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].

restrict our attention to symmetric games, as in the previous sections of this chapter, where the roles of the two players are interchangeable.

The payoff matrix for a completely general two-player, two-strategy game that is symmetric can be written as in Figure 7.3.

		Organism 2	
		S	T
Organism 1	S	a, a	b, c
	T	c, b	d, d

Figure 7.3: General Symmetric Game

Let's check how to write the condition that S is evolutionarily stable in terms of the four variables a, b, c, and d. As before, we start by supposing that for some very small positive number x, a $1-x$ fraction of the population uses S and an x fraction of the population uses T.

- What is the expected payoff to an organism playing S in a random interaction in this population? With probability $1-x$, it meets another player of S, receiving a payoff of a, while with probability x, it meets a player of T, receiving a payoff of b. Therefore its expected payoff is
$$a(1-x) + bx.$$
- What is the expected payoff to an organism playing T in a random interaction in this population? With probability $1-x$, it meets a player of S, receiving a payoff of c, while with probability x, it meets another player of T, receiving a payoff of d. Therefore its expected payoff is
$$c(1-x) + dx.$$

Therefore, S is evolutionarily stable if for all sufficiently small values of $x > 0$, the inequality
$$a(1-x) + bx > c(1-x) + dx$$
holds. As x goes to 0, the left-hand side becomes a and the right-hand side becomes c. Hence, if $a > c$, then the left-hand side is larger once x is sufficiently small, while if $a < c$ then the left-hand side is smaller once x is sufficiently small. Finally, if $a = c$, then the left-hand side is larger precisely when $b > d$. Therefore, we have a simple way to express the condition that S is evolutionarily stable:

> *In a two-player, two-strategy, symmetric game, S is evolutionarily stable precisely when either (i) $a > c$, or (ii) $a = c$ and $b > d$.*

It is easy to see the intuition behind our calculations that translates into this condition, as follows.

- First, in order for S to be evolutionarily stable, the payoff to using strategy S against S must be at least as large as the payoff to using strategy T against S. Otherwise, an invader who uses T would have a higher fitness than the rest of population, and the fraction of the population who are invaders would have a good probability of growing over time.

- Second, if S and T are equally good responses to S, then in order for S to be evolutionarily stable, players of S must do better in their interactions with T than players of T do with each other. Otherwise, players of T would do as well as against the S part of the population as players of S, and at least as well against the T part of the population, so their overall fitness would be at least as good as the fitness of players of S.

7.4 Relationship Between Evolutionary and Nash Equilibria

Using our general way of characterizing evolutionarily stable strategies, we can now understand how they relate to Nash equilibria. If we go back to the General Symmetric Game from the previous section, we can write down the condition for (S, S) (i.e. the choice of S by both players) to be a Nash equilibrium. (S, S) is a Nash equilibrium when S is a best response to the choice of S by the other player: this translates into the simple condition

$$a \geq c.$$

If we compare this to the condition for S to be evolutionarily stable,

$$\text{(i) } a > c, \text{ or (ii) } a = c \text{ and } b > d,$$

we immediately get the conclusion that

If strategy S is evolutionarily stable, then (S, S) is a Nash equilibrium.

We can also see that the other direction does not hold: it is possible to have a game where (S, S) is a Nash equilibrium, but S is not evolutionarily stable. The difference in the two conditions above tells us how to construct such a game: we should have $a = c$ and $b < d$.

To get a sense for where such a game might come from, let's recall the Stag Hunt Game from Chapter 6. Here, each player can hunt stag or hunt hare; hunting hare successfully just requires your own effort, while hunting the more valuable stag requires that you both do so. This produces payoffs as shown in Figure 7.4.

		Hunter 2	
		Hunt Stag	*Hunt Hare*
Hunter 1	*Hunt Stag*	$4, 4$	$0, 3$
	Hunt Hare	$3, 0$	$3, 3$

Figure 7.4: Stag Hunt

In this game, as written, *Hunt Stag* and *Hunt Hare* are both evolutionarily stable, as we can check from the conditions on a, b, c, and d. (To check the condition for *Hunt Hare*, we simply need to interchange the rows and columns of the payoff matrix, to put *Hunt Hare* in the first row and first column.)

However, suppose we make up a modification of the Stag Hunt Game, by shifting the payoffs as follows. In this new version, when the players mis-coordinate, so that one hunts stag while the other hunts hare, then the hare-hunter gets an extra benefit due to the lack of competition for hare. In this way, we get a payoff matrix as in Figure 7.5.

		Hunter 2	
		Hunt Stag	*Hunt Hare*
Hunter 1	*Hunt Stag*	$4, 4$	$0, 4$
	Hunt Hare	$4, 0$	$3, 3$

Figure 7.5: Stag Hunt: A version with added benefit from hunting hare alone

In this case, the choice of strategies *(Hunt Stag, Hunt Stag)* is still a Nash equilibrium: if each player expects the other to hunt stag, then hunting stag is a best response. But *Hunt Stag* is not an evolutionarily stable strategy for this version of the game, because (in the notation from our General Symmetric Game) we have $a = c$ and $b < d$. Informally, the problem is that a hare-hunter and a stag-hunter do equally well when each is paired with a stag-hunter; but hare-hunters do better than stag-hunters when each is paired with a hare-hunter.

There is also a relationship between evolutionarily stable strategies and the concept of a *strict Nash equilibrium*. We say that a choice of strategies is a strict Nash equilibrium if each player is using the unique best response to what the other player is doing. We can check that for symmetric two-player, two-strategy games, the condition for (S, S) to be a strict Nash equilibrium is that $a > c$. So we see that in fact these different notions of equilibrium naturally *refine* each other. The concept of an evolutionarily stable strategy can be viewed as a refinement of the concept of a Nash equilibrium: the set of evolutionarily stable strategies S is a subset of the set of strategies S for which (S, S) is a Nash equilibrium. Similarly, the concept of a strict Nash equilibrium (when the players use the same strategy) is a refinement of evolutionary stability: if (S, S) is a strict Nash equilibrium, then S is evolutionarily stable.

It is intriguing that, despite the extremely close similarities between the conclusions of evolutionary stability and Nash equilibrium, they are built on very different underlying stories. In a Nash equilibrium, we consider players choosing mutual best responses to each other's strategy. This equilibrium concept places great demands on the ability of the players to choose optimally and to coordinate on strategies that are best responses to each other. Evolutionary stability, on the other hand, supposes no intelligence or coordination on the part of the players. Instead, strategies are viewed as being hard-wired into the players, perhaps because their behavior is encoded in their genes. According to this concept, strategies which are more successful in producing offspring are selected for.

Although this evolutionary approach to analyzing games originated in biology, it can be applied in many other contexts. For example, suppose a large group of people are being matched repeatedly over time to play the General Symmetric Game from Figure 7.3. Now the payoffs should be interpreted as reflecting the welfare of the players, and not their number of offspring. If any player can look back at how others have played and can observe their payoffs, then imitation of the strategies that have been most successful may induce an evolutionary dynamic. Alternatively, if a player can observe his own past successes and failures then his learning may induce an evolutionary dynamic. In either case, strategies that have done relatively well in the past will tend to be used by more people in the future. This can lead to the same behavior that underlies the concept of evolutionarily stable strategies, and hence can promote the play of such strategies.

7.5 Evolutionarily Stable Mixed Strategies

As a further step in developing an evolutionary theory of games, we now consider how to handle cases in which no strategy is evolutionarily stable.

In fact, it is not hard to see how this can happen, even in two-player games that have pure-strategy Nash equilibria.[2] Perhaps the most natural example is the Hawk-Dove Game from Chapter 6, and we use this to introduce the basic ideas of this section. Recall that in the Hawk-Dove Game, two animals compete for a piece of food; an animal that plays the strategy *Hawk* (H) behaves aggressively, while an animal that plays the strategy *Dove* (D) behaves passively. If one animal is aggressive while the other is passive, then the aggressive animal benefits by getting most of the food; but if both animals are aggressive, then they risk destroying the food and injuring each other. This leads to a payoff matrix as shown in Figure 7.6.

In Chapter 6, we considered this game in contexts where the two players were making choices about how to behave. Now let's consider the same game in a setting where each

[2] Recall that a player is using a *pure strategy* if she always plays a particular one of the strategies in the game, as opposed to a *mixed strategy* in which she chooses at random from among several possible strategies.

		Animal 2	
		D	H
Animal 1	D	$3, 3$	$1, 5$
	H	$5, 1$	$0, 0$

Figure 7.6: Hawk-Dove Game

animal is genetically hard-wired to play a particular strategy. How does it look from this perspective, when we consider evolutionary stability?

Neither D nor H is a best response to itself, and so using the general principles from the last two sections, we see that neither is evolutionarily stable. Intuitively, a hawk will do very well in a population consisting of doves — but in a population of all hawks, a dove will actually do better by staying out of the way while the hawks fight with each other.

As a two-player game in which players are actually choosing strategies, the Hawk-Dove Game has two pure Nash equilibria: (D, H) and (H, D). But this doesn't directly help us identify an evolutionarily stable strategy, since thus far our definition of evolutionary stability has been restricted to populations in which (almost) all members play the same pure strategy. To reason about what will happen in the Hawk-Dove Game under evolutionary forces, we need to generalize the notion of evolutionary stability by allowing some notion of "mixing" between strategies.

Defining Mixed Strategies in Evolutionary Game Theory. There are at least two natural ways to introduce the idea of mixing into the evolutionary framework. First, it could be that each individual is hard-wired to play a pure strategy, but some portion of the population plays one strategy while the rest of the population plays another. If the fitness of individuals in each part of the population is the same, and if invaders eventually die off, then this could be considered to exhibit a kind of evolutionary stability. Second, it could be that each individual is hard-wired to play a particular mixed strategy — that is, they are genetically configured to choose randomly from among certain options with certain probabilities. If invaders using any other mixed strategy eventually die off, then this too could be considered a kind of evolutionary stability. We will see that for our purposes here, these two concepts are actually equivalent to each other, and we will focus initially on the second idea, in which individuals use mixed strategies. Essentially, we will find that in situations like the Hawk-Dove game, the individuals or the population as a whole must display a mixture of the two behaviors in order to have any chance of being stable against invasion by other forms of behavior.

The definition of an evolutionarily stable mixed strategy is in fact completely parallel to the definition of evolutionary stability we have seen thus far — it is simply that we now greatly enlarge the set of possible strategies, so that each strategy corresponds to a particular randomized choice over pure strategies.

Specifically, let's consider the General Symmetric Game from Figure 7.3. A mixed strategy here corresponds to a probability p between 0 and 1, indicating that the organism plays S with probability p and plays T with probability $1-p$. As in our discussion of mixed strategies from Chapter 6, this includes the possibility of playing the pure strategies S or T by simply setting $p = 1$ or $p = 0$. When Organism 1 uses the mixed strategy p and Organism 2 uses the mixed strategy q, the expected payoff to Organism 1 can be computed as follows. There is a probability pq of an (X, X) pairing, yielding a for the first player; there is a probability $p(1-q)$ of an (X, Y) pairing, yielding b for the first player; there is a probability $(1-p)q$ of a (Y, X) pairing, yielding c for the first player; and there is a probability $(1-p)(1-q)$ of a (Y, Y) pairing, yielding d for the first player. So the expected payoff for this first player is

$$V(p, q) = pqa + p(1-q)b + (1-p)qc + (1-p)(1-q)d.$$

As before, the *fitness* of an organism is its expected payoff in an interaction with a random member of the population. We can now give the precise definition of an evolutionarily stable mixed strategy.

> *In the General Symmetric Game, p is an evolutionarily stable mixed strategy if there is a (small) positive number y such that when any other mixed strategy q invades p at any level $x < y$, the fitness of an organism playing p is strictly greater than the fitness of an organism playing q.*

This is just like our previous definition of evolutionarily stable (pure) strategies, except that we allow the strategy to be mixed, *and* we allow the invaders to use a mixed strategy. An evolutionarily stable mixed strategy with $p = 1$ or $p = 0$ is evolutionarily stable under our original definition for pure strategies as well. However, note the subtle point that even if S were an evolutionarily stable strategy under our previous definition, it is not necessarily an evolutionarily stable mixed strategy under this new definition with $p = 1$. The problem is that it is possible to construct games in which no pure strategy can successfully invade a population playing S, but a mixed strategy can. As a result, it will be important to be clear in any discussion of evolutionary stability on what kinds of behavior an invader can employ.

Directly from the definition, we can write the condition for p to be an evolutionarily stable mixed strategy as follows: for some y and any $x < y$, the following inequality holds for all mixed strategies $q \neq p$:

$$(1-x)V(p, p) + xV(p, q) > (1-x)V(q, p) + xV(q, q). \tag{7.1}$$

This inequality also makes it clear that there is a relationship between mixed Nash equilibria and evolutionarily stable mixed strategies, and this relationship parallels the one we saw earlier for pure strategies. In particular, if p is an evolutionarily stable mixed strategy, then we must have $V(p,p) \geq V(q,p)$, and so p is a best response to p. As a result, the pair of strategies (p,p) is a mixed Nash equilibrium. However, because of the strict inequality in Equation (7.1), it is possible for (p,p) to be a mixed Nash equilibrium without p being evolutionarily stable. So again, evolutionary stability serves as a refinement of the concept of mixed Nash equilibrium.

Evolutionarily Stable Mixed Strategies in the Hawk-Dove Game. Now let's see how to apply these ideas to the Hawk-Dove Game. First, since any evolutionarily stable mixed strategy must correspond to a mixed Nash equilibrium of the game, this gives us a way to search for possible evolutionarily stable strategies: we first work out the mixed Nash equilibria for the Hawk-Dove, and then we check if they are evolutionarily stable.

As we saw in Chapter 6, in order for (p,p) to be a mixed Nash equilibrium, it must make the two players indifferent between their two pure strategies. When the other player is using the strategy p, the expected payoff from playing D is $3p+(1-p) = 1+2p$, while the expected payoff from playing H is $5p$. Setting these two quantities equal (to capture the indifference between the two strategies), we get $p = 1/3$. So $(1/3, 1/3)$ is a mixed Nash equilibrium. In this case, both pure strategies, as well as any mixture between them, produce an expected payoff of $5/3$ when played against the strategy $p = 1/3$.

Now, to see whether $p = 1/3$ is evolutionarily stable, we must check Inequality (7.1) when some other mixed strategy q invades at a small level x. Here is a first observation that makes evaluating this inequality a bit easier. Since (p,p) is a mixed Nash equilibrium that uses both pure strategies, we have just argued that all mixed strategies q have the same payoff when played against p. As a result, we have $V(p,p) = V(q,p)$ for all q. Subtracting these terms from the left and right of Inequality (7.1), and then dividing by x, we get the following inequality to check:

$$V(p,q) > V(q,q). \tag{7.2}$$

The point is that since (p,p) is a mixed equilibrium, the strategy p can't be a strict best response to itself — all other mixed strategies are just as good against it. Therefore, in order for p to be evolutionarily stable, it must be a strictly better response to every other mixed strategy q than q is to itself. That is what will cause it to have higher fitness when q invades.

In fact, it is true that $V(p,q) > V(q,q)$ for all mixed strategies $q \neq p$, and we can check this as follows. Using the fact that $p = 1/3$, we have

$$V(p,q) = 1/3 \cdot q \cdot 3 + 1/3(1-q) \cdot 1 + 2/3 \cdot q \cdot 5 = 4q + 1/3$$

while

$$V(q,q) = q^2 \cdot 3 + q(1-q) \cdot 1 + (1-q) \cdot q \cdot 5 = 6q - 3q^2.$$

Now we have

$$V(p,q) - V(q,q) = 3q^2 - 2q + 1/3 = \frac{1}{3}(9q^2 - 6q + 1) = \frac{1}{3}(3q-1)^2.$$

This last way of writing $V(p,q)-V(q,q)$ shows that it is a perfect square, and so it is positive whenever $q \neq 1/3$. This is just what we want for showing that $V(p,q) > V(q,q)$ whenever $q \neq p$, and so it follows that p is indeed an evolutionarily stable mixed strategy.

Interpretations of Evolutionarily Stable Mixed Strategies. The kind of mixed equilibrium that we see here in the Hawk-Dove Game is typical of biological situations in which organisms must break the symmetry between two distinct behaviors, when consistently adopting just one of these behaviors is evolutionarily unsustainable.

We can interpret the result of this example in two possible ways. First, all participants in the population may actually be mixing over the two possible pure strategies with the given probability. In this case, all members of the population are genetically the same, but whenever two of them are matched up to play, any combination of D and H could potentially be played. We know the empirical frequency with which any pair of strategies will be played, but not what any two animals will actually do. A second interpretation is that the mixture is taking place at the population level: it could be that 1/3 of the animals are hard-wired to always play D, and 2/3 are hard-wired to always play H. In this case, no individual is actually mixing, but as long as it is not possible to tell in advance which animal will play D and which will play H, the interaction of two randomly selected animals results in the same frequency of outcomes that we see when each animal is actually mixing. Notice also that in this case, the fitnesses of both kinds of animals are the same, since both D and H are best responses to the mixed strategy $p = 1/3$. Thus, these two different interpretation of the evolutionarily stable mixed strategy lead to the same calculations, and the same observed behavior in the population.

There are a number of other settings in which this type of mixing between pure strategies has been discussed in biology. A common scenario is that there is an undesirable, fitness-lowering role in a population of organisms — but if some organisms don't choose to play this role, then everyone suffers considerably. For example, let's think back to the Virus Game in Figure 7.2 and suppose (purely hypothetically, for the sake of this example) that the payoff when both viruses use the strategy ΦH2 were (0.50, 0.50), as shown in Figure 7.7.

		Virus 2: Φ6	Virus 2: ΦH2
Virus 1	Φ6	1.00, 1.00	0.65, 1.99
	ΦH2	1.99, 0.65	0.50, 0.50

Figure 7.7: The Virus Game: Hypothetical payoffs with stronger fitness penalties to ΦH2.

In this event, rather than having a Prisoner's Dilemma type of payoff structure, we'd have a Hawk-Dove payoff structure: having both viruses play ΦH2 is sufficiently bad that one of them needs to play the role of Φ6. The two pure equilibria of the resulting two-player game

— viewed as a game among rational players, rather than a biological interaction — would be (Φ6,ΦH2) and (ΦH2,Φ6). In a virus population we'd expect to find an evolutionarily stable mixed strategy in which both kinds of virus behavior were observed.

This example, like the examples from our earlier discussion of the Hawk-Dove Game in Section 6.6, suggests the delicate boundary that exists between Prisoner's Dilemma and Hawk-Dove. In both cases, a player can choose to be "helpful" to the other player or "selfish". In Prisoner's Dilemma, however, the payoff penalties from selfishness are mild enough that selfishness by both players is the unique equilibrium — while in Hawk-Dove, selfishness is sufficiently harmful that at least one player should try to avoid it.

There has been research into how this boundary between the two games manifests itself in other biological settings as well. One example is the implicit game played by female lions in defending their territory [218, 327]. When two female lions encounter an attacker on the edge of their territory, each can choose to play the strategy *Confront*, in which she confronts the attacker, or *Lag*, in which she lags behind and tries to let the other lion confront the attacker first. If you're one of the lions, and your fellow defender chooses the strategy *Confront*, then you get a higher payoff by choosing *Lag*, since you're less likely to get injured. What's harder to determine in empirical studies is what a lion's best response should be to a play of *Lag* by her partner. Choosing *Confront* risks injury, but joining your partner in *Lag* risks a successful assault on the territory by the attacker. Understanding which is the best response is important for understanding whether this game is more like Prisoner's Dilemma or Hawk-Dove, and what the evolutionary consequences might be for the observed behavior within a lion population.

In this, as in many examples from evolutionary game theory, it is beyond the power of current empirical studies to work out detailed fitness values for particular strategies. However, even in situations where exact payoffs are not known, the evolutionary framework can provide an illuminating perspective on the interactions between different forms of behavior in an underlying population, and how these interactions shape the composition of the population.

7.6 Exercises

1. In the payoff matrix below the rows correspond to player A's strategies and the columns correspond to player B's strategies. The first entry in each box is player A's payoff and the second entry is player B's payoff.

		Player B	
		x	y
Player A	x	$2,2$	$0,0$
	y	$0,0$	$1,1$

(a) Find all pure strategy Nash equilibria.

(b) Find all Evolutionarily Stable strategies. Give a brief explanation for your answer.

(c) Briefly explain how the sets of predicted outcomes relate to each other.

2. In the payoff matrix below the rows correspond to player A's strategies and the columns correspond to player B's strategies. The first entry in each box is player A's payoff and the second entry is player B's payoff.

		Player B	
		x	y
Player A	x	$4, 4$	$3, 5$
	y	$5, 3$	$5, 5$

(a) Find all pure strategy Nash equilibria.

(b) Find all Evolutionarily Stable strategies. Give a brief explanation for your answer.

(c) Briefly explain how the answers in parts (2a) and (2b) relate to each other.

3. In this problem we will consider the relationship between Nash equilibria and evolutionarily stable strategies for games with a strictly dominant strategy. First, let's define what we mean by *strictly dominant*. In a two-player game, strategy, X is said to be a strictly dominant strategy for a player i if, no matter what strategy the other player j uses, player i's payoff from using strategy X is strictly greater than his payoff from any other strategy. Consider the following game in which a, b, c, and d are non-negative numbers.

		Player B	
		X	Y
Player A	X	a, a	b, c
	Y	c, b	d, d

Suppose that strategy X is a strictly dominant strategy for each player, i.e. $a > c$ and $b > d$.

(a) Find all of the pure strategy Nash equilibria of this game.

(b) Find all of the evolutionarily stable strategies of this game.

(c) How would your answers to parts (a) and (b) change if we change the assumption on payoffs to: $a > c$ and $b = d$?

		Player B X	Player B Y
Player A	X	$1, 1$	$2, x$
	Y	$x, 2$	$3, 3$

4. Consider following the two-player, symmetric game where x can be 0, 1, or 2.

   (a) For each of the possible values of x find all (pure strategy) Nash equilibria and all evolutionarily stable strategies.

   (b) Your answers to part (a) should suggest that the difference between the predictions of evolutionary stability and Nash equilibrium arises when a Nash equilibrium uses a *weakly dominated strategy*. We say that a strategy s_i^* is weakly dominated if player i has another strategy s_i' with the property that:

   (a) No matter what the other player does, player i's payoff from s_i' is at least as large as the payoff from s_i^*, and

   (b) There is some strategy for the other player so that player i's payoff from s_i' is strictly greater than the payoff from s_i^*.

   Now, consider the following claim that makes a connection between evolutionarily stable strategies and weakly dominated strategies.

   > *Claim:* Suppose that in the game below, (X, X) is a Nash equilibrium and that strategy X is weakly dominated. Then X is not an evolutionarily stable strategy.

		Player B X	Player B Y
Player A	X	a, a	b, c
	Y	c, b	d, d

   Explain why this claim is true. (You do not need to write a formal proof; a careful explanation is fine.)