soc. Lond. B*, 27, S25–S26.

Hamburger, H. (1973). N-person Prisoners Dilemma. *J. Math. Sociol.*, 3, 27–48.

Hamilton, W.D. (1964). The genetical evolution of social behaviour. *J. Theor. Biol.*, 7, 1–52.

Hamilton, W.D. & Zuk, M. (1982). Heritable true fitness and bright birds, a role for parasites? *Science*, 218, 384–387.

Hammerstein, P. & Hagen, E. (2005). The second wave of evolutionary economics in biology. *Trends Ecol. Evol.*, 20, 604–609.

Hardin, J. (1968). The tragedy of the commons. *Science*, 162, 1243–1248.

Hauert, C., Michor, F., Nowak, M.A. & Doebeli, M. (2006). Synergy and discounting in social dilemmas. *J. Theor. Biol.*, 239, 195–202.

Heath, K.D. & Tiffin, P. (2009). Stabilizing mechanisms in a legume-rhizobium mutualism. *Evolution*, 63, 652–662.

Heil, M., González-Teuber, M., Clement, L., Kautz, S., Verhaagh, M. & Silva-Bueno, J. (2009). Divergent investment strategies of *Acacia* myrmecophytes and the coexistence of mutualists and exploiters. *PNAS USA*, 106, 18091–18096.

Hemker, H.C. & Hemker, P.W. (1969). General kinetics of enzyme cascades. *Proc. R. Soc. B*, 173, 411–420.

Holmström, B. (1979). Moral hazard and observability. *Bell J. Econ.*, 10, 74–91.

Huth, C.J. & Pellmyr, O. (2000). Pollen-mediated selective abortion in yuccas and its consequences for the plant-pollinator mutualism. *Ecology*, 81, 1100–1107.

Kaltenpoth, M. (2009). Actinobacteria as mutualists: general healthcare for insects? *Trends Microbiol.*, 17, 529–535.

Kiers, E.T., Rousseau, R.A., West, S.A. & Denison, R.F. (2003). Host sanctions and the legume-rhizobium mutualism. *Nature*, 425, 78–81.

Kiers, E.T., Rousseau, R.A. & Denison, R.F. (2006). Measured sanctions: legume hosts detect quantitative variation in rhizobium cooperation and punish accordingly. *Evol. Ecol. Res.*, 8, 1077–1086.

Kollock, P. (1998). Social dilemmas, the anatomy of cooperation. *Annu. Rev. Sociol.*, 24, 183–214.

Krugman, K. (1996). WHAT ECONOMISTS CAN LEARN FROM EVOLUTIONARY THEORISTS, in talk given to the European Association for Evolutionary Political Economy. Available at: http://web.mit.edu/krugman/www/evolute.html. Last accessed 1 Aug 2011.

Lee, H.H., Molla, M.N., Cantor, C.R. & Collins, J.J. (2010). Bacterial charity work leads to population-wide resistance. *Nature*, 467, 82–86.

Lehmann, L., Keller, L. & Sumpter, D. (2007). The evolution of helping and harming on graphs, the return of the inclusive fitness effect. *J. Evol. Biol.*, 20, 2284–2295.

MacLean, R.C., Fuentes-Hernandez, A., Greig, D., Hurst, L.D. & Gudelj, I. (2010). A mixture of "cheats" and "co-operators" can enable maximal group benefit. *PLoS Biol.*, 8, e1000486.

Mailath, G. & Samuelson, L. (2006). *Repeated Games and Reputations: Long-run Relationships*. Oxford University Press, New York.

Marr, D.L. & Pellmyr, O. (2003). Effect of pollinator-inflicted ovule damage on floral abscission in the yucca-yucca moth mutualism: the role of mechanical and chemical factors. *Oecologia*, 136, 236–243.

Matessi, C. & Jayakar, S.D. (1976). Conditions for the evolution of altruism under Darwinian selection. *Theor. Popul. Biol.*, 9, 360–387.

Maynard Smith, J. & Harper, D. (2003). *Animal Signals*. Oxford University Press, Oxford, UK.

Maynard Smith, J. & Price, G.R. (1973). The logic of animal conflict. *Nature*, 246, 15–18.

Mendes, P. (1997). Biochemistry by numbers: simulation of biochemical pathways with Gepasi 3. *TIBS* 22, 361–363.

Motro, U. (1991). Co-operation and defection, playing the field and ESS. *J. Theor. Biol.*, 151, 145–154.

Myatt, D.P. & Wallace, C. (2009). Evolution, teamwork and collective action, production targets in the private provision of public goods. *Econ. J.*, 119, 61–90.

Noë, R. & Hammerstein, P. (1994). Biological markets: supply and demand determine the effect of partner choice in cooperation, mutualism and mating. *Behav. Ecol. Sociobiol.*, 35, 1–11.

Nöldeke, G. & Samuelson, L. (1999). How costly is the honest signalling of need? *J. Theor. Biol.*, 197, 527–539.

Nöldeke, G. & Samuelson, L. (2003). Strategic choice handicaps when females seek high male net viability. *J. Theor. Biol.*, 221, 53–59.

Nyholm, S.V. & McFall-Ngai, M.J. (2004). The winnowing: establishing the squid-vibrio symbiosis. *Nat. Rev. Microbiol.*, 2, 632–642.

Nyholm, S.V., Stewart, J.J., Ruby, E.G. & McFall-Ngai, M.J. (2009). Recognition between symbiotic Vibrio fischeri and the haemocytes of Euprymna scolopes. *Environmental Microbiology*, 11, 483–493.

Olson, M. (1965). *The Logic of Collective Action, Public Goods and the Theory of Groups*. Harvard University Press, Cambridge, MA.

Pacheco, J.M., Santos, F.C., Souza, M.O. & Skyrms, B. (2009). Evolutionary dynamics of collective action in N-person stag hunt dilemmas. *Proc. R. Soc. Lond. B*, 276, 315–321.

Packer, C., Scheel, D. & Pusey, A.E. (1990). Why lions form groups: food is not enough. *Am. Nat.*, 136, 1–19.

Palfrey, T.R. & Rosenthal, H. (1984). Participation and the provision of public goods, a strategic analysis. *J. Public Econ.*, 24, 171–193.

Pellmyr, O. & Huth, C.J. (1994). Evolutionary stability of mutualism between yuccas and yucca moths. *Nature*, 372, 257–260.

Pellmyr, O. & Huth, C.J. (1995). Differential abortion in the yucca: reply. *Nature*, 376, 558.

Rabenold, K.N. (1984). Cooperative enhancement of reproductive success in tropical wren societies. *Ecology*, 65, 871–885.

Rapoport, A. & Chammah, A.M. (1966). The game of chicken. *Am. Behav. Sci.*, 10, 10–28.

Ricard, J. & Noat, G. (1986). Catalytic efficiency, kinetic co-operativity of oligometric enzymes and evolution. *J. Theor. Biol.*, 123, 431–451.

Ritchie, K. (2006). Regulation of microbial populations by coral surface mucus and mucus-associated bacteria. *Mar. Ecol. Prog. Ser.*, 322, 1–14.

Richter, K.S. & Weis, A.E. (1995). Differential abortion in the yucca. *Nature*, 376, 557–558.

Richter, K.S. & Weis, A.E. (1998). Inbreeding and outcrossing in Yucca whipplei: consequences for the reproductive success of plant and pollinator. *Ecol Lett*, 1, 21–24.

Rosvall, M. & Bergstrom, C.T. (2008). Maps of random walks on complex networks reveal community structure. *PNAS USA*, 105, 1118–1123.

Rothschild, M. & Stiglitz, J.E. (1976). Equilibrium in competitive insurance markets: an essay on the economics of imperfect information. *Q. J. Econ.*, 90, 626–649.

Ruby, E.G. & McFall-Ngai, M.J. (1999). Oxygen-utilizing reactions and symbiotic colonization of the squid light organ by *Vibrio fischeri*. *Trends Microbiol.*, 7, 414–420.

Sachs, J.L., Mueller, U.G., Wilcox, T.P. & Bull, J.J. (2004). The evolution of cooperation. *Q. Rev. Biol.*, 79, 135–160.

Samuelson, P.A. (1954). The pure theory of public expenditure. *Rev. Econ. Stat.*, 36, 387–389.

Samuelson, L. (1997). *Evolutionary Games and Equilibrium Selection*. MIT Press, Cambridge, MA, USA.

Schelling, T.C. (1960). *The Strategy of Conflict*. Harvard University Press, Cambridge, MA.

Selva, L., Viana, D., Regev-Yochay, G., Trzcinski, K., Corpa, J., Lasa, I. *et al.* (2009). Killing niche competitors by remote-control bacteriophage induction. *PNAS USA*, 106, 1234–1238.

Soto, M.J., Dominguez-Ferreras, A., Perez-Mendoza, D., Sanjuan, J. & Olivares, J. (2009). Mutualism versus pathogenesis: the give-and-take in plant-bacteria interactions. *Cell Microbiol.*, 11, 381–388.

Souza, M.O., Pacheco, J.M. & Santos, F.C. (2009). Evolution of cooperation under N-person snowdrift games. *J. Theor. Biol.*, 260, 581–588.

Spence, M. (1973). Job market signaling. *Q. J. Econ.*, 87, 355–374.

Stander, P.E. (1991). Foraging dynamics of lions in semi-arid environment. *Can. J. Zool.*, 70, 8–21.

Sui, S., Fedynak, A., Hsiao, W., Langille, M. & Brinkman, F. (2009). The association of virulence factors with genomic islands. *PLoS ONE*, 4, e8094.

Számadó, S. (2008). How threat displays work: species-specific fighting techniques, weaponry and proximity risk. *Anim. Behav.*, 76, 1455–1463.

Számadó, S. (2011). Long-term commitment promotes honest status signalling. *Anim. Behav.*, 82, 295–302.

Szilágyi, A., Scheuring, I., Orivel, J., Edwards, D.P. & Yu, D.W. (2009). The evolution of intermediate castration virulence and ant coexistence in a spatially structured environment. *Ecol. Lett.*, 12, 1306–1316.

Taylor, M. & Ward, H. (1982). Chickens, whales, and lumpy goods: alternative models of public-goods provision. *Polit. Stud.*, 30, 350–370.

Tong, D., Rozas, N.S., Oakley, T.H., Mitchell, J., Colley, N.J. & McFall-Ngai, M.J. (2009). Evidence for light perception in a bioluminescent organ. *PNAS USA*, 106, 9836–9841.

West, S.A., Kiers, E.T., Pen, I. & Denison, R.F. (2002a). Sanctions and mutualism stability: when should less beneficial mutualists be tolerated? *J. Evol. Biol.*, 15, 830–837.

West, S.A., Kiers, E.T., Simms, E.L. & Denison, R.F. (2002b). Sanctions and mutualism stability: why do rhizobia fix nitrogen? *Proc. R. Soc. Lond. B*, 269, 685–694.

Weyl, E.G., Frederickson, M.E., Yu, D.W. & Pierce, N.E. (2010). Economic contract theory tests models of mutualism. *PNAS USA*, 107, 15712–15716.

Yip, E.C., Powers, K.S. & Aviles, L. (2008). Cooperative capture of large prey solves scaling challenge faced by spider societies. *PNAS USA*, 105, 11818–11822.

Zivkovic, A.M., German, J., Lebrilla, C. & Mills, D. (2010). Human milk glycobiome and its impact on the infant gastrointestinal microbiota. *PNAS USA*, 108, 4653–4658.