

Games Microbes Play

R. E. LENSKI* and G. J. VELICER

Center for Microbial Ecology, Michigan State University, East Lansing, USA

(Received: ...)

John Maynard Smith has addressed an incredible range of evolutionary problems in his remarkable career, including the development of evolutionary game theory and population-genetic analyses of bacteria in nature, among many others. This paper reviews three studies of complex dynamical interactions among competing microorganisms (including viruses as well as bacteria), each of which can be usefully cast in the game-theoretical framework. These studies demonstrate some of the most interesting outcomes of evolutionary game theory, including multiple stable equilibria, prisoner's dilemma, and balanced polymorphisms. Thus, even simple organisms can exhibit complex behaviors, our understanding of which benefits from game-theoretical analysis. Moreover, by virtue of their simple genetic systems, microbes are especially well suited for performing rigorous tests of evolutionary game theory. Finally, all three of these studies were performed in the laboratory, where the factors that shaped the interactions were further shown to depend on the genetic structure of the microbial populations. Thus, the evolutionary significance of these complex behaviors and interactions in nature remains poorly known. Future progress in understanding these fascinating systems will benefit from integrating evolutionary game theory with fine-scale analyses of population structure aimed at elucidating the patterns of genetic relationship among interacting individuals.

Keywords: Altruism, bacteria, conflict, cooperation, fitness, game theory, Maynard Smith, population genetics, prisoner's dilemma, viruses.

1. Introduction

John Maynard Smith (JMS) has worked on a remarkable diversity of evolutionary problems in his career. No doubt two of his many lasting contributions will be his application of game theory to animal behavior (Maynard Smith and Price, 1973; Maynard Smith and Parker, 1976; Maynard Smith, 1982) and his more recent work applying population genetics to bacteria (Maynard Smith, 1990; Maynard Smith et al., 1991, 1993). As far as we know, JMS has not drawn any direct link between these two interests. However, we believe that the answers to some fascinating biological questions will require a fusion between these areas.

When discussing game theoretical problems in evolutionary biology, one usually thinks first of the behaviors of complex organisms with brains: humans, birds, fish, and perhaps the occasional dung fly. Indeed, the very names of the games and their

players – for example, “prisoner's dilemma” and “hawks and doves” – reflect that tendency. But while evolutionary game theory has flourished as a mathematical theory, it has been difficult to test the models using complex organisms with brains, especially given the problems of characterizing and isolating genetically influenced differences in the behaviors of interest. Such organisms tend to be hard to manipulate, they have long generations, their behaviors are subject to complex environmental and genetic influences, and so on. Or as JMS once told one of us, “It's bloody hell easier to get good data from bacteria than from birds.” So while JMS continues to *watch* the birds that have intrigued him since his youth, he now spends much of his time *analyzing* bacteria.

The bacterial data analyzed by JMS so far have been DNA sequences and other molecular markers, despite his deeper interest in organismal phenotypes. Sequences can sometimes be interesting (especially in the hands of a master), but the evolutionary successes and failures of organisms are played out at the phenotypic level. The question is: Do

* Corresponding author: R. E. Lenski, Center for Microbial Ecology, Michigan State University, East Lansing, MI 48824, USA. E-mail: lenski@msu.edu

microorganisms have any interesting phenotypes, in particular the sorts of complex behaviors that produce the tension between conflict and cooperation that is the essence of so much game theory? (This tension arises in certain non-zero-sum games because cooperation can, at least in principle, be rewarded. No such tension exists in zero-sum games.) Indeed, many microbes have fascinating behaviors. And because microbes are genetically simple and easy to manipulate, one can see more clearly the range of game theoretical outcomes and elucidate the factors that produce different outcomes.

In this paper, we review three studies with microbes that examine altruism, cheating, and the tension between conflict and cooperation. In each case, we describe the organisms, their alternative behaviors, the frequency-dependent interactions that ensue, and the resulting payoff (fitness) matrices. Each study nicely illustrates a fundamentally different payoff matrix. But each study also raises important questions about how these laboratory-based interactions play out in nature. To answer these questions will require linking our game-theoretical understanding of the behaviors with knowledge of the microorganism's fine-scale population-genetic structure.

2. "Suicide bombers" in *E. coli*

Different strains of the bacterium *Escherichia coli* engage in fierce competition with one another over limiting resources. Among the weapons that some genotypes deploy are colicins, which are particles that kill most other *E. coli* strains. Other cells belonging to the colicin-producing genotype are unharmed by the particles, but the genotype pays another price: producing colicins is lethal to the individual cell that does so. In effect, some individuals of the producing genotype become suicidal bombers that kill themselves, while potentially taking out many individuals of other genotypes, thereby freeing up resources for their own kind. Of course, if all individuals of the producing population actually expressed the bomber phenotype, then the genotype would become extinct. What happens instead is that only a small fraction of cells of the producing genotype, say 1%, are stochastically induced to express that phenotype, even though the genetic *potential* to express the phenotype is inherited clonally.

Chao and Levin (1981) performed an elegant set of experiments to identify the conditions in which a colicin-producing genotype would have a selective advantage over a non-producing, sensitive strain. They did not explicitly cast their results in a game-theoretical context (nor did they use the term suicide bomber). But they did cast their study in the context of understanding the evolution of an altruistic trait, and their findings can readily be examined in a game-theoretical framework. In the context of their paper, the production of colicins – that is, the expression of the suicide-bomber phenotype – is an altruistic behavior.

In humans, one may not usually associate the actions of a suicide bomber with altruism. Yet, in an evolutionary context, an altruistic act is one that harms the future reproduction of the individual who performs the act while benefiting the group to which that individual belongs. Thus, a worker honeybee who dies while giving a sting in the defense of her hive is behaving altruistically. The action of a bomber in a genetically heterogeneous *E. coli* population may also fulfill both criteria: the action clearly harms the bomber's own future success, while providing additional resources for its own type by eliminating competitors. (Indeed, even in humans, suicide bombers are sometimes revered as martyrs who sacrifice their lives for the benefit of the group to which they belong.)

Chao and Levin (1981) examined the interaction between producers and sensitive genotypes in two distinct environments, a mass-action liquid environment and a physically structured surface environment. The interaction in the liquid medium is such that each genotype has the advantage when it is common, and neither type can invade when it is initially rare. This positive frequency-dependence means that the system has two stable and monomorphic equilibria (as well as a polymorphic equilibrium that is unstable). The reasons for these curious dynamics are readily understood in a game-theoretical context.

In the payoff matrix shown below, B denotes the genotype that occasionally expresses the suicide-bomber phenotype, and in which all individuals are protected against a bomb's effects. N is a genotype that never expresses this phenotype but is vulnerable to killing by a bomb. Each element in the matrix is the relative fitness that accrues, on average, to an individual belonging to the genotype listed on the left

when the vast majority of individuals in the population have the genotype indicated at the top. [Note: This matrix representation is not the conventional one based on pairwise games, and it cannot be used to solve precise equilibria as though it were. However, it can be used to understand invasion conditions, and hence the opportunity for coexistence, in a game theoretical context. Also, in this and later matrices, the precise values in the four cells are not important, only their relative ranks are. The numbers that we show preserve the ranks in the original studies, but not the exact values, which varied depending on experimental details. Among the three studies we review, only the one by Turner and Chao (1999) explicitly derived the fitness matrix; in the other studies, relative fitness ranks were readily interpreted from data in the papers.]

	B	N
B	0.99	0.99
N	0	1

In a population in which every cell is N, the average fitness is set to unity. In a pure population of B, the average fitness is somewhat lower owing to the 1% or so of the cells in each generation that self-destruct, but which cause no harm to their surviving relatives. Next consider when N is rare and tries to invade a dense population containing millions of B cells in a tiny volume. Although only a small fraction of the B become suicide bombers, the overall effect is such a large amount of “shrapnel” in the medium that each N is killed before it can reproduce.

Finally, when B is rare, it bears the same cost of the suicidal bombings, but it does not receive any appreciable benefit. Why is this so? In the mass-action environment – and when the B genotype is rare – the resources that become available as a consequence of the occasional death of an N-type cell are distributed randomly (by diffusion) to all individuals in the population, regardless of their genotype. In other words, the B genotype pays the cost but does not receive any disproportionate benefit. But, you may ask, what about the extra deaths of N caused by B? In principle, this creates a slight asymmetry in the fitness matrix, because the average payoff to N when competing against all N is 1, whereas the payoff to N when competing with rare B (mostly N) will be slightly less than 1 owing to bombing deaths. In

practice, however, this effect on the death rate of N is trivial in the limit as B becomes rare. Imagine a population of 2×10^7 cells, the vast majority of which are N. A small minority in the population are 100 B cells, one of which expresses the bomber phenotype, which causes the death of 200 of the N genotype. N suffers many more casualties than does B, but the average reduction in fitness of B is 1% whereas the fitness loss to N is only 0.001%. These illustrative calculations depend on certain parameters, which could be varied to give B an advantage even when very rare, for example, by increasing the kill ratio. However, the experimental results obtained by Chao and Levin (1981) confirm that B cannot invade when it is initially very rare in the liquid environment that they studied.

Thus, the fitness matrix for the mass-action environment shows that it pays to be B when everyone else is also B, whereas it pays to be N when all the other cells are N, hence the two stable equilibria. By contrast, in a structured surface environment, the B genotype has a selective advantage at all initial frequencies and spreads to fixation. This advantage accrues to the B genotype because bacteria grow as colonies – and thus more often adjacent to other members of their own type – on a surface. In that case, resources that become available upon the death of N-type cells are not distributed randomly across the entire surface, but instead the resources tend to remain near the B killer and thus near the killer’s kin who are also B. Therefore, even a tiny number of B cells can invade a population of N growing on a surface (provided the bomb does not go off before a few relatives have accumulated in a patch). As a consequence, the fitness matrix on a surface has the following form:

	B	N
B	0.99	1.1
N	0	1

In the physically-structured surface environment, it pays to be the B genotype no matter what the rest of the population is doing, hence its eventual spread to fixation across all initial frequencies. This study by Chao and Levin (1981) demonstrated, in the laboratory, the potential for both kinds of dynamic, one in which B prevails in any case and the other in which the eventual winner depends on the initial densities of B and N. Which dynamic is more important in na-

ture? The answer clearly must depend on the physical structure of the environment, which governs not only the diffusion of the resources but also the population-genetic structure of the bacteria. To extend our understanding of colicin-producers from the laboratory to nature, it will become necessary to analyze the population-genetic structure of colicin-producing and sensitive bacteria at a fine level of resolution, one that corresponds to the physical scale of their interactions.

The interaction between B and N is very interesting, but there also exist other strategies. For example, some genotypes are immune to the effects of the bombs (like B) but never engage in their detonation (like N). Chao and Levin (1981) studied a resistant mutant (derived from N) that was altered for the bomb's target, as opposed to one altered for the bomb's detonation (which would derive from B). The mutant they chose was rather uninteresting from a game-theoretical standpoint, because the fitness cost of resistance to the mutant was greater than the combined costs of detonation and intrinsic resistance to the bomber. Thus, the mutant could invade neither B nor N. One can imagine, however, a more interesting B-derived mutant that pays the cost of intrinsic resistance but avoids the cost of stochastic detonation. Such a mutant, M, would imply non-transitive pairwise interactions in a structured environment: (i) rare M invades B owing to M's lower cost and its resistance to B; (ii) rare N invades M owing to N's even lower cost and no need for any resistance; and (iii) rare B invades N because the benefit, on a surface, to B of killing N more than offsets the costs of doing so (cf. Getty, 1979; Durrett and Levin, 1997).

3. Viral "prisoners" in host cells

Even simpler organisms than bacteria can exhibit interactions between genotypes that may be understood using game theory. A recent study by Turner and Chao (1999) showed that $\phi 6$, a small RNA virus that infects certain bacteria, can become caught in the prisoner's dilemma. The prisoner's dilemma refers to those interactions in which the payoffs are structured such that the stable strategy is to defect, D, rather than cooperate, C, even though the payoff to each individual is lower when they both defect than when they both cooperate.

The name comes from a puzzle involving the behavior of two prisoners, when they are each asked by the jailer (out of earshot of the other) which one of them did some misdeed, and they are both knowledgeable of the payoffs. If neither accuses the other – that is, they cooperate with one another – then their prison terms will be unchanged. If they both accuse one another – that is, they defect – then the jailer will think they are both liars and add five more years to the sentence of each. If one prisoner accuses the other (defects), while the other does not accuse the first (cooperates), then the jailer assumes both are truthful; he adds ten years to the sentence of the accused and takes off five years from the sentence of the accuser. No matter what one prisoner may do, it is better for the other prisoner to defect: if the first cooperates, the second stands to reduce his own sentence by five years if he defects; and if the first also defects, the second protects himself by defecting because his sentence will be lengthened by five years instead of ten. Both prisoners, being rational, will defect; and each will receive a longer sentence than if they had both cooperated, hence the dilemma.

So how do viruses find themselves in a prisoner's dilemma inside their host's cells? Turner and Chao (1999) sought to study the effect of recombination on virus evolution. Viruses can recombine with one another only when they co-infect the same individual cell, and the authors manipulated the extent of viral recombination by varying the multiplicity of infection (MOI). At high MOI, many viruses infect each cell and can recombine with one another; at low MOI, each virus infects a cell that no other virus has found and recombination cannot occur. The authors were surprised to find that the viruses evolved lower fitness relative to the ancestral state in the high MOI treatment, whereas viral fitness improved in the low MOI treatment.

The reason for this declining fitness may lie in the fact that, when two or more viruses infect the same cell, they end up sharing the same intracellular pool of products needed for their replication. This intracellular pool comprises virus products as well as those materials intrinsic to the host cell. Thus, in addition to scramble competition between viruses for host products, there may also be more complex interactions. For example, (i) viruses may actively interfere with one another; or (ii) they may manipulate their own contributions to the pool of viral products in a way that is beneficial to the individual virus but

harmful to the group of co-infecting viruses as a whole. In effect, selection at high MOI may favor viruses that attempt to monopolize intracellular resources and thereby push the system toward lower productivity.

To test this hypothesis, Turner and Chao (1999) performed a set of careful experiments. They first examined whether competition at high MOI between the genotypes evolved at high MOI and their progenitor was frequency dependent. Indeed, the evolved genotypes had a fitness advantage across the entire range of frequencies, but the magnitude of their relative advantage declined significantly with their own frequency. In other words, as shown by the payoff matrix below, it was always better to defect at high MOI, but the magnitude of the benefit of defecting was less when the other viruses were also trying to take more than their fair share of the pool of building materials.

	C	D
C	1	0.6
D	2	0.8

Finally, to establish that the net effect of viral evolution was a decline in absolute fitness, and that this paradox reflected interactions inside the host cell, Turner and Chao (1999) performed another experiment in which the two viral types were mixed only after they had already infected the host cells (again at high MOI to avoid confounding effects of MOI on virus production). Hence, all intracellular competition was between identical viruses, and the evolved defectors could not exploit their more cooperative ancestor. In this experiment, it was observed that the ancestral virus was indeed more productive ($1 > 0.8$) than the defector genotype that evolved at high MOI, thereby establishing this essential feature of the prisoner's dilemma. In other words, the defecting strategy invaded and took over the evolving population, even though this process led to a lower average fitness than if the viruses had found some way to maintain cooperation.

(Note: The payoff matrix for bacterial suicide bombers and non-bombers in the physically structured environment also has the properties of a prisoner's dilemma, except that B was defined as the altruist in that game and in fact wins. However, this distinction may be semantic, as the evolutionary dynamic is similar. That is, a particular strategy in-

vades and takes over while reducing the overall productivity of the group.)

The study by Turner and Chao (1999) demonstrated a complex game-theoretical outcome in a very simple biological system. Moreover, by observing that this outcome was obtained when viruses evolved at high, but not low, MOI, their work points again to the need to understand the population-genetic structure of these viruses in nature. If multiple viruses frequently co-infect the same cell, then the prisoner's dilemma may be important for their evolution in nature; but if single infections are the rule, then such interactions may be inconsequential.

Turner and Chao (1999) examined two genetic strategies that viruses can follow, but others may also come into play. Defective interfering (DI) particles exist in many viral species. In effect, DI particles are more extreme forms of the defectors in f6. DI particles not only benefit from co-infection with a cooperator, but in fact they absolutely require a cooperator's presence for their own replication. As a consequence, DI particles can invade a population of cooperators, but they cannot spread to fixation. The expected equilibrium at high MOI is therefore a genetically mixed population (Szathmáry, 1992), similar to that in our third case study below. Adding more game-theoretical complexity, there exist other viral genotypes that are resistant to exploitation by DI particles (Szathmáry, 1992).

4. Developmental "cheaters" in myxobacteria

The final study involves myxobacteria, which exhibit several remarkable behaviors that are more obviously social than those described in either of the previous examples. Myxobacteria are soil-dwelling prokaryotes that share several interesting behaviors with eukaryotic slime molds. In response to starvation, both slime molds and myxobacteria will physically aggregate and then undergo multicellular development, involving many intercellular signals, to produce complex fruiting bodies in which only a fraction of the cells become hardy spores that survive to germinate later (or elsewhere) under more favorable circumstances. Also, both slime molds and myxobacteria are highly motile predators that mount collective attacks on prey, which they digest by secreting extracellular enzymes and then transporting

the nutrients back inside their cells. (This phenotypic convergence in complex behaviors is especially remarkable given that, in the tree of life, myxobacteria are near *E. coli* while slime molds are our own close cousins.)

These social behaviors indicate a high degree of cooperation. However, such cooperation seems vulnerable to exploitation by social “parasites” that derive the benefits of belonging to the group while avoiding the associated costs. Imagine, for example, a mutant that produces less than its share of extracellular enzyme for digesting prey and uses the savings to produce more enzyme used for taking up the resulting nutrients. Or imagine another mutant that somehow finds a way, during fruiting-body development, to become a surviving spore (rather than a dying stalk cell) more often than its neighbors. Do such cheaters exist? And if so, what are their evolutionary fates?

Velicer et al. (2000) studied, in the laboratory, the potential for developmental cheating to occur in *Myxococcus xanthus*. They examined two different classes of potential cheaters, with each class being defective in development when grown alone. Some potential cheaters were lines that evolved for 1000 generations in an asocial regime (a nutrient-rich and physically unstructured environment in which fruiting-body development was never needed). The other potential cheaters were mutants defective in molecular signaling pathways, which were generated by geneticists seeking to understand fruiting-body development. In both cases, potential cheaters were mixed at various initial ratios with their developmentally proficient progenitors, and their respective contributions to surviving spores were measured.

A demonstration of developmental cheating would require two outcomes in this system. First, when introduced at low frequency in a mixture, a cheater should obtain disproportionate success in becoming a spore relative to its progenitor. Second, when a cheater is present at sufficiently high frequency in the mixture, the overall yield of the mixture should be lower than that of a population composed entirely of the progenitor. In fact, Velicer et al. (2000) found that these criteria were met by three of six evolved lines and two of three mutants they examined. And the effects were not subtle, as all the cheaters were over-represented among the spores

produced by tenfold (or more) when introduced at an initial frequency of 1%.

Clearly then, developmental cheaters are able to invade, but what is their fate? Can they sweep to fixation, as did the viruses at high MOI? Or do the cheaters lose their advantage within the mixture when they become too common? Velicer et al. (2000) examined these questions in two cheaters, and in both cases the developmentally-proficient genotype regained its advantage when a cheater became numerically dominant. The payoff matrix below summarizes the relative fitness values based on spore production during development of the multicellular fruiting-body.

	P	C
P	1	0.5
C	10	0.1

P indicates the developmentally proficient genotype, while C denotes a developmental cheater. As before, each entry is the mean fitness of the genotype listed on the left when the vast majority of individuals with which it competes have the genotype shown at the top. This matrix clearly corresponds to a balanced polymorphism, because each type has a selective advantage when rare.

This study demonstrates the existence of genetically-based cheating behavior during the development of myxobacteria fruiting bodies. But whether cheaters are rare or common in nature depends on the population-genetic structure of the bacteria, which in turn may depend on whether spores are dispersed as individuals or as clumps (Velicer et al. 2000). If spores disperse as individuals, then subsequent fruiting bodies may often develop from progeny of a single founder. In that case, fruiting bodies will be genetically homogeneous, and competition will occur between (rather than within) fruiting bodies. The P genotype would be favored ($1 > 0.1$), while C would depend on new mutations in each developmental cycle for its existence. Alternatively, if spores disperse in clumps, then there should be frequent competition within fruiting bodies, a situation in which C has the advantage when rare ($10 > 1$) and can persist indefinitely as a social parasite.

In fact, the spores of myxobacteria are quite sticky and it is difficult to separate them, suggesting that the latter scenario is correct. One may ask why

the spores should be so sticky, if that promotes cheating. The growth of myxobacteria on nonhydrolyzed media in the laboratory is density-dependent (Rosenberg et al., 1977), suggesting that spores in nature may germinate faster in clumps than they would in isolation. Also, after germination in a favorable environment, the myxobacteria may require many cells to forage successfully and to develop into fruiting bodies during the next bout of starvation. In other words, it may be less costly to support a few cheaters than for these social organisms to go it alone even in a favorable environment. We emphasize that these adaptive hypotheses for sticky spores are speculative, but they do serve to illustrate that there may be costs to opposing cheating, as well as benefits.

5. Summary

We have reviewed three laboratory studies in which simple microorganisms exhibit surprisingly complex behaviors during competition. In each case, the resulting dynamics can be readily understood using evolutionary game theory. Yet, questions also remain as to which strategies prevail in nature. Finding the answers will require elucidating the population-genetic structure of the microorganisms to determine whether competitive interactions are mostly within or between genetically distinct clones, whether the genes influencing the relevant behaviors are transmitted horizontally as well as vertically, and so on. JMS has led the way in evolutionary game theory, and in population-genetic analyses of microbes. Now we should combine these two approaches to better understand the evolutionary forces that shape fascinating behaviors, both cooperative and selfish, that exist throughout the biological realm.

Acknowledgements

This paper is dedicated to JMS, a great thinker and teacher of us all. We thank L. Chao, T. Getty, L. Kroos, D. Lenski, B. Levin, E. Szathmary, and P. Turner for interesting and helpful conversations, collaborations, and comments on this paper. Our research on bacterial evolution has been supported by Michigan State University and the U.S. National Science Foundation.

References

- CHAO, L. and LEVIN, B. R. (1981): Structured habitats and the evolution of anticompeteritor toxins in bacteria. *Proc. Natl. Acad. Sci. USA* **78**:6324–6328.
- DURRETT, R. and LEVIN, S. A. (1997): Allelopathy in spatially distributed populations. *J. Theor. Biol.* **185**:165–171.
- GETTY, T. (1979): On the benefits of aggression: the adaptiveness of inhibition and super territories. *Amer. Nat.* **114**:605–609.
- MAYNARD SMITH, J. (1982): *Evolution and the Theory of Games*. Cambridge Univ. Press, Cambridge.
- MAYNARD SMITH, J. (1990): The evolution of prokaryotes: does sex matter? *Annu. Rev. Ecol. Syst.* **21**:1–12.
- MAYNARD SMITH, J. and PARKER, G. A. (1976): The logic of asymmetric contests. *Anim. Behav.* **25**:1–9.
- MAYNARD SMITH, J. and PRICE, G. R. (1973): The logic of animal conflict. *Nature* **246**:15–18.
- MAYNARD SMITH, J., DOWSON, C. G. and SPRATT, B. G. (1991): Localized sex in bacteria. *Nature* **349**:29–31.
- MAYNARD SMITH, J., SMITH, N. H., O’ROURKE, M. and SPRATT, B. G. (1993): How clonal are bacteria? *Proc. Natl. Acad. Sci. USA* **90**:4384–4388.
- ROSENBERG, E., KELLER, K. H. and DWORKIN, M. (1977): Cell density-dependent growth of *Myxococcus xanthus* on casein. *J. Bacteriol.* **129**:770–777.
- SZATHMARY, E. (1992): Natural selection and dynamical coexistence of defective and complementing virus segments. *J. Theor. Biol.* **157**:383–406.
- TURNER, P. E. and CHAO, L. (1999): Prisoner’s dilemma in an RNA virus. *Nature* **398**:441–443.
- VELICER, G. J., KROOS, L. and LENSKI, R. E. (2000): Developmental cheating in the social bacterium *Myxococcus xanthus*. *Nature* **403**, in press.