

Evolutionary models of mutualism^{*}

Erol Akçay

Department of Biology, University of Pennsylvania

Mutualisms are among the most important ecological interactions, yet disciplinary folklore holds that they were long neglected in favor of competition and prey-predator interactions. Fortunately, this situation changed dramatically in the last three decades and the evolutionary theory of mutualisms has been attracting increasing interest. As a result, we now have a large enough collection of evolutionary models of mutualisms to warrant its own survey. Most of the present chapter focuses on the question that has been the dominant one in mutualism theory: how cooperation between species can evolve and be maintained through natural selection.

In many mutualisms, helping an individual of the partner species (e.g., providing protection, housing or nutrients) is costly, i.e., all else being equal, it decreases the fitness of the individual providing it. Such costly helping behaviors (or phenotypes) are termed investments. Mutualisms based on investments are liable to be taken over by what are generally called cheaters (see Chapter 7): individuals that enjoy the benefit from the mutualism, but do not provide any benefits in return. An example would be a plant-ant that feeds on the extra-floral nectaries of its host tree, but does not provide any protection in return (see Box 7.1 and Chapter 10). In the absence of a means to control cheating, trees hosting such ants would get selected to reduce the amount of extra-floral nectar they produce, until cooperation disappears (or remains at vestigial levels). Many evolutionary models of mutualisms are concerned with how mutualisms can persist in the face of cheaters, or equivalently, how mutualistic interactions resolve the inherent conflicts of interests between species. Accordingly, most of this chapter will deal with the modeling approaches used in this effort.

The twin goals of this chapter are to provide an overview of the rich theory of mutualism evolution as well as an introduction to the mathematical methods used in it. These methods are diverse, from adaptive dynamics to quantitative genetics, and a whole host of diverse approaches inspired by economics. Many of the modeling approaches in mutualism theory parallel those that deal with cooperation within species, which is hardly surprising given the conceptual closeness of the questions and the fact that many theorists contribute to both literatures. Yet, there are also some differences in emphasis. Although selection in structured populations (i.e., kin- or multi-level selection) is the dominant explanation for within-species cooperation, its action in mutualisms is much less studied (see section 4.2). Conversely, mutualism theory puts a lot more emphasis on flexible investment phenotypes (e.g., sanctions, partner choice; see section 4.4) than within-species theory. On balance, there is scope for much better integration of within- and between-species cooperation theories and modeling approaches.

^{*}This is the draft of a chapter that has been accepted for publication by Oxford University Press in the forthcoming book *Mutualism*, edited by Judith Bronstein, due for publication in 2015.

I organized the chapter around the biological concepts, discussing the relevant modeling methods in due order. The first four sections, making up the bulk of the chapter are devoted to the different biological mechanisms that can resolve the cheater problem. Section 4.5 discusses the evolution of diversity and section 4.6 some of the exciting future directions for mutualism theory.

4.1 Vertical transmission of symbioses

In some mutualisms, such as those between the leaf-cutter ants and their fungal symbionts, the hosts acquire their symbionts (see Box 6.1) from their parents. Such “vertical” transmission of symbionts (as opposed to horizontal, or environmentally acquired) is expected to favor mutualisms (Fine, 1975; Ewald, 1987): helping the host produce more progeny will also produce more symbiont progeny if symbionts are transferred to the offspring. Vertical transmission may have led to the evolution of two of the most important symbioses in the history of life: chloroplasts and mitochondria. On a less grand scale, maternally transmitted *Wolbachia* spreading through populations of *Drosophila simulans* in California has rapidly evolved from imposing a fecundity cost to infected females to conferring a fecundity benefit (Weeks et al., 2007). Thus, perfect (or high-fidelity) vertical transmission can in theory and in the real world resolve the conflict of interest between the host and symbiont and make the cheater problem disappear. But how do separate organisms evolve vertical transmission in the first place?

This question has been studied using the theory of virulence evolution in disease causing organisms. Virulence is defined as an index of the harm a symbiont causes to the host, so a mutualist is simply a symbiont with negative virulence. Many virulence evolution models have to deal with the coupled ecological and evolutionary dynamics of host and symbiont populations. Generally, this is an intractable task, but in the last two decades, a framework called adaptive dynamics (AD) was developed that can yield analytical insight under some simplifying assumptions (Dieckmann and Law, 1996). The basics of AD are described in Box 1.

Box 1: Adaptive Dynamics Erol Akçay

Adaptive dynamics achieves tractability on the problem of eco-evolutionary dynamics by assuming that the ecological dynamics of populations are very fast compared to the rate mutations occur. The typical AD model starts with a monomorphic population (or multiple monomorphic populations for coevolutionary models) at its ecological equilibrium, introduces a rare mutant to this population, and calculates its “invasion fitness” (its rate of increase when rare) at this ecological equilibrium. In the case of epidemiological models, the ecological dynamics usually keep track of different class of individuals separately (e.g., infected, susceptible, etc.). If the invasion fitness, denoted by $f(x_m, x_r)$ where x_r is the resident genotype’s trait value and x_m the mutant’s, is higher than the resident’s increase rate, the mutant invades and takes over the population, and the procedure starts again. Assuming mutants differ very little from the residents, one can approximate this process by a simple dynamics that climbs the gradient of the invasion fitness function, i.e., $\frac{dx_r}{dt} \propto \frac{\partial f(x_m, x_r)}{\partial x_m}$, evaluated at $x_m = x_r$. If, for some trait value x_r , $\frac{\partial f(x_m, x_r)}{\partial x_m} = 0$, that value is said to be an evolutionarily singular point. If,

further, $\frac{\partial^2 f(x_m, x_r)}{\partial x_m^2} < 0$, the trait value is evolutionarily stable: mutants to either direction of the singularity will have lower fitness than the resident. Conversely, if $\frac{\partial^2 f(x_m, x_r)}{\partial x_m^2} > 0$, the resident trait value is a fitness minimum and mutants with both higher and lower trait values can invade the resident population. With sufficient mutational variability, the trait value will branch into two and start diverging (Metz et al., 1996). Finally, whether or not a population will evolve towards an evolutionarily singular point is determined by the cross derivative $\frac{\partial^2 f(x, x_r)}{\partial x_m \partial x_r} < 0$, which means for values below the singular point, mutants with greater than resident trait value have the advantage and invade, and above the singular point the converse is true. Such singular points are called convergent stable. A convergent stable singular point is called an evolutionarily stable strategy (ESS) if it is a fitness maximum, and an evolutionary branching point (EBP) if it is a fitness minimum. EBPs serve as a model for investigating the conditions under which selection can lead to diversification and adaptive speciation (Doebeli and Dieckmann, 2000), though their relevance to speciation dynamics in nature is debated (Doebeli et al., 2005; Gavrillets, 2005).

References

- Doebeli, M. and Dieckmann, U. (2000). Evolutionary branching and sympatric speciation caused by different types of ecological interactions. *The american naturalist*, 156(S4):S77–S101.
- Doebeli, M., Dieckmann, U., Metz, J. A., and Tautz, D. (2005). What we have also learned: Adaptive speciation is theoretically plausible. *Evolution*, 59(3):691–695.
- Gavrillets, S. (2005). “Adaptive speciation”—it is not that easy: Reply to Doebeli et al. *Evolution*, 59(3):696–699.
- Metz, J. A. J., Geritz, S. A. H., Meszéna, G., Jacobs, F. J., and Van Heerwaarden, J. (1996). Adaptive dynamics, a geometrical study of the consequences of nearly faithful reproduction. *Stochastic and spatial structures of dynamical systems*, 45:183–231.

Yamamura (1993) applies the AD methodology to a model where a symbiont can evolve to harm or help its host, with harm increasing the symbiont’s own offspring production. The model confirms that high vertical transmission rates favor helping traits. However, how the transmission rate evolves depends on who controls it. The symbiont is always selected for higher vertical transmission rates, because there is no tradeoff between vertical and horizontal transmission in the model. The host on the other hand is selected to lower the transmission rate when it is low, and increase it when it is already above a threshold. This model suggests that vertical transmission can evolve at first by a parasite overcoming host defenses and increasing its vertical transmission rate, becoming less virulent as this happens. When the threshold transmission rate is crossed both the host and the symbiont will be selected to ensure that vertical transmission happens with high fidelity, which should lead to maximum mutualistic benefit to be obtained. The concordance of interest between the host and symbiont at the later stage of this process suggests that once a vertically

transmitted mutualism evolves, it would be very hard to break up. This prediction is in agreement with recent data on the evolutionary transitions in symbioses that show there are more transitions from parasitism to mutualism than the other way around (Sachs et al., 2011).

More recent work by Ferdy and Godelle (2005) considers the tradeoff between vertical and horizontal transmission. Such a trade-off is expected if traits that favor higher horizontal transmission (infection rates) reduce the fitness of the host and thus vertical transmission rates and is supported in many empirical studies (e.g., Stewart et al., 2005). Ferdy and Godelle allow the shape of the trade-off to be convex (i.e., decreasing marginal fecundity cost of horizontal transmission to the host) or concave (increasing marginal costs). With convex costs, they find evolutionary branching into divergent strategies of high horizontal transmission-high virulence and low horizontal transmission-low virulence (or helping) can occur. This is a general phenomenon in AD models, where the shape of costs and benefits and the fitness components they affect (e.g., fecundity of survivorship) determine the evolutionary outcome. In particular, evolutionary branching frequently occurs when costs are decelerating (are convex) or benefits are accelerating (concave) over some range of traits, as we will also see in the section on evolutionary diversification of mutualists in section 4.5.

4.2 Population structure and genetic associations

One of the most important advances in evolutionary theory after the modern synthesis was the emergence of theories of selection under population structure, i.e., when individuals do not interact or disperse in a completely mixed manner. Hamilton (1964) famously showed that if individuals interact with other individuals that carry similar genotypes, helping others even at a personal cost to oneself can evolve, provided that the coefficient of relatedness r (appropriately defined) is greater than the ratio of the personal cost, c to the benefit b supplied, or $rb > c$. For mutualisms, population structure operates in two ways: first, population structure within a species can promote helping the other species if the return benefits are likely to go related individuals. Second, population structure and limited dispersal can create association between mutualist genotypes, which also promotes cooperation.

There are many approaches to modeling evolution in structured populations (and to deriving versions of the Hamilton's rule). The most rigorous models explicitly specify the population structure (most commonly a population structured in discrete demes), population regulation (most commonly fixed deme size) and the life-cycle of the organisms (most commonly a Wright-Fisher or Moran process), and write down recursion equations for the probabilities of individuals being identical by descent (IBD) that can then be solved for determining relatedness. (Lehmann and Rousset, 2010, provides a good overview of this approach). Typically, such models assume that selection is weak, in which case relatedness can be approximated by the IBD probabilities under neutrality, and the demographical effects (such as local competition with related individuals) can be separated from the fertility effects of the traits (see Akçay and Van Cleve, 2012, for an exposition of how this can be achieved). Practically no model of mutualisms has been built using this approach. Instead, models of spatial population structure in mutualism tend to use either a phenomenological approach (as in Bever and Simms (2000), who use IBD probabilities as an exogenous parameter), or use individual-based simulations on a lattice (Doebeli and Knowlton, 1998; Yamamura et al., 2004). Another powerful approach that has been very influential in social evolution theory is the

Price equation, which I discuss here in some detail.

4.2.1 The Price equation

The Price equation (Price, 1970, 1972) expresses the total change in the mean value of a trait in a population at two time points, say, times t_1 and t_2 as:

$$\Delta z = \text{cov}(w, z) + E(w\bar{\delta}) . \quad (1)$$

Here, the first term is a covariance between an individual's trait at t_1 , z , and its fitness, w , defined as how many of the individuals at t_2 can be assigned to that particular individual at t_1 . This term is usually interpreted as giving the effects of selection on the mean trait value. The second term is the expectation of fitness w and the difference between the mean trait of the descendants (alive at t_2) of a focal individual and that of the focal individual itself (measured at t_1), given by $\bar{\delta}$. This second term gives the effects of other processes, such as transmission bias, phenotypic change, and selection at different levels on the change in mean trait value. The Price equation is an exact identity, meaning that it holds anytime when one can identify an ancestral and a descendant population, measure traits of all individuals in them, and assign individuals from the descendant population to individuals at the ancestral population. In other words, the Price equation allows an elegant summary of the change in a population, if one already knows every single thing necessary to calculate the change. Therefore, in its exact form, the Price equation is not actually a model of evolution or a prediction of the change in the mean trait. Nonetheless, one can sacrifice the exactness of the Price equation, and use its partitioning to organize and interpret the different processes influencing fitness when we do make models for those processes, such as one based on quantitative genetics (Frank, 1997).

One of the cases the Price equation becomes most fruitful (and the reason I discuss it in this section) is when individual fitness is a function of the phenotypes of other individuals, and individuals interact non-randomly with each other. Suppose, for example, a population with pairwise interactions, in which fitness of a focal individual is given by an additive function of an individual's own phenotype and that of a partner, i.e.,

$$w_i = \beta_{w,z_i} z_i + \beta_{w,z_j} z_j + e , \quad (2)$$

where the β 's give the regression coefficients of the fitness on z_i and z_j , respectively, and e is an error term independent of either phenotype. The last assumption makes this statistical model of fitness less general than the exact Price equation, so like many models we have sacrificed some generality in order to gain simplicity. Let us further assume that there is no transmission bias, i.e., the second term in equation (1) vanishes. Substituting (2) into (1), we obtain:

$$\Delta z = \text{cov}(w, z) = \text{cov}(\beta_{w,z_i} z_i + \beta_{w,z_j} z_j + e, z_i) = \beta_{w,z_i} \text{var}(z) + \beta_{w,z_j} \text{cov}(z_j, z_i) . \quad (3)$$

Supposing that the phenotypes correspond directly to the genotypes of the individuals, or $z_i = g_i$, we can replace the phenotypic variance and covariance by the genotypic ones, and divide by the variance term to obtain:

$$\frac{\Delta z}{\text{var}(g)} = \beta_{w,z_i} + \beta_{w,z_j} R , \quad (4)$$

where we have defined $R = \text{cov}(g_j, g_i)/\text{var}(g)$, or the regression of the partner’s genotype on the focal individual’s, and is one of the canonical definitions of relatedness (Frank, 1998). If the phenotype z is a costly helping behavior, β_{w,z_i} and β_{w,z_j} correspond to the cost $-c$ and benefit b , respectively, so the condition for the increase of such a behavior can be seen to be

$$Rb - c > 0 ,$$

which of course is the celebrated Hamilton’s rule (Hamilton, 1964). Another useful special case is when the phenotypes z are continuous, fitness is a smooth function of them, expressed as $w(z_i, z_j)$, and the amount of genetic variation around the population mean is small (which would happen with small effect mutants and low mutation rate). In that case, we can use the following approximations: $\beta_{w,z_i} \sim \frac{\partial w}{\partial z_i}$ and $\beta_{w,z_j} \sim \frac{\partial w}{\partial z_j}$ (Taylor and Frank, 1996). This is one of the easiest ways to model phenotypic evolution in structured populations, provided the assumptions are palatable for the particular system.

In one of the very few applications of the Price equation to mutualisms, Frank (1994) uses it to derive an extension of Hamilton’s rule, where a cooperative investment of an individual of, say, Species 1 benefits not only conspecific neighbors, but also individuals of a mutualist species (Species 2). In that case, the cooperative investment by Species 1 will be selected for if the following condition holds:

$$b_{1,1}R_{1,1} + b_{2,1}R_{2,1} - c_1 > 0 , \tag{5}$$

where $b_{1,1}$ and $b_{2,1}$ are the benefit from a Species 1 individual’s investment to conspecifics, and the benefit from Species 2’s investment to a Species 1 individual, respectively. The coefficients $R_{1,1}$ and $R_{2,1}$ are again defined as regression coefficients between Species 1 individuals’ investment genotypes and those of their conspecific or heterospecific interaction partners, respectively. Inequality (5) says that cooperation between species can evolve either if the mutualistic investment has “spillover” benefits to conspecifics, and relatedness is high enough (the first term in (5)), or if the investment from the partner species is beneficial enough for the focal species, and the cooperative genotypes from each species are non-randomly associated with each other (the second term in (5)). Thus, for the evolution of mutualism in structured populations, both within- and between-species genetic associations are potentially important. Furthermore, mutualistic interactions are inherently expected to generate between-species genetic associations, as Frank (1994) showed: when more mutualistic genotypes of both species co-occur in a patch, they will increase in abundance while populations in less mutualistic patches will decline. Such local build-up of highly mutualist populations is inherent in the population dynamics of mutualisms, but its effects have not been as widely studied, apart from a few individual-based simulation studies (e.g. Yamamura et al., 2004).

As in intraspecific models of social evolution, there is a tension between explicitly quantifying genetic association as a function of demographic quantities such as dispersal and fertility on one hand, and being able to derive analytically tractable models focusing on mechanisms that operate at the individual level. Where one chooses to be along this trade-off is a matter of what the model is primarily intended to do. As it stands, there is much work left to be done using either approach. For example, focusing on the build up of genetic associations in ecological time-scales and their evolutionary effects in diverse communities would require explicitly modeling the population dynamics and demography of partner species. Agent-based simulations can yield important insight to this question. On the other hand, the integration of genetic association with other mechanisms

for fitness feedbacks also remains to be fully explored (see section 4.6.1) which can be done with models that treat population structure implicitly (McGlothlin et al., 2010; Akçay and Van Cleve, 2012; Van Cleve and Akçay, 2014).

4.3 Modeling phenotypic feedbacks

In many mutualisms, the rewards from investing into a partner arise from the partner changing its investment in return. One classic example (described in more detail in Section 4.4.3, and in Boxes 5.2 and 6.2) is the legume-rhizobium symbiosis, where high nitrogen fixation by a nodule increases the resource allocation of the plant to that nodule, and low (or no) fixation elicits reduction in the plant investment (Kiers et al., 2003; Simms et al., 2006). I use the term *phenotypic feedbacks* to denote feedbacks between the investments that happen within the time-scale of the interaction, reflecting the fact that investments in such mutualisms can be thought of as flexible phenotypes of individuals. Phenotypic feedbacks are distinct from genetic associations and inter-generational feedbacks due to differential reproduction of partners. A huge amount of research effort in mutualism theory is focused on modeling phenotypic feedbacks. In this section, I will focus on the simplest cases, feedbacks in pairwise interactions, or that are otherwise undirected to particular individuals. Section 4.4 will deal with phenotypic feedbacks directed at particular partners, e.g., as a function of their performance.

The theory of iterated games (Maynard Smith and Price, 1973; Axelrod and Hamilton, 1981) has historically been a very popular way of modeling phenotypic feedbacks, starting from the first explicitly game theoretic study in biology, Maynard Smith and Price (1973). Phenotypic feedbacks can be modeled as conditional strategies in iterated games, where behavior in one round depends on the partner's behavior in the previous round. This approach had a meteoric rise after the landmark paper by Axelrod and Hamilton (1981), who showed that a very simple conditional strategy that simply copies the opponent's previous move can outcompete a multitude of sophisticated strategies. This set off an explosion of studies that pitched a multitude of simple conditional strategies against in each other in various agent-based simulations. While the modeling enterprise soon took on a life of its own with relatively little connection to biological questions, there have been several studies that have provided useful insight into biological questions and mutualisms in particular.

In one such study, Doebeli and Knowlton (1998) ask how cooperation between species can evolve when an individual's investment in the mutualism is responsive to its partner's investment. Doebeli and Knowlton assume pairwise interactions on a lattice structured population. Each pair plays a constant, finite number of rounds with each other before reproducing. At each round, each individual's investment to the mutualism is determined by the the outcome of the previous round. More precisely, if a focal individual invests x and its partner y in a given round, the payoff to the focal individual is given by $B(y) - C(x)$ (benefit from partner's investment minus the cost from own investment). The investment of the focal individual in the next round is then given by $a + b(B(y) - C(x))$, where a and b are assumed to be genetically encoded strategies (i.e., they are constant during the time-scale of the interaction). The investments of the partner species are determined likewise using a linear response rule. The total fitness of an individual is the sum of its payoffs in each round; this total fitness determines the evolution of the investment strategies a and b .

Using individual-based simulations, Doebeli and Knowlton find that in a well-mixed population

(with global dispersal), the only evolutionarily stable investment strategy is no investment at all ($a = b = 0$). On the other hand, if dispersal is localized (so that genetic associations start building up within and between species), unconditional and conditional investments can evolve even in finitely-repeated games. The intuition behind this result is that although a positive response slope is costly to a focal individual at the end of a finite interaction (when there is no possibility of positive feedbacks), population structure allows some of this cost to be “recouped” through effects on kin: offspring of a focal individual are more likely to interact with other mutualists, as a consequence of the focal individual’s positive response. Population structure and phenotypic feedback in general interact synergistically in supporting cooperation (see also section 4.6.1).

One of the key assumptions of Doebeli and Knowlton (1998) is that the interaction lasts for a finite number of rounds. This differentiates their model from a set of related studies for within-species cooperation (e.g. McNamara et al., 1999; Taylor and Day, 2004) that instead assume that the number of rounds is very large and therefore concentrate on the equilibrium of the phenotypic dynamics (i.e., when from one round to the next, the investment levels do not change). Among these, Taylor and Day (2004) find that conditional investment and cooperative investment equilibria can indeed evolve even in unstructured populations. The intuition behind this result is sometimes called partner fidelity feedback (Bull and Rice, 1991; Sachs et al., 2004): if one stays with a given partner long enough and the partner employs a conditional investment strategy, it does not pay to cheat it for short term benefit.

My coauthors and I (Akçay et al., 2009; Akçay and Van Cleve, 2012; Van Cleve and Akçay, 2014) generalized the behavioral equilibrium-based approach and found that the effects of such feedbacks in continuous, symmetric games can in general be characterized by a local response coefficient ρ , which gives how much a partner of a focal individual increases (or decreases their investments) in response to a small change in the partner’s investment. For within-species interactions with population structure, we used the Price equation approach to show that the response coefficient ρ plays a role symmetric to relatedness r in its effect on selection Akçay and Van Cleve (2012). Moreover, relatedness and responsiveness interact with each other synergistically in promoting more investment, such that populations with both genetic structure and phenotypic responses can maintain significantly more cooperation than expected by their separate effects. Van Cleve and Akçay (2014) extend such analysis to a simple model of symbiotic mutualism. They find that the synergistic interaction between phenotypic response coefficients on the one hand and within- and between-species genetic associations on the other is a general prediction of these models.

Phenotypic feedbacks can also be modeled using the methodology of quantitative genetics, which refers to phenotypic feedbacks as indirect genetic effects (IGEs) (Moore et al., 1997; McGlothlin et al., 2010). The IGE framework starts with a standard quantitative genetics framework but then incorporates the fact that in a social interaction, an individual’s phenotype (a behavior or a morphological trait) is a function of not just its own genetic make-up, but those of its partners’ as well. In the IGE framework, each individual’s phenotype is a linear, additive function of its own breeding value and the phenotypes of others it interacts with multiplied with some IGE coefficient (usually denoted by Ψ). In this way, IGE models bear a close resemblance to the linear response rule models, except that IGE coefficient is usually taken to be a constant (with some exceptions, e.g., Kazancıoğlu et al., 2012). Akçay and Van Cleve (2012) demonstrate that for a single phenotype, there is a one-to-one mapping between IGE models in structured populations (McGlothlin et al., 2010) and the behavioral response formulation, with ρ and Ψ invertible functions of each other. The

advantage of IGE models is that they use the machinery of quantitative genetics, which can handle multiple traits and their phenotypic and genetic covariances, and provides a robust framework for empirical measurement of such responses (McGlothlin and Brodie III, 2009). The intersection of IGE models and models of behavioral and proximate mechanisms mediating the phenotypic feedbacks is an exciting area of future research (Akçay and Simms, 2011; Akçay and Van Cleve, 2012).

4.4 Partner choice and beyond: the economic theory mutualism

Cooperation between species can also evolve if individuals preferentially interact with a cooperative subset of partners, or selectively terminate interactions with non-cooperative ones; in other words, if they exert partner choice. The idea of choice by organisms as a selective force goes back at least to Darwin's sexual selection theory, and has also played a central role in social evolution theory (Axelrod and Hamilton, 1981; Bull and Rice, 1991). It also has close parallels with consumer choice and related economic concepts and thus has been a focus for application of ideas and techniques from economics to biological problems (Hammerstein and Hagen, 2005). There are many important biological questions to ask about partner choice, such as when to choose, what do base partner choices on, how stringent to partner choice should be, etc. Answering this diverse set of questions requires a somewhat eclectic set of approaches, which are the focus on these section.

4.4.1 Supply and demand: biological markets

One prominent application of an economic idea to biology has been the concept of biological markets (Noe and Hammerstein, 1994, 1995). Biological markets theory asserts that just like in human markets, the supply and demand for partners of high quality will determine the evolutionary dynamics of a mutualism through partner choice. This idea was first modeled by Noe and Hammerstein (1994), who consider two species that trade in some "commodities" (e.g., resources such as nitrogen or carbon, or services such as pollination). One of the species (say, species B) can invest at into producing commodities of high or low value to the other species, low-value commodities being less costly to produce. Furthermore, there is some noise in production of commodities by species B, such that a high-investment might produce a low value, and vice-versa. This last feature ensures that there is phenotypic variability in the population that makes partner choice worthwhile, even if the population is fixed for one investment type. Species A, on the other hand, has a fixed amount of investment, but it has the ability to choose its partner from the B species, but that comes at a cost, in terms of the effort spent sampling different partners. The model has two different kinds equilibria depending on the relative abundance of the species: if species A is abundant, A individuals will accept low-value partners, and B species invest into low-value commodities. On the other hand, if B individuals are abundant, A individuals only trade with high-value partners and consequently the B population is fixed for high investment. These results make sense from a supply-demand perspective: when A individuals are relatively rare, B individuals are willing to pay the cost of high-investment to be chosen by a partner. Interestingly, there is a range of relative abundances where both equilibria exist at the same time, meaning that two different populations with the same abundances can inhabit either regime, depending on the initial conditions.

Applying the consumer-choice metaphor at the individual decision-making level, Kummel and

Salant (2006) model a plant that can associate with multiple mycorrhizal partners. Each partner is characterized by an exchange function, which denotes how much nutrient it supplies for a given amount of carbon it receives. The plant then has to decide how much nutrient it “buys” from each “supplier”. Kummel and Salant show that the optimal allocation of carbon will equalize the marginal cost of nutrient acquisition from each partner. Interestingly, the exchange function characterization admits a much richer space of variation in partner quality than a dichotomy between mutualists vs. cheaters: a partner can be very beneficial at low exchange rates (exchange function goes up steeply at first) but not at higher exchange rates (exchange function levels off), and vice versa. Kummel and Salant show that depending on set of exchange functions available to the plant, complex patterns of partner choice can emerge. An extension of this model that incorporates mycorrhizal fitness shows that it can also generate context dependency the benefits to the mycorrhizae (Akçay, unpublished results).

Yet another approach based on the idea of markets was proposed by Schwartz and Hoeksema (1998; Hoeksema and Schwartz 2003), who ask when trade between two species is beneficial to begin with. In their model, both species need two distinct resources (say, carbon and phosphorus). The two species can either take both resources up from the environment or they can each take up only one and trade with each other for the other. This situation is directly analogous to comparative advantage models from international trade going back to Adam Smith and David Ricardo (Ricardo, 1891), The basic idea is that countries might do best if they focus on producing goods for which they have a comparative advantage (e.g., Portuguese wine is better than British), and rely on trade to obtain other goods that they need. Schwartz and Hoeksema (1998) show that under some “exchange rates” of the two resources, it is optimal for each species to specialize in taking up only one resource and trade with each other for the other. The existence of such exchange rates depends on the ratio at which each species optimally consumes the two resources (i.e., the stoichiometry of their growth), and their relative efficiency in obtaining each resource directly from the environment. Specialization with trade is optimal when the two species are relatively more efficient in taking up different resources from the environment.

4.4.2 The paradox of the market

The mechanisms of partner choice and biological markets as means to select against cheaters create something of a paradox: if partner choice is effective, cooperative partners will increase in frequency, leading most individuals to interact only with cooperative partners, leading to lower benefits from choice. Eventually, the cost of partner choice will create selection against it, which opens the door for cheaters to invade again. This dynamic is well-known in the study of sexual selection as the “paradox of the lek” (Kirkpatrick and Ryan, 1991), and has attracted a persistent research effort over more than 30 years. Foster and Kokko (2006) tackle this issue with a model in which a host interacts with symbionts that are drawn from a distribution of mutualistic quality, but chooses the more beneficial ones. This choice depletes available variation in mutualist quality locally, but Foster and Kokko propose that it is continually replenished by migration from a source population in which variation is held constant. With enough such migration, costly partner choice and a high level of average cooperation can be maintained. In a fundamental sense, though, this model begs the question, since it assumes some unspecified mechanism that maintains the variation in the source population. As Foster and Kokko point out, the benefit from choice in mutualisms does not require additive genetic variation in the partner quality (as it does in the paradox of the lek),

so purely environmental or non-additive genetic variation (both of which are easy to maintain even under selection) can make partner choice beneficial in the absence of additive genetic variation in partner quality. This is important, because empirical evidence suggests that most variation in partner quality is not additive (with unconditionally beneficial and non-beneficial partners); rather, the costs and benefits are highly context dependent (see Chapters 7 and 11, as well as section 4.6.2). Such non-additive variation and mechanisms that generate and maintain variation in mutualist quality remain relatively unexplored in theory (Heath and Stinchcombe, 2014, see also section 4.5).

4.4.3 Sanctions: partner choice in the legume-rhizobium symbiosis

One of the most important model empirical systems for studying partner choice has been the symbiosis between legume plants and the nitrogen-fixing soil bacteria collectively called rhizobia (see Boxes 5.2 and 6.2). Rhizobia are taken up from the soil and inhabit specialized plant root organs called nodules. A single plant can have hundreds of nodules, each inhabited by one or more rhizobium strains. Since rhizobia are not vertically transmitted, there is a conflict of interest between the plant and the rhizobia, and hence the prospect of cheating.

Partner choice in the legume-rhizobium system is frequently referred to as “sanctions”, a term introduced by Denison (2000) who proposed that legume plants control cheating by selectively withholding resources to nodules that do not fix nitrogen. West et al. (2002b,a) have formalized Denison’s ideas in two separate models. West et al. (2002b) consider a kin-selection model with local competition among a population of rhizobia that occupy the same plant, using the approach of Taylor and Frank (1996, see section 4.2.1). To model partner choice by the plant, West et al. (2002b) assume that a focal rhizobium strain’s fitness is given by a continuous, increasing function of the focal strain’s nitrogen fixation rate. They consider two different functional forms for this function: one where it depends only on the focal strain’s fixation rate (which they term fixed-rule sanctions) and one where it depends both on the focal strain’s fixation rate and the average of all nodules (termed relative-rule sanctions). Both types of partner choice select for more nitrogen fixation compared to no partner choice. Increasing relatedness between rhizobia on different nodules increases the evolutionarily stable fixation rate under fixed-rule sanctions, but interestingly, it has the opposite effect under relative-rule sanctions. This is because high relatedness increases the average nitrogen fixation rate on plants that contain a highly-fixing focal strain, which decreases the benefit the focal strain obtains from the plant under relative-rule sanctions.

In the companion paper, West et al. (2002a) use the same modeling approach to ask when it is beneficial for the plant to engage in sanctions when interacting with rhizobium strains that vary in their nitrogen fixation rate. They model plant fitness as proportional to $R * N$, where R stands for the carbon the plant has available for its own growth and reproduction (after investing into the production of nodules) and N is the total nitrogen the plant obtains from the nodules and the soil. The amount of resources available for growth, R , is a function of how many nodules a plant makes initially, how many it sanctions, and what proportion of the resources invested into sanctioned nodule the plant can reclaim. Likewise, N is a function of how many nodules there are after sanctions, their average fixation rate, and the availability of nitrogen from the soil. The crucial assumption in the model is that the average fixation rate of the non-sanctioned nodules is increasing in the fraction of nodules sanctioned, although the mechanism by which this happens is not modeled. Under these assumptions, stricter sanctions are favored when more resources are recouped by sanctioning, higher mean fixation rate among rhizobia before sanctions, and higher

variance in fixation rates.

The two West et al. models confirm our intuition that sanctions might favor cooperation and evolve through direct benefits. However, since they treat rhizobia and plants separately, they do not fully elucidate the patterns of fitness variation that sanctions would generate between plants and rhizobia. Further, West et al. (2002a) sidestep two important facts: first, that the distribution of rhizobium strains of varying quality on a plant will determine how much benefit the plants gain by sanctioning, and second, the possibility of mutual control over partner choice. These factors are treated in the models I discuss next.

4.4.4 Control, negotiation and outside options in mutualisms

In the models in sections 4.4.1, 4.4.2, and 4.4.3, partner choice is a unilateral affair: one party has complete control over who it interacts with and how much. When both parties have some control, richer dynamics ensue. Johnstone and Bshary (2002) consider a model inspired by the cleaner fish mutualism, where small cleaner fish remove ectoparasites from larger fish that visit their territories, but prefer to feed on healthy mucus tissue (see Box 7.2 for more on this mutualism). In the model of Johnstone and Bshary (2002), client-cleaner pairs gain mutual benefits as an increasing function of time they spend together. The cleaners can however exploit the clients by biting healthy mucus tissue rather than removing parasites (Bshary and Grutter, 2002), gaining an additional benefit per unit time at a cost to their clients. The evolutionarily stable exploitation level will be determined by how the duration of the interaction at different exploitation levels, which in turn depends on which party controls the duration of the interaction. Johnstone and Bshary (2002) assume that the interaction ends when the marginal cost of continuation to the cleaners exceed the marginal benefit multiplied by a coefficient signifying the “power” cleaners have over the clients. The more power cleaners have, the closer the duration will be to their optimum. With increasing power to the cleaners, evolutionarily stable exploitation levels increase. This shows that the effectiveness of partner choice depends on whether it can be exercised free from interference from the partners.

In their initial model, Johnstone and Bshary (2002) do not account for the alternatives clients or cleaners have once they leave an interaction, in other words, their outside options. Johnstone and Bshary (2008) model outside options arising from population level availability of cleaners and clients, which is a function of their relative abundances, as well as how long they stay in each pairwise-interaction. Johnstone and Bshary (2008) use a time-in/time-out approach, where individuals transition from being free-living (time in the “market”) to being in an interaction (time out), at a rate that is proportional to the availability of partners. The time out duration is determined as in Johnstone and Bshary (2002), but with the additional consideration of the opportunity cost of leaving, which depends on how fast a new partner can be found. This opportunity cost emerges as a property of equilibrium, with only the overall abundance of clients and cleaners specified exogenously. The results support the biological market hypothesis: when clients are rare, low or no exploitation is expected at equilibrium, but when cleaners are rare, exploitation levels are expected to be high. In the former case, cleaners face a high opportunity cost of exploiting a client and causing it to leave, since replacing a partner takes a long time; in the latter case clients are facing the high opportunity cost. Thus, Johnstone and Bshary (2008) show the importance of outside options in determining evolutionary outcomes of biological markets.

In interactions where fast partner switching is not an option, as in the legume-rhizobium mutualism, other mechanisms will determine the outcome of the mutualism. The models of West et al.

(2002a,b) assume that the nitrogen fixation rates are “sealed bids”, determined by the rhizobium genotype only, while the plant responds to these bids. We (Akçay and Roughgarden, 2007) investigated an alternative where the fixation rates are the result of a “negotiation” process. The key idea is that once a nodule is established, there will be a range of carbon and nitrogen exchange rates with a focal nodule that make both the plant and the rhizobia in that nodule better off (grow more) relative to no exchange. But some of these exchange rates will favor the plant more, and some the rhizobia. Where along that continuum the outcome ends up can have important consequences for how partner choice works at the longer time-scale. In the negotiation between bacteria and plants, the “offers” come in the form of stochastic fluctuations in resource fluxes. When these fluctuations increase one party’s growth rate while decreasing the other’s, we conjectured that there would be a period of costly conflict between the parties (modeled as a war of attrition), in which the party with more to gain (relative to giving in) wins. This simple process converges to the maximum of the product of the plant’s and nodule’s growth rates. This stable equilibrium point is called the Nash bargaining solution (Nash, 1950; Rubinstein, 1982) in economics and provides a basis for quantitatively predicting the outcome of a negotiation process.

Akçay and Simms (2011) later extended the pairwise negotiation model to negotiation with two nodules, each providing an outside option to the plant vis-a-vis the other and allowing the ‘negotiation strategy’ of the rhizobia to evolve. In particular, we considered rhizobia that have two evolving traits. The first one is what we anthropomorphically termed “stubbornness”, i.e., how hard it is to get a rhizobium strain to increase its nitrogen fixation or decrease carbon intake. A more stubborn strain fixes less nitrogen for more carbon, and thus is less cooperative. Second, a focal nodule could provide a good or bad outside option to the plant, by increasing its nitrogen fixation temporarily when the plant is negotiating with the other nodule. This outside option trait is nearly neutral for the focal nodule, since it is only a temporary increase in nitrogen fixation and does not affect equilibrium fixation rate for the focal nodule directly. When the more stubborn (less cooperative) nodule has a lower growth rate than the more cooperative nodule, we can say the plant is exercising effective partner choice (or sanctions). We show that this requires that the less stubborn strain also provides a better outside option to the plant, and not otherwise. This finding provides a potential explanation for why some studies of the legume-rhizobium symbiosis show effective partner choice (Kiers et al., 2003; Simms et al., 2006; Kiers et al., 2006), but others do not (Marco et al., 2009; Gubry-Rangin et al., 2010). Strikingly, regardless of whether partner choice is effective, selection for stubbornness (i.e., against cooperation) is always positive. This is because the plant has incomplete control over partner choice, and all else being equal, a more stubborn rhizobium strain will always extract more resources from the plant. To counteract this selection, a positive feedback at the whole-plant level is required, so that nodules on plants that grow better do better altogether. This finding shows that phenotypic feedbacks and selection at multiple levels might be required to overcome the limitations of partner choice when both parties have control over the interaction (see Box 5.2).

Grman et al. (2012) recently combined the Nash bargaining approach of Akçay and Roughgarden (2007) with the comparative advantage model of Schwartz and Hoeksema (1998) to ask how a negotiated exchange rate affects the evolution of specialization in a plant-mycorrhizal mutualism. Mycorrhizal fungi infect the roots of a wide variety of plants, and trade nutrients they acquire from the soil with the plant for carbon (Smith and Read, 2008). Grman et al. (2012) find that specialization is somewhat less likely with negotiated exchange, since it makes a party completely dependent on the partner (i.e., reduces the outside options) and means a worse bargaining outcome. Nev-

ertheless, for empirically derived parameter regimes, specialization (especially of the fungus into nutrient uptake) does occur at the equilibrium. This study and Johnstone and Bshary (2008) together demonstrate the utility of combining models that apply at pair- and population-levels (see Box 2, and section 4.6.1).

The general message from the set of models in this subsection is that the outcome of a mutualism depends strongly on the mechanisms of control each party has over partners and the relative outside options (e.g., alternative partners, ability to take up resources from the environment, etc.) each party has. This fact sets the stage for rich evolutionary and ecological dynamics that have not been fully explored yet.

4.4.5 Screening

An alternative to trying to stop a partner's exploitation during the interaction is to not associate with an exploiter in the first place. In the cleaner fish mutualism, for example, clients choose to interact with only non-exploiting partners, as determined by own experience or observing interactions with others (Bshary and Grutter, 2002, 2006). As discussed in 4.4.1, such "market choice" can preclude the evolution of exploitation, but it requires that clients get reliable information in the face of incentives by the cleaners to misrepresent.

This is a problem that has direct analogies in economics: when a business wants to hire new employees, it will have imperfect information about the quality of the applicants. Obviously, applicants for the vacancy have an incentive to exaggerate their abilities. One way to resolve this problem is to make the applicants reliably signal their quality before the interaction, which usually involves costly signals (Spence 1973; the same idea was proposed independently in biology as the handicap principle by Zahavi 1975). But an entirely different way of separating high and low-quality applicants is to let them sort themselves out. This is termed "screening" in economics (Rothschild and Stiglitz, 1976; Riley, 2001). For example, a firm can offer a job contract to a candidate that entails a low salary with a big bonus that is only paid if a certain performance criterion (e.g., sales volume) is satisfied. Candidates that know they cannot satisfy the criterion will have no incentive to accept the offer, and those that do accept it will automatically have the correct incentives to work hard, without the need for constant monitoring by their employer. The result mimics effective partner choice, but there is no explicit signaling of quality, or evaluation of partners by the chooser.

A series of recent papers proposed that screening theory can be applied to understand mutualisms (Archetti et al., 2011; Scheuring and Yu, 2012). Archetti et al. (2011) use the symbiosis between the Hawaiian bobtail squid (*Euprymna scolopes*) and *Vibrio fischeri* bacteria (see Box 6.3) as an example. The bacteria colonize the light organ of the squid, where they emit light that reduces detection of the squid by predators. The luciferase enzyme that is responsible for bioluminescence also happens to remove reactive oxygen species from the interior of the light organ where it can damage the bacterial cells. Archetti et al. interpret this as a screening mechanism: only those bacterial strains that can produce light can survive in this hostile environment. Therefore, for would-be cheaters (strains that don't produce light), colonizing the squid is a fitness dead-end, and only light-producing strains will evolve the traits necessary to colonize the squid. This obviates the need to control cheating after the initial infection, similar to the job offer with bonus example in the previous paragraph.

The second paper, by Scheuring and Yu (2012), proposes a different screening mechanism, based on competition between antibiotic producing bacterial strains. Using a model originally developed

by Mao-Jones et al. (2010) for the establishment of coral diseases, Scheuring and Yu argue that the host can screen for a beneficial microbiome by letting the antibiotics produced by beneficial symbionts create the hostile environment for the non-beneficial ones. Following Archetti et al. (2011), they call this mechanism “competitive screening”. The crucial divergence of this model from economic screening models is that the symbiont dynamics within a host are bistable: if non-beneficial microbes initially colonize a host at high frequency, they persist at equilibrium. Thus, unlike standard screening models, the selection on cheaters for not colonizing the host requires some pre-infection partner choice mechanism (or vertical transmission) on the part of the host that results in high initial frequency of beneficial microbes.

4.4.6 Economics and mutualisms: a bright future

This section illustrated how fruitful economic analogies and techniques have been in understanding partner choice and related phenomena in mutualisms. This is not a coincidence of course: mutualisms are exchanges of resources and services between all organisms, whereas economics studies the exchange of resources and services among humans. The future of economic theory of mutualisms seems very bright, with a great deal of work that remains to be done in exploring strategic questions related to mutualisms using game theory, market modeling, and related theories such as signaling and screening.

4.5 Evolution and maintenance of diversity in mutualisms

In nature, most mutualisms involve tremendous variation in how beneficial the different species or strains are to their partners. However, relatively few models have specifically focused on the emergence and maintenance of diversity (more specifically, variation in cooperative phenotypes) in mutualisms.

One of the most sustained theoretical contributions to this question has come from the adaptive dynamics (AD) approach. As described in Box 5.1, AD can be used to find conditions under which selection will cause a population to branch phenotypically, evolving in two different directions. Doebeli and Dieckmann (2000) analyze an instance of such evolutionarily branching in mutualisms as part of a larger compendium of models for adaptive speciation. They consider standard Lotka-Volterra dynamics of two mutualists, but let the interaction coefficients between the species evolve, with more benefits from the interaction if the trait values (e.g., flowering time for a flower and emergence time for its insect pollinator) match, with the trade-off that each species also has an independent (e.g., abiotic) optimum value for the trait. Doebeli and Dieckmann find that if these trait optima are far enough apart evolutionary branching might occur, resulting in two species pairs that match each other well.

One potential objection to Doebeli and Dieckmann’s approach is that in their model, the cost of mutualism is due to deviation from some external optimum, rather than costly investments that are beneficial to the partner. This issue is addressed by Ferrière et al. (2002) who consider the evolution of costly provisioning of some resource by partners in an obligate mutualism. Their model incorporates intraspecific competition in each species for both the resources provided by the partner species as well as for some other, mutualism-independent resource. The ecological dynamics of the two populations are stable when rates of resource provision is positive but not too high. Cheat-

ing can destabilize mutualism evolutionarily (and lead populations to extinction) when individuals compete equally well for partner-provided resources regardless of their own provision. A competitive advantage for better mutualists –not surprisingly– counteracts cheating, but there is also the potential for evolutionary branching in the provision rate when the competitive advantage of mutualists is accelerating faster than the cost of providing commodities. Biologically, this means that mechanisms such as partner choice that create steep differences between the fitness of good and bad mutualist partners are more likely to create diversity in mutualistic quality.

Ferrière et al. (2007) extend this model to include a third, unrelated species that only exploits its partner, without providing any benefits. They find that if an exploiter invades a system after high levels of cooperation have evolved, it can take over and cause the collapse of the mutualism. In contrast, if the exploiter invades the system early, the provision rates evolve to lower stable values, but are also immune to further invasion by exploiters, similar to what was found by Golubski and Klausmeier (2010) in an “ultimatum game” model. In the same three-species setting, Jones et al. (2009) further find that the intensity of intraspecific competition among the mutualists and exploiters mediates their coexistence, with high intraspecific competition in both promoting coexistence, and low competition in one tending to drive the other to extinction.

Borrowing an approach from ecological coexistence theory, Hoeksema and Kummel (2003) use a patch occupancy model to explore how cooperative genotypes (in their case, of mycorrhizae) can coexist with non-cooperative genotypes. In contrast to the AD models, they assume the cooperative genotypes are actually competitively inferior to the non-cooperative types, due to their costly resource provisioning. Nonetheless, both can survive if cooperative mycorrhizae is better able to colonize empty root tips as they appear. This is one of the few models that has investigated the trade-offs between different traits affecting mutualist performance, which is likely to be of importance in maintaining mutualism variation.

4.6 Future directions

4.6.1 Integrating mechanisms and modeling approaches

The main message of this chapter is that there are many different mechanisms and selective pathways can resolve conflicts in mutualisms. Existing models tend to consider each selective pathway in isolation, but it is likely that multiple mechanisms frequently operate at the same time, and at different levels. For example, the model of (Akçay and Simms, 2011) for the legume-rhizobium mutualism suggest that the interaction of feedbacks at the nodule and whole plant level determines whether there will be selection for cooperation or not. Box 5.2 by Simms and Porter discuss in more detail the multiple levels of selection in a hierarchically organized symbiosis. Using the legume-rhizobium mutualism as an example, they argue that phenotypic feedbacks and population structure at different scales generate different selection pressures at multiple levels of selection. The net outcome of selection on symbiosis traits is determined by the interaction of these feedbacks and levels. So far, there has been relatively little theoretical work that tries to integrate feedbacks and selection pressures at multiple levels.

In one of the only attempts to integrate multiple mechanisms, Foster and Wenseleers (2006) consider a model where the fitness of a mutualist of species A depends on the mutualist’s own investment, a “passive” phenotypic feedback (akin to by-products) from the partner in response to

investment, active partner choice, and the help received from the partner. They then decompose the latter three feedbacks as a product of a total potential benefit, the realized fraction of the potential (due to the degree of fidelity between the partners), and the relatedness of species A individuals that benefit from each feedback, and write the first-order condition for increased investment (i.e., the total derivative of fitness with respect to A's investment) in this expanded form. Interpreting this equation, Foster and Wenseleers argue that passive feedbacks and active partner choice should play a more important role in promoting mutualisms than genetic associations between species. Importantly, Foster and Wenseleers make explicit that different processes might operate at different scales, and therefore be subject to different population structures, summarized by their separate relatedness coefficients. On the other hand, the way they model phenotypic feedbacks precludes any interaction between these processes; in particular, each species' investment phenotype is assumed to be independent from these feedbacks. This would be inconsistent with what most people regard as active or passive phenotypic feedbacks, in which, say, a positive feedback to a good mutualist would most likely be higher mutualistic investment into that partner (e.g. Doebeli and Knowlton, 1998). When feedbacks occur over investments, a variety of models find that the genetic associations and phenotypic feedback coefficients interact both multiplicatively and additively (Lehmann and Keller, 2006; McGlothlin et al., 2010; Akçay and Van Cleve, 2012; Van Cleve and Akçay, 2014, section 4.3), and Foster and Wenseleer's conclusions will likely change.

Box 2: Phenotypic feedbacks and multilevel selection in mutualisms

Ellen Simms and Stephanie Porter

Costly cooperative investments trade off with individual fitness and directly select for cheating (earlier in chapter; Fig 1A), but positive host responses to symbiont cooperation can counter selection for cheating (Fig 1B; sections 4.3, 4.4, and 4.6.3). These phenotypic feedbacks can align partners' fitness (Sachs et al., 2004) and produce positive net fitness covariance between partners (Fig 1C), despite the potential benefits of cheating.




Figure 1: Direct symbiont fitness cost, W_S^d , as a function of resources allocated to cooperation phenotype, z_s (A). Indirect symbiont fitness benefit W_S^i , as a function of symbiont cooperation phenotype, resulting from feedback from the host (B). Net fitness alignment of host, W_H , and symbiont fitness W_S in stable mutualism (C).

Many mutualisms exhibit hierarchical spatial organization delimited by modules. Examples include flowers in the yucca-moth interaction, inflorescences in the wasp-fig interaction and nodules in the legume-rhizobium interaction. Hierarchical spatial organization of symbionts creates the potential for multi-level selection (Okasha, 2006). Therefore, exclusively

measuring selection and feedbacks at only one level - among modules, for example - would be inadequate for predicting the evolution of cooperation if multiple symbiont genotypes occur within a module or feedbacks and selection can occur at levels above and/or below the module (e.g., see discussion of Akçay and Simms, 2011, in section 4.4.4).

The legume-rhizobium mutualism offers an especially interesting case for examining multi-level selection on symbionts. Soil-dwelling rhizobium bacteria infect legume roots, stimulate nodule formation, proliferate, and may differentiate to fix nitrogen in exchange for plant-derived photosynthates (Oldroyd et al., 2011). Nodules can contain one or more rhizobium genotypes (Friesen and Mathias, 2010). Reproductive rhizobium cells return to soil populations upon nodule senescence and symbiosis is re-initiated with each new nodule. Thus, the scales of selection relevant to the maintenance of cooperative rhizobia are: 1) among free-living rhizobium cells, 2) among rhizobium genotypes within a mixed nodule, 3) among nodules, and 4) among plants (Fig 2). Further, phenotypic feedbacks by hosts can modulate selection within each host generation.




Figure 2: Levels of selection and phenotypic feedbacks relevant to the maintenance of cooperative rhizobia. Dark circles, cooperative rhizobia; light circles, uncooperative rhizobia.

Selection on cooperation in the legume-rhizobium symbiosis may differ in sign and magnitude among levels. Less cooperative rhizobia might better survive in soil or compete for nodulation sites (Denison and Kiers, 2011; Oldroyd et al., 2011), leading to selection against rhizobium cooperation, but phenotypic feedback via host choice prior to nodule formation can favor more cooperative rhizobium cells (Heath and Tiffin, 2009). After infection, sanctions (Kiers et al., 2003) or partner fidelity feedback (Bull and Rice, 1991; Sachs et al., 2004) can favor nodules with rhizobium subpopulations that express greater cooperation, leading to positive selection at the nodule level. If less-cooperative rhizobia could manipulate host responses or free-ride within mixed nodules then direct selection would favor less cooperative rhizobia

within nodule subpopulations (Friesen and Mathias, 2010). However, rhizobium cheating might reduce host fitness and thereby favor cooperation via selection at the level of the entire host. Phenotypic feedback by hosts could also create cell-level selection by favoring more-cooperative rhizobium cells within a nodule. Thus, understanding the maintenance of cooperation in this symbiosis requires accounting for the levels, signs, and relative magnitudes of selection in the soil, in nodules, and among plants.

This accounting will be complicated by environmental variation that could dynamically shift the relative magnitudes of selection and host responses across scales. Increasing soil rhizobium density increases the number of mixed infection nodules on a host (Friesen and Mathias, 2010; Sachs et al., 2010), which would increase within-nodule selection for cheating if uncooperative rhizobia were more fit than cooperative genotypes within mixed infection nodules. Host ontogeny affects the magnitude of selection at the level of the whole host. Seedlings host few nodules and if early viability selection on seedlings dominates selection among plants, then plant-level selection for rhizobium cooperation could be strong. However, the number of nodules, and thus the potential diversity of symbionts sharing a host, increases as plants mature, which reduces whole-plant level selection.

The availability of traded resources could also impact selection (Heath and Tiffin, 2009; Akçay and Simms, 2011). For example, shading and high external nitrogen both increase carbon limitation, leading plants to drive a harder bargain with rhizobia and thus favor “more-stubborn” rhizobia that are poorer nitrogen fixers. However, if shading resulted from inter-plant competition, it could strengthen viability selection among plants, thereby increasing whole-plant level selection in favor of rhizobial cooperation. High soil nitrate availability reduces nodulation rates via autoregulation of nodulation, which could also increase selection for cooperation if it strengthened pre-infection partner choice. While we have highlighted levels of selection that we expect to be sensitive to external contexts, further development of this hypothesis requires more explicit theoretical models to clarify the implications of such selection, as well as empirical measurements of selection on symbionts at multiple scales in contrasting environmental conditions. This research will improve predictions about the evolution of cooperation within realistically rich patterns of environmental heterogeneity.

References

- Akçay, E. and Simms, E. L. (2011). Negotiation, sanctions, and context dependency in the legume-Rhizobium mutualism. *The American naturalist*, 178(1):1–14.
- Bull, J. J. and Rice, W. R. (1991). Distinguishing mechanisms for the evolution of co-operation. *Journal of Theoretical Biology*, 149:63–74.
- Denison, R. F. and Kiers, E. T. (2011). Life histories of symbiotic rhizobia and mycorrhizal fungi. *Current biology : CB*, 21(18):R775–85.
- Friesen, M. L. and Mathias, a. (2010). Mixed infections may promote diversification of mutualistic symbionts: why are there ineffective rhizobia? *Journal of Evolutionary Biology*, 23:323–334.

- Heath, K. D. and Tiffin, P. (2009). Stabilizing mechanisms in a legume–rhizobium mutualism. *Evolution*, 63(3):652–662.
- Kiers, E. T., Rousseau, R. A., West, S. A., and Denison, R. F. (2003). Host sanctions and the legume–rhizobium mutualism. *Nature*, 425:78–81.
- Okasha, S. (2006). *Evolution and the levels of selection*. Oxford University Press.
- Oldroyd, G. E., Murray, J. D., Poole, P. S., and Downie, J. A. (2011). The rules of engagement in the legume–rhizobial symbiosis. *Annual review of genetics*, 45:119–144.
- Sachs, J. L., Mueller, U. G., Wilcox, T. P., and Bull, J. J. (2004). The evolution of cooperation. *Quarterly Review of Biology*, 79(2):135–160.
- Sachs, J. L., Russell, J. E., Lii, Y. E., Black, K. C., Lopez, G., and Patil, a. S. (2010). Host control over infection and proliferation of a cheater symbiont. *Journal of Evolutionary Biology*, 23(9):1919–27.

4.6.2 Context dependency

One of the major lessons from the past couple of decades in mutualism research is that the benefits from most mutualisms (or even whether the interaction is mutualistic or not) depend on both the biotic and abiotic context (see Chapter 11). Increasingly it appears that most variation we see in the costs and benefits of mutualisms in nature are due to these context-dependent effects. Identifying cheaters unequivocally in nature has been challenging (see Chapter 7). One important source of context dependency is the evolved phenotypic feedbacks (partner fidelity or partner choice) that stabilize mutualisms. For example, in the relative sanctions model of West et al. (2002b), rhizobium fitness depends on the other strains connecting the same plant, because the plant’s partner choice is relative to the mean nitrogen fixation rate. Models that explicitly focus on the causes and consequences of context dependency either due to environment (Pringle et al., 2013), life-history variation (Palmer et al., 2010), or phenotypic responses (Akçay and Simms, 2011) will be useful in understanding the variation in the costs and benefits of mutualisms in the environment. Coupling these models with the ecological and resource dynamics of the mutualisms (e.g., soil nitrogen and plant interspecific competition in the legume–rhizobium case) will also help predict the responses of mutualistic interactions to global change (Kiers et al., 2010).

4.6.3 Alignment of evolutionary interests

With the focus on the cheater problem in the literature (see Sachs, Chapter 7), it is easy to conclude mutualisms as a constant battleground. However, empirical research shows broad alignment of interest between parties in most, if not all, mutualisms (e.g. Friesen, 2012, see also Chapter 7). Yet, most theory is based on the assumption that mutualisms must have an underlying conflict of interests. This assumption is certainly supported by what we know about “first principles” in many mutualisms, for example, the fact nitrogen fixation *is* energetically costly for rhizobia. At face value, we seem to have a mismatch between theory and empirical findings. But a more nuanced reading of

the theory suggests a different interpretation. All of the mechanisms discussed in this chapter that increase mutualistic genotypes' fitness relative to cheaters as implicitly aligning the evolutionary interests of the parties. In this view, mutualisms are rife with *potential* for conflicts of interests, but these are resolved through one or more mechanisms. Equivalently, one might say that mutualisms evolve only when the underlying conflicts of interests are sufficiently resolved.

Alignment of interests therefore provides a complementary perspective to the cheating problem. However, theoretical literature only recently started to focus explicitly on resolving conflicts of interests. Along these lines, Worden and Levin (2007) consider a population that starts out playing a prisoner's dilemma game, but is subject to mutations that change the payoff matrix of the game. They show that eventually, the population evolves to play a game of perfectly aligned interests. The evolution towards mutual interest is unconstrained because Worden and Levin assume no cost to helping (or not hurting) one's partner. Later, we (Akçay and Simms, 2011) analyzed a related model where the mutation that gives an incentive to the partner to cooperate is costly to the focal individual. We showed that despite costs, such incentives can invade a population if they are sufficiently big to induce the partner to put up with cheating. Interestingly, under a diploid population genetic model, a stable polymorphism arises between the different alleles that modify the game. The result is a rich diversity of game structures that coexist in a population. Behaviorally, even though defecting individuals on average have higher fitness than cooperators, the system is nonetheless stable, because the genetic polymorphism is not in the behavior itself, but the games individuals play. These models have not yet been extended to interactions between species, which represents a fruitful avenue for future research. It will also be important to develop a formal framework to measure and quantify the alignment of fitness interests. Quantitative genetic theory is likely to provide an good basis for such a theory.

Alignment of interests is a harder problem to solve when there is asymmetric or private information (i.e., one party can condition its behavior on some information while the other party cannot), as in the signaling or screening models discussed in section 4.4.5. The theory of mechanism design in economics (Myerson, 1979; Myerson and Satterthwaite, 1983; Akçay et al., 2012) deals with these situations and concerns itself with finding games that achieve the best outcome (defined according to some performance measure) given private information and conflicts of interests. In general, it is not always possible to achieve full efficiency (i.e., taking advantage of all the opportunities for mutual benefit) when parties are privately informed (Myerson and Satterthwaite, 1983; Akçay et al., 2012). However, mechanism design provides techniques to compute game structures that achieve a second-best outcome given the constraints of private information. If one makes the working assumption that natural selection modifies games to achieve the most mutual benefit that is compatible with individual-level selection, these second-best mechanisms provide testable hypotheses about what we might expect to see in nature. More importantly, an evolutionary version of mechanism design theory can open the door to a deeper understanding of how natural selection structures biological organization for mutual benefit in both mutualisms and within-species cooperation (Akçay et al., 2010).

4.6.4 Conclusion

This chapter presented a survey of the main ideas that have been influential in mutualism research. It is of course impossible to be all-inclusive (the existing literature is already too big for that), but I have tried to highlight the principle ideas in evolutionary theory of mutualisms and the mod-

eling approaches used to address them. We are beginning to see a new generation of theoretical approaches to social evolution and mutualisms, and these hold the promise of both uncovering new theoretical principles but also bringing models closer to data. When this book is updated in a few decades we will surely to see the fruits of that effort.

References

- Akçay, E., Meirowitz, A., Ramsay, K. W., and Levin, S. A. (2012). Evolution of cooperation and skew under imperfect information. *Proceedings of the National Academy of Sciences*, 109(37):14936–14941.
- Akçay, E., Roughgarden, J., Fearon, J. D., Ferejohn, J. A., and Weingast, B. (2010). Biological institutions: the political science of animal cooperation. working paper.
- Akçay, E. and Roughgarden, J. (2007). Negotiation of mutualism: rhizobia and legumes. *Proc. R. Soc. B*, 274(1606):25–32.
- Akçay, E. and Simms, E. L. (2011). Negotiation, sanctions, and context dependency in the legume-Rhizobium mutualism. *The American naturalist*, 178(1):1–14.
- Akçay, E. and Van Cleve, J. (2012). Behavioral responses in structured populations pave the way to group optimality. *The American Naturalist*, 179(2):257–269.
- Akçay, E., Van Cleve, J., Feldman, M. W., and Roughgarden, J. (2009). A theory for the evolution of other-regard integrating proximate and ultimate perspectives. *Proc. Natl. Acad. Sci. USA*, 106(45):19061–6.
- Archetti, M., Ubeda, F., Fudenberg, D., Green, J., Pierce, N. E., and Yu, D. W. (2011). Let the right one in: a microeconomic approach to partner choice in mutualisms. *The American naturalist*, 177(1):75–85.
- Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211:1390–1396.
- Bever, J. D. and Simms, E. L. (2000). Evolution of nitrogen fixation in spatially structured populations of rhizobium. *Heredity*, 85(4):366–372.
- Bshary, R. and Grutter, A. S. (2002). Asymmetric cheating opportunities and partner control in a cleaner fish mutualism. *Anim. Behav.*, 63:547–555.
- Bshary, R. and Grutter, A. S. (2006). Image scoring and cooperation in a cleaner fish mutualism. *Nature*, 441(7096):975–978.
- Bull, J. J. and Rice, W. R. (1991). Distinguishing mechanisms for the evolution of co-operation. *Journal of Theoretical Biology*, 149:63–74.
- Denison, R. F. (2000). Legume sanctions and the evolution of symbiotic cooperation by rhizobia. *American Naturalist*, 156:567–576.

- Dieckmann, U. and Law, R. (1996). The dynamical theory of coevolution: a derivation from stochastic ecological processes. *Journal of Mathematical Biology*, 34(5):579–612.
- Doebeli, M. and Dieckmann, U. (2000). Evolutionary branching and sympatric speciation caused by different types of ecological interactions. *The american naturalist*, 156(S4):S77–S101.
- Doebeli, M. and Knowlton, N. (1998). The evolution of interspecific mutualisms. *Proc. Natl. Acad. Sci. USA*, 95(15):8676–8680.
- Ewald, P. W. (1987). Transmission modes and evolution of the parasitism-mutualism continuum. *Annals of the New York Academy of Sciences*, 503:295–306.
- Ferdy, J.-B. and Godelle, B. (2005). Diversification of transmission modes and the evolution of mutualism. *American Naturalist*, 166(5):613–27.
- Ferrière, R., Bronstein, J. L., Rinaldi, S., Law, R., and Gauduchon, M. (2002). Cheating and the evolutionary stability of mutualisms. *Proc. R. Soc. B*, 269(1493):773–80.
- Ferrière, R., Gauduchon, M., and Bronstein, J. L. (2007). Evolution and persistence of obligate mutualists and exploiters: competition for partners and evolutionary immunization. *Ecology letters*, 10(2):115–26.
- Fine, P. E. (1975). Vectors and vertical transmission: an epidemiologic perspective. *Annals of the New York Academy of Sciences*, 266(1):173–194.
- Foster, K. R. and Kokko, H. (2006). Cheating can stabilize cooperation in mutualisms. *Proc. R. Soc. B*, 273(1598):2233–9.
- Foster, K. R. and Wenseleers, T. (2006). A general model for the evolution of mutualisms. *Journal of Evolutionary Biology*, 19:1283–1293.
- Frank, S. (1994). Genetics of mutualism: the evolution of altruism between species. *Journal of Theoretical Biology*, 170(4):393–400.
- Frank, S. (1998). *Foundations of Social Evolution*. Princeton Univ Press, Princeton.
- Frank, S. A. (1997). The price equation, fisher’s fundamental theorem, kin selection, and causal analysis. *Evolution*, pages 1712–1729.
- Friesen, M. L. (2012). Widespread fitness alignment in the legume–rhizobium symbiosis. *New Phytologist*, 194(4):1096–1111.
- Golubski, A. J. and Klausmeier, C. A. (2010). Control in mutualisms: combined implications of partner choice and bargaining roles. *Journal of theoretical biology*, 267(4):535–545.
- Grman, E., Robinson, T. M. P., and Klausmeier, C. A. (2012). Ecological specialization and trade affect the outcome of negotiations in mutualism. *The American Naturalist*, 179(5):567–581.
- Gubry-Rangin, C., Garcia, M., and Béna, G. (2010). Partner choice in *Medicago truncatula*–*Sinorhizobium* symbiosis. *Proc. R. Soc. B*, 277(1690):1947–51.

- Hamilton, W. D. (1964). The genetical evolution of social behaviour. *Journal of Theoretical Biology*, 7(1):1–16.
- Hammerstein, P. and Hagen, E. H. (2005). The second wave of evolutionary economics in biology. *Trends in Ecology & Evolution*, 20(11):604–9.
- Heath, K. D. and Stinchcombe, J. R. (2014). Explaining mutualism variation: A new evolutionary paradox? *Evolution*, 68:309–317.
- Hoeksema, J. D. and Kummel, M. (2003). Ecological persistence of the plant-mycorrhizal mutualism: a hypothesis from species coexistence theory. *American Naturalist*.
- Hoeksema, J. D. and Schwartz, M. W. (2003). Expanding comparative–advantage biological market models: contingency of mutualism on partner’s resource requirements and acquisition trade-offs. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(1518):913–919.
- Johnstone, R. A. and Bshary, R. (2002). From parasitism to mutualism: partner control in asymmetric interactions. *Ecology Letters*, 5(5):634–639.
- Johnstone, R. A. and Bshary, R. (2008). Mutualism, market effects and partner control. *Journal of Evolutionary Biology*, 21(3):879–88.
- Jones, E. I., Ferrière, R., and Bronstein, J. L. (2009). Eco-evolutionary dynamics of mutualists and exploiters. *The American Naturalist*, 174(6):780–794.
- Kazancıoğlu, E., Klug, H., and Alonzo, S. H. (2012). The evolution of social interactions changes predictions about interacting phenotypes. *Evolution*, 66(7):2056–2064.
- Kiers, E. T., Palmer, T. M., Ives, A. R., Bruno, J. F., and Bronstein, J. L. (2010). Mutualisms in a changing world: an evolutionary perspective. *Ecology Letters*, 13(12):1459–1474.
- Kiers, E. T., Rousseau, R. A., and Denison, R. F. (2006). Measured sanctions: legume hosts detect quantitative variation in rhizobium cooperation and punish accordingly. *Evolutionary Ecology Research*, 8(6):1077–1086.
- Kiers, E. T., Rousseau, R. A., West, S. A., and Denison, R. F. (2003). Host sanctions and the legume-rhizobium mutualism. *Nature*, 425:78–81.
- Kirkpatrick, M. and Ryan, M. J. (1991). The evolution of mating preferences and the paradox of the lek. *Nature*, 350(6313):33–38.
- Kummel, M. and Salant, S. W. (2006). The economics of mutualisms: Optimal utilization of mycorrhizal mutualistic partners by plants. *Ecology*, 87(4):892–902.
- Lehmann, L. and Keller, L. (2006). The evolution of cooperation and altruism—a general framework and a classification of models. *Journal of Evolutionary Biology*, 19(5):1365–1376.
- Lehmann, L. and Rousset, F. (2010). How life history and demography promote or inhibit the evolution of helping behaviours. *Phil. Trans. R. Soc. B*, 365(1553):2599–617.

- Mao-Jones, J., Ritchie, K. B., Jones, L. E., and Ellner, S. P. (2010). How microbial community composition regulates coral disease development. *PLoS biology*, 8(3):e1000345.
- Marco, D. E., Carbajal, J. P., Cannas, S., Pérez-Arnedo, R., Hidalgo-Perea, A., Olivares, J., Ruiz-Sainz, J. E., and Sanjuán, J. (2009). An experimental and modelling exploration of the host-sanction hypothesis in legume-rhizobia mutualism. *Journal of Theoretical Biology*, 259(3):423–33.
- Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246(5427):15–18.
- McGlothlin, J. W. and Brodie III, E. D. (2009). How to measure indirect genetic effects: the congruence of trait-based and variance-partitioning approaches. *Evolution; international journal of organic evolution*, 63(7):1785–95.
- McGlothlin, J. W., Moore, A. J., Wolf, J. B., and Brodie, E. D. (2010). Interacting Phenotypes and the Evolutionary Process. Iii. Social Evolution. *Evolution*, pages 1–17.
- McNamara, J. M., Gasson, C., and Houston, A. I. (1999). Incorporating rules for responding into evolutionary games. *Nature*, 401(6751):368–371.
- Moore, A. J., Brodie III, E. D., and Wolf, J. B. (1997). Interacting phenotypes and the evolutionary process: I. direct and indirect genetic effects of social interactions. *Evolution*, pages 1352–1362.
- Myerson, R. (1979). Incentive compatibility and the bargaining problem. *Econometrica*, 47(1):61–74.
- Myerson, R. and Satterthwaite, M. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281.
- Nash, J. (1950). The bargaining problem. *Econometrica*, 18(2):155–162.
- Noe, R. and Hammerstein, P. (1994). Biological markets: supply and demand determine the effect of partner choice in cooperation, mutualism and mating. *Behavioral Ecology and Sociobiology*, 35:1–11.
- Noe, R. and Hammerstein, P. (1995). Biological markets. *Trends in Ecology & Evolution*, 10:336–339.
- Palmer, T. M., Doak, D. F., Stanton, M. L., Bronstein, J. L., Kiers, E. T., Young, T. P., Goheen, J. R., and Pringle, R. M. (2010). Synergy of multiple partners, including freeloaders, increases host fitness in a multispecies mutualism. *Proc. Natl. Acad. Sci. USA*, 107(40).
- Price, G. R. (1970). Selection and Covariance. *Nature*, 227:520–521.
- Price, G. R. (1972). Extension of covariance selection mathematics. *Annals of Human Genetics*, 35:485–490.
- Pringle, E. G., Akçay, E., Raab, T. K., Dirzo, R., and Gordon, D. M. (2013). Water stress strengthens mutualism among ants, trees and scale insects. *PLoS Biology*, 11(11):e1001705.
- Ricardo, D. (1891). *Principles of political economy and taxation*. G. Bell and sons.

- Riley, J. G. (2001). Silver Signals: Twenty-Five Years of Screening and Signaling. *Journal of Economic Literature*, 39:432–478.
- Rothschild, M. and Stiglitz, J. (1976). Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *The Quarterly Journal of Economics*, 90(4):629.
- Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica*, 50:97–109.
- Sachs, J. L., Mueller, U. G., Wilcox, T. P., and Bull, J. J. (2004). The evolution of cooperation. *Quarterly Review of Biology*, 79(2):135–160.
- Sachs, J. L., Skophammer, R. G., and Regus, J. U. (2011). Evolutionary transitions in bacterial symbiosis. *Proceedings of the National Academy of Sciences*, 108(Supplement 2):10800–10807.
- Scheuring, I. and Yu, D. W. (2012). How to assemble a beneficial microbiome in three easy steps. *Ecology letters*, 15(11):1300–1307.
- Schwartz, M. W. and Hoeksema, J. D. (1998). Specialization and resource trade: biological markets as a model of mutualisms. *Ecology*, 79(3):1029–1038.
- Simms, E. L., Taylor, D. L., Povich, J., Shefferson, R. P., Sachs, J. L., Urbina, M., and Tausczik, Y. (2006). An empirical test of partner choice mechanisms in a wild legume-rhizobium interaction. *Proc. R. Soc. B*, 273(1582):77–81.
- Smith, S. E. and Read, D. J. (2008). *Mycorrhizal symbiosis*. Academic Press.
- Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374.
- Stewart, A. D., Logsdon, J. M., and Kelley, S. E. (2005). An empirical study of the evolution of virulence under both horizontal and vertical transmission. *Evolution*, 59(4):730–739.
- Taylor, P. D. and Day, T. (2004). Stability in negotiation games and the emergence of cooperation. *Proceedings of the Royal Society of London, Series B: Biological Sciences*, 271(1540):669–674.
- Taylor, P. D. and Frank, S. A. (1996). How to make a kin selection model. *Journal of Theoretical Biology*, 180:27–37.
- Van Cleve, J. and Akçay, E. (2014). Pathways to social evolution: reciprocity, relatedness, and synergy. *Evolution*, 68:2245–2258.
- Weeks, A. R., Turelli, M., Harcombe, W. R., Reynolds, K. T., and Hoffmann, A. A. (2007). From parasite to mutualist: rapid evolution of wolbachia in natural populations of drosophila. *PLoS biology*, 5(5):e114.
- West, S. A., Kiers, E. T., Pen, I., and Denison, R. F. (2002a). Sanctions and mutualism stability: when should less beneficial mutualists be tolerated. *Journal of Evolutionary Biology*, 15(5):830–837.
- West, S. A., Kiers, E. T., Simms, E. L., and Denison, R. F. (2002b). Sanctions and mutualism stability: why do rhizobia fix nitrogen? *Proc. R. Soc. B*, 269(1492):685.

- Worden, L. and Levin, S. A. (2007). Evolutionary escape from the prisoner's dilemma. *Journal of Theoretical Biology*, 245:411–422.
- Yamamura, N. (1993). Vertical transmission and evolution of mutualism from parasitism. *Theoretical Population Biology*, 44(1):95–109.
- Yamamura, N., Higashi, M., Behera, N., and Yuichiro Wakano, J. (2004). Evolution of mutualism through spatial effects. *Journal of Theoretical Biology*, 226(4):421–428.
- Zahavi, A. (1975). Mate selection: a selection for a handicap. *Journal of Theoretical Biology*, 53:205–214.