From the book *Networks, Crowds, and Markets: Reasoning about a Highly Connected World.*
By David Easley and Jon Kleinberg. Cambridge University Press, 2010.
Complete preprint on-line at http://www.cs.cornell.edu/home/kleinber/networks-book/

Chapter 7

Evolutionary Game Theory

In Chapter 6, we developed the basic ideas of game theory, in which individual players make decisions, and the payoff to each player depends on the decisions made by all. As we saw there, a key question in game theory is to reason about the behavior we should expect to see when players take part in a given game.

The discussion in Chapter 6 was based on considering how players simultaneously reason about what the other players may do. In this chapter, on the other hand, we explore the notion of *evolutionary game theory*, which shows that the basic ideas of game theory can be applied even to situations in which no individual is overtly reasoning, or even making explicit decisions. Rather, game-theoretic analysis will be applied to settings in which individuals can exhibit different forms of behavior (including those that may not be the result of conscious choices), and we will consider which forms of behavior have the ability to persist in the population, and which forms of behavior have a tendency to be driven out by others.

As its name suggests, this approach has been applied most widely in the area of evolutionary biology, the domain in which the idea was first articulated by John Maynard Smith and G. R. Price [375, 376]. Evolutionary biology is based on the idea that an organism's genes largely determine its observable characteristics, and hence its *fitness* in a given environment. Organisms that are more fit will tend to produce more offspring, causing genes that provide greater fitness to increase their representation in the population. In this way, fitter genes tend to win over time, because they provide higher rates of reproduction.

The key insight of evolutionary game theory is that many behaviors involve the *interaction* of multiple organisms in a population, and the success of any one of these organisms depends on how its behavior interacts with that of others. So the fitness of an individual organism can't be measured in isolation; rather it has to be evaluated in the context of the full population in which it lives. This opens the door to a natural game-theoretic analogy:

Draft version: June 10, 2010

an organism's genetically-determined characteristics and behaviors are like its strategy in a game, its fitness is like its payoff, and this payoff depends on the strategies (characteristics) of the organisms with which it interacts. Written this way, it is hard to tell in advance whether this will turn out to be a superficial analogy or a deep one, but in fact the connections turn out to run very deeply: game-theoretic ideas like equilibrium will prove to be a useful way to make predictions about the results of evolution on a population.

7.1 Fitness as a Result of Interaction

To make this concrete, we now describe a first simple example of how game-theoretic ideas can be applied in evolutionary settings. This example will be designed for ease of explanation rather than perfect fidelity to the underlying biology; but after this we will discuss examples where the phenomenon at the heart of the example has been empirically observed in a variety of natural settings.

For the example, let's consider a particular species of beetle, and suppose that each beetle's fitness in a given environment is determined largely by the extent to which it can find food and use the nutrients from the food effectively. Now, suppose a particular mutation is introduced into the population, causing beetles with the mutation to grow a significantly larger body size. Thus, we now have two distinct kinds of beetles in the population — small ones and large ones. It is actually difficult for the large beetles to maintain the metabolic requirements of their larger body size — it requires diverting more nutrients from the food they eat — and so this has a negative effect on fitness.

If this were the full story, we'd conclude that the large-body-size mutation is fitness-decreasing, and so it will likely be driven out of the population over time, through multiple generations. But in fact, there's more to the story, as we'll now see.

Interaction Among Organisms. The beetles in this population compete with each other for food – when they come upon a food source, there's crowding among the beetles as they each try to get as much of the food as they can. And, not surprisingly, the beetles with large body sizes are more effective at claiming an above-average share of the food.

Let's assume for simplicity that food competition in this population involves two beetles interacting with each other at any given point in time. (This will make the ideas easier to describe, but the principles we develop can also be applied to interactions among many individuals simultaneously.) When two beetles compete for some food, we have the following possible outcomes.

- When beetles of the same size compete, they get equal shares of the food.
- When a large beetle competes with a small beetle, the large beetle gets the majority of the food.

- In all cases, large beetles experience less of a fitness benefit from a given quantity of food, since some of it is diverted into maintaining their expensive metabolism.

Thus, the fitness that each beetle gets from a given food-related interaction can be thought of as a numerical payoff in a two-player game between a first beetle and a second beetle, as follows. The first beetle plays one of the two strategies *Small* or *Large*, depending on its body size, and the second beetle plays one of these two strategies as well. Based on the two strategies used, the payoffs to the beetles are described by Figure 7.1.

		Beetle 2	
		Small	*Large*
Beetle 1	*Small*	5, 5	1, 8
	Large	8, 1	3, 3

Figure 7.1: The Body-Size Game

Notice how the numerical payoffs satisfy the principles just outlined: when two small beetles meet, they share the fitness from the food source equally; large beetles do well at the expense of small beetles; but large beetles cannot extract the full amount of fitness from the food source. (In this payoff matrix, the reduced fitness when two large beetles meet is particularly pronounced, since a large beetle has to expend extra energy in competing with another large beetle.)

This payoff matrix is a nice way to summarize what happens when two beetles meet, but compared with the game in Chapter 6, there's something fundamentally different in what's being described here. The beetles in this game aren't asking themselves, "What do I want my body size to be in this interaction?" Rather, each is genetically hard-wired to play one of these two strategies through its whole lifetime. Given this important difference, the idea of choosing strategies — which was central to our formulation of game theory — is missing from the biological side of the analogy. As a result, in place of the idea of Nash equilibrium — which was based fundamentally on the relative benefit of changing one's own personal strategy — we will need to think about strategy changes that operate over longer time scales, taking place as shifts in a population under evolutionary forces. We develop the fundamental definitions for this in the next section.

7.2 Evolutionarily Stable Strategies

In Chapter 6, the notion of Nash equilibrium was central in reasoning about the outcome of a game. In a Nash equilibrium for a two-player game, neither player has an incentive to deviate from the strategy they are currently using — the equilibrium is a choice of strategies that tends to persist once the players are using it. The analogous notion for evolutionary settings will be that of an *evolutionarily stable strategy* — a genetically-determined strategy that tends to persist once it is prevalent in a population.

We formulate this as follows. Suppose, in our example, that each beetles is repeatedly paired off with other beetles in food competitions over the course of its lifetime. We will assume the population is large enough that no two particular beetles have a significant probability of interacting with each other repeatedly. A beetle's overall fitness will be equal to the average fitness it experiences from each of its many pairwise interactions with others, and this overall fitness determines its reproductive success — the number of offspring that carry its genes (and hence its strategy) into the next generation.

In this setting, we say that a given strategy is *evolutionarily stable* if, when the whole population is using this strategy, any small group of invaders using a different strategy will eventually die off over multiple generations. (We can think of these invaders either as migrants who move to join the population, or as mutants who were born with the new behavior directly into the population.) We capture this idea in terms of numerical payoffs by saying that when the whole population is using a strategy S, then a small group of invaders using any alternate strategy T should have strictly lower fitness than the users of the majority strategy S. Since fitness translates into reproductive success, evolutionary principles posit that strictly lower fitness is the condition that causes a sub-population (like the users of strategy T) to shrink over time, through multiple generations, and eventually die off with high probability.

More formally, we will phrase the basic definitions as follows.

- We say the *fitness* of an organism in a population is the expected payoff it receives from an interaction with a random member of the population.

- We say that a strategy T *invades* a strategy S at level x, for some small positive number x, if an x fraction of the underlying population uses T and a $1 - x$ fraction of the underlying population uses S.

- Finally, we say that a strategy S is *evolutionarily stable* if there is a (small) positive number y such that when any other strategy T invades S at any level $x < y$, the fitness of an organism playing S is strictly greater than the fitness of an organism playing T.

Evolutionarily Stable Strategies in our First Example. Let's see what happens when we apply this definition to our example involving beetles competing for food. We will first check whether the strategy *Small* is evolutionarily stable, and then we will do the same for the strategy *Large*.

Following the definition, let's suppose that for some small positive number x, a $1 - x$ fraction of the population uses *Small* and an x fraction of the population uses *Large*. (This is what the picture would look like just after a small invader population of large beetles arrives.)

- What is the expected payoff to a small beetle in a random interaction in this population? With probability $1 - x$, it meets another small beetle, receiving a payoff of 5, while with probability x, it meets a large beetle, receiving a payoff of 1. Therefore its expected payoff is
$$5(1 - x) + 1 \cdot x = 5 - 4x.$$

- What is the expected payoff to a large beetle in a random interaction in this population? With probability $1 - x$, it meets a small beetle, receiving a payoff of 8, while with probability x, it meets another large beetle, receiving a payoff of 3. Therefore its expected payoff is
$$8(1 - x) + 3 \cdot x = 8 - 5x.$$

It's easy to check that for small enough values of x (and even for reasonably large ones in this case), the expected fitness of large beetles in this population exceeds the expected fitness of small beetles. Therefore *Small* is not evolutionarily stable.

Now let's check whether *Large* is evolutionarily stable. For this, we suppose that for some very small positive number x, a $1-x$ fraction of the population uses *Large* and an x fraction of the population uses *Small*.

- What is the expected payoff to a large beetle in a random interaction in this population? With probability $1-x$, it meets another large beetle, receiving a payoff of 3, while with probability x, it meets a small beetle, receiving a payoff of 8. Therefore its expected payoff is
$$3(1 - x) + 8 \cdot x = 3 + 5x.$$

- What is the expected payoff to a small beetle in a random interaction in this population? With probability $1 - x$, it meets a large beetle, receiving a payoff of 1, while with probability x, it meets another small beetle, receiving a payoff of 5. Therefore its expected payoff is
$$(1 - x) + 5 \cdot x = 1 + 4x.$$

In this case, the expected fitness of large beetles in this population exceeds the expected fitness of small beetles, and so *Large* is evolutionarily stable.

Interpreting the Evolutionarily Stable Strategy in our Example. Intuitively, this analysis can be summarized by saying that if a few large beetles are introduced into a population consisting of small beetles, then the large beetles do extremely well — since they rarely meet each other, they get most of the food in almost every competition they experience. As a result, the population of small beetles cannot drive out the large ones, and so *Small* is not evolutionarily stable.

On the other hand, in a population of large beetles, a few small beetles will do very badly, losing almost every competition for food. As a result, the population of large beetles resists the invasion of small beetles, and so *Large* is evolutionarily stable.

Therefore, if we know that the large-body-size mutation is possible, we should expect to see populations of large beetles in the wild, rather than populations of small ones. In this way, our notion of evolutionary stability has predicted a strategy for the population — as we predicted outcomes for games among rational players in Chapter 6, but by different means.

What's striking about this particular predicted outcome, though, is the fact that the fitness of each organism in a population of small beetles is 5, which is larger than the fitness of each organism in a population of large beetles. In fact, the game between small and large beetles has precisely the structure of a Prisoner's Dilemma game; the motivating scenario based on competition for food makes it clear that the beetles are engaged in an arms race, like the game from Chapter 6 in which two competing athletes need to decide whether to use performance-enhancing drugs. There it was a dominant strategy to use drugs, even though both athletes understand that they are better off in an outcome where neither of them uses drugs — it's simply that this mutually better joint outcome is not sustainable. In the present case, the beetles individually don't understand anything, nor could they change their body sizes even if they wanted to. Nevertheless, evolutionary forces over multiple generations are achieving a completely analogous effect, as the large beetles benefit at the expense of the small ones. Later in this chapter, we will see that this similarity in the conclusions of two different styles of analysis is in fact part of a broader principle.

Here is a different way to summarize the striking feature of our example: Starting from a population of small beetles, evolution by natural selection is causing the fitness of the organisms to decrease over time. This might seem troubling initially, since we think of natural selection as being fitness-increasing. But in fact, it's not hard to reconcile what's happening with this general principle of natural selection. Natural selection increases the fitness of individual organisms in a fixed environment — if the environment changes to become more hostile to the organisms, then clearly this could cause their fitness to go down. This is what is happening to the population of beetles. Each beetle's environment includes all the other beetles, since these other beetles determine its success in food competitions; therefore the increasing fraction of large beetles can be viewed, in a sense, as a shift to an environment that is more hostile for everyone.

Empirical Evidence for Evolutionary Arms Races. Biologists have offered recent evidence for the presence of evolutionary games in nature with the Prisoner's-Dilemma structure we've just seen. It is very difficult to truly determine payoffs in any real-world setting, and so all of these studies are the subject of ongoing investigation and debate. For our purposes in this discussion, they are perhaps most usefully phrased as deliberately streamlined examples, illustrating how game-theoretic reasoning can help provide qualitative insight into different forms of biological interaction.

It has been argued that the heights of trees can obey Prisoner's-Dilemma payoffs [156, 226]. If two neighboring trees both grow short, then they share the sunlight equally. They also share the sunlight equally if they both grow tall, but in this case their payoffs are each lower because they have to invest a lot of resources in achieving the additional height. The trouble is that if one tree is short while its neighbor is tall, then the tall tree gets most of the sunlight. As a result, we can easily end up with payoffs just like the Body-Size Game among beetles, with the trees' evolutionary strategies *Short* and *Tall* serving as analogues to the beetles' strategies *Small* and *Large*. Of course, the real situation is more complex than this, since genetic variation among trees can lead to a wide range of different heights and hence a range of different strategies (rather than just two strategies labeled *Short* and *Tall*). Within this continuum, Prisoner's-Dilemma payoffs can only apply to a certain range of tree heights: there is some height beyond which further height-increasing mutations no longer provide the same payoff structure, because the additional sunlight is more than offset by the fitness downside of sustaining an enormous height.

Similar kinds of competition take place in the root systems of plants [181]. Suppose you grow two soybean plants at opposite ends of a large pot of soil; then their root systems will each fill out the available soil and intermingle with each other as they try to claim as many resources as they can. In doing so, they divide the resources in the soil equally. Now, suppose that instead you partition the same quantity of soil using a wall down the middle, so that the two plants are on opposite sides of the wall. Then each still gets half the resources present in the soil, but each invests less of its energy in producing roots and consequently has greater reproductive success through seed production.

This observation has implications for the following simplified evolutionary game involving root systems. Imagine that instead of a wall, we had two kinds of root-development strategies available to soybean plants: *Conserve*, where a plant's roots only grow into its own share of the soil, and *Explore*, where the roots grow everywhere they can reach. Then we again have the scenario and payoffs from the Body-Size Game, with the same conclusion: all plants are better off in a population where everyone plays *Conserve*, but only *Explore* is evolutionarily stable.

As a third example, there was recent excitement over the discovery that virus populations can also play an evolutionary version of the Prisoner's Dilemma [326, 392]. Turner and Chao studied a virus called Phage Φ6, which infects bacteria and manufactures products needed for its own replication. A mutational variant of this virus called Phage ΦH2 is also able to replicate in bacterial hosts, though less effectively on its own. However, ΦH2 is able to take advantage of chemical products produced by Φ6, which gives ΦH2 a fitness advantage when it is in the presence of Φ6. This turns out to yield the structure of the Prisoner's Dilemma: viruses have the two evolutionary strategies Φ6 and ΦH2; viruses in a pure Φ6 population all do better than viruses in a pure ΦH2 population; and regardless of what the other viruses are doing, you (as a virus) are better off playing ΦH2. Thus only ΦH2 is evolutionarily stable.

The virus system under study was so simple that Turner and Chao were able to infer an actual payoff matrix based on measuring the relative rates at which the two viral variants were able to replicate under different conditions. Using an estimation procedure derived from these measurements, they obtained the payoffs in Figure 7.2. The payoffs are re-scaled so that the upper-left box has the value $1.00, 1.00$.[1]

		Virus 2	
		Φ6	ΦH2
Virus 1	Φ6	$1.00, 1.00$	$0.65, 1.99$
	ΦH2	$1.99, 0.65$	$0.83, 0.83$

Figure 7.2: The Virus Game

Whereas our earlier examples had an underlying story very much like the use of performance-enhancing drugs, this game among phages is actually reminiscent of a different story that also motivates the Prisoner's Dilemma payoff structure: the scenario behind the Exam-or-Presentation game with which we began Chapter 6. There, two college students would both be better off if they jointly prepared for a presentation, but the payoffs led them to each think selfishly and study for an exam instead. What the Virus Game here shows is that shirking a shared responsibility isn't just something that rational decision-makers do; evolutionary forces can induce viruses to play this strategy as well.

7.3 A General Description of Evolutionarily Stable Strategies

The connections between evolutionary games and games played by rational participants are suggestive enough that it makes sense to understand how the relationship works in general. We will focus here, as we have thus far, on two-player two-strategy games. We will also

---

[1] It should be noted that even in a system this simple, there are many other biological factors at work, and hence this payoff matrix is still just an approximation to the performance of Φ6 and ΦH2 populations under real experimental and natural conditions. Other factors appear to affect these populations, including the density of the population and the potential presence of additional mutant forms of the virus [393].