

Introduction to Game Theory

Part 2: Evolution

2 A Primer in Evolutionary Game Theory

Remark 2.1 (A recapitulation)

So far we have looked at some basic concepts of classical game theory. First we have discussed how we can translate strategic decision-making scenarios into formal games. The respective process required us to identify the five elements of a game: the players, their actions, the order of moves, the information the players have, and their payoffs. Then we looked at how we can define ‘solutions’ for certain types of games (*normal-form games*, in which players only have to make one decision, and where all aspects of the game are known to the players). We realized that even in these simple types of games, the respective solution concepts seemed to make rather strong assumptions on the players’ cognitive abilities. In some cases, players did not only need to be rational (i.e., they did not only need to fully understand the game and aim to maximize their payoffs). They also needed to know that their co-player is rational, they needed to know their co-player knows they know their co-player is rational, etc. (This infinite chain of knowledge is sometimes called *common knowledge of rationality*). Those strong requirements seem to undercut the value of the entire theory. The question we want to address in the following is: Can we get similar or alternative solutions with weaker assumptions on the individuals’ cognition? To address this question, we first make a little detour, and consider a seemingly unrelated problem.

Example 2.2 (Ritual fighting in deer)

The problem. In biology, there are often situations in which two males engage in a combat in order to gain access to a female, or to defend their territories, etc. If evolution favors the survival of those individuals with better fighting techniques, or with more deadly weapons, one would expect escalating evolution, towards ever more effective weapons. Instead, many combats in nature seem to be of ‘limited war type’. Individuals would use ineffective weapons or ritualized fighting techniques that rarely result in serious injuries. Why is that? One hypothesis would be that individuals consider what is good for their species. According to that hypothesis, individuals would realize that by constantly escalating a fight, they diminish their own numbers. However, that account seems to contradict Darwin’s theory: there, individuals aim to maximize their own fitness, not some abstract fitness of the species.

The model. Maynard Smith and Price (1973) proposed a possible resolution of this puzzle. They consider a species with two types of males, those that escalate a fight (‘hawks’) and those that adhere to ritual fighting (‘doves’). When a hawk encounters another, the two of them fight until one of them is seriously

injured. In that case, the winner gains access to the territory (which yields a benefit b), whereas the other suffers a considerable fitness cost $c > b$. When a hawk encounters a dove, the dove escapes and the hawk wins the conflict. Here, the hawk gains a fitness of b (the value of the territory), whereas the effect on the dove's fitness is zero. Finally when two doves interact, one of them wins the ritual fight, without doing any harm to the other (yielding a fitness gain of b to the winner and of zero to the other). To model this scenario, suppose the current fraction of hawks in the entire population is x_H . Moreover, assume for simplicity that animals have the same general health and fighting ability, such that if a fight occurs, both contestants are equally likely to win it. In that case, we can compute the expected fitness of the two types. For example, if f_H is the average fitness of a hawk and f_D the fitness of a dove, we obtain:

$$\begin{aligned} f_H &= \frac{b-c}{2} \cdot x_H + b \cdot (1-x_H), \\ f_D &= 0 \cdot x_H + \frac{b}{2} \cdot (1-x_H). \end{aligned} \tag{1}$$

An analysis. To get an intuition for the long-run dynamics of the resulting model, it is instructive to look at two limiting cases. First, suppose hawks are rare, such that $x_H \approx 0$. In that case, the two fitness functions simplify to

$$f_H = b \quad \text{and} \quad f_D = b/2.$$

Therefore, a rare hawk would have a higher fitness than the residents. This hawk would be expected to have more offspring and to invade the resident population. On the other hand, now assume that hawks are very common, such that $x_H \approx 1$. Now, the fitnesses are

$$f_H = (b-c)/2 \quad \text{and} \quad f_D = 0.$$

Because we assumed $c > b$ (the cost of a serious injury outweighs the value of the territory), it is now the doves who have a higher fitness. Hence, now doves should increase in frequency.

Overall, we conclude that neither a pure population of hawks nor a pure population of doves is stable. Over time, we would expect that there is a change in the proportion x in hawks until the two types have equal fitness. By solving $f_D = f_H$, we obtain an equilibrium at

$$x_H^* = b/c.$$

In particular, especially in species in which the cost of a fight would be considerable (where c is large, as in deer), we would expect that most individuals are doves, and engage in ritual fighting.

Connection to classical game theory. Instead of this evolutionary approach, we could have analyzed a very similar game among humans. Again, we could have assumed that people need to decide whether they want to act like a hawk or like a dove, and that the resulting payoff consequences are similar to before. In

that case we would have ended up with the payoff matrix

	Hawk	Dove
Hawk	$\frac{b-c}{2}, \frac{b-c}{2}$	$b, 0$
Dove	$0, b$	$\frac{b}{2}, \frac{b}{2}$

We could analyze this game the same way as we did in the first part. It turns out that the game has no dominated actions. However, there are three Nash equilibria. Two of the equilibria are asymmetric (the two players coordinate on choosing the opposite action: one player uses Hawk and the other player uses Dove). However, there is also a symmetric equilibrium. Here both players use a mixed strategy, where they play Hawk with probability $x_H^* = b/c$ and Dove with probability $x_D^* = 1 - b/c$.

Remark 2.3 (On the parallels between evolutionary and classical models)

- In the above example we observe the same overall outcome with both modeling approaches (by considering an evolving population, or by considering a normal-form game). In both cases, on average x_H^* of the observed actions are Hawk and x_D^* are Dove. In the classical viewpoint, it is the individuals who use mixed strategies; in the evolutionary viewpoint it is the mixed population.
- One aspect that is noteworthy: to make sense of the Nash equilibrium earlier today, we had to make rather stringent assumptions on people's cognitive abilities. In particular they needed to be aware they are part of a game, they needed to understand the game's rules, try to optimize their behavior, and they needed to assume the co-player would do the same. In contrast, when following the evolutionary approach, no cognitive abilities are required at all. Here, it is the evolutionary process that optimizes behavior, not each single individual.
- Instead of a biological evolutionary process (where fitter individuals reproduce more often), we could have also imagined a process where successful traits or behaviors spread through cultural evolution. For example, players with a high payoff may be imitated more often than others. It is worth to formalize this alternative interpretation in the following.

Remark 2.4 (Deriving a model of pairwise comparisons)

To derive a simple model of cultural evolution, consider a population of size N . Suppose the members of this population randomly meet each other to interact in a hawk-dove game with each other. Suppose the current number of hawks is n_H and the number of doves is n_D , such that $n_H + n_D = N$. Then we can again compute the players' expected payoffs (or 'fitnesses'), similar to Eq. (1),

$$\begin{aligned} f_D &= \frac{b-c}{2} \cdot \frac{n_H-1}{N-1} + b \cdot \frac{n_D}{N-1}, \\ f_H &= 0 \cdot \frac{n_H}{N-1} + \frac{b}{2} \cdot \frac{n_D-1}{N-1}. \end{aligned} \tag{2}$$

In addition to these interactions, we assume that occasionally, individuals consider updating their strategy (of either playing hawk or dove). For example, we might assume that every once in a while, a random population member compares its own payoff f to the payoff of a randomly chosen role model f' . If the




Figure 1: A realization of the pairwise comparison dynamics process for the hawk-dove game. The figure shows one representative run of the code in Table 1.

role model's payoff is higher, the focal player switches its strategy. This elementary updating process can then be iterated for many time steps. To explore the dynamics of this process, we look at simulations (for some Matlab code, see Table 1). According to these simulations, individuals eventually indeed end up in a state that reflects the game's Nash equilibrium, see Figure 1.

Group Exercise 2.5 (Possible exercises)

- Beginner's version: Try to recreate the above process with your own favorite programming language. Run it for different payoff matrices (e.g., hawk-dove games with different parameters, or the split-and-steal game, or a rock-scissors-paper game). What do you observe?
- More advanced version: In the above model, we have considered a *well-mixed population*. This means that every individual interacted with everyone else to obtain their payoffs; moreover, when individuals looked for a role-model they picked some random other individual, taken from the entire population. Instead one could also imagine a structured population, where individuals are placed on some network. They would only play the game with their immediate neighbors. Moreover, when updating strategies, they would only compare their own payoff to the payoff of a randomly chosen neighbor. How does the evolutionary dynamics look there (e.g. how do strategies spread across the network)? Again, you could apply your code to different games, or different networks!
- An alternative advanced option: Some might be familiar with reinforcement learning. What happens if there are two agents who update their strategies (hawk or dove) with reinforcement learning?

Remark 2.6 (A summary)

In this second part of this course, you should have learned:

- that there is an alternative interpretation of game theory. This interpretation does not require conscious rational players who reason about their own options and about their co-player's reasoning. Instead, this approach only requires that individuals tend to adopt better strategies over time (either due to biological

evolution, or due to learning).

- how to implement such a learning process with an individual-based simulation.
- that the results of such simulations often recover (and hence justify) classical Nash equilibrium predictions (however, there are exceptions, see e.g., Hofbauer and Sandholm, 2011)
- that the cognitive requirements under such a learning interpretation seem much weaker. Instead of fully understanding the game, the associated payoffs, and the possible considerations of the co-player, individuals merely need to see what strategies other people are using and what payoffs they have. If strategies are inherited (rather than learnt) no information about the underlying game is required at all!

References

- Hofbauer, J. and Sandholm, W. H. (2011). Survival of dominated strategies under evolutionary dynamics. *Theoretical economics*, 6:341–377.
- Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246:15–18.

```

function [time,proportionhawks]=pairwisecomparisondynamics();

%% Explanation:
% [time,proportionhawks]=pairwisecomparisondynamics();
% Simulates the dynamics of a population that interacts in a hawk-dove game.
% Population members update their strategies according to some pairwise comparison process.
% Output "time" contains at which time steps we have recorded the state of the population.
% Output "proportionhawks" records the proportion of hawks at that time.

%% Setting the parameters of the game
benefit=2; cost=3; % benefit of winning a fight, cost of an injury
populationsize=200; numberhawks=20; numberdoves=populationsize-numberhawks;
% population size, initial proportion of hawks and doves
totalruntime = 1000; % Number of times we allow individuals to update their strategies
time=zeros(1,totalruntime); proportionhawks=zeros(1,totalruntime); % Initializing the output

%% Doing the actual simulation
for t=1:totalruntime
    % Computing payoffs according to formula
    payoffhawk = (benefit-cost)/2*(numberhawks-1)/(populationsize-1) + ...
        benefit*numberdoves/(populationsize-1);
    payoffdove = benefit/2*(numberdoves-1)/(populationsize-1);

    % Randomly drawing numbers to determine the identity of the learner
    learner = randi(populationsize);
    rolemodel = randi(populationsize);

    if learner <= numberhawks & rolemodel > numberhawks % learner is hawk, role model a dove
        if payoffdove > payoffhawk
            numberhawks = numberhawks -1;
            numberdoves = numberdoves +1;
        end
        elseif learner > numberhawks & rolemodel <=numberhawks % learner is dove, role model a hawk
            if payoffhawk > payoffdove
                numberhawks = numberhawks +1;
                numberdoves = numberdoves -1;
            end
        end
    end

    % Updating the output vectors
    time(t)=t;
    proportionhawks(t) = numberhawks / populationsize;
end

%% Plotting the outcome
colorblue = [0,0.5,0.9];
ax1=axes('Position',[0.15,0.15,0.8,0.8],'XTick',0:totalruntime/5:totalruntime,'YTick',0:0.2:1,...
    'YTickLabel',{'0.0','0.2','0.4','0.6','0.8','1.0'},'FontSize',14,'FontName','Arial');
hold on
axis([-totalruntime*0.05, totalruntime*1.05, -0.05 1.05])
plot(time,proportionhawks,'Color',colorblue,'LineWidth',4)
plot([0,totalruntime], benefit/cost*[1 1],'k--','LineWidth',2);
xlabel('Time','FontName','Arial','FontSize',14)
ylabel('Proportion of hawks','FontName','Arial','FontSize',14)
text(0,benefit/cost*1.07,'Nash equilibrium b/c','FontName','Arial','FontSize',14)
text(totalruntime,benefit/cost*0.92,'Simulation Result','Color',colorblue,'FontSize',14,...
    'HorizontalAlignment','right');
text(totalruntime/2,1,'Population's learning dynamics','FontName','Arial','FontSize',16,...
    'HorizontalAlignment','center')
end

```

Table 1: Some Matlab Code to implement the pairwise comparison process for the hawk-dove game.