



Game Theory and Evolution

Steve Phelps, *University of Essex*

Michael Wooldridge, *University of Oxford*

Game theory and evolution might seem unlikely bedfellows, and yet over the past three decades, game theoretic ideas have shed much light on a range of problems in biology, while ideas from evolution have begun to be applied with great success in economics. Here, we describe the key concepts from the area now known as *evolutionary game theory*, and introduce some important applications of these concepts.

Hawks and Doves

Evolutionary game theory is usually introduced by way of a simple scenario, known as the *Hawk-Dove Game*.¹ The game is as follows. Imagine we have a very large population of individuals, which come in two variants: Hawks and Doves (we aren't thinking of the species of birds that go by these names, by the way—these terms refer to innate predispositions towards aggressive or passive behavior). These two variants play the role of *strategies* in conventional game theory. Note that, in a biological context, we don't imagine an individual consciously deciding whether to be a Hawk or a Dove: each variant reflects a genetically hard-wired behavior.

Individuals will reproduce over time, but in the Hawk-Dove game, reproduction is *asexual*—an individual doesn't need a partner to reproduce, and if an individual reproduces, then it begets offspring of the same type. Thus, when a Hawk reproduces it creates a Hawk, and when a Dove reproduces it creates a Dove (in a very large population, we can consider the effect of random mutation to be negligible). The key attribute of an individual that determines how likely they are to reproduce is a numeric value that we'll call their *fitness*. In evolutionary terms, the fitness of an individual is a measure of how likely that individual is to be able to reproduce and pass on their

genes: an individual with a high fitness value will be likely to be able to reproduce and pass on their genes, while an individual with low fitness is unlikely to do so.

In the Hawk-Dove Game, individuals can *increase* their fitness by obtaining a particular *resource* from the environment. For example, if we think of the resource as being food or water in the animal kingdom, then this is a natural interpretation: animals that succeed in obtaining food and water are surely more likely to survive and reproduce than those that don't.

In our game, individuals are in competition with others to obtain resources. We think of the Hawks as being fierce and unafraid to fight for resources, while the Doves are timid and reluctant to fight. The rules defining what happens when Hawks and Doves compete are then as follows:

- When a Hawk competes with a Dove, the Hawk takes the whole of the resource. (The Dove is afraid to fight and backs down, leaving the whole of the resource to the Hawk.)
- When a Dove competes with a Dove, they share the resource equally.
- When a Hawk competes with a Hawk, then they fight, and have an equal chance of obtaining the resource or being injured.

Let V denote the value of the resource: this is the increase in fitness that an individual would gain by obtaining the resource. Let C denote the *cost of injury*—this is the amount by which fitness would *decrease* if an individual fought for the resource and lost. So, the changes in fitness resulting from an encounter in the game are as follows:

- When a Hawk meets a Hawk: they fight, and have an equal chance of increasing their fitness

Replicator Dynamics

The replicator dynamics of a scenario in evolutionary game theory describe how the proportion, or frequency, of different variants changes over time in response to their fitness. The subtlety here is that a variant's fitness also depends on its frequency (in the terminology of mathematical biology, we say that fitness is *frequency-dependent*). The model we present here is the standard model for the Hawk-Dove game, as developed by John Maynard Smith (1920–2004).¹

Recall that V is the value that would be gained by obtaining the resource, C is the cost that would be incurred by a Hawk if it fought another Hawk and lost, and P is the initial frequency of Hawks in the population. Thus, if $P = 0$, then there are no Hawks in the initial population, while if $P = 1$, then the population is initially composed entirely of Hawks. If we imagine that reproduction occurs at discrete time steps, then at each step we compute the average fitness of Hawks and Doves. Let $W(H)$ and $W(D)$ denote the average fitness of Hawks and Doves. Taking Hawks first, we have

$$W(H) = P(V - C)/2 + (1 - P)V.$$

The first term is the fitness that Hawks obtain by competing with Hawks; the second term is the increase in fitness resulting from encounters with Doves.

The corresponding equation for Doves is simpler, because we don't need a term describing the increase in fitness that a

Dove would obtain from encountering a Hawk: we know this will be 0, and all we need is to describe the increase in fitness obtained by Doves encountering Doves:

$$W(D) = (1 - P)(V/2).$$

The final element we need in our replicator dynamics describes how the frequencies of different variants change over time in response to these fitness values. The idea is that the frequency of a variant in a population will increase at a rate that's proportional to their fitness compared with the average fitness of the population as a whole. Thus, if Hawks increase their average fitness by 2, and Doves also increase their average fitness by 2, then the frequency of Hawks and Doves in the next time step will remain the same. However, if Hawks increase their fitness more than Doves, then there will be proportionally more Hawks. Formally, we let P' denote the frequency of Hawks present in the next time period, and let $W = PW(H) + (1 - P)W(D)$ denote the average fitness of an individual in the population. The new frequency of Hawks is then given by

$$P' = P \frac{W(H)}{W} = \frac{PW(H)}{PW(H) + (1 - P)W(D)}.$$

Reference

1. J.M. Smith, *Evolution and the Theory of Games*, Cambridge Univ. Press, 1982.

by V or decreasing their fitness by C ; on average, this will result in an increase of fitness by $(V - C)/2$.

- When a Dove meets a Dove, they share the resource equally, each obtaining an increase of fitness of $V/2$.
- When a Hawk meets a Dove, the Hawk takes the whole of the resource, resulting in an increase of fitness of V for the Hawk, while the Dove has no change in fitness.

To calculate the fitness of the Hawks and the Doves in our population, we imagine that the individuals in the population encounter each other several times before they reproduce. We model this by “mixing” the population; randomly chosen pairs of individuals are picked from the population, and we repeat this several times, choosing a new pair on each iteration. On each encounter, the two individuals compete with each other, obtaining an improvement in fitness

for each individual as described by the aforementioned rules. We then calculate the average fitness improvement of all of the Hawks, and correspondingly for the Doves.

When the population reproduces, the fraction of each variant increases at a rate proportional to its current fitness relative to the rest of the population. The exact rules defining the proportions of Hawks and Doves are known as the *replicator dynamics* of the game: see the related sidebar for a detailed explanation of the replicator dynamics for the Hawk-Dove scenario. The basic idea is that if, for example, Hawks have improved their average fitness more than Doves, then the proportion of Hawks in the next time step will be increased accordingly. Similarly, if Doves increase their average fitness more than Hawks, then the proportion of Doves will increase. This process proceeds iteratively, and can continue indefinitely unless the population reaches

an equilibrium in which the proportions are unchanging.

Figure 1 shows how the population of Hawks and Doves changes over time for a typical instance of the game. The proportion of Hawks in the population starts as only 1/10 of a percent of the population, but grows over successive generations, and eventually Hawks completely eliminate the Dove population. To see why this is, suppose there was just one Hawk in the population, and 999 Doves. Initially, that Hawk would be guaranteed to compete with a Dove, taking the whole of the value V , while every other Dove would compete with a Dove, sharing the value V equally. The fitness of Hawks increases dramatically compared to Doves, and the proportion of Hawks in the next generation increases accordingly, while the proportion of Doves decreases. This growth continues, with Hawks essentially preying on the Dove population. But observe

Evolutionarily Stable Strategies

Informally, a strategy S is said to be evolutionarily stable if a population consisting almost entirely of S can't be invaded over time by another strategy. It turns out that we can characterize exactly the conditions required for S to be evolutionarily stable. There are two requirements.

The first and most basic property is that for S to be an evolutionarily stable strategy, there can be no other strategy that does better when played against S than S does itself. If S didn't have this property, then some other strategy—call it T —would be able to invade a population of S individuals. Imagine on the first generation of population of 999 S individuals and one T . The T is guaranteed to meet with an S and do better, while all the other S individuals meet with other S individuals. The one T individual will thus have an improvement in fitness greater than the S individuals, and will begin to invade the S population, just as Hawks invade a population of Doves in the Hawk-Dove game (see Figure 1 in the main text).

This first condition is, however, not enough to guarantee that S will survive and prosper over successive generations. This is because there might be other strategies that do no better when played against S than S does, but do better when playing against themselves than S does when played against them. They thus benefit compared to S when they play against themselves. So, the second condition required for an evolutionarily stable strategy is that S does better when played against every alternative strategy T than T does when played against itself.

In the terminology of conventional game theory, the first condition amounts to saying that S playing against S is a conventional Nash equilibrium (in fact, a symmetric Nash equilibrium). The second condition ensures that it's a refinement of the ordinary notion of a Nash equilibrium. It ensures that other strategies can't prosper, compared to S , when they meet themselves.

Table 1. Payoff matrix.

		Alex	
		Defect	Cooperate
Bob	Defect	2, 2	4, 1
	Cooperate	1, 4	3, 3

Axelrod's Tournament

The Hawk-Dove game nicely illustrates the basic ideas of evolutionary game theory, but it's a little abstract, so let's consider a more concrete (and much more famous) scenario. In the late 1970s, a political scientist called Robert Axelrod was pondering the question of how cooperation could emerge in a society of purely self-interested individuals.² To shed some light on this question, he organized a tournament in which interested parties from around the world were invited to submit computer programs to play a game called the *Iterated Prisoner's Dilemma*. As the name suggests, the tournament is based on the well-known game called the Prisoner's Dilemma. To recap (and with apologies to regular readers of this column, who will have seen this explanation before), the Prisoner's Dilemma game has two players: Alex and Bob. Each player must choose between two actions: *cooperation* or *defection*. Depending on the combination of choices made, the players receive *payoffs*, as defined in the following *payoff matrix* (see Table 1).

We read this payoff matrix as follows. Bob is the *row player*, so called because his possible choices correspond to the matrix rows, while

the shape of the graph in Figure 1: at around the eleventh generation, the rate of increase in the proportion of Hawks starts to decrease sharply. This is because at this point Hawks increasingly encounter other Hawks in the population, rather than Doves, and have to fight for resources.

In the terminology of evolutionary game theory, we say that in the scenario illustrated in Figure 1, the Hawk population *invades* the Dove population. Now, for this scenario, Hawks are able to invade the Dove population even starting from a low initial proportion of the population. In fact, as long as the value of the resource is greater than the cost of conflict (that is, as long as $V > C$), then having just one Hawk in the initial population is sufficient to guarantee that eventually the Doves will be wiped out. Of course, in this scenario, Doves won't be able to invade a Hawk population.

This brings us to the main analytical concept in evolutionary game

theory, which plays a role analogous to that played by Nash equilibrium in classical game theory. Let's say a strategy is *evolutionarily stable* if a population consisting of almost entirely that strategy can't be invaded by any other strategy. In the Hawk-Dove game, Hawk is an evolutionary stable strategy if $V > C$. It's not hard to see this: if $V > C$, then being a Hawk is always going to be more successful than being a Dove: in the worst case, if you meet another Hawk, you'll still get some positive improvement in fitness. A related sidebar presents some more discussion of evolutionary stable strategies.

So, when we analyze a scenario in terms of evolutionary game theory, instead of asking "What are the Nash equilibrium strategies?" we ask "What are the evolutionarily stable strategies?" In informal terms, this means asking *which strategies will prosper over time, in an evolutionary setting*, taking into account the other strategies present in the environment.

Alex is the *column player*, so called because his choices correspond to the matrix columns. Each cell in the matrix is a possible *outcome* of the game, corresponding to the combination of choices made by the players. The numbers in a matrix cell are the payoffs that the players receive in that outcome: Bob's is listed first, then the payoff for Alex. Thus, if Alex cooperates while Bob defects, we obtain the outcome in the matrix's top right cell: Bob gets a payoff of 4, while Alex gets a payoff of 1. Players prefer higher payoffs, so this outcome is the best possible outcome for Bob, and the worst possible outcome for Alex.

It's easy to see that, no matter what Bob does, Alex's best response would be to defect, while no matter what Alex does, Bob's best response would be to defect. Thus both players choose to defect, each receiving a payoff of 2 (in the matrix's top left cell), although in fact *both* players would have done better to cooperate, which would give each of them a payoff of 3 (bottom right cell).

In the Iterated Prisoner's Dilemma, the same two players will meet and play a certain number of rounds of the Prisoner's Dilemma; their payoffs on each round are summed to find their overall payoff. At each round, both players are able to see what their counterpart did on the previous round. It was this game that was played by the computer programs submitted to Axelrod's tournament. Before proceeding, let's see some of the 63 strategies that were submitted to the tournament:

- ALLD: Defect on every round, no matter what your counterpart does.
- ALLC: Cooperate on every round, no matter what your counterpart does.
- TIT-FOR-TAT: On the first round, cooperate; on subsequent rounds,

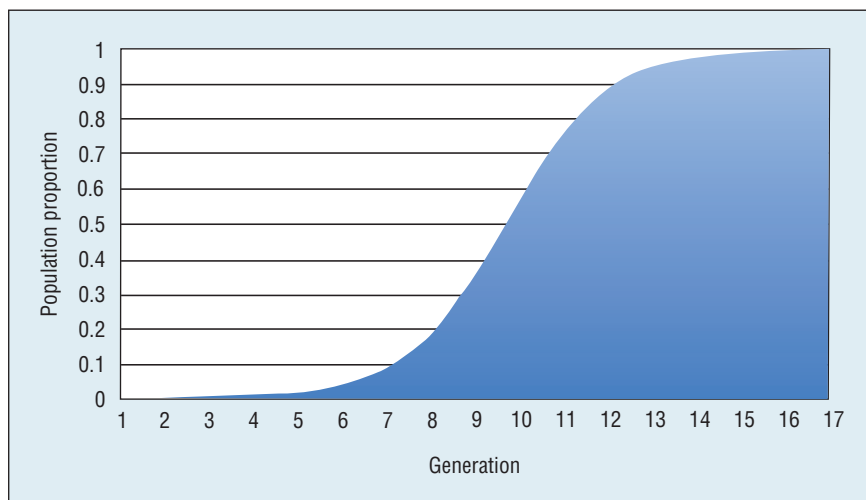


Figure 1. Population dynamics in a Hawk-Dove game. Hawks invade a population of Doves, starting from a low proportion of the population. The x-axis shows the generation, and the y-axis shows the proportion of the population that are Hawks at the corresponding time. As the proportion of Hawks begins to dominate, the rate of increase in the Hawk population decreases; this is because as Doves are eliminated from the population, Hawks find themselves increasingly competing with Hawks, reaping smaller improvements in fitness than when they were competing with Doves. The initial proportion of Hawks for the scenario in this graph was 0.001. The initial fitness for both Hawks and Doves was 2, with $V = 4$ and $C = 3$.

simply copy whatever your counterpart did on the preceding round.

- JOSS: As TIT-FOR-TAT, except that it defects 10 percent of the time instead of cooperating.
- FRIEDMAN: Start by cooperating, and repeat cooperation as long as your counterpart does; but if ever your counterpart defects, then defect forever after. (In the literature, this strategy is sometimes called GRIM or GRIM TRIGGER.)
- RANDOM: Choose between cooperation and defection randomly.

In Axelrod's tournament, every submitted strategy got to play the Iterated Prisoner's Dilemma for 200 rounds against every other submitted entry. The overall score for a strategy was then the sum of its scores over each individual contest. Thus, ALLD played the Iterated Prisoner's Dilemma for 200 rounds against ALLC, 200 rounds against TIT-FOR-TAT, and so on. It's important to remember in what follows that a strategy got to compete against *all* other strategies, and that its overall score took

into account how it performed in all of these individual contests.

The outcome of Axelrod's tournament was that the clear overall winner was TIT-FOR-TAT. Axelrod gives a compelling analysis of how and why TIT-FOR-TAT did so well in his eminently readable 1984 book describing the tournament (on p. 54):

What accounts for TIT-FOR-TAT's robust success is its combination of being nice, retaliatory, forgiving and clear. Its niceness [the fact that it starts by cooperating] prevents it from getting into unnecessary trouble. Its retaliation [the fact that it punishes defection by immediate defection] discourages the other side from persisting whenever defection is tried. Its forgiveness [the fact that it reciprocates cooperation even after previous defections] helps restore mutual co-operation. And its clarity makes it intelligible to the other player, thereby eliciting long-term co-operation.²

Axelrod's work deservedly received widespread acclaim. Many readers interpreted the results as "solving"

the Prisoner's Dilemma, and TIT-FOR-TAT was even promoted as a model for cooperation between nations. However, some commentators were critical: game theorists in particular pointed out that the fact that cooperation can be rationally sustained in some iterated versions of the Prisoner's Dilemma had been well-known for several decades previously, and highlighted a range of methodological issues with the tournament, which spurred much future research.

Important though this debate is, our present purpose is to focus on Axelrod's tournament in the context of evolutionary game theory. Notice that a key reason why TIT-FOR-TAT did well is that it was able to meet a *range* of other strategy types. In particular, when it met other "nice" strategies (those that start by cooperating, and reciprocate cooperation) it was able to benefit from this. Now, given our earlier discussion, a very natural evolutionary interpretation of Axelrod's tournament suggests itself. Instead of having just two types of individuals (Hawk and Dove) in the population, we have 63—the strategies that were submitted to the tournament. Let's suppose these are initially present in the population in equal proportions. We then compute fitness as being the score that the strategy achieves in the tournament played with those individuals currently in the environment. We can then use a similar model of replicator dynamics to those presented in the "Replicator Dynamics" sidebar to compute the populations of strategies in the next generation, and then we can look to see which strategies prosper and which don't, over successive generations. The dynamics of this process for Axelrod's tournament are fascinating (see pp. 50–51 of his book):

The first thing that happens is that the lowest-ranking eleven entries fall to half their initial size by the fifth generation while the middle-ranking entries tend to hold their own and the top-ranking entries gradually grow in size. By the fiftieth generation, the [strategies] that ranked in the bottom third of the tournament have virtually disappeared, while most of those in the middle third have started to shrink, and those in the top third are continuing to grow. The process simulates survival of the fittest. A [strategy] that is successful on average with the current distribution of [strategies] in the population will become an even larger proportion of the environment ... in the next generation. At first, a rule that is successful with all sorts of rules will proliferate, but later as the unsuccessful rules disappear, success requires success with other successful rules.²

As in the overall tournament, TIT-FOR-TAT was the best-performing strategy in the evolutionary model. Although it isn't technically an evolutionarily stable strategy, over time, TIT-FOR-TAT came to be the dominant strategy in the environment. For Axelrod, one of the crucial aspects of this result is the following: Starting with the initial distribution of strategies, and then simply applying some easily motivated reproduction rules, *we eventually end up with a population of cooperating individuals*. And note that no guiding hand, outside force, or even deep reasoning was required to get there: *we get to mutually sustained cooperation just by applying basic evolutionary principles*.

Lest we get overly excited about the success of TIT-FOR-TAT, it's important to note that its success greatly depends on its initial environment. For example, if we consider a population consisting of all possible strategies that can be represented by finite automata with just two states

(TIT-FOR-TAT is such a strategy), then the FRIEDMAN strategy does better.

Learning and Equilibrium Selection

Although evolutionary game theory was originally formulated to understand problems in biology, it has found applications in many other fields. This is because the underlying mathematics also applies in scenarios where individuals *learn* from one another.

For example, let's revisit the Hawk versus Dove scenario—however, this time we'll think of the individuals in the population as people who are allowed to switch between two different strategies (and correspondingly replace the fitness values with monetary payoffs). Now also imagine that these people are somewhat pushed for time, and rather than sitting down to work out the optimal behavior deductively, instead they use a heuristic rule of thumb which says "if someone else appears to be doing better than you, then copy their strategy." If we also assume that the probability of actually noticing that somebody else is doing better is proportional to the relative payoff between the two strategies, then the changes in the proportion of people using each strategy over time will be governed in exactly the same way as for the evolutionary case. We can therefore think of the replicator dynamics not just as a model of *genetic* evolution, but more generally as a model of *social learning*, or alternatively as *cultural* evolution.³

Remarkably, it also turns out that replicator dynamics are a good approximation of learning even if we drop the assumption that people imitate each other's strategies, and instead assume that they learn by trial and error by gradually reinforcing behaviors,

which yield good payoff to themselves. Thus, ideas from evolutionary game theory have been used in AI research to understand how systems comprised of multiple autonomous systems can learn to perform tasks with or against each other. For example, this conceptualization of evolutionary game theory provides an answer to the *equilibrium selection problem*. The equilibrium selection problem can be illustrated by the following scenario.

Alice and Bob like each other but are too shy to speak to one another. They know that both like to go to pubs in the evening, and there are two pubs in town, the Turf and the Anchor. Ann and Bob each have to independently decide which pub to go to—they can't go to both. If they go to the same pub, both will be happy; if they go to different pubs, both will be miserable.

Problems like this, where both players must *independently* make *mutually consistent* choices, are known as *coordination problems*. The difficulty is that there are two equilibria in this game: one in which both go to

the *Turf*, and one in which both go to the *Anchor*. But if there's nothing to distinguish these equilibria, the players don't communicate, and nobody dictates the outcome, how should they choose? Evolutionary game theory can provide an answer. The idea is that rather than deliberating and rationalizing between strategies, the players apply a trial-and-error process over time, being more likely to favor strategies that prove successful. By learning over time, they can ultimately coordinate.

Evolutionary game theory provides a fundamentally different view on strategy selection to that proposed by classical game theory. In evolutionary game theory, we don't think of individuals engaged in complex deductive reasoning while they try to decide what to do. Instead, strategies are selected over time by the inexorable process of evolution.

These ideas were originally applied to understanding biological evolution, but they apply more generally to scenarios where players reason

inductively by preferring strategies that have performed well over those that haven't. Thus, evolutionary game theory can be used to understand not only how intelligent systems evolve, but also how they learn. ■

References

1. J.M. Smith, *Evolution and the Theory of Games*, Cambridge Univ. Press, 1982.
2. R. Axelrod, *The Evolution of Cooperation*, Basic Books, 1984.
3. R. Boyd and P.J. Richerson, *Culture and the Evolutionary Process*, Univ. Chicago Press, 1988.

Steve Phelps is a lecturer in the Centre for Computational Finance and Economic Agents (CCFEA) at the University of Essex. Contact him at phelps.sg@gmail.com.

Michael Wooldridge is a professor of computer science at the University of Oxford. Contact him at mjw@cs.ox.ac.uk.

cn Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.

continued from page 75

Standards Edge: Future Generations, Sheridan Books, 2005.

3. L. Ding, V. Peristeras, and M. Hausenblas, "Linked Open Government Data," *IEEE Intelligent Systems*, vol. 27, no. 3, 2012, pp. 11–15.
4. Interoperability Solutions for European Public Administrations (ISA), *Towards Open Government Metadata*, European Commission, Sept. 2011; https://joinup.ec.europa.eu/sites/default/files/towards_open_government_metadata_0.pdf

5. ISA, *eGovernment Core Vocabularies*, European Commission, Feb. 2012; https://joinup.ec.europa.eu/sites/default/files/ISA_eGovernment-Core-Vocabularies_February2012.pdf.

Vassilios Peristeras is a program manager in the Interoperability Solutions for European Public Administration Unit at the European Commission in Brussels, and an elected assistant professor on Web Information Systems at the Interna-

tional Hellenic University in Thessaloniki, Greece. His research interests include e-government, semantic interoperability, enterprise architecture, open and linked data, and metadata management. Contact him at vassilios.peristeras@ec.europa.eu.

cn Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.