Prepared for submission to JHEP

Mathematical models for animals and humans

Marcos Mariño

Département de Physique Théorique et Section de Mathématiques
Université de Genève, Genève, CH-1211 Switzerland

E-mail: marcos.marino@unige.ch

Abstract: An introduction to game-theory based mathematical modelling.

Contents

1 Conflicts and games. Nash equilibrium

> The invention of *game theory* might come to be seen as the most important single advance of the social sciences in the twentieth century.
>
> Jon Elster, *Explaining social behavior*, Cambridge University Press, 2015, p. 308

> All of social science is just a branch of game theory.
>
> Ken Binmore, interview, in D. Colander, R. P. F. Holt and J.B. Rosser, Jr. *The changing face of economics*, University of Michigan Press, 2004, p. 74

1.1 Social interactions as games

In this course we want to develop mathematical models describing social interactions among animals or among human beings. By following standard usage, we will refer to such entities as *agents* or *players*. In a typical interaction between agents, we have two ingredients: the *strategy* of each agent, and her *payoff* or *utility*, i.e. a real number measuring the benefit or gain (with a sign!) obtained from the interaction. Note that, if I interact with you, my payoff depends on both your strategy and my strategy. We will call such an interaction, with strategies and payoffs, a *game* (a precise, mathematical definition will be given later in this Chapter). For the moment being we will assume that players play simultaneously, like in the rock-scissors-paper (RSP) game.

Let us start with some simple examples of games. For simplicity, we will consider games with two players and with a discrete set of "pure" strategies (we will discuss "mixed" strategies later on). Then, the game can be represented as a square box. The rows represent strategies of the first player, and the columns represent strategies of the second player. Inside each box, we put the payoffs of the two players according to the strategies that have been used. Our first example is probably the most famous of all games in town, the *Prisoner's Dilemma*, or PD for short. In this game, there are two strategies, which we will denote by C and D (cooperate and defect). If we both cooperate, we have to make some effort which decreases our payoff by a cost $-c$, but we get in exchange a benefit b. If I cooperate but you defect, I suffer the cost $-c$, while you get the whole benefit b. If we both defect we don't get anywhere and we both get 0. This game is represented in Fig. 1 below.

	C	D
C	$(b-c, b-c)$	$(-c, b)$
D	$(b, -c)$	$(0, 0)$

Figure 1. The payoffs in the Prisoner's dilemma

Another important game which is played all time in human life is known sometimes as *Bach or Stravinsky*[1]. Two players, Gillian and Bob, want to spend the evening together and attend a concert. There are two interesting concerts in town, one playing Bach's *Goldberg Variations*, and the other Stravinsky's *Le sacre du printemps*. Bob likes Stravinsky better (S), while Gillian prefers Bach (B). However, they would prefer to stay together. In contrast to the previous game, this one is asymmetric. The payoff structure can be represented as in Fig. 2.

Our last example is closely related to the PD, and we will call it *Joker's social experiment*. In the film *The dark knight*, the Joker, in his insatiable pursuit of evil, puts bombs in two different ships, which can be regarded as the players of the game. Each ship has a detonator for the bomb in the *other* ship. There are two strategies: to detonate (D) or not to detonate (C, as in cooperation). If both ships detonate, they destroy each other. If one detonates and the other doesn't, the Joker will spare the defector. If both decide not to detonate, the Joker will explode them anyway, but there is a chance that Batman will save them. The payoffs of this game are represented in Fig. 3.

[1]The original name of the game is "battle of the sexes".

		Bob	
		B	S
Gillian	B	(3,2)	(1,1)
	S	(0,0)	(2,3)

Figure 2. Bach or Stravinsky

	C	D
C	$(-\epsilon,-\epsilon)$	(-1,0)
D	(0,-1)	(-1,-1)

Figure 3. Joker's social experiment, with $1 > \epsilon > 0$.

After these elementary examples, we are now ready for some formal definitions.

Definition 1.1. A *game* G in *normal or strategic form* consists of a set $\mathcal{I} = \{1, \cdots, I\}$ of I players, a set of *strategies* S_i for each player i, and a payoff function

$$u_i : S_1 \times \cdots \times S_I \to \mathbb{R}, \qquad i = 1, \cdots, I, \tag{1.1}$$

for each player. We will denote the game G as $G = [\mathcal{I}, \{S_i\}_{i\in\mathcal{I}}, \{u_i(\cdot)\}_{i\in\mathcal{I}}]$. An element $s_i \in S_i$ will be called a *pure* strategy for the i-th player.

So far we have assumed that players choose *pure* strategies $s_i \in S_i$. However, a player in a game might choose a strategy according to a probability distribution. For example, if you play rock-scissors-paper (RSP) many times, you might change your choice in each tournament according to some probability law. You might even randomize completely your choice of strategy (this is actually the best strategy when playing RSP). This leads to the important concept of *mixed strategy.*

Definition 1.2. Let us assume that the strategy set for the i-th player, S_i, has M_i pure strategies, i.e.

$$S_i = \{s_{1i}, \cdots, s_{M_i i}\}. \tag{1.2}$$

A *mixed strategy* is a function

$$\sigma_i : S_i \to [0, 1] \tag{1.3}$$

which assigns to each pure strategy $s_{mi} \in S_i$, $m = 1, \cdots, M_i$, a probability $p_{mi} \geq 0$, such that

$$\sum_{m=1}^{M_i} p_{mi} = 1. \tag{1.4}$$

The set of mixed strategies can be identified with the points of the simplex

$$\Delta(S_i) = \left\{ (p_{1i}, \cdots, p_{M_i i}) \in \mathbb{R}^{M_i} : p_{mi} \geq 0, \text{ for all } m = 1, \cdots, M_i \text{ and } \sum_{m=1}^{M_i} p_{mi} = 1 \right\}. \tag{1.5}$$

We will sometimes write a mixed strategy as a formal linear combination of the pure strategies in S_i, whose coefficients are the probabilities, i.e. we will write

$$\sigma_j = \sum_{k=1}^{M_j} p_{kj} s_{kj}, \qquad j = 1, \cdots, I. \tag{1.6}$$

Once we allow mixed strategies, we have to extend the payoff functions $u_i(\cdot)$ to the full set of mixed-strategy profiles

$$\Delta(S_1) \times \cdots \times \Delta(S_I). \tag{1.7}$$

The standard way to do this is to require multi-linearity. The resulting function is the *Von Neumann–Morgenstern expected utility* or payoff. If σ_j is given by (1.6), the payoff of the i-th player is given by

$$u_i(\sigma_1, \cdots, \sigma_I) = \sum_{k_1=1}^{M_1} \cdots \sum_{k_I=1}^{M_I} p_{k_1 1} \cdots p_{k_I I} u_i\left(s_{k_1 1}, \cdots, s_{k_I I}\right). \tag{1.8}$$

The concept of mixed strategy is very useful, both mathematically and in terms of modeling interactions. As we will eventually explain, when considering games in animal or human populations, the probabilities p_{s_i} will be interpreted as the fraction of the population adopting the strategy s_i, or possessing the trait s_i relevant to the interaction described by the game.

Remark 1.3. A notational remark: we will often denote a profile of pure strategies for the player i's opponents by

$$s_{-i} = (s_1, \cdots, s_{i-1}, s_{i+1}, \cdots, s_I) \in S_{-i} \tag{1.9}$$

where

$$S_{-i} = S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_I. \tag{1.10}$$

In this case, we will denote

$$s = (s_i, s_{-i}). \tag{1.11}$$

A similar notation applies to mixed strategies.

1.2 Dominated strategies

The obvious question when presented with a game is the following; what is the strategy to follow, in order to get the best possible payoff? For the moment being, we will focus on pure strategies. Mixed strategies will be addressed later. The easiest thing to guess is how *not* to play: you don't want to use a strategy which is worst than another one for *every* possible strategic choice of the other players. The simplest example of this is the "cooperate" strategy in PD. Clearly, whatever the other player does, you better defect, since

$$b = u_1(D, C) > u_1(C, C) = b - c \tag{1.12}$$

when the other player cooperates, and

$$0 = u_1(D, D) > u_1(C, D) = -c. \tag{1.13}$$

This type of strategy is called a *strictly dominated strategy.* Let us write down the formal definition.

Definition 1.4. Let $G = [\mathcal{I}, \{S_i\}_{i\in\mathcal{I}}, \{u_i(\cdot)\}_{i\in\mathcal{I}}]$ be a game in normal form. A strategy $s_i \in S_i$ is *strictly dominated* for player i if there exists another strategy $t_i \in S_i$ such that, for *any* profile $s_{-i} \in S_{-i}$ of strategies for the other players, one has

$$u_i(t_i, s_{-i}) > u_i(s_i, s_{-i}). \tag{1.14}$$

We say that t_i *strictly dominates* s_i, or that s_i is strictly dominated by t_i.

In some cases, we are lucky and there is a strategy that strictly dominates any other. You should clearly play that strategy. This also suggests the following definition.

Definition 1.5. A strategy $t_i \in S_i$ is a *strictly dominant strategy* for player i in game G if it strictly dominates any other strategy in S_i.

Example 1.6. In the PD, the strategy D strictly dominates C, so it is strictly dominant. $\square$

We note that, if we relax the strict inequality in definition (1.4) to just an inequality, we obtain the definition of a so-called *weakly dominated strategy*:

Definition 1.7. A strategy $s_i \in S_i$ is *weakly dominated* for player i if there exists another strategy $t_i \in S_i$ such that, for any profile $s_{-i} \in S_{-i}$ of strategies for the other players,

$$u_i(t_i, s_{-i}) \geq u_i(s_i, s_{-i}). \tag{1.15}$$

We say that t_i *weakly dominates* s_i.

	L	R
U	$(1, -1)$	$(-1, 1)$
M	$(-1, 1)$	$(1, -1)$
D	$(-2, 5)$	$(-3, 2)$

Figure 4. A game with a strictly dominated strategy

Example 1.8. In the game represented in Fig. 4, the strategy D is strictly dominated by both U and M. However, there is no strictly dominant strategy.

In some cases, we can find the most "reasonable" strategy for a game G by eliminating strictly dominated strategies in an iterative way. This leads to a *reduced game* H with less strategies. It can be shown that the resulting game H is independent of the order in which strictly dominated strategies are eliminated. If the game H consists of a single strategy for each player, we say that the game G is *solvable by iterated elimination of strictly dominated strategies.*

Example 1.9. Let us consider for example the game shown in Fig. 5. Clearly, the strategy "Right" for player 1 is strictly dominated by "Middle", so it is not rational for her to play it. We can then *eliminate it* from the game, to obtain another game in which player 1 has only two strategies, as shown in Fig. 6. In this new game, the strategy "Down" is now strictly dominated by "Up" for player 2. After eliminating Down, we have the game shown in Fig. 7. In this final game, player 1 has the dominant strategy "Middle". We conclude that the best move for the players (or the "solution" to the game) is that player 1 plays Middle, and player 2 plays Up. $\square$

	Up	Down
Left	$(0,1)$	$(3,0)$
Middle	$(2,1)$	$(1,0)$
Right	$(1,0)$	$(0,2)$

Figure 5. A game which can be solved by iterative elimination of dominated strategies

	Up	Down
Left	$(0,1)$	$(3,0)$
Middle	$(2,1)$	$(1,0)$

Figure 6. The game obtained after eliminating the strategy "Right".

	Up
Left	$(0,1)$
Middle	$(2,1)$

Figure 7. A game which can be solved by iterative elimination of dominated strategies

1.3 Nash equilibrium

The method of solving a game by iterated deletion of strictly dominated strategies does not work in many examples. For example, in Bach or Stravinsky, no player has a strictly dominated strategy. What can be done in those situations? We will now introduce a very fruitful concept of "solution" of a game: the Nash equilibrium. This is a strictly stronger solution concept than the solution obtained by iteration of strongly dominated strategies, in the sense that a Nash equilibrium always survives the iteration process described above, but the converse is not true. Let us first focus on pure strategies.

Definition 1.10. Let $G = [\mathcal{I}, \{S_i\}_{i\in\mathcal{I}}, \{u_i(\cdot)\}_{i\in\mathcal{I}}]$ be a game. A strategy profile $(s_1^*, \cdots, s_I^*)$ is a *Nash equilibrium* of the game if for every $i = 1, \cdots, I$, we have

$$u_i(s_i^*, s_{-i}^*) \geq u_i(t_i, s_{-i}^*), \tag{1.16}$$

for all $t_i \in S_i$.

There is a very insightful way of rephrasing the concept Nash equilibrium which is also useful in practice: in a Nash equilibrium, each player provides the best response to the strategies played by her rivals. To formalize this, we introduce the concept of best-response correspondence.

Definition 1.11. i-th's player *best-response correspondence*

$$BR_i : S_{-i} \to S_i \tag{1.17}$$

in the game $G = [\mathcal{I}, \{S_i\}_{i\in\mathcal{I}}, \{u_i(\cdot)\}_{i\in\mathcal{I}}]$, is a correspondence which assigns to each $s_{-i} \in S_{-i}$ the set

$$BR_i(s_{-i}) = \{s_i \in S_i : \ u_i(s_i, s_{-i}) \geq u_i(t_i, s_{-i}), \ \forall t_i \in S_i\}. \tag{1.18}$$

We now have the following immediate result.

Proposition 1.12. A strategy profile $(s_1^*, \cdots, s_I^*)$ is a Nash equilibrium of the game $G = [\mathcal{I}, \{S_i\}_{i\in\mathcal{I}}, \{u_i(\cdot)\}_{i\in\mathcal{I}}]$ if and only if, for every $i = 1, \cdots, I$,

$$s_i^* \in BR_i(s_{-i}^*). \tag{1.19}$$

We can now go back to games where the iterated elimination of iterated strategies doesn't work, and try to find Nash equilibria.

	l	m	r
U	$(5,3)$	$(0,4)$	$(3,5)$
M	$(4,0)$	$(5,5)$	$(4,0)$
D	$(3,5)$	$(0,4)$	$(5,3)$

Figure 8. A game with no striclty dominated strategies and a single Nash equilibrium.

Example 1.13. Consider the two-player game with three strategies presented in Fig. 8. It is easy to see that this game has no strictly dominated strategies. Let us try to find Nash equilibria. The best-response correspondence of the first player is given by

$$BR_1(l) = \{U\}, \quad BR_1(m) = \{M\}, \quad BR_1(r) = \{D\}, \tag{1.20}$$

while the best-response correspondence of the second player is

$$BR_2(U) = \{r\}, \quad BR_2(M) = \{m\}, \quad BR_2(D) = \{l\}, \tag{1.21}$$

and we conclude that (M, m) is the only Nash equilibrium of this game, since $M \in BR_1(m)$ and $m \in BR_2(M)$. □

Example 1.14. In the game Bach or Stravinsky, Gillian has the following best-response correspondence

$$BR_1(B) = \{B\}, \qquad BR_1(S) = \{S\}. \tag{1.22}$$

For Bob we have

$$BR_2(B) = \{B\}, \qquad BR_2(S) = \{S\}. \tag{1.23}$$

This means that the game has *two* Nash equilibria: (B, B) and (S, S). □

The relationship between iterated elimination of strictly dominated strategies and Nash equilibria is summarized in the following propositions.

Proposition 1.15. If the strategy profile $s^* = (s_1^*, \cdots, s_I^*)$ is a Nash equilibrium of the game $G = [I, \{S_i\}, \{u_i(\cdot)\}]$, it survives iterated elimination of strictly dominated strategies.

Proof: The proof is by contradiction. Let us assume that the Nash equilibrium does not survive the iterated elimination process, and let s_i^* be the first strategy in $s^* = (s_1^*, \cdots, s_I^*)$ that is eliminated. This means that there is a strategy t_i, which has not been eliminated yet, which strictly dominates s_i^*, i.e. such that

$$u_i(t_i, r_{-i}) > u_i(s_i^*, r_{-i}), \tag{1.24}$$

for any strategy profile r_{-i} of the other players. One possible profile is s_{-i}^*, so we should have

$$u_i(t_i, s_{i-}^*) > u_i(s_i, s_{-i}^*) = u(s^*). \tag{1.25}$$

This obviously contradicts the characterizing property (1.16) of a Nash equilibrium. □

This proposition shows that a Nash equilibrium is a stronger solution to a game than the iterated elimination of strictly dominated strategies: a Nash equilibrium always survives the process of iterated elimination. However, the reciprocal is not true: there are strategies that survive iterated elimination which are *not* Nash equilibria, as we have seen in for example the game of Fig. 8, where the full game survives.

Remark 1.16. Nash equilibria do not necessarily survive the iterated elimination of *weakly* dominated strategies. However, as the next proposition shows, if this process ends up in a single strategy profile, then this strategy profile is a Nash equilibrium.

Proposition 1.17. If the iterated elimination of weakly dominated strategies eliminates all but the strategy profile $s^* = (s_1^*, \cdots, s_I^*)$, then this is a Nash equilibrium of the game.

Proof: We will argue by contradiction. Let us suppose that the iterated elimination eliminates all but the strategy profile s^*, and that s^* is not a Nash equilibrium. Therefore, for some player i, there is a strategy $s_i \in S_i$ such that

$$u_i(s_i, s_{-i}^*) > u_i(s_i^*, s_{-i}^*). \tag{1.26}$$

At the same time, s_i was dominated at some point of the iteration process, i.e. there exists a s_i' in the set of strategies available at some stage of the process such that

$$u_i(s_i', r_{-i}) \geq u_i(s_i, r_{-i}), \tag{1.27}$$

where r_{-i} is a profile for the remaining players made out of strategies available at that stage. However, since the profile s^* is never eliminated, by assumption, we can set $r_{-i} = s_{-i}^*$, and (1.27) reads

$$u_i(s_i', s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \tag{1.28}$$

If $s_i' = s_i^*$, then we find immediately a contradiction with (1.26). If $s_i' \neq s_i^*$, then some other strategy s_i'' dominates later s_i', since s_i' does not survive the process of iterative elimination.

Therefore, there is an inequality like (1.28) with s_i' and s_i'' replacing s_i and s_i', respectively, i.e. we have

$$u_i(s_i'', s_{-i}^*) \geq u_i(s_i', s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \tag{1.29}$$

where we used (1.28). Once again, if $s_i'' = s_i^*$, we contradict (1.26), and if not we should go on. Since the number of strategies is finite, we must eventually reach s_i^*. The survival profile is then a Nash equilibrium. □

Corollaire 1.1. If the iterated elimination of strongly dominated strategies eliminates all but the strategy profile $s^* = (s_1^*, \cdots, s_I^*)$, then this is the *unique* Nash equilibrium of the game.

Proof: by the previous proposition, s^* is a Nash equilibrium. By proposition 1.15, all Nash equilibria survive the iteration process, therefore s^* is the unique Nash equilibrium of the game. □

We conclude in particular that the Nash equilibrium of the PD is the strategy (D, D). One disturbing point of this equilibrium is that it is clearly not optimal: *both* players could strictly improve their sort by switching simultaneously to (C, C). The standard way to formalize this is the important notion of Pareto optimality or Pareto efficiency.

Definition 1.18. A strategy profile σ is *Pareto optimal* (or *Pareto efficient*) if there is no other strategy profile σ' such that $u_i(\sigma') \geq u_i(\sigma)$ for all $i = 1, \cdots, I$, and $u_i(\sigma') > u_i(\sigma)$ for at least one player i.

If a set of strategies is Pareto optimal, one cannot change it without making at least one player worse off. Conversely, if a set of strategies is not Pareto optimal, one can make at least one player better off without making any player worse off. It follows that the Nash equilibrium in PD is not Pareto-optimal. This is an important difference between Nash equilibrium and for example the Walrasian equilibrium in a perfectly competitive market. The first fundamental theorem of welfare economics states that the market Walrasian equilibrium is Pareto optimal. This theorem captures Adam Smith's intuition that, in a competitive market, each agent pursues her interests but the overall result is an efficient allocation in which the only way to improve one's lot is at the expense of somebody else. However, Nash equilibria are not necessarily Pareto optimal. That's why, in the movie *A beautiful mind*, John Nash asserts that "Adam Smith needs revision."

Not every game has a Nash equilibrium in pure strategies. A simple example is RSP, shown in Fig. 9. Anybody who has played this game knows that there is no "winning strategy". To see this formally, consdier the best-response correspondence for the first player:

$$BR_1(R) = \{P\}, \qquad BR_1(S) = \{R\}, \qquad BR_1(P) = \{S\}. \tag{1.30}$$

We have the same situation for the second player,

$$BR_2(R) = \{P\}, \qquad BR_2(S) = \{R\}, \qquad BR_2(P) = \{S\}, \tag{1.31}$$

and it is easy to see that there is no Nash equilibrium in pure strategies. In fact, in RSP, the best guarantee not to be badly beaten is to play at random. This means that in order to find some sort of equilibrium we should consider *mixed* strategies. Let us then extend the notion of Nash equilibrium to allow for mixed strategies. In particular, the set of strategies of the i-th player will now be the simplex $\Delta(S_i)$.

	R	S	P
R	$(0,0)$	$(1,-1)$	$(-1,1)$
S	$(-1,1)$	$(0,0)$	$(1,-1)$
P	$(1,-1)$	$(-1,1)$	$(0,0)$

Figure 9. Rock-scissors-paper.

Definition 1.19. A mixed strategy profile $\sigma^* = (\sigma_1^*, \cdots, \sigma_I^*)$ is a Nash equilibrium of the game $G = [\mathcal{I}, \{\Delta(S_i)\}_{i\in\mathcal{I}}, \{u_i(\cdot)\}_{i\in\mathcal{I}}]$ if for every $i = 1, \cdots, I$, we have

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\tau_i, \sigma_{-i}^*), \tag{1.32}$$

for all $\tau_i \in \Delta(S_i)$.

In order to check that a mixed strategy profile is a Nash equilibrium, it is sufficient to check it against pure strategies, since we have the following

Proposition 1.20. A mixed strategy profile $\sigma^* = (\sigma_1^*, \cdots, \sigma_I^*)$ is a Nash equilibrium of the game $G = [\mathcal{I}, \{\Delta(S_i)\}_{i\in\mathcal{I}}, \{u_i(\cdot)\}_{i\in\mathcal{I}}]$ if and only if

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(t_i, \sigma_{-i}^*), \tag{1.33}$$

for all $t_i \in S_i$ and for all $i = 1, \cdots, I$.

Proof: The proof is elementary. Clearly, if σ^* is a Nash equilibrium, (1.33) holds, since pure strategies t_i are particular cases of mixed strategies τ_i. To prove the converse, let $t_{ki} \in S_i$, $k = 1, \cdots, M_i$ be the pure strategies in S_i. Then, by assumption, we have

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(t_{ki}, \sigma_{-i}^*), \qquad k = 1, \cdots, M_i. \tag{1.34}$$

Let us consider the mixed strategy

$$\tau_i = \sum_{k=1}^{M_i} p_{ki} t_{ki}, \qquad \sum_{k=1}^{M_i} p_{ki} = 1. \tag{1.35}$$

It follows that

$$\left(\sum_{k=1}^{M_i} p_{ki}\right) u_i(\sigma_i^*, \sigma_{-i}^*) \geq \sum_{k=1}^{M_i} p_{ki} u_i(t_{ki}, \sigma_{-i}^*) = u_i(\tau_i, \sigma_i^*), \tag{1.36}$$

which is the definition of Nash equilibrium. □

Given a game in normal form, we often want to calculate its Nash equilibria in mixed strategies. One possibility is to extend the idea of best-response correspondence to mixed strategies. This goes as follows. Consider the correspondence

$$BR_i : \prod_{j\neq i} \Delta(S_j) \to \Delta(S_i) \tag{1.37}$$

defined as in (1.18),

$$BR_i(\sigma_{-i}) = \{\sigma_i \in \Delta(S_i) : \ u_i(\sigma_i, \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}), \ \forall \tau_i \in \Delta(S_i)\}. \tag{1.38}$$

Then, a profile $\sigma^* = (\sigma_1^*, \cdots, \sigma_I^*)$ is a Nash equilibrium if $\sigma_i^* \in BR_i(\sigma_{-i}^*)$ for all $i = 1, \cdots, I$. This can be used to find mixed Nash equilibria, but it is more practical to use the "fundamental theorem" of mixed equilibria, which we now state. Before doing that, we need the following definition.

Definition 1.21. In a finite game, the support of a mixed strategy σ_i, supp σ_i, is the set of pure strategies to which σ_i assigns positive probability.

Theorem 1.22. *(Fundamental theorem of mixed strategy Nash equilibrium). A strategy profile $(\sigma_1^*, \cdots, \sigma_I^*)$ is a Nash equilibrium of the game $G = [\mathcal{I}, \{\Delta(S_i)\}_{i\in\mathcal{I}}, \{u_i(\cdot)\}_{i\in\mathcal{I}}]$ if and only if, for every $i = 1, \cdots, I$,*

1. *$u_i(s_i, \sigma_{-i}^*) = u_i(s_i', \sigma_{-i}^*)$, for all $s_i, s_i' \in \operatorname{supp}\sigma_i^*$.*
2. *$u_i(s_i, \sigma_{-i}^*) \geq u_i(t_i, \sigma_{-i}^*)$, for all $s_i \in \operatorname{supp}\sigma_i^*$ and $t_i \notin \operatorname{supp}\sigma_i^*$*

Proof: To see that the conditions are necessary, suppose that they do not hold for some player i. Let us first suppose that (1) does not hold. This means that there are strategies $s_i, s_i' \in \operatorname{supp}\sigma_i^*$ such that

$$u_i(s_i', \sigma_{-i}^*) > u_i(s_i, \sigma_{-i}^*) \qquad \text{or} \qquad u_i(s_i, \sigma_{-i}^*) > u_i(s_i', \sigma_{-i}^*) \tag{1.39}$$

If the first inequality is the case, σ_i^* cannot be a Nash equilibrium, since player i can increase her payoff by playing strategy s_i' whenever he would have played strategy s_i. A similar argument can be made for the second inequality. Let us now suppose that (2) does not hold. This means that there is a strategy $s_i \in \operatorname{supp}\sigma_i^*$, $s_i' \notin \operatorname{supp}\sigma_i^*$ and such that

$$u_i(s_i', \sigma_{-i}^*) > u_i(s_i, \sigma_{-i}^*). \tag{1.40}$$

Again, player i could then increase her payoff by switching. We conclude that (1) and (2) should hold in a mixed Nash equilibrium.

To see that conditions (1) and (2) are sufficient to characterize a mixed Nash equilibrium, let us reason by contradiction. Let us suppose that (1) and (2) hold but that σ^* is not a Nash equilibrium. By Proposition 1.20, there is a player i and a pure strategy t_i such that

$$u_i(t_i, \sigma_{-i}^*) > u_i(\sigma_i^*, \sigma_{-i}^*). \tag{1.41}$$

There are two possibilities: either $t_i \in \operatorname{supp}\sigma_i^*$, or $t_i \notin \operatorname{supp}\sigma_i^*$. In the first case, we contradict (1), since according to (1)

$$u_i(s_i, \sigma_{-i}^*) = u_i(t_i, \sigma_{-i}^*), \tag{1.42}$$

for all $s_i \in \operatorname{supp}\sigma_i^*$. By summing this equality over all elements s_i in supp σ_i^*, weighted with the probabilities with which they appear in σ_i^*, we find

$$u_i(\sigma_i^*, \sigma_{-i}^*) = u_i(t_i, \sigma_{-i}^*), \tag{1.43}$$

which contradicts (1.41). If $t_i \notin \operatorname{supp}\sigma_i^*$ we contradict (2), since according to (2)

$$u_i(s_i, \sigma_{-i}^*) \geq u_i(t_i, \sigma_{-i}^*), \tag{1.44}$$

for all $s_i \in \text{supp } \sigma_i^*$. Again, by summing this inequality over all elements in supp σ_i^*, weighted with the probabilities appearing in σ_i^*, we find

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(t_i, \sigma_{-i}^*), \tag{1.45}$$

which contradicts (1.41) again. □

The Fundamental Theorem is extremely useful in order to obtain mixed Nash equilibria. In particular, condition (1) says that, in a mixed Nash equilibrium, players should be indifferent between pure strategies in the support of the equilibrium. To see this principle at work, let us consider the following example.

Example 1.23. *Nash equilibrium in RSP*. We know that RSP has no Nash equilibrium in pure strategies. Let us look for a mixed Nash equilibrium, with

$$\begin{aligned} \sigma_1^* &= q_1 R + q_2 S + (1 - q_1 - q_2) P, \\ \sigma_2^* &= p_1 R + p_2 S + (1 - p_1 - p_2) P. \end{aligned} \tag{1.46}$$

By the fundamental theorem, we have

$$u_1(R, \sigma_2^*) = u_1(S, \sigma_2^*) = u_1(P, \sigma_2^*). \tag{1.47}$$

Since

$$\begin{aligned} u_1(R, \sigma_2^*) &= p_2 - (1 - p_1 - p_2) = 2p_2 + p_1 - 1, \\ u_1(S, \sigma_2^*) &= -p_1 + (1 - p_1 - p_2) = 1 - 2p_1 - p_2, \\ u_1(P, \sigma_2^*) &= p_1 - p_2, \end{aligned} \tag{1.48}$$

the solution is

$$p_1 = p_2 = p_3 = \frac{1}{3}. \tag{1.49}$$

The same argument shows that $\sigma_1^* = \sigma_2^*$, and the NE is $\sigma^* = (\sigma_1^*, \sigma_1^*)$. Such an equilibrium, in which all players play the same strategy, is called a *symmetric* Nash equilibrium. □

Example 1.24. As we saw in example 1.14, Bach or Stravinsky has two Nash equilibria in pure strategies. To look for mixed strategies, we use the Fundamental Theorem. Let us consider the strategy profile

$$\sigma^* = (\sigma_1^*, \sigma_2^*), \tag{1.50}$$

where

$$\begin{aligned} \sigma_1^* &= \alpha B + (1 - \alpha) S, \qquad 0 < \alpha < 1, \\ \sigma_2^* &= \beta B + (1 - \beta) S, \qquad 0 < \beta < 1. \end{aligned} \tag{1.51}$$

We want to solve for α, β. To solve for α, we note that player 2 has to be indifferent between B and S, therefore,

$$u_2(\sigma_1^*, B) = u_2(\sigma_1^*, S), \tag{1.52}$$

i.e.

$$\alpha u_2(B, B) + (1 - \alpha) u_2(S, B) = \alpha u_2(B, S) + (1 - \alpha) u_2(S, S), \tag{1.53}$$

which gives

$$2\alpha = \alpha + 3(1 - \alpha). \tag{1.54}$$

This is solved by

$$\alpha = \frac{3}{4}. \tag{1.55}$$

Player 1 has to be indifferent between B and S, therefore

$$u_1(B, \sigma_2^*) = u_1(S, \sigma_2^*). \tag{1.56}$$

This can be rewritten as

$$\beta u_1(B, B) + (1 - \beta) u_1(B, S) = \beta u_1(S, B) + (1 - \beta) u_1(S, S). \tag{1.57}$$

Using the payoffs of player 1, we find

$$3\beta + 1 - \beta = 2(1 - \beta) \Rightarrow \beta = \frac{1}{4}. \tag{1.58}$$

The resulting profile is $\sigma^* = (\sigma_1^*, \sigma_2^*)$, with

$$\sigma_1^* = \frac{3}{4} B + \frac{1}{4} S, \qquad \sigma_2^* = \frac{1}{4} B + \frac{3}{4} S, \tag{1.59}$$

and it satisfies by construction condition (1) in the Fundamental Theorem. Since there are only two strategies for each player, and both are in the support of σ_i^*, $i = 1, 2$, we conclude that (2) is also satisfied, and σ^* is a Nash equilibrium in mixed strategies. □

One important property of Nash equilibrium is that it always exists, at least in mixed strategies. This is the content of the famous theorem proved by John Nash in his Ph. D. thesis.

Theorem 1.25. *(Nash). Every game $G = [\mathcal{I}, \{\Delta(S_i)\}_{i \in \mathcal{I}}, \{u_i(\cdot)\}_{i \in \mathcal{I}}]$ in which the sets S_i, $i = 1, \cdots, I$ have a finite number of elements, has a mixed strategy Nash equilibrium.*

The proof of the theorem is simple once we assume a technical result known as Kakutani's fixed point theorem. The idea of the proof is to consider the best-response correspondences (1.37), and to consider an application

$$BR : \Delta(S_1) \times \cdots \Delta(S_I) \to \Delta(S_1) \times \cdots \Delta(S_I), \tag{1.60}$$

defined by

$$BR(\sigma_1, \cdots, \sigma_I) = BR_1(\sigma_{-1}) \times \cdots BR_I(\sigma_{-I}). \tag{1.61}$$

A Nash equilibrium is clearly a fixed point of this application. The existence of such fixed point is guaranteed in turn by Kakutani's theorem (we need a more sophisticated fixed point theorem than e.g. Brouwer's theorem since BR is not an application, but a correspondence, i.e. each point can have more than one image). Nash' theorem can be extended to other situations in which the strategy sets are not necessarily of the form $\Delta(S_i)$ for a finite set S_i.

The existence theorem of Nash is very important, since it guarantees that games will have at least one equilibrium, or "solution." Therefore, the problem with Nash equilibria is not existence, but uniqueness, since there are often many equilibria and we don't know in which one of them we will end up. We will address this embarrassment of richness at various points in this course.

2 Applications of Nash equilibrium

2.1 Cournot oligopoly

We will now discuss a classical application of Nash equilibrium to economic theory. In this game, the players are I firms competing in a market and producing the same good. A strategy for the firm i is the quantity $q_i \geq 0$ of the good that it produces, therefore the strategy space is

$$S_i = \mathbb{R}_+, \qquad i = 1, \cdots, I. \tag{2.1}$$

A strategy profile in this game is then a vector

$$\mathbf{q} = (q_1, \cdots, q_I) \in S_1 \times \cdots S_I. \tag{2.2}$$

Each firm has a cost function for producing the good, i.e. a function

$$C_i : \mathbb{R}_+ \to \mathbb{R}_+. \tag{2.3}$$

The price of the good depends on the aggregate quantity produced by all the firms,

$$Q = \sum_{i=1}^{I} q_i, \tag{2.4}$$

through a given *price function* $P(Q)$ (this can be regarded as the inverse of a demand function $Q = F(P)$, which specifies the total demand for a good as a function of its price.) We assume that all the outputs of all the firms are sold in the market. Therefore, the payoff function of each firm is

$$u_i(\mathbf{q}) = P(Q)q_i - C_i(q_i). \tag{2.5}$$

We can now look for a Nash equilibrium in this game. This will be a profile

$$\mathbf{q}^* = (q_1^*, \cdots, q_I^*) \in S_1 \times \cdots S_I, \tag{2.6}$$

such that

$$u_i(\mathbf{q}^*) \geq u_i(q_i, q_{-i}^*) \tag{2.7}$$

for all $q_i \in S_i$. One way of proceeding is to compute best-response functions (these are similar to the best-response correspondences introduced above, but since in this case the image of the correspondence is a single point, the correspondence reduces to a conventional function). Assume that the firm i is confronted to outputs q_j^*, $j \neq i$, from the other firms. Its payoff is

$$u_i(q_i, q_{-i}^*) = P\left(q_i + \sum_{j\neq i} q_j^*\right) q_i - C_i(q_i). \tag{2.8}$$

The best-response function gives the set of points maximizing this payoff. In particular, if the points are in the interior of S_i, they satisfy the first-order condition:

$$\frac{\partial u_i(q_i, q_{-i}^*)}{\partial q_i} = P'\left(q_i + \sum_{j\neq i} q_j^*\right) q_i + P\left(q_i + \sum_{j\neq i} q_j^*\right) - C_i'(q_i) = 0. \tag{2.9}$$

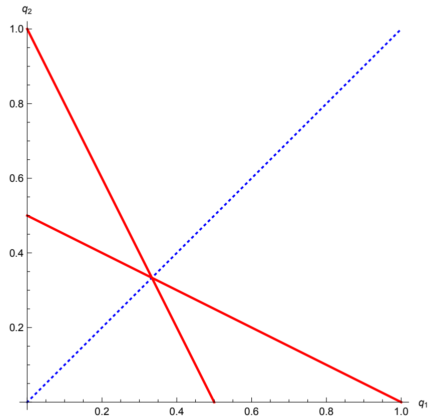

Figure 10. The Nash equilibrium in the (q_1, q_2) plane for the payoff functions given in (2.37), with $a = 2$, $b = c = 1$. It is given by the intersection of the best-response functions (2.15), (2.17).

Let us assume that $\mathbf{q}^*$ is an (interior) Nash equilibrium. Then, it has to satisfy the above first order condition with $q_i = q_i^*$, and for all $i = 1, \cdots, I$. We then find I extremization conditions,

$$P'(Q^*)q_i^* + P(Q^*) - C_i'(q_i^*) = 0, \qquad i = 1, \cdots, I. \tag{2.10}$$

The solution to these conditions is the Nash–Cournot equlibrium of this game.

A concrete application will help to understand this game. Let us assume that there are only two firms, with identical cost functions,

$$c_1(q) = c_2(q) = cq. \tag{2.11}$$

The inverse demand function is

$$P(Q) = a - bQ. \tag{2.12}$$

The payoff functions are

$$\begin{aligned} u_1(q_1, q_2) &= P(Q)q_1 - cq_1 = (a-c)q_1 - bq_1^2 - bq_1q_2, \\ u_2(q_1, q_2) &= P(Q)q_2 - cq_2 = (a-c)q_2 - bq_2^2 - bq_1q_2. \end{aligned} \tag{2.13}$$

The best-response function $B_1(q_2)$ of the first firm is determined by

$$\frac{\partial u_1(q_1, q_2)}{\partial q_1} = 0 \Rightarrow a - c - 2bq_1 - bq_2 = 0, \tag{2.14}$$

i.e.

$$B_1(q_2) = \begin{cases} \frac{a-c}{2b} - \frac{1}{2}q_2, & \text{if } q_2 \leq \frac{a-c}{b}, \\ 0, & \text{otherwise.} \end{cases} \tag{2.15}$$

Note that

$$\frac{\partial^2 u(q_1, q_2)}{\partial q_1^2} = -2b < 0, \tag{2.16}$$

so this is indeed a local maximum. Since the situation for the firm 2 is completely symmetric, we have

$$B_2(q_1) = \begin{cases} \frac{a-c}{2b} - \frac{1}{2}q_1, & \text{if } q_1 \leq \frac{a-c}{b}, \\ 0, & \text{otherwise.} \end{cases} \tag{2.17}$$

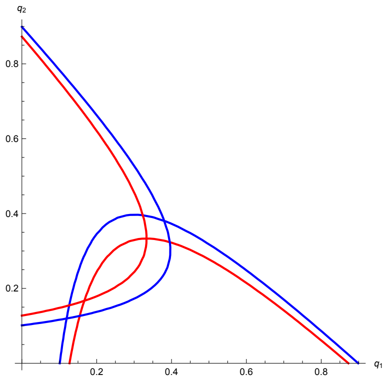

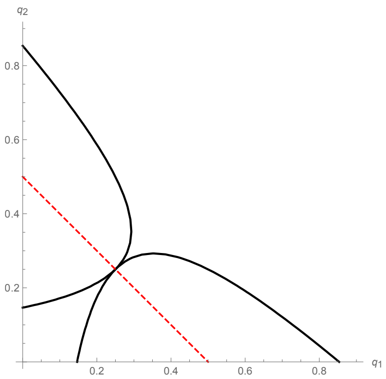

Figure 11. (Left) The level curves of the payoff functions of the two firms, for $a = 2$, $b = c = 1$. In blue we represent $u_1(q_1, q_2) = 1/11$ and $u_2(q_1, q_2) = 1/11$, and in red we represent $u_1(q_1, q_2) = 1/9$ and $u_2(q_1, q_2) = 1/9$. The region inside the intersection of the two curves contains points which are Pareto superior to the intersection points of the curves. (Right) The Pareto optimal outcomes correspond to the points of tangency of the two level curves. This defines the contract curve, which is shown here as a dashed line. We also show the tangency point when the payoffs are equal: $u_1(q_1, q_2) = u_2(q_1, q_2) = 1/8$.

The Nash equilibrium is the point (q_1^*, q_2^*) satisfying

$$B_1(q_2^*) = q_1^*, \qquad B_1(q_1^*) = q_2^*. \tag{2.18}$$

Since by symmetry we must have $q_1^* = q_2^* = q_*$, we find

$$q_* = \frac{a - c}{3b}. \tag{2.19}$$

This point can be easily understood graphically: it is the point of intersection of the two best-response functions, see Fig. 10. At this point, the payoffs of the producers are

$$u_{1,2}(q_*, q_*) = \frac{(a - c)^2}{9b}. \tag{2.20}$$

It is easy to see that the Nash–Cournot equilibrium is not Pareto optimal. Let us do the analysis for the particular example discussed above. We look at the curves of constant payoffs of the two firms, in the plane (q_1, q_2). Due to the nature of the payoff function, higher utility for the firm 1 (2) means curves which are closer to the q_1 axis (q_2 axis, respectively). In Fig. 11, on the left, we show two examples of level curves for payoff functions. In blue we represent $u_1(q_1, q_2) = 1/11$ and $u_2(q_1, q_2) = 1/11$, and in red we represent $u_1(q_1, q_2) = 1/9$ and $u_2(q_1, q_2) = 1/9$, when $a = 2$, $b = c = 1$. Clearly, the points inside the regions obtained by intersection of the two curves give higher payoffs for *both* agents, so they are Pareto improvements over the intersection points. In order to find a Pareto optimal outcome, both curves should be tangent at their common point, as shown in the figure on the right in Fig. 11. Analytically, this point must satisfy

$$\frac{\partial_{q_1} u_1(q_1, q_2)}{\partial_{q_2} u_1(q_1, q_2)} = \frac{\partial_{q_1} u_2(q_1, q_2)}{\partial_{q_2} u_2(q_1, q_2)}. \tag{2.21}$$

Equivalently, one can obtain this condition by maximizing one of the payoff functions, while the other remains constant, i.e. one maximizes

$$u_i(q_1, q_2), \qquad i = 1, 2, \tag{2.22}$$

to find a point $(q_1^{\mathrm{P}}, q_2^{\mathrm{P}})$, with the condition that

$$u_j(q_1, q_2) = u_j(q_1^{\mathrm{P}}, q_2^{\mathrm{P}}), \qquad j \neq i. \tag{2.23}$$

We can solve this extremum problem by using Lagrange multipliers. Let us for example maximize the payoff of the first player. We consider the following function

$$\mathcal{L}_1(q_1, q_2) = u_1(q_1, q_2) + \lambda_1 \left(u_2(q_1, q_2) - u_2(q_1^{\mathrm{P}}, q_2^{\mathrm{P}}) \right), \tag{2.24}$$

where λ_1 is a Lagrange multiplier. The extremum conditions are

$$\begin{aligned} \frac{\partial \mathcal{L}_1}{\partial q_1} &= \frac{\partial u_1}{\partial q_1} + \lambda_1 \frac{\partial u_2}{\partial q_1} = 0, \\ \frac{\partial \mathcal{L}_1}{\partial q_2} &= \frac{\partial u_1}{\partial q_2} + \lambda_1 \frac{\partial u_2}{\partial q_2} = 0. \end{aligned} \tag{2.25}$$

The second equation gives

$$\lambda_1 = -\frac{\partial_{q_2} u_1}{\partial_{q_2} u_2}, \tag{2.26}$$

and replacing this value in the first equation we obtain (2.21). It is easy to check that the same condition is obtained if we maximize the payoff of the second player.

The geometric locus defined by (2.21), which is the geometric locus of Pareto efficient points, is called the *contract curve*. In this example, the contract curve can be found analytically. By calculating derivatives, we find

$$\frac{a - c - 2bq_1 - bq_2}{-bq_1} = \frac{-bq_2}{a - c - 2bq_2 - q_1}. \tag{2.27}$$

Let us denote $\xi = (a - c)/b$. We find the quadratic equation

$$2Q^2 - 3\xi Q + \xi^2 = 0, \qquad Q = q_1 + q_2. \tag{2.28}$$

This has two solutions. The first one is $Q = \xi/2$, or

$$q_1 + q_2 = \frac{a - c}{2b}. \tag{2.29}$$

The second one is $Q = \xi$. It can be seen that this second solution is spurious since $u_1(q_1, q_2) = u_2(q_1, q_2) = 0$ along this curve. Therefore, the contract curve is given in this case by (2.29).

We can now go back to the Nash equilibrium and check that it cannot be Pareto optimal. This is easy to see analytically: at the Nash equilibrium, we have

$$\frac{\partial u_1}{\partial q_1} = \frac{\partial u_2}{\partial q_2} = 0, \tag{2.30}$$

which does not agree with (2.21). This is also easy to see graphically, by representing the payoff level curves which pass through the Nash equilibrium: as shown in Fig. 12, the Nash equilibrium can be Pareto improved by going to the points inside the intersection of the two curves. In the example above, the Pareto optimum with equal payoffs (and equal production) corresponds to an output

$$q_1^{\mathrm{P}} = q_2^{\mathrm{P}} = q^{\mathrm{P}} = \frac{a - c}{4b}, \tag{2.31}$$

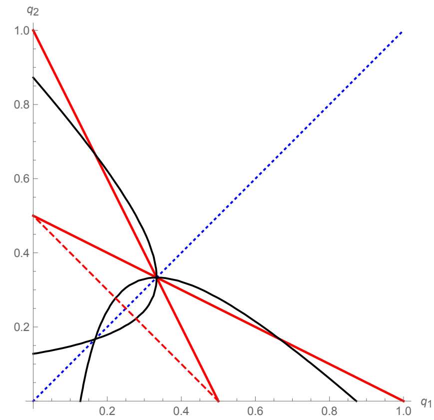

Figure 12. The Nash equilibrium in a Cournot duopoly is not Pareto optimal, as it is not on the contract curve.

with payoffs

$$u_{1,2}(q^{\mathrm{P}}, q^{\mathrm{P}}) = \frac{(a-c)^2}{8b}. \tag{2.32}$$

If both producers could commit themselves to this choice, they would improve their situation. However, in the absence of commitment, the Pareto optimum above is not an equilibrium, since it is in the interest of each producer to defect. Suppose for example that the first producer decides to betray his partner, who is producing at the Pareto optimal level. Assuming that $q_2 = q_2^{\mathrm{P}}$, it is in the interest of the first producer to increase the production to his best-response function

$$B_1(q^{\mathrm{P}}) = \frac{3(a-c)}{8b}. \tag{2.33}$$

With this output, his payoff increases to

$$u_1(B_1(q^{\mathrm{P}}), q^{\mathrm{P}}) = \frac{9(a-c)^2}{64b}, \tag{2.34}$$

while the second producer worsens his lot since his payoff is now

$$u_2(q^{\mathrm{P}}, B_1(q^{\mathrm{P}})) = \frac{3(a-c)^2}{32b}. \tag{2.35}$$

The structure of this game is then very similar to a PD.

Example 2.1. *Cournot duopoly with different production costs.* We can consider the slightly more general situation in which the two producers have different production costs satisfying

$$c_1 < c_2 < \frac{a}{2}. \tag{2.36}$$

The utilities are given by

$$\begin{aligned} u_1(q_1, q_2) &= (a - c_1)q_1 - bq_1^2 - bq_1q_2, \\ u_2(q_1, q_2) &= (a - c_2)q_2 - bq_2^2 - bq_1q_2. \end{aligned} \tag{2.37}$$

It is easy to see that, in this case, the Nash equilibrium corresponds to the utilities

$$u_1^{\mathrm{NE}} = \frac{(a - 2c_1 + c_2)^2}{9b}, \qquad u_2^{\mathrm{NE}} = \frac{(a - 2c_2 + c_1)^2}{9b}. \tag{2.38}$$

The contract curve is more complicated in this case, and it is described by the hyperbola

$$(a - c_1 - b(2q_1 + q_2))\,(a - c_2 - b(2q_2 + q_1)) - b^2 q_1 q_2 = 0. \tag{2.39}$$

□

2.2 Common property resources

Nash equilibria which are Pareto inefficient can be regarded as *coordination failures*: both players could improve their outcomes if they were able to coordinate their strategies. This type of failures are widespread in society. We will now analyze the Nash equilibrium arising in common property resources, which leads to a similar type of coordination failure. A *common property resource* is a good where users can not be excluded (non-excludability), but where the use of the resource by one user diminishes the benefits available to other users (rivalness). Examples of common property resources are public roads and internet access (with their traffic congestion problems), or fisheries and water resources (with the corresponding depletion problems). For concreteness, we will illustrate the problem with a particular example: the "tragedy of the fishers" (taken from [4]).

Let us consider two fishers, Upper and Lower, who fish in the same lake. The amount of fish each one catches correlates positively with the effort of each one, and negatively with the effort of the other. In other words, since fish is a rival commodity, the more fish one of them catches, the less fish is left for the other. To be specific, let y, Y be the amount of fish caught by Lower (respectively, Upper), and e, E the amount of time or effort they devote to fishing. We assume the following functional form:

$$\begin{aligned} y(e, E) &= \alpha(1 - \beta E)e, \\ Y(e, E) &= \alpha(1 - \beta e)E. \end{aligned} \tag{2.40}$$

Here, α and β are positive constants. The payoffs for the fishers are

$$\begin{aligned} u(y, e) &= y - e^2, \\ U(Y, E) &= Y - E^2, \end{aligned} \tag{2.41}$$

which correlates positively with the amount of fish they catch, and negatively with their effort. Note that, mathematically, this is almost identical to the model of Cournot duopoly discussed in the previous section. When written in terms of e, E, we have

$$\begin{aligned} u(e, E) &= \alpha e - e^2 - \alpha\beta e E, \\ U(e, E) &= \alpha E - E^2 - \alpha\beta e E, \end{aligned} \tag{2.42}$$

which are identical to the payoff functions in (2.37), up to an overall rescaling of the variables and the function. We can immediately solve for the Nash equilibrium for the efforts,

$$e^* = E^* = \frac{\alpha}{2 + \alpha\beta}. \tag{2.43}$$

As in the case of the Cournot duopoly, this is not a Pareto optimal allocation, and the effort level of the fishers is too big. In other words, *the common resource is overused*. This is a precise mathematical statement of the famous "tragedy of the commons" pointed out by Garrett Hardin in [18]. It is clearly another version of a PD: each player would like to fish less (i.e. to cooperate), but if she did so, the other player might take advantage of such a decision by fishing more (i.e. by defecting), so they both fish too much (i.e. they both defect).

There have been many proposals to avoid the tragedy of the commons, which can be mathematically modeled. For an insightful survey, see Chapter 4 of [4].

2.3 Sex allocation

In many animal species (including ours), one observes that the proportion of males to females is very close to one, i.e. both sexes are equally represented in the population (at least in terms of births). From an evolutionary point of view, this is paradoxical: as any farmer knows, a male can inseminate many females, so an equal share of both sexes is really unnecessary from the point of view of the reproductive needs of the species as a whole. This evolutionary puzzle, first noted by Darwin himself, was solved by Fisher in 1930.

Fisher's argument can be reformulated as a Darwinian game. In this game, the players are females, and their reproductive strategy is the ratio of sons to daughters in their brood. The payoff is reproductive success, or *fitness*, which can be measured in terms of grandchildren. Suppose now that there are less males than females in the population. Since every individual needs a mother and a father, when males are rare, they have more reproductive success than females, comparatively. Producing sons instead of daughters is then evolutionarily advantageous to a female, since this will lead to more grandchildren. In this way, the strategy "producing more males" is favored. When the number of males in the population surpasses the number of females, the relative advantages get reversed, and the winning strategy is "producing more females." Clearly, producing as many males as females (i.e. the ratio of one to one, or *Fisherian sex ratio*) is an equilibrium of the game.

We will now show, formally, that the Fisherian sex ratio is a Nash equilibrium in a certain game. This is doubly interesting because, when the rules of the game change, the Fisherian sex ratio is not necessarily the best strategy. This was noted by William Hamilton in 1967, in a remarkable paper on extraordinary sex ratios [17]. Biased or extraordinary sex ratios are common in certain animals, and they arise as Nash strategies in the Darwinian game. The problem of "choosing" gender ratios as reproductive strategies is known in evolutionary biology as the problem of *sex allocation*.

Our model of sex allocation is as follows. There are n players in the game, labeled by $i = 1, \cdots, n$. Each player is a female, and her strategy is the production of m_i males and f_i females in the next generation. These two variables are not independent, and we will assume that there is a constraint of the form

$$f_i = F_i(m_i), \qquad i = 1, \cdots, n. \tag{2.44}$$

This constraint is due to the fact that a female has a limited biological budget to make children. If sons and daughters are equally costly, and the total number of children made by the i-th female is b_i, we will have the constraint

$$f_i + m_i = b_i, \tag{2.45}$$

so that

$$F_i(m) = b_i - m. \tag{2.46}$$

Since m_i and f_i are related by (2.44), the strategy of the i-th player is specified by giving the value of, say, m_i. The payoff of each player is the number of sons and daughters, weighted by the reproductive value of males and females in the population. The reproductive value, or reproductive success, is the expected contribution of an individual to the genetic pool of the population. We will denote by v_f, v_m the reproductive successes of females and males. These quantities are related to the total number of males and females in the population, n_m, n_f. Indeed, let us assume that the species is *diploid*, i.e. that mothers and fathers contribute equally to the number of genes of an individual. Then, the total genetic contribution of males must be equal to the total genetic contribution of females, and we have

$$v_f n_f = v_m n_m, \tag{2.47}$$

On the other hand, we have that

$$\frac{n_f}{n_m} = \frac{\sum_{j=1}^n f_j}{\sum_{j=1}^n m_j}. \tag{2.48}$$

We will assume now that the reproductive success of females is a constant v_f, while the reproductive success of males v_m is obtained from (2.47), and is given by

$$v_m = v_f \frac{n_f}{n_m} = v_f \frac{\sum_{j=1}^n f_j}{\sum_{j=1}^n m_j}. \tag{2.49}$$

Note that v_m depends on the relative proportion of males and females in the population. Therefore, the payoff to the i-th player is

$$u_i(m_1, \cdots, m_n) = v_f f_i + v_m m_i = v_f \left(f_i + m_i \frac{n_f}{n_m} \right) = v_f \left(f_i + m_i \frac{\sum_{j=1}^n f_j}{\sum_{j=1}^n m_j} \right), \qquad i = 1, \cdots, n. \tag{2.50}$$

This payoff can be also interpreted as the contribution of the i-th player to the genetic pool of the second generation. We have the contribution of the daughters, f_i, and then we multiply the contribution of the sons, m_i, by the probability that each son mates with a female in the population, which is n_f/n_m.

We will assume that all the players are identical (in particular, the constraint functions are the same for all players: $F_i(m) = F(m)$ for all $i = 1, \cdots, n$). We will search for a Nash equilibrium in which $m_i = m_*$ for all $i = 1, \cdots, n$[2]. What is the best response of the i-th player, assuming that the other players choose $m_j = m_*$, $j \neq i$, and the corresponding $f_* = F(m_*)$? We have to find the extremum of

$$u_i(m_i, m^*_{-i}) = v_f \left(F(m_i) + m_i \frac{(n-1)f_* + f_i}{(n-1)m_* + m_i} \right) \tag{2.51}$$

w.r.t. m_i, and set $m_i = m_*$. It is easy to write down the general condition for this, but we will consider two particular cases.

In the first case, we take $n \to \infty$, and we write $m_i = m$, $f_i = f$, so that the payoff function reads

$$u^\infty(m, m_*) = v_f \left(F(m) + m \frac{f_*}{m_*} \right). \tag{2.52}$$

Extremizing w.r.t. m and setting $m = m_*$ afterwards we find,

$$F'(m_*) + \frac{f_*}{m_*} = 0. \tag{2.53}$$

This is sometimes called the *Shaw–Mohler theorem*. In the simplest case in which $F(m) = b - m$, we find

$$f_* = m_*. \tag{2.54}$$

[2]This is called a *symmetric* Nash equilibrium, for obvious reasons. We will study in detail symmetric Nash equilibria in games of two players in Chapter 4.

This is precisely the Fisherian sex ratio, which we have then derived as a Nash equilibrium.

Sometimes producing males is more costly than producing females, or viceversa[3]. Suppose for example than producing males is twice more costly than producing females. Then, the constraint would be

$$f + 2m = b, \tag{2.55}$$

and the Shaw–Mohler theorem gives

$$f_* = 2m_*, \tag{2.56}$$

i.e. sex ratio is biased in favor of females, due to the different costs of each sex.

Let us now consider the case in which n is finite, and sons and daughters are equally costly, so that $F(m) = b - m$. Let us express everything in terms of the variable r_i, defined by

$$m_i = br_i, \tag{2.57}$$

so that r_i is the proportion of sons in the brood of the i-th player. We then have,

$$u_i(r_i, r^*_{-i}) = v_f b \left(1 - r_i + r_i \frac{(n-1)(1-r_*) + 1 - r_i}{(n-1)r_* + r_i} \right). \tag{2.58}$$

Taking the derivative w.r.t. r_i we find

$$\frac{\partial u_i}{\partial r_i} = v_f b \left\{ -1 + \frac{(n-1)(1-r_*) + 1 - r_i}{(n-1)r_* + r_i} - r_i \frac{n}{((n-1)r_* + r_i)^2} \right\}, \tag{2.59}$$

and the condition

$$\left. \frac{\partial u_i}{\partial r_i} \right|_{r_i = r_*} = 0 \tag{2.60}$$

gives

$$r_* = \frac{n-1}{2n}. \tag{2.61}$$

This is Hamilton's famous result for the sex ratio in a diploid population in which n females reproduce. Note that, as $n \to \infty$, we recover Fisher's ratio $1/2$. However, for finite n, this is smaller than $1/2$: it is more convenient to produce fewer males and more females. The reason is that, in a small population, there is competition for partners, so it is better to have more daughters. For this reason, this type of models of sex allocation are known as models of *local mate competition*.

Hamilton's result has been applied in certain species where a small number of females called the "foundresses" lay their eggs in a closed environment, so that their descendants have to reproduce among themselves. One example are fig wasps. In this species, a group of foundresses lay their eggs in the same fig fruit. Their sons and daughters mate inside the fig. Males die inside the fruit, and inseminated females leave to find another fig. These are precisely the conditions for local mate competition. A similar behavior occurs in the parasitoid wasps *Ichneumonidae*. These wasps paralyze a host (e.g. a caterpillar) and lay their eggs inside their body[4]. Since

[3] Among mammals, males tend to be bigger, therefore more costly. However, in most species, females are bigger, see for example Chapter 3 of [15].

[4] The hosts are eaten alive by the larvae of the wasp. This animal behavior led Darwin to write: "There seems to me too much misery in the world. I cannot persuade myself that a beneficent and omnipotent God would have designedly created the *Ichneumonidae* with the express intention of their feeding within the living bodies of caterpillars." (letter to Asa Gray, May the 22nd, 1860)

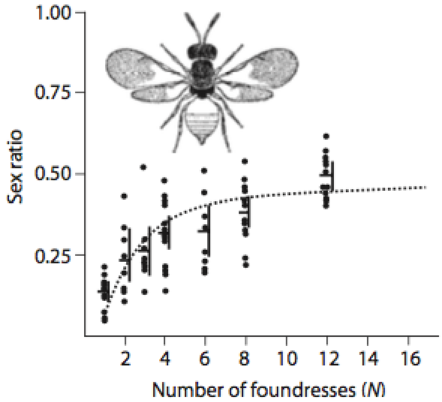

Figure 13. The sex ratio as a function of the number of foundresses in a parasitoid wasp. The dots represent experimental data, while the dotted line is Hamilton's prediction (2.62). Figure taken from [28].

wasps are haplodiploid, the formula (2.61), which was derived for the diploid case, gets modified to (see for example [5, 28])

$$r_* = \frac{2n-1}{4n-1}\frac{n-1}{n}. \tag{2.62}$$

This is due to the peculiar genetics of haplodiploidy. However, this value of r_* has the same qualitative behavior shown by (2.61): it gives a biased sex ratio in favor of females for small n, and it approaches $1/2$ as n becomes large. Experimental data on the dependence of r_* with n, as well as the theoretical prediction (2.62), are shown in Fig. 13.

A extreme case occurs when $n = 1$, i.e. there is a single reproducing mother. In this case, Hamilton's result predicts no males, which of course can not be strictly true, since at least one male is needed. However, the above discussion indicates that we shouldn't expect more than one male, and this is indeed what is found in some species of mites: each brood involves a single male who inseminates his sisters and dies shortly afterwards. In some extreme cases, the male reproduces *inside* his mother, and it dies before it is even born (see Chapter 6 of [14] for a fascinating account). This is maleness reduced to its absolute minimum!

3 Bargaining

3.1 The bargaining problem

In a *bargaining problem*, two or more agents have to reach an agreement. Different agreements give different utilities or payoffs to the agents. The simplest example consist of two agents which have to agree on how divide a fixed amount of money, say 100 Swiss francs. Another example is a couple, living together, who has to bargain over the amount of housework that each one will assume.

However, people do not really bargain over things (even money) but on the *payoffs* or utilities that they derive from them. Therefore, we will first introduce the set of possible utilities that agents can obtain in a bargain, by agreeing to cooperate. This set has to satisfy certain properties, so we will start by recalling the definition of a convex set.

Definition 3.1. A *convex set* $X \in \mathbb{R}^I$ is a set that satisfies the property that, if $x, y \in X$, then the points in the line segment connecting x and y are also in X, i.e.

$$\alpha x + (1-\alpha)y \in X, \qquad \alpha \in [0,1]. \tag{3.1}$$

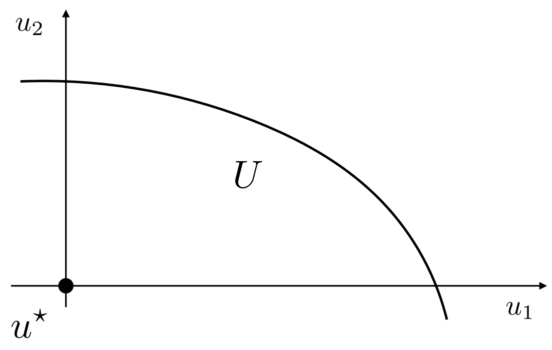

Figure 14. A utility possibility set in $\mathbb{R}^2$

X is *strictly convex* if every point on the line segment connecting x and y other than the endpoints is inside the interior of X.

We will also introduce a partial order relation in $\mathbb{R}^I$.

Definition 3.2. If $u, u' \in \mathbb{R}^I$, we will say that $u \geq u'$ if $u_i \geq u'_i$ for all $i = 1, \cdots, I$.

We can now define our basic object.

Definition 3.3. Given I agents, a *utility possibility set* (UPS) $U \subset \mathbb{R}^I$ is a *closed* and *convex* set which is bounded from above (i.e. there is a $b \in \mathbb{R}^I$ such that $u \leq b$ for all $u \in U$). It also satisfies the *free disposal property*: if $u' \leq u$ and $u \in U$, then $u' \in U$.

The free disposal property means that, if u is a possible agreement, then any point in which all agents lose utility is also a possible agreement (there is no limit to bad possibilities).

There is also a notion of Pareto optimality for UPS.

Definition 3.4. The *Pareto frontier* of the UPS U is the set of points $u = (u_1, \cdots, u_I) \in U$ for which there is no other $u' = (u'_1, \cdots, u'_I) \in U$ with $u'_i \geq u_i$ for every i and $u'_i > u_i$ for some i.

An example of a UPS is shown in Fig. 14. Its Pareto frontier is the upper curve bounding it from above.

In many situations, an UPS for I agents is obtained by considering a set of alternatives X together with utility functions for the agents

$$u_i : X \to \mathbb{R}, \qquad i = 1, \cdots, I. \tag{3.2}$$

Given some conditions on the utility functions and X one can show that $u(X)$ is a UPS, as the next example shows. We first recall the definition of a concave function, which will be useful later.

Definition 3.5. The function $f : A \to \mathbb{R}$ defined on the convex set A is *concave* if

$$f(\alpha x' + (1-\alpha)x) \geq \alpha f(x') + (1-\alpha) f(x), \qquad x, x' \in A, \quad \alpha \in [0,1]. \tag{3.3}$$

If the inequality is strict for all $x' \neq x$ and all $\alpha \in (0,1)$ we say that the function is *strictly concave*.

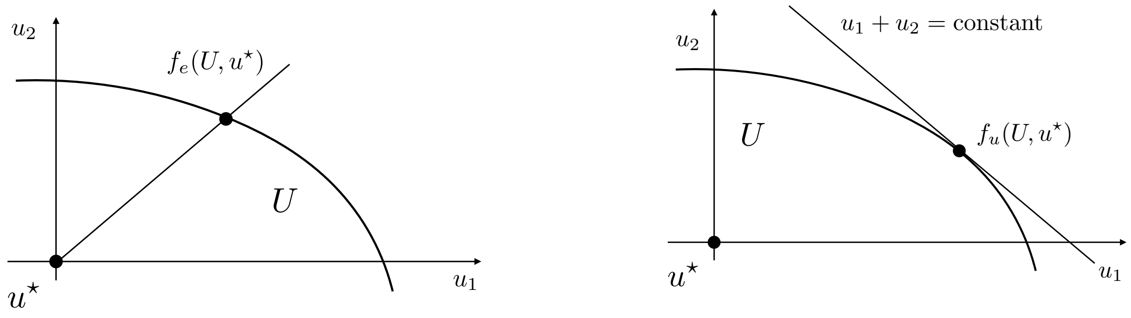

Figure 15. The egalitarian (left) and utilitarian (right) bargaining solutions.

Example 3.6. *Divide a dollar, I.* Let us consider the problem in which one divides a unit of money between two agents, so that the set of alternatives is

$$X = \{(x_1, x_2) \in \mathbb{R}^2 : x_1 + x_2 \leq 1\}. \tag{3.4}$$

x_i is the money that the i-th agent gets after the bargaining. Note that, in accordance with the free disposal property, we have assumed that the x_i can be negative. We now assume that the utility functions $u_i : X \to \mathbb{R}$ are of the form

$$u_i(x_1, x_2) = v_i(x_i), \qquad i = 1, 2, \tag{3.5}$$

where $v_i(x_i)$ is the utility function of the i-th agent. Note that we assume that the utility of the i-th agent does not depend on what the other agent gets. We also assume that $v_i(x_i)$ are continuous, strictly increasing, bounded above, and concave functions. These are standard assumptions for utility functions. With these assumptions, it can be shown that $u(X)$ is a UPS. □

In order to have a bargaining problem, we also need a *threat* or *status-quo* point $u^\star \in U$. This point represents the utility that agents will obtain if the bargaining does not succeed and there is a breakdown of cooperation. We will also assume $u^\star$ is in the interior of U, that U contains some point $u > u^\star$, and that the set of points

$$U_\star = \{u \in U : u \geq u^\star\} \tag{3.6}$$

is bounded. Note that the set $U_\star$ is convex and compact (verify it!).

Definition 3.7. A *bargaining problem* for I agents consists of a UPS $U \subset \mathbb{R}^I$ and a *threat*, or *status-quo* point $u^\star \in U$. A *bargaining solution* is a rule that assigns an element of U, $f(U, u^\star) \in U$, to every bargaining problem $(U, u^\star)$.

Example 3.8. *The egalitarian and utilitarian solutions.* In the *egalitarian solution*, all agents get the same utility, and one picks the vector in the boundary of U with all its coordinates equal:

$$f_e(U, u^\star) = (u, \cdots, u) \in \partial U. \tag{3.7}$$

In the *utilitarian solution* $f_u(U, u^\star)$, we choose the point $u \in U$ which maximizes the function

$$\sum_{i=1}^{I} u_i. \tag{3.8}$$

There is a unique such point if the set U is strictly convex (if U is just convex, and its boundary contains a straight line with slope 1, then the solution is not unique). The egalitarian and utilitarian solutions are shown in Fig. 15. □

3.2 The Nash bargaining solution

In 1950, Nash found a bargaining solution based on a set of *axioms* that it should satisfy. The axioms are the following.

1. *IUO* (independence of utility origins). The bargaining solution is *independent of utility origins* if, for any $\alpha = (\alpha_1, \cdots, \alpha_I) \in \mathbb{R}^I$, we have

$$f_i(U', u^\star + \alpha) = f_i(U, u^\star) + \alpha_i, \qquad i = 1, \cdots, I, \tag{3.9}$$

   for $U' = \{u + \alpha : u \in U\}$. This means that the solution does not depend on absolute scales of utility. Note that, if this property is satisfied, we can set $u^\star = 0$ without loss of generality.

2. *IUU* (independence of utility units). The bargaining solution is *independent of utility units* if, for any $\beta = (\beta_1, \cdots, \beta_I) \in \mathbb{R}^I_{>0}$, we have

$$f(\beta \cdot U, \beta \cdot u^\star) = \beta \cdot f(U, u^\star), \tag{3.10}$$

   where $\beta \cdot U = \{\beta \cdot u : u \in U\}$.

3. *P* (Pareto property). The bargaining solution satisfies the *Pareto property* if, for every U, $f(U, u^\star)$ is in the Pareto frontier of U.

4. *IIA* (independence of irrelevant alternatives). The bargaining solution is *independent of irrelevant alternatives* if, whenever $U' \subset U$ and $f(U, u^\star) \in U'$, we have that

$$f(U', u^\star) = f(U, u^\star), \tag{3.11}$$

   see Fig. 16 for an illustration of this axiom.

5. *S* (symmetry). The bargaining solution is symmetric if whenever U is a symmetric set (i.e. it is invariant under permutations of the axes) and $u^\star$ has all its entries identical, $f(U, u^\star)$ also has all its entries identical.

The axioms are all well motivated. IUO and IUU mean that utilities are equivalent up to linear transformations of the form $u \mapsto \beta u + \alpha$, $\beta > 0$. This equivalence property is part of Von Neumann–Morgenstern theory of expected utility: since utilities represent *preferences*, and at the same time they should be linear in probabilities, they are equivalent precisely under this type of transformations. The P axiom means that the bargaining solution cannot be improved. Finally, IIA means that the alternatives which were not optimal when they were available will be irrelevant when they are not available anymore, as long as you still have the optimal alternative

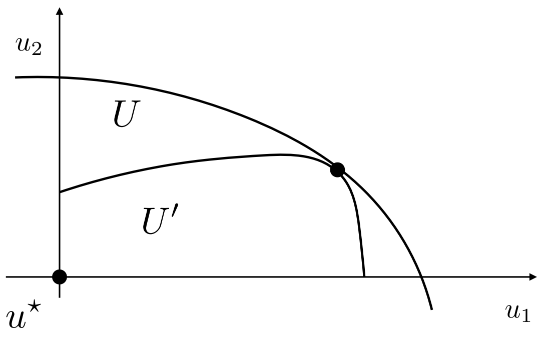

Figure 16. IIA: if the bargaining solution for $(U, u^\star)$ is the point in the boundary, then the bargaining solution for $(U', u^\star)$ should be the same point.

available. However, the axioms have also been criticized. Let us first develop the theory based on the axioms, and then we will come to the criticisms.

We note that axiom IIA is satisfied whenever the bargaining solution is obtained by maximizing a function on U, since if the maximum on U is attained on a point in U', the maximum on U' is obtained at the same point.

Definition 3.9. The *Nash bargaining solution* $f_{\mathrm{N}}(U, u^\star)$ is obtained by maximizing on $U_\star$ the function

$$N(u; u^\star) = \prod_{i=1}^{I} (u_i - u_i^\star), \tag{3.12}$$

or equivalently its logarithm

$$n(u; u^\star) = \sum_{i=1}^{I} \log(u_i - u_i^\star). \tag{3.13}$$

Let $\nu = (\nu_1, \cdots, \nu_I)$ be a vector of positive real numbers. We can normalize it in such a way that

$$\sum_{i=1}^{I} \nu_i = I. \tag{3.14}$$

The *generalized Nash bargaining solution* is obtained by maximizing on $U_\star$ the function

$$N_\nu(u; u^\star) = \prod_{i=1}^{I} (u_i - u_i^\star)^{\nu_i} \tag{3.15}$$

or its logarithm

$$n_\nu(u; u^\star) = \sum_{i=1}^{I} \nu_i \log(u_i - u_i^\star). \tag{3.16}$$

The Nash bargaining solution corresponds to the case $\nu_i = 1$, for all $i = 1, \cdots, I$.

We will now prove existence and uniqueness of the generalized Nash bargaining solution. To do so, we need some additional ingredients from the theory of (quasi) concave functions.

Definition 3.10. A function $f : A \to \mathbb{R}$ defined on the convex set $A \subset \mathbb{R}^n$ is quasiconcave if

$$f(\alpha x' + (1-\alpha)x) \geq \min\{f(x), f(x')\}, \qquad x, x' \in A, \quad \alpha \in (0,1). \tag{3.17}$$

The function is strictly quasiconcave if (3.17) holds with strict inequality whenever $x \neq x'$.

The following result is immediate.

Lemma 3.11. *If $f : A \to \mathbb{R}$ defined on a convex set is (strictly) concave, then it is (strictly) quasi-concave.*

We also have the following uniqueness result.

Lemma 3.12. *If $f : A \to \mathbb{R}$ defined on a convex set is strictly quasi-concave, it attains its maximum at a unique point of A.*

Proof: Indeed, if there were two points x, x' where it attains its maximum, then by (3.17) the point $\alpha x' + (1-\alpha)x$, which is also in A by convexity, would give an even larger value, a contradiction. □

We now need a criterium for strict quasi-concavity. We have the following

Theorem 3.13. *Let $f : A \to \mathbb{R}$ be twice differentiable. If, for every $x \in A$, the Hessian matrix $D^2 f(x)$ is negative definite in the subspace*

$$\{z \in \mathbb{R}^N,\ z \neq 0 : \nabla f(x) \cdot z = 0\}, \tag{3.18}$$

then f is strictly quasi-concave.

Proof: Suppose by contradiction that f is not strictly quasi-concave. This means that there are two points $x \neq y \in A$, with, say, $f(y) \geq f(x)$ and

$$f(\lambda^0 x + (1-\lambda^0)y) \leq f(x), \qquad \lambda^0 \in (0,1). \tag{3.19}$$

Consider the vector $z = x - y$, and the function:

$$\begin{aligned} g : [0,1] &\to \mathbb{R} \\ \lambda &\mapsto g(\lambda) = f(y + \lambda z). \end{aligned} \tag{3.20}$$

By assumption g is continuous and twice differentiable. We also have

$$g(1) = f(x) \leq g(0) = f(y), \qquad g(\lambda^0) \leq g(1) \tag{3.21}$$

Therefore, there must be an $\alpha \in (0,1)$ such that $g(\alpha) \leq g(\lambda)$ for all $\lambda \in [0,1]$. This due to the existence of a minimum of g on $[0,1]$, and the fact that there is at least an interior point λ^0 where g takes values which are bounded above by both $g(0)$ and $g(1)$. At this point we have $g'(\alpha) = 0$, $g''(\alpha) \geq 0$. But this means that there is an h such that

$$z \cdot \nabla f(x + \alpha z) = 0, \qquad z \cdot D^2 f(x + \alpha z) \cdot z \geq 0, \tag{3.22}$$

contradicting the assumptions. □

Lemma 3.14. *The generalized Nash bargaining solution exists and is unique.*

Proof: The existence follows from the fact that the function $N_\nu(u; u^\star)$ is continuous, therefore it achieves its maximum inside the compact set $U_\star$. We now show that $N_\nu(u; u^\star)$ is strictly quasi-concave, by using the criterium in theorem 3.13. We calculate the derivatives:

$$\frac{\partial}{\partial u_i} N_\nu = \nu_i (u_i - u_i^\star)^{\nu_i - 1} \prod_{j \neq i} (u_j - u_j^\star)^{\nu_j} = N_\nu \frac{\nu_i}{u_i - u_i^\star}. \tag{3.23}$$

We also have,

$$\begin{aligned} \frac{\partial^2}{\partial u_i^2} N_\nu &= \frac{\nu_i(\nu_i - 1)}{(u_i - u_i^\star)^2} N_\nu, \\ \frac{\partial^2}{\partial u_i \partial u_j} N_\nu &= \frac{\nu_i \nu_j}{(u_i - u_i^\star)(u_j - u_j^\star)} N_\nu, \qquad i \neq j. \end{aligned} \tag{3.24}$$

Therefore,

$$\begin{aligned} z \cdot D^2 N_\nu \cdot z &= N_\nu \left(\sum_{i=1}^N \frac{z_i^2 \nu_i (\nu_i - 1)}{(u_i - u_i^\star)^2} + \sum_{i \neq j} \frac{z_i z_j \nu_i \nu_j}{(u_i - u_i^\star)(u_j - u_j^\star)} \right) \\ &= N_\nu \left[\left(\sum_{i=1}^N \frac{z_i \nu_i}{(u_i - u_i^\star)} \right)^2 - \sum_{i=1}^N \frac{z_i^2 \nu_i}{(u_i - u_i^\star)^2} \right] \\ &= \frac{1}{N_\nu} (z \cdot \nabla N_\nu)^2 - N_\nu \sum_{i=1}^N \frac{z_i^2 \nu_i}{(u_i - u_i^\star)^2} \end{aligned} \tag{3.25}$$

Clearly, this is manifestly negative definite whenever $z \cdot \nabla N_\nu = 0$. We conclude that N_ν is strictly quasi-concave, and by lemma 3.12, its maximum is unique. □

Geometrically, the (generalized) Nash bargaining solution can be characterized as follows. Let $\hat{u} = f_N(U, u^\star)$, and let us consider the hypersurface

$$n_\nu(u; u^\star) = n_\nu(\hat{u}; u^\star) \tag{3.26}$$

passing through that point. The gradient $\nabla n_\nu(\hat{u}; u^\star)$ is orthogonal to the hypersurface (3.26). The plane defined by

$$\nabla n_\nu(\hat{u}; u^\star)(y - \hat{u}) = 0 \tag{3.27}$$

is the tangent plane to (3.26) at $\hat{u}$, and divides $\mathbb{R}^I$ in two halves. The *hyperplane theorem* (which we will not prove) says that the convex set U is contained in the lower half. Due to the strict concavity of $n_\nu(u; u^\star)$, the set

$$n_\nu(u; u^\star) \geq n_\nu(\hat{u}; u^\star) \tag{3.28}$$

is also convex, and it is contained in the upper half. The tangent plane is then a hyperplane separating the two sets, which have only one point in common, and it is called a *separating hyperplane*. We will use this result in order to prove Nash theorem. These geometric ingredients are shown in Fig. 17.

It is clear that the generalized Nash bargaining solution satisfies the axiom IIA, since it is obtained by maximizing a function on U. Since

$$n_\nu(\beta u + \alpha; \beta u^\star + \alpha) = n_\nu(u; u^\star) + \sum_{i=1}^I \nu_i \log(\beta_i) \tag{3.29}$$

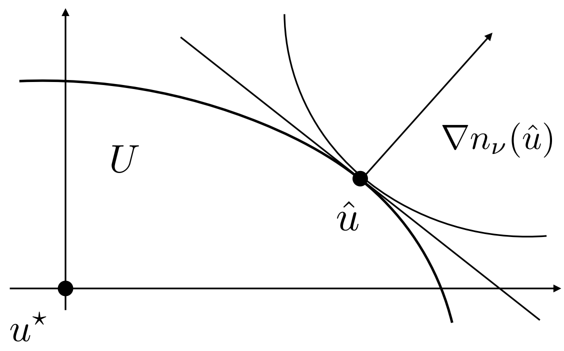

Figure 17. The geometry of the Nash bargaining solution.

which has the same maximum than $n_\nu(u;u^\star)$, the Nash bargaining solution satisfies IUO and IUU. It also satisfies the Pareto property (argue by contradiction: if there was a $u \in U$ with $u > f(U,u^\star)$, clearly u would lead to a higher value for the generalized Nash function $n_\nu(u;u^\star)$). Therefore, the generalized Nash bargaining solution satisfies P, IIA, IUO and IUU. Surprisingly, the converse is also true, and we have the following

Theorem 3.15. *(Nash) If a bargaining solution satisfies the properties P, IIA, IUO and IUU, then it is the generalized Nash bargaining solution.*

Proof. Since the solution satisfies IUO, we can set $u^\star = 0$, and we will simply denote $(U,0)$ by U. Let $f(\cdot)$ be a bargaining solution satisfying the conditions stated in the theorem. Let us first consider the UPS given by

$$U'' = \{u \in \mathbb{R}^I : \sum_{i=1}^{I} u_i \le I\}. \tag{3.30}$$

By P, $f(U'')$ must be in the boundary of this set, and we will denote

$$f(U'') = (\nu_1, \cdots, \nu_I), \qquad \sum_{i=1}^{I} \nu_i = I. \tag{3.31}$$

Let us now consider the UPS U, and let $\hat{u} = f_{\mathrm{N}}(U) \in U$ be the generalized Nash bargaining solution with exponents ν. Consider now the set

$$U' = \{u \in \mathbb{R}^I : \sum_{i=1}^{I} \nu_i \frac{u_i}{\hat{u}_i} \le I\}. \tag{3.32}$$

The gradient of the Nash function

$$\nabla n_\nu(\hat{u}) = \left(\frac{\nu_1}{\hat{u}_1}, \cdots, \frac{\nu_I}{\hat{u}_I}\right) \tag{3.33}$$

is orthogonal to the separating hyperplane (3.27) passing through $\hat{u}$. The points $u \in U$ satisfy then,

$$(u - \hat{u}) \cdot \nabla n_\nu(\hat{u}) \le 0, \tag{3.34}$$

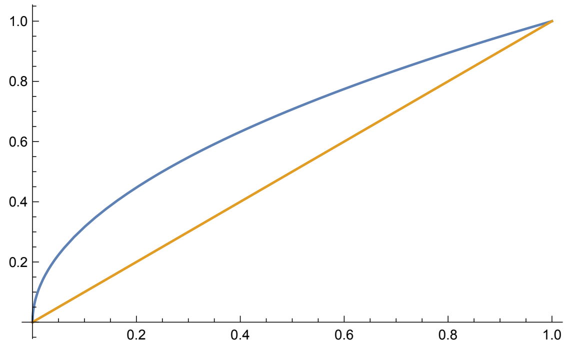

Figure 18. Two different utility functions: the upper line is the function $u(x) = x^{1/2}$, while the lower line is the linear utility function $u(x) = x$. In the first case, a small increase in x gives a higher increase of utility as compared to the second case. It corresponds to a "poorer" or more-risk averse agent.

since points under the plane form an angle $\pi > \theta > \pi/2$ with the gradient. It follows that $U \subset U'$.

Consider now the affine transformation $\tau : \mathbb{R}^I \to \mathbb{R}^I$ defined by

$$\tau(x_1, \cdots, x_I) = \left(\frac{\hat{u}_1}{\nu_1} x_1, \cdots, \frac{\hat{u}_I}{\nu_I} x_I \right). \tag{3.35}$$

It maps U'' into U':

$$\tau(U'') = U'. \tag{3.36}$$

Therefore, by IIU, we must have

$$f(U') = f(\tau(U'')) = \tau(f(U'')) = \hat{u} = f_{\mathrm{N}}(U). \tag{3.37}$$

On the other hand, since $U \subset U'$ and $\hat{u} \in U$, we have, by IIA

$$f(U) = f(U') = f_{\mathrm{N}}(U), \tag{3.38}$$

which is what we wanted to prove. □

Corollaire 3.1. If a bargaining solution satisfies the properties P, IIA, IUO, IUU and S, then it is the Nash bargaining solution, i.e. $\nu_i = 1$ for all $i = 1, \cdots, I$.

3.3 Examples

We now discuss various examples of Nash' bargaining solution.

Example 3.16. *Divide a dollar, II.* We consider again the problem of dividing a dollar between two persons. Let us assume that the utilities are given by

$$u_i(x) = x^{\nu_i}, \qquad i = 1, 2, \tag{3.39}$$

where $0 < \nu_i \leq 1$ and $x \in [0, 1]$. These functions are concave, and we note that ν_i characterizes the preferences of the agent w.r.t. money. In the linear case, $\nu = 1$, we are dealing with an agent for which money has always the same utility (a "rich" agent). In contrast, an agent with $\nu < 1$ will value more a small increase in the amount of money (a "poor" agent). The two different utility functions can be also regarded as two different attitudes towards risk, with smaller ν

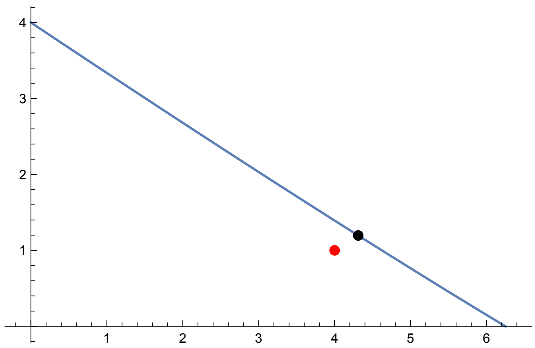

Figure 19. The UPS for a Cournot duopoly with $a = 6$ and $b = 1$, and different production costs $c_1 = 1$, $c_2 = 2$. The point in the frontier is the Nash bargaining solution for this problem. The point in the interior is $u^\star$, which corresponds to the Nash equilibrium of the duopoly.

indicating greater risk aversion, see Fig. 18 for an illustration. The generalized Nash solution is obtained by maximizing

$$x_1^{\nu_1 \alpha} x_2^{\nu_2 \beta} \tag{3.40}$$

satisfying the constraint $x_1 + x_2 = 1$. A simple calculation gives

$$x_1 = \frac{\nu_1 \alpha}{\nu_1 \alpha + \nu_2 \beta}, \qquad x_2 = \frac{\nu_2 \beta}{\nu_1 \alpha + \nu_2 \beta}. \tag{3.41}$$

An interesting property of this solution is the following: even in the case of a symmetric Nash solution with $\alpha = \beta = 1$, the outcome is *not* egalitarian, due to the different preferences of the agents: the agent who is more risk-averse (or "poorer") will have a smaller share. Therefore, asymmetric results in a bargaining are not necessarily the result of different bargaining power, but also of different preferences. □

Example 3.17. *Collusion in a Cournot duopoly.* As we saw earlier in the course, the Nash equilibrium in a duopoly is not Pareto efficient. We will now analyze bargaining between two Cournot duopolists, who want to reach a Pareto-efficient agreement. We will assume that they reach a symmetric Nash solution.

Let us first consider the situation in which they have equal production costs. The boundary of the UPS is the contract curve, but in utility space. By using (2.29), it is easy to see that the Pareto frontier of the UPS is given by the straight line

$$u_1 + u_2 = \frac{(a-c)^2}{4b}. \tag{3.42}$$

The threat point is the Nash equilibrium solution, given by

$$u^\star = \left(\frac{(a-c)^2}{9b}, \frac{(a-c)^2}{9b} \right). \tag{3.43}$$

It is clear that, in this case, the Nash bargaining solution is the midpoint of the Pareto frontier (as measured w.r.t. $u^\star$), i.e.

$$f_{\mathrm{N}}(U, u^\star) = \left(\frac{(a-c)^2}{8b}, \frac{(a-c)^2}{8b} \right). \tag{3.44}$$

We will now consider the more interesting situation in which the duopolists have different production costs. The threat point is in this case

$$u^\star = \left(\frac{(a - 2c_1 + c_2)^2}{9b}, \frac{(a - 2c_2 + c_1)^2}{9b} \right). \tag{3.45}$$

The contract curve is given by (2.39), and it has two branches in the first quadrant. The lowest brach intersects the axes in the (q_1, q_2) plane at the points

$$\left(\frac{a - c_1}{2b}, 0 \right) \quad \text{and} \quad \left(0, \frac{a - c_2}{2b} \right), \tag{3.46}$$

while the frontier of the UPS intersects the axes at the points

$$\left(\frac{(a - c_1)^2}{4b}, 0 \right) \quad \text{and} \quad \left(0, \frac{(a - c_2)^2}{4b} \right), \tag{3.47}$$

The Nash bargaining solution can be obtained by finding the extrema of $n(u; u^\star)$ as a function of (q_1, q_2), i.e. by solving

$$\frac{\partial n(u; u^\star)}{\partial q_1} = \frac{\partial n(u; u^\star)}{\partial q_2} = 0. \tag{3.48}$$

The analytic solution to this equation is too complicated to write down, but it is easy to obtain numerical values. As an example, let us consider a duopoly with $a = 6$, $b = 1$, $c_1 = 1$ and $c_2 = 2$. The UPS is shown in Fig. 19, together with the Nash bargaining solution, which occurs at

$$(u_1, u_2) \approx (4.31, 1.20). \tag{3.49}$$

Note that it is the agent with a lower cost who gets the better share, as one could have expected.

There is however an even better deal for the two Cournot duopolists: the agent with a lower cost can become a monopolist and pay a compensation to the second agent for not entering into the market. To understand this possibility, we have to enlarge the possibility set by adding points of the form $(u_1 - r, r_2 + r)$, where r represents a monetary transfer between the agents (a compensation, or a "bribe" if the procedure is illegal). When this is done, we obtain a new UPS which is given by the half-plane below a straight line of slope -1. This line passes through the most outward point of the original UPS.

In the case at hand, the frontier of the new UPS is the line of slope -1 going through the points

$$\left(\frac{(a - c_1)^2}{4b}, 0 \right) \quad \text{and} \quad \left(0, \frac{(a - c_1)^2}{4b} \right). \tag{3.50}$$

The Nash bargaining solution for the new UPS is

$$(u_1, u_2) = \left(\frac{(a - c_1)^2}{8b} + \frac{u_1^\star - u_2^\star}{2}, \frac{(a - c_1)^2}{8b} + \frac{u_2^\star - u_1^\star}{2} \right). \tag{3.51}$$

Note that this agreement means that the first agent has effectively a monopoly, with a monopolistic profit of

$$\frac{(a - c_1)^2}{4b}. \tag{3.52}$$

Part of this profit, namely,

$$\frac{(a - c_1)^2}{8b} + \frac{u_2^\star - u_1^\star}{2} \tag{3.53}$$

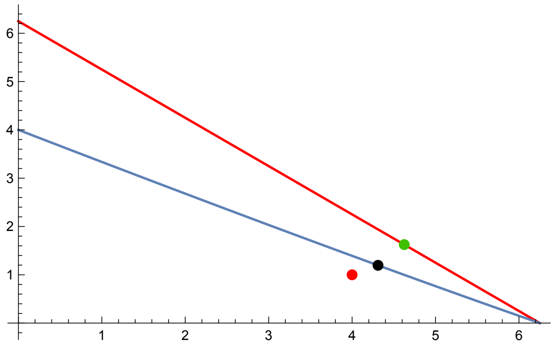

Figure 20. If we allow for monetary compensations (or "bribes"), the UPS becomes bigger. Its boundary is the straight line at the top, which we draw here for $a = 6$, $b = 1$, $c_1 = 1$ and $c_2 = 2$. The green point in the new frontier is the Nash solution (3.54), in which the first agent is a monopolistic producer and pays a compensation to the second agent to stay out of the market.

is paid as a compensation to the second agent. In the numerical example we have been considering, this corresponds to

$$(u_1, u_2) = (4.625, 1.125). \tag{3.54}$$

Note that this is clearly an improvement over the previous arrangement, as shown in Fig. 20.

Example 3.18. *Bargaining over the housework.* One of the persistent gender gaps in modern society concerns the time devoted to housework. Many studies have found, surprisingly, that the gap turns out to depend very weakly on other sociological variables (for example, it is largely independent of social class). In average, men only do one third of housework, while women do two-thirds (see chapter 8 of [23] for such a study, which discusses various hypothesis in order to explain the persistence of the gap).

It is natural to model the division of housework as a result of a bargaining process, and it has been suggested that the asymmetries observed in society are due to different bargaining powers. Here we will present a simple model, due to K. Binmore [2, 3], which uses the Nash bargaining solution to explain the unequal division of housework as a result of different preferences.

We will then consider two agents, Misha and Willy, who live together and have to bargain over how much time each of them is going to dedicate to the housework. They turn out to have different views about the amount of hours that have to be invested on housework per week. For Misha, $\mathcal{M}$ hours per week are enough, while Willy thinks that $\mathcal{W}$ hours per week are needed, and $\mathcal{M} < \mathcal{W}$. Let us now denote by t_m, t_w the number of hours per week that they devote to housework. We normalize units so that, when all the housework has been done, they enjoy a payoff of one unit. The payoffs of Misha and Willy are then

$$u_M = \begin{cases} 1 - m t_m, & \text{if } t_m + t_w \geq \mathcal{M}, \\ -m t_m, & \text{otherwise}, \end{cases}, \qquad u_W = \begin{cases} 1 - w t_w, & \text{if } t_m + t_w \geq \mathcal{W}, \\ -w t_w, & \text{otherwise}, \end{cases} \tag{3.55}$$

where $m > w$, i.e. Misha suffers more by doing housework than Willy. We will also assume that $m\mathcal{M} < 1$, $w\mathcal{W} < 1$. Note that, if not enough time is devoted to housework, they do not achieve their payoff of one unit.

What is the status quo? Before living together, they were devoting $\mathcal{M}$ and $\mathcal{W}$ hours to the housework, respectively, so that the status quo is

$$(u_M^\star, u_W^\star) = (1 - m\mathcal{M}, 1 - w\mathcal{W}). \tag{3.56}$$

Let us now assume that both agree that the total number of hours should be $\mathcal{W}$, as proposed by Willy. We will see in a moment that working less hours than that is less efficient. Therefore, we have

$$t_m + t_w = \mathcal{W}. \tag{3.57}$$

What is the frontier of the utility set? Since

$$u_M = 1 - mt_m, \qquad u_W = 1 - w(\mathcal{W} - t_m), \tag{3.58}$$

we can eliminate t_m from these equations and find

$$mu_W + wu_M = w + m(1 - w\mathcal{W}), \tag{3.59}$$

which is a straight line. Let us now note that, if the number of hours that are devoted to housework is $\mathcal{H}$ instead of $\mathcal{W}$, with $\mathcal{M} < \mathcal{H} < \mathcal{W}$, then the payoff of Willy is $-wt_w$, and the Pareto frontier is at

$$mu_W + wu_M = w - mw\mathcal{H}. \tag{3.60}$$

But

$$w + m(1 - w\mathcal{W}) > w - mw\mathcal{H}, \tag{3.61}$$

since

$$m(1 - w\mathcal{W} + w\mathcal{H}) > 0. \tag{3.62}$$

Therefore, the Pareto frontier (3.60) is lower than the one in (3.59), and the corresponding utility set is smaller. However, in a bargaining problem, efficiency requires that we look for a solution in the Pareto frontier of the largest possible utility set, and Willy and Misha will agree to work $\mathcal{W}$ hours in total.

We will assume that Misha and Willy have equal bargaining power. The Nash bargaining solution is obtained by maximizing the function

$$N(u_M, u_W) = (u_W - 1 + w\mathcal{W})(u_M - 1 + m\mathcal{M}) \tag{3.63}$$

on the frontier of the utility set. We can e.g. express everything in terms of u_m, and obtain the function

$$f(u_M) = \frac{w}{m}(1 - u_M)(u_M - 1 + m\mathcal{M}). \tag{3.64}$$

Its critical point is

$$f'(u_M) = 0 \Rightarrow u_M^{\mathrm{N}} = 1 - \frac{m\mathcal{M}}{2}, \tag{3.65}$$

so Misha works $\mathcal{M}/2$ hours, while Willy works the remaining hours. For example, if $\mathcal{W} = 2\mathcal{M}$ (i.e. Willy thinks one should devote double time to housework than Misha), then Misha works $\mathcal{M}/2 = \mathcal{W}/4$ hours, while Willy works

$$\mathcal{W} - \frac{\mathcal{W}}{4} = \frac{3}{4}\mathcal{W}. \tag{3.66}$$

In other words, Willy will end up doing three quarters of the housework, and Misha only one quarter.

What is the moral of this? Note that we have not assumed that Misha and Willy have different bargaining power. What has introduced the asymmetry in the outcome are the different *preferences* that Misha and Willy have concerning the number of hours that should be devoted to housework. It is similar to the divide-the-dollar example: even if the two agents have the same bargaining power, the one which is more risk-averse will have smaller share. □

3.4 Other bargaining solutions

The Nash bargaining solution (and the axioms that determine it uniquely) have been criticized on various grounds.

The first criticism is that the solution has no "microfoundations", i.e. it is not clear that it emerges as the solution of a non-cooperative game. This was noted by Nash himself, and has led to the so-called "Nash program", in which bargaining solutions are not justified by set of axioms, but rather by being solutions of a game. We will come back to this when we study dynamic games

Criticizing Nash solution means criticizing the underlying axioms. We will now consider problems with these axioms.

Example 3.19. *The Rolls–Royce paradox.* Let us consider two agents bargaining to agree on two numbers p, q with $p + q \leq 1$. If they agree, the first agent wins a prize in a lottery with probability p, and the second agent wins a prize in a *different* lottery with probability q. If they do not agree, they do not get anything. Let us consider a first bargaining problem in which the two lotteries have the same prize, say a bicycle. Then, it is clear that the result of a symmetric bargaining is

$$p = q = \frac{1}{2}. \tag{3.67}$$

Let us consider now a second bargaining problem, in which the prize for the first agent is a Rolls–Royce, while the one for the second agent is a bicycle. Intuitively, it seems clear that, since the expected utility in the second situation is bigger, one should have $p < 1/2$. However, IUU tells us that the second bargaining problem should have the same outcome than the first one. Namely, let u, $v < u$ be the utilities of a Rolls–Royce and a bicycle for the first agent, respectively. The expected utility of the first agent in the second bargaining, for the outcome (p, q), est

$$pu = \frac{u}{v} \cdot pv, \tag{3.68}$$

which is a linear transformation of the utility in the first bargaining problem. By IUU, the outcome of the second bargaining problem should be the same as for the first one. In actual, experimental bargainings, the result of the second one is different from the first one, with $p < 1/2$[5]. □

Another axiom which has been criticized is the IIA. One counterintuitive aspect of IIA is the following. Let us consider the two UPS U and U' shown in Fig. 21: U has the shape of a triangle, and U' the shape of a trapezoid, and $U' \subset U$. In the bargaining problem $(U, u^\star)$, the symmetric Nash bargaining solution is indicated by the black dot. If we assume the IIA axiom, the bargaining problem $(U', u^\star)$ has the same bargaining solution thatn $(U, u^\star)$. However, this is somewhat counter-intuitive, since in the second problem the second agent has less bargaining possibilities. In fact, in the second problem the second agent gets the maximum of what she can get. One would expect that, in the second problem, the second agent does worst than in the first one.

Kalai and Smorodinsky introduced a bargaining solution which incorporates this objection, and therefore gives up axiom IIA.

Definition 3.20. Let $(U, u^\star)$ be a bargaining problem. For each agent $i = 1, \cdots, I$, let

$$u_i^\flat = \max\{u_i : u \in U, u \geq u^\star\}. \tag{3.69}$$

[5] The Rolls–Royce paradox is further discussed in [10], pp. 62-63, and [25], pp. 78-80.

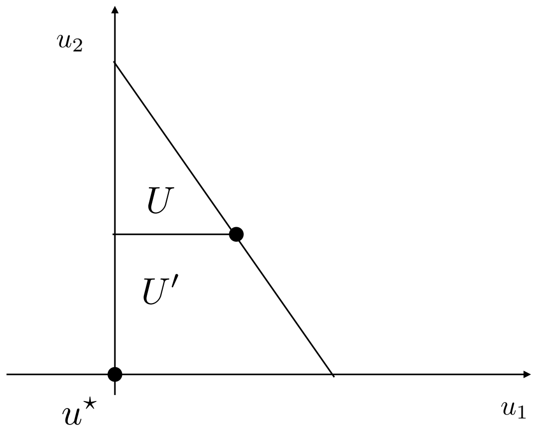

Figure 21. Two UPS with the same Nash bargaining solution.

This is the maximal utility that the i-th agent can get in the bargaining problem. Consider the point,

$$u^\flat = (u_1^\flat, \cdots, u_I^\flat). \tag{3.70}$$

This point is typically not even in the UPS. It is sometimes called the *ideal point* of the bargaining problem $(U, u^\star)$ (Ken Binmore calls it the "Utopian point" [2]).

Definition 3.21. Let $(U, u^\star)$ be a bargaining problem, and let $u^\flat$ be its ideal point. Let us consider the straight line joining $u^\star$ and $u^\flat$, $L(u^\star, u^\flat)$. Then, the *Kalai–Smorodinsky solution* to the bargaining problem $f^{\mathrm{KS}}(U, u^\star)$ is the maximal element of U on this line, with respect to the partial order relation on $\mathbb{R}^I$.

The properties of this solution will be studied in the exercises.

4 Evolutionary games and stable strategies

In *evolutionary game theory*, we use the tools of game theory to model the Darwinian evolution of traits in populations. Conversely, ideas and tools of dynamical systems (which are used in the study of population dynamics) will shed new light on the concept of equilibrium in "classical" game theory.

4.1 Evolutionary stability

We will now focus on two-player *symmetric* games.

Definition 4.1. A game G in normal form with two players is called *symmetric* if both players have the same set of strategies, denoted by

$$S = \{s_1, \cdots, s_n\}, \tag{4.1}$$

and in addition the payoffs satisfy

$$u_2(s_i, s_j) = u_1(s_j, s_i), \qquad i, j = 1, \cdots, n. \tag{4.2}$$

We will denote the payoffs of the first player in a symmetric game by

$$u_1(s_i, s_j) = u_{ij}, \tag{4.3}$$

so that the different entries form a matrix

$$U = (u_{ij})_{i,j=1,\cdots,n}. \tag{4.4}$$

The payoffs of the second player are then given by the transpose of U. In view of symmetry, we only have to consider one payoff function, say for the first player, and we will denote

$$u(\sigma_1, \sigma_2) = u_1(\sigma_1, \sigma_2), \tag{4.5}$$

where $\sigma_{1,2}$ are arbitrary mixed strategies.

Definition 4.2. Let G a symmetric game of two players in normal form. A Nash equilibrium of G is called *symmetric* if both players play the same (in general mixed) strategy at equilibrium.

We have the following result, due to Nash:

Theorem 4.3. *Let G a symmetric game of two players in normal form. Then, it has a symmetric Nash equilibrium.*

We want to use symmetric, two-player games to describe interactions in populations. We assume that in the population there are various possible strategies $s_1, \cdots, s_n$. We will consider mixed strategies of the form,

$$\sigma = \sum_{i=1}^{n} p_i s_i. \tag{4.6}$$

Remark 4.4. In "classical" game theory, a mixed strategy represents a randomization of strategies in the same individual. In evolutionary game theory, this can represent various things. It can represent a population where all individuals are identical and randomize in the way prescribed by the probabilities p_i. But it can also represent a mixed population in which p_i is the proportion of members adopting the pure strategy s_i. Finally, it can also represent an intermediate situation between these two. In fact, as emphasized by Richard Dawkins in Chapter 7 of [9], this is not that important: when doing evolutionary game theory, we want to think in terms of strategies, not in terms of individuals.

We will sometimes identify a mixed strategy with an element in the simplex

$$\Delta^{n-1} = \{(p_1, \cdots, p_n) \in \mathbb{R}^n : \sum_{i=1}^{n} p_i = 1\}. \tag{4.7}$$

If an individual with an strategy s_i interacts with a random member of the population characterized by the mixed strategy σ, its payoff will be

$$u(s_i, \sigma) = \sum_{j=1}^{n} u_{ij} p_j. \tag{4.8}$$

Similarly, the payoff of a mixed strategy σ playing against a mixed strategy τ is

$$u(\sigma, \tau) = \sigma^T U \tau = \sum_{i,j=1}^{n} \sigma_i u_{ij} \tau_j. \tag{4.9}$$

Definition 4.5. A mixed strategy $\sigma^* \in \Delta^{n-1}$ is an *evolutionarily stable strategy* (ESS) if, for any $\sigma \neq \sigma^*$, there is an $\bar{\epsilon} > 0$ such that, for any $0 < \epsilon \leq \bar{\epsilon}$,

$$(\sigma^*)^T U \left[(1-\epsilon)\sigma^* + \epsilon\sigma\right] > \sigma^T U \left[(1-\epsilon)\sigma^* + \epsilon\sigma\right]. \tag{4.10}$$

The condition for an ESS has an easy interpretation. Let us replace a (small) fraction ϵ of the population σ^* by a "mutant" type σ, to obtain the new population

$$\Sigma_\epsilon = (1-\epsilon)\sigma^* + \epsilon\sigma. \tag{4.11}$$

Can the population σ^* be "invaded" by a small number of mutants, adopting the strategy σ? To answer this, we have to compare payoffs. The payoff of the mutant in the population is precisely the r.h.s. of (4.10),

$$u(\sigma, \Sigma_\epsilon) = \sigma^T U \left[(1-\epsilon)\sigma^* + \epsilon\sigma\right]. \tag{4.12}$$

The payoff of a non-mutant in the population is the l.h.s.

$$u(\sigma^*, \Sigma_\epsilon) = (\sigma^*)^T U \left[(1-\epsilon)\sigma + \epsilon\sigma'\right]. \tag{4.13}$$

Therefore, σ^* is an ESS if a mutant in the population does worse than a non-mutant, in terms of payoffs.

A slightly easier characterization of an ESS is given by the following proposition.

Proposition 4.6. A strategy profile σ^* in a symmetric game G is an ESS if and only if the following conditions hold.

1. *Equilibrium condition.* For all $\sigma \in \Delta^{n-1}$,

$$u(\sigma^*, \sigma^*) \geq u(\sigma, \sigma^*). \tag{4.14}$$

2. *Stability condition.* For all $\sigma \neq \sigma^* \in \Delta^{n-1}$, satisfying $u(\sigma, \sigma^*) = u(\sigma^*, \sigma^*)$, we have

$$u(\sigma^*, \sigma) > u(\sigma, \sigma). \tag{4.15}$$

Proof: first of all, we rewrite the condition for an ESS as follows:

$$(1-\epsilon)\left\{u(\sigma^*, \sigma^*) - u(\sigma, \sigma^*)\right\} + \epsilon\left\{u(\sigma^*, \sigma) - u(\sigma, \sigma)\right\} > 0. \tag{4.16}$$

In view of this, the proof is elementary. Let us prove the implication $\Rightarrow$. Assume σ^* is an ESS. If (1) is not satisfied, there is some $\sigma \neq \sigma^*$ such that

$$u(\sigma^*, \sigma^*) < u(\sigma, \sigma^*). \tag{4.17}$$

Then, there is some $\hat{\epsilon} > 0$ such that, for $0 < \epsilon \leq \hat{\epsilon}$,

$$(1-\epsilon)u(\sigma^*, \sigma^*) + \epsilon u(\sigma^*, \sigma) < (1-\epsilon)u(\sigma, \sigma^*) + \epsilon u(\sigma, \sigma), \tag{4.18}$$

or, equivalently,

$$(1-\epsilon)\left\{u(\sigma^*, \sigma^*) - u(\sigma, \sigma^*)\right\} + \epsilon\left\{u(\sigma^*, \sigma) - u(\sigma, \sigma)\right\} < 0, \tag{4.19}$$

which contradicts the definition of ESS. To check (2), let us assume that $u(\sigma^*, \sigma^*) = u(\sigma, \sigma^*)$ but

$$u(\sigma^*, \sigma) \leq u(\sigma, \sigma). \tag{4.20}$$

It follows that

$$(1-\epsilon)\left\{u(\sigma^*, \sigma^*) - u(\sigma, \sigma^*)\right\} + \epsilon\left\{u(\sigma^*, \sigma) - u(\sigma, \sigma)\right\} \leq 0, \tag{4.21}$$

which contradicts again the definition of ESS.

We now prove the opposite implication $\Leftarrow$. Assume that (1) and (2) are true. Let $\sigma \neq \sigma^*$. If (1) holds with a strict inequality, then (4.16) clearly holds if ϵ is small enough. If (1) holds with equality, then we must have the inequality $u(\sigma^*, \sigma) > u(\sigma, \sigma)$, and then (4.16) holds for all $\epsilon > 0$. $\square$

A very important consequence of the above proposition is that, due to (1), an ESS *must be* a symmetric Nash equilibrium of the symmetric game. However, the reciprocal is not true, since an ESS has to satisfy in addition the stability condition (2). Therefore, in the case of symmetric two-player games, the concept of ESS is a strict *refinement* of symmetric Nash equilibria, which takes into account stability issues.

A strategy profile σ^* which is not an ESS can be "invaded" by a mutant population σ. This happens for example if condition (1) is violated and

$$u(\sigma, \sigma^*) > u(\sigma^*, \sigma^*). \tag{4.22}$$

In this case, σ^* is not even a symmetric Nash equilibrium, since (4.22) means that σ is a better response to σ^* than σ^* itself. It might also happen that condition (1) holds but condition (2) is not satisfied. As we will see in examples, in many symmetric games with a symmetric Nash equilibrium (σ^*, σ^*), there exists an invading mutant strategy $\sigma \neq \sigma^*$ violating (2), i.e. a strategy σ with $u(\sigma, \sigma^*) = u(\sigma^*, \sigma^*)$ but such that

$$u(\sigma, \sigma) \geq u(\sigma^*, \sigma). \tag{4.23}$$

	Hawk	Dove
Hawk	$\left(\frac{v-c}{2}, \frac{v-c}{2}\right)$	$(v, 0)$
Dove	$(0, v)$	$\left(\frac{v}{2}, \frac{v}{2}\right)$

Figure 22. The payoffs in the Hawk-Dove game

Example 4.7. *The Hawk-Dove game.* The most famous game in evolutionary game theory is the Hawk-Dove game, which was introduced to model animal conflict, and in particular to address the following question: why is it that many animal conflicts are of limited intensity, instead of being of a "total war" type? In the Hawk-Dove game, there are two strategies: the Hawk strategy, which we will denote by H, involves fighting your rival for a resource with value v. The cost of a fight (due for example to injuries) will be denoted by c. The Dove strategy, which we will denote by D, is to look tough but flee when a fight is about to start. When two Hawks meet, they fight, and there is a probability 1/2 for each of them to win or to lose. The average payoff is

$$u(H,H) = \frac{v-c}{2}, \tag{4.24}$$

When a Hawk meets a Dove, the latter escapes and the Hawk gets everything, i.e.

$$u(H,D) = v, \qquad u(D,H) = 0. \tag{4.25}$$

When two Doves meet, they share the resource in equal parts, so that

$$u(D,D) = \frac{v}{2}. \tag{4.26}$$

The payoffs in this game are represented in Fig. 22.

Let us analyze the ESS in this model. There are three different cases. If $v > c$, D is a strictly dominated strategy, and the Nash equilibrium is (H,H). It is easy to see that condition (2) is always satisfied in this case, and H is the only ESS.

When $v < c$, there are no symmetric Nash equilbria in pure strategies. Indeed,

$$BR_1(H) = \{D\}, \qquad BR_1(D) = \{H\}, \tag{4.27}$$

and the same holds for the second player. Therefore, the Nash equilibria are (H,D), (D,H), none of which are symmetric. There is however a symmetric Nash equilibrium in *mixed* strategies. Let us consider the mixed strategy

$$\sigma^* = p^* H + (1-p^*)D. \tag{4.28}$$

Applying the Fundamental Theorem, we get

$$u(H,\sigma^*) = u(D,\sigma^*), \tag{4.29}$$

which reads

$$p^* \frac{v-c}{2} + (1-p^*)v = (1-p)\frac{v}{2}. \tag{4.30}$$

Solving, we obtain

$$p^* = \frac{v}{c}. \tag{4.31}$$

Let $\sigma = qH + (1-q)D$. Then,

$$u(\sigma,\sigma^*) = qu(H,\sigma^*) + (1-q)u(D,\sigma^*) = u(H,\sigma^*) = u(\sigma^*,\sigma^*). \tag{4.32}$$

For σ^* to be an ESS, we must verify in addition condition (2) in Proposition 4.6, i.e. we must have

$$u(\sigma^*,\sigma) > u(\sigma,\sigma) \tag{4.33}$$

for all $\sigma \neq \sigma^*$. We calculate

$$\begin{aligned} u(\sigma^*,\sigma) &= u(p^*H + (1-p^*)D, qH + (1-q)D) = -qv + \frac{v}{2c}(v+c), \\ u(\sigma,\sigma) &= u(qH + (1-q)D, qH + (1-q)D) = -\frac{cq^2}{2} + \frac{v}{2}. \end{aligned} \tag{4.34}$$

Therefore,

$$u(\sigma^*,\sigma) - u(\sigma,\sigma) = \frac{1}{2c}(v-qc)^2 > 0. \tag{4.35}$$

We conclude that (4.28), with $p^* = v/c$, is a mixed ESS. This is not a very nice world to live in. Indeed, the average payoff for

$$\sigma = pH + (1-p)D \tag{4.36}$$

is

$$u(\sigma,\sigma) = p^2 u(H,H) + p(1-p)\left(u(H,D) + u(D,H)\right) + (1-p)^2 u(D,D) = \frac{v}{2} - \frac{c}{2}p^2. \tag{4.37}$$

This is maximized when $p = 0$, not when $p = p^*$. Life would be much better in a society with no Hawks, but this society would not be in equlibrium. A society with only doves could be easily invaded by Hawks. In a dynamical situation, we would expect the fraction of Hawks in the population to grow from $p = 0$ to $p = p^*$. Conversely, a population of Hawks can be invaded by a small group of Doves, and we would expect the number of Doves to grow until they reach the equilibrium proportion $1-p^*$ in the population. We will see that this dynamical expectations are indeed fulfilled when the evolution of the populations is described by the replicator equation. □

An important property of Nash equilibria is that they always exist. In fact, very often there are too many. Moreover, it can be shown that, in a symmetric two-player game, there is always a symmetric Nash equilibrium (this is also a result due to Nash). The concept of ESS refines the concept of symmetric Nash equilibrium, but ESS do not always exist. A classical example is the RSP game, as we will now see.

Example 4.8. *Absence of ESS in the RSP game.* The only Nash equilibrium in the RSP game is the mixed strategy

$$\sigma^* = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right), \tag{4.38}$$

which is also symmetric. Therefore this is our only candidate for an ESS. Note that

$$u(\sigma^*, \sigma^*) = 0. \tag{4.39}$$

It this strategy was an ESS, we would have, by (4.15) that for every strategy $\sigma \neq \sigma^*$ with

$$u(\sigma, \sigma^*) = 0, \tag{4.40}$$

the following inequality should hold

$$u(\sigma^*, \sigma) > u(\sigma, \sigma). \tag{4.41}$$

However, the strategy

$$\sigma = (1, 0, 0) \tag{4.42}$$

satisfies (4.40), but

$$u(\sigma^*, \sigma) = u(\sigma, \sigma) = 0. \tag{4.43}$$

Therefore, it violates the stability condition (4.15). □

4.2 Replicator dynamics

The concept of ESS introduces a notion of stability, and therefore suggests an underlying dynamical origin. Roughly speaking, an equilibrium is stable if, when the system is perturbed away from it, it comes back in finite time. We will now introduce a natural dynamical system for the frequencies of strategies. Let

$$x(t) = (x_1(t), \cdots, x_n(t)) \in \Delta^{n-1} \tag{4.44}$$

be the vector of strategy frequencies in a population at a time t, where $x_q(t)$ is the frequency of the strategy s_q. If $\nu(t)$ is the total size of the population,

$$\nu_q(t) = x_q(t)\nu(t) \tag{4.45}$$

is the total number of individuals adopting the strategy $x_q(t)$ (or the total number of times that the strategy s_q has been adopted). We will now interpret strategies as inherited genetic traits, and the payoffs of the game in terms of Darwinian fitness: the offspring of an individual inherits the strategy of the parent, and the size of the offspring is proportional to the relative payoff of the strategy. Mathematically, we postulate that the *per capita* growth rate of a given strategy,

$$\frac{\dot{x}_q(t)}{x_q(t)} \tag{4.46}$$

is proportional to the relative fitness, i.e. to the payoff of the strategy q at time t, minus the average fitness of the population. The evolution of $x_q(t)$ is then given by

$$\frac{\dot{x}_q(t)}{x_q(t)} = u(s_q, x(t)) - u(x(t), x(t)) = \sum_{r=1}^{n} u_{qr} x_r(t) - x(t)Ux(t), \qquad t \geq 0. \tag{4.47}$$

This dynamical system is called the *replicator dynamics*.

Note that the dynamical system is compatible with the fact that $x(t)$ lives in a simplex for all t. Indeed, let

$$S(t) = \sum_{q=1}^{n} x_q(t). \tag{4.48}$$

Then,

$$\dot{S}(t) = \frac{\mathrm{d}}{\mathrm{d}t} \sum_{q=1}^{n} x_q(t) = \sum_{q=1}^{n} x_q(t) \left(\sum_{r=1}^{n} u_{qr} x_r(t) - x(t)Ux(t) \right) = (1 - S(t))\, x(t)Ux(t). \tag{4.49}$$

Clearly, $S(t) = 1$, $\forall t \geq 0$, is a solution to this equation. Therefore, if we start with a vector of population frequencies at $t = 0$, such that $S(0) = 1$, we will have $S(t) = 1$ for all further evolution.

Example 4.9. *Replicator dynamics for a general* 2×2 *game.* Let us consider a general symmetric 2×2 game with payoff matrix

$$U = \begin{pmatrix} a & b \\ c & d \end{pmatrix}. \tag{4.50}$$

The corresponding two-player game is shown in Fig. 23. Since $x_1(t) + x_2(t) = 1$, we can focus

	s_1	s_2
s_1	(a, a)	(b, c)
s_2	(c, b)	(d, d)

Figure 23. The payoffs in a generic symmetric game with two players and two strategies.

on the dynamics for $x_1(t)$. Let us denote $x(t) = x_1(t)$. Simple algebra shows that the replicator dynamics leads to the non-linear ODE

$$\dot{x} = f(x) \equiv x(1-x)\left(w_1 - (w_1 + w_2)x\right), \tag{4.51}$$

where

$$w_1 = b - d, \qquad w_2 = c - a. \tag{4.52}$$

4.3 Review of dynamical systems

Let us now discuss the equilibrium and stability properties of the replicator dynamics. To do this, we review some elementary aspects of dynamical systems. Let $f : A \subset \mathbb{R}^n \to \mathbb{R}^n$ a C^∞ function defined on an open set A, and let us consider the following (autonomous) system of ordinary differential equations:

$$\dot{x} = f(x). \tag{4.53}$$

This is of course a vectorial equation in $\mathbb{R}^n$ and one has, in components,

$$\dot{x}_i(t) = f_i\left(x_1(t), \cdots, x_n(t)\right), \qquad i = 1, \cdots, n. \tag{4.54}$$

Definition 4.10. x^* is a *rest point* or an *equilibrium* of the system if $f(x^*) = 0$.

Note that, if the system is originally at $t = 0$ at the rest point, it will remain so for $t \geq 0$.

Definition 4.11. We say that a rest point is *stable* (also called *neutrally stable* or *Lyapunov stable*) if, for any neigbourhood U of x^* there exists a neigbourhood $V \subset U$ of x^* such that, if $x(0) \in V$, $x(t) \in U$ for all $t \geq 0$. If a rest point is not stable, we will call it *unstable*.

Definition 4.12. We say that x_* is *asymptotically stable* if it is stable and in addition there is a neighbourhood B of x_* such that, if $x(0) \in B$,

$$\lim_{t\to\infty} x(t) = x^*. \tag{4.55}$$

In general, determining the nature of an equilibrium point in a general dynamical system is a very difficult problem, and often one has to resort to numerical methods. A tractable system is the general n-dimensional linear system:

$$\dot{x} = Ax, \tag{4.56}$$

where A is an $n \times n$ dimensional matrix with real entries. The stability properties of the rest point $x = 0$ are completely determined by the eigenvalues of A, which will be denoted by λ_i, $i = 1, \cdots, n$:

1. $x = 0$ is asympotically stable if and only if

$$\mathrm{Re}(\lambda_i) < 0, \qquad i = 1, \cdots, n, \tag{4.57}$$

2. If there is at least one eigenvalue λ_j with

$$\mathrm{Re}(\lambda_j) > 0, \tag{4.58}$$

then $x = 0$ is unstable.

This result follows from the general solution of the linear system (4.56). As an illustration, let us analyze the linear dynamical system in one and two dimensions.

Example 4.13. In one dimension, we simply have the equation

$$\dot{x} = ax, \qquad a \in \mathbb{R}\backslash\{0\}. \tag{4.59}$$

Its solution is

$$x(t) = x(0)\mathrm{e}^{at}, \qquad t \in \mathbb{R}. \tag{4.60}$$

Therefore, it is clear that $x = 0$ is stable if $a > 0$ and unstable if $a < 0$. □

Example 4.14. Let us now consider a general linear dynamical system in two dimensions:

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = A \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}, \qquad A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}. \tag{4.61}$$

The eigenvalue equation for A is

$$(\lambda - a_{11})(\lambda - a_{22}) - a_{12}a_{22} = 0, \tag{4.62}$$

which can be written as

$$\lambda^2 - 2\alpha\lambda + \beta = 0 \tag{4.63}$$

where

$$\alpha = \frac{1}{2}\mathrm{Tr}(A), \qquad \beta = \det(A). \tag{4.64}$$

The eigenvalues are then

$$\lambda_{1,2} = \alpha + \sqrt{\gamma}, \qquad \gamma = \alpha^2 - \beta. \tag{4.65}$$

We will assume that $\beta \neq 0$. There are three different situations:

1. If $\gamma > 0$, the eigenvalues are real. The solution of the system is of the form

$$\begin{aligned} x_1(t) &= a\mathrm{e}^{\lambda_1 t} + b\mathrm{e}^{\lambda_2 t}, \\ x_2(t) &= c\mathrm{e}^{\lambda_1 t} + d\mathrm{e}^{\lambda_2 t}, \end{aligned} \tag{4.66}$$

for appropriate constants a, b, c, d. If $\lambda_{1,2} < 0$, $x = 0$ is asymptotically stable and is called a *stable node*. Otherwise, it is unstable. If $\lambda_{1,2} > 0$, we have a *unstable node*, and if one eigenvalue is positive and the other is negative, we have a *saddle point*

2. f $\gamma < 0$, we set $\omega = \sqrt{-\gamma}$. The eigenvalues are complex conjugate and of the form

$$\lambda_{1,2} = \alpha \pm \mathrm{i}\omega. \tag{4.67}$$

The solution of the system is of the form

$$\begin{aligned} x_1(t) &= \mathrm{e}^{\alpha t} \left(a \cos(\omega t) + b \sin(\omega t) \right), \\ x_2(t) &= \mathrm{e}^{\alpha t} \left(c \cos(\omega t) + d \sin(\omega t) \right), \end{aligned} \tag{4.68}$$

for appropriate constants a, b, c, d. If $\alpha > 0$, $x = 0$ is stable and is called a *stable focus*. If $\alpha > 0$, $x = 0$ is unstable and is called a *unstable focus*. If $\alpha = 0$, $x = 0$ is stable but not asymptotically stable, and is called a *center*.

3. If $\gamma = 0$, the solution of the system is of the form

$$\begin{aligned} x_1(t) &= \mathrm{e}^{\alpha t} \left(at + b \right), \\ x_2(t) &= \mathrm{e}^{\alpha t} \left(ct + d \right), \end{aligned} \tag{4.69}$$

for appropriate constants a, b, c, d. If $\alpha < 0$, the origin is a stable node, and if $\alpha \geq 0$, it is an unstable node.

□

Let us now consider a general non-linear dynamical system (4.53) with a rest point $x = x^*$. In order to determine its stability properties, the first step is to linearize the system. Let us consider the Jacobian matrix of f at the equilibrium point x^*, defined by

$$\mathcal{J}_{ij} = \left. \frac{\partial f_i}{\partial x_j} \right|_{x = x^*}. \tag{4.70}$$

The linear system associated to the fixed point x^* is given by

$$\dot{y} = \mathcal{J} y. \tag{4.71}$$

Definition 4.15. We say that the rest point x^* of (4.53) is *hyperbolic* if every eigenvalue of the Jacobian matrix $\mathcal{J}$ has a non-zero real part.

Theorem 4.16. *Let x^* be an hyperbolic equilibrium point of (4.53). Then, if all the eigenvalues of $\mathcal{J}$ have a strictly negative real part, x^* is asymptotically stable. If there is an eigenvalue with a negative real part, x^* is unstable.*

This theorem is extremely useful, since, in favorable cases, it reduces the study of stability to the diagonalization of the matrix giving the linear approximation. For a proof, see for example [1].

4.4 Replicator dynamics: general properties and examples

We note that the replicator dynamics is an autonomous system of ODEs with $A = \Delta^{n-1}$ and

$$\left(f(x) \right)_q = x_q \left(u(s_q, x) - u(x, x) \right) \tag{4.72}$$

We will now provide some relations between equilibria in the replicator dynamics, and the equilibria in the underlying game.

Proposition 4.17. Let σ^* be a Nash equilibrium of the symmetric game G. Then, $x^* = \sigma^*$ is an equilibrium or rest point of the replicator dynamics.

Proof: The proof is elementary. By the fundamental theorem of mixed Nash equilibria, if s_q is a strategy played with a positive weight in σ^*, we have

$$u(s_q, \sigma^*) = u(\sigma^*, \sigma^*), \tag{4.73}$$

and $(f(\sigma^*))_q = 0$. If s_q is played with zero weight, then $x_q^* = 0$ and $(f(\sigma^*))_q = 0$ as well.

□

Proposition 4.18. Let σ^* be an asymptotically stable equilibrium for the replicator dynamics. Then, the strategy σ^* is a symmetric Nash equilibrium of the underlying game.

Proof: Let us assume that σ^* is an asymptotically stable equilibrium. This implies in particular that there is an open neigbourhood of σ^*, V, such that, if $x(0) \in V$,

$$\lim_{t\to\infty} x(t) = \sigma^*. \tag{4.74}$$

Let us now suppose that σ^* does not define a Nash equilibrium of the underlying game. Therefore, by Proposition 1.20, there must exist a strategy s_a giving a higher payoff when played against σ^* than σ^*. In terms of the matrix u_{ij} specifying the game, we must have

$$\sum_{r=1}^{n} u_{ar}\sigma_r^* - \sigma^* U \sigma^* > 0. \tag{4.75}$$

By assumption, we have (4.74), i.e. at large times $x(t)$ is very closed to σ^*. Since the function

$$f(y) = \sum_{r=1}^{n} u_{ar} y_r - y^T U y, \qquad y \in \mathbb{R}^n, \tag{4.76}$$

is continuous, at large times $f(x(t))$ is very close to $f(\sigma^*)$, which is positive by (4.75). Therefore, there exists a $T > 0$ such that, if $t > T$,

$$\sum_{r=1}^{n} u_{ar} x_r(t) - x(t) U x(t) > 0. \tag{4.77}$$

But in view of the replicator equation

$$\frac{\dot{x}_a(t)}{x_a(t)} = \sum_{r=1}^{n} u_{ar} x_r(t) - x(t) U x(t) \tag{4.78}$$

we have

$$\frac{\dot{x}_a(t)}{x_a(t)} > 0, \qquad t > T. \tag{4.79}$$

This contradicts (4.74), which implies that

$$\lim_{t\to\infty} \dot{x}_a(t) = 0. \tag{4.80}$$

□

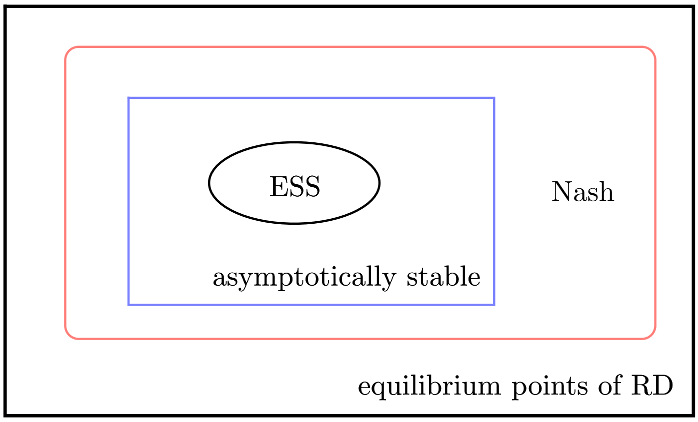

Figure 24. Different types of equlbria in the replicator dynamics (RD) and in game theory.

Proposition 4.18 is very interesting, since it shows that asymptotic stability in the replicator dynamics leads to a refinement of the concept of Nash equilibrium. An asymptotically stable equilibrium must be a Nash equilibrium, but the converse is not true: a Nash equilibrium is not necessarily an asymptotically stable equilibrium of the replicator dynamics.

We can also ask what is the relationship between asymptotic stability and ESS. This is answered by the following theorem. Its proof uses techniques of Lyapunov stability and we will not give the details here, see for example [19].

Theorem 4.19. *Let $\sigma^* \in \Delta^{n-1}$ be an ESS. Then, σ^* is an asymptotically stable equilibrium of the replicator dynamics.*

The converse of this theorem is not true, as we will see in explicit examples in the Exercises. In Fig. 24, we summarize the relationships between the types of equilibrium relevant to the replicator dynamics, and the concepts of equilibrium in classical and evolutionary game theory.

Example 4.20. Let us study the stability properties of equilibrium points in the simplest case, namely the one-dimensional replicator dynamics described by (4.51). Clearly, $x = 0, 1$ are equilibrium points for any value of the parameters. A candidate for a third equilibrium point is

$$x^* = \frac{w_1}{w_1 + w_2} = \frac{b-d}{b-d+c-a}. \tag{4.81}$$

For this to be an admissible solution, we must have

$$0 < \frac{w_1}{w_1 + w_2} < 1. \tag{4.82}$$

The local stability of the equilibrium points is easy to establish. We have

$$f'(0) = w_1, \qquad f'(1) = w_2, \tag{4.83}$$

and

$$f'(x^*) = -\frac{w_1 w_2}{w_1 + w_2}. \tag{4.84}$$

There are four different types of equilibrium.

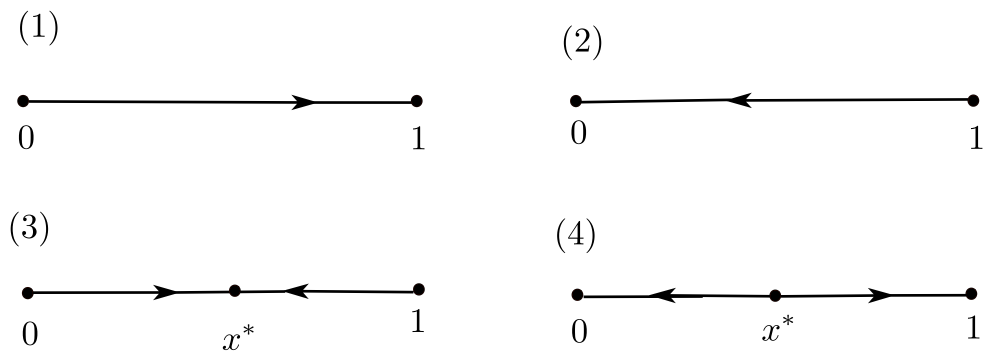

Figure 25. The four different cases in Example 4.20.

1. If $w_1 > 0 > w_2$, the only equilibrium points are $x = 0$, $x = 1$. The point $x = 1$ is an asymptotically stable equilibrium of the replicator dynamics. The other equilibrium, $x = 0$, is unstable. Correspondingly, (s_1, s_1) is a Nash equilibrium of the underlying game. In fact, since $a > c$ and $b > d$, s_1 strictly dominates s_2.

2. If $w_1 < 0 < w_2$, we have the reverse situation and $x = 0$ is asymptotically stable. The only other equilibrium, $x = 1$, is unstable. The strategy s_2 strictly dominates s_1 and (s_2, s_2) is Nash equilibrium.

3. If $w_1, w_2 > 0$, there is an asymptotically stable equilibrium at x^*. $x = 0, 1$ are unstable.

4. If $w_1, w_2 < 0$, $x = 0, 1$ are asymptotically stable equilibria. The equilibrium at x^* is unstable. The basins of attraction of $x = 0$, $x = 1$ are

$$[0, x^*), \qquad (x^*, 1], \tag{4.85}$$

respectively.

The different flow diagrams associated to these four situations are shown in Fig. 25.

As an application of this classification, let us consider the Hawk-Dove game of Example 4.7. In this case, we have that

$$w_1 = \frac{v}{2}, \qquad w_2 = \frac{c - v}{2}. \tag{4.86}$$

When $v > c$, we are in the case (1), and the symmetric Nash equilibrium (H, H) is asymptotically stable. When $c > v$, we are in case (3), and the asymptotically stable equilibrium at x^* corresponds to the ESS σ^* found in Example 4.7.

□

Example 4.21. Let us now consider the example of RSP, and let us write the RD. The matrix U is

$$U = \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}, \tag{4.87}$$

and we have that

$$xUx = x_1(x_2 - x_3) + x_2(x_3 - x_1) + x_3(x_1 - x_2) = 0. \tag{4.88}$$

Therefore, the replicator dynamics reduces to

$$\begin{aligned}\dot{x}_1 &= x_1(x_2 - x_3),\\ \dot{x}_2 &= x_2(x_3 - x_1),\\ \dot{x}_3 &= x_3(x_1 - x_2),\end{aligned} \tag{4.89}$$

which we can write compactly as

$$\dot{x}_q = x_q(x_{q+1} - x_{q+2}), \qquad q = 1, 2, 3, \tag{4.90}$$

where the understanding that the subindices are periodic modulo 3.

As usual, we can focus on the first two variables $x_{1,2}(t)$, since $x_3(t) = 1 - x_1(t) - x_2(t)$. The resulting two-dimensional dynamical system is

$$\begin{aligned}\dot{x}_1 &= f_1(x_1, x_2) = x_1(x_1 + 2x_1 - 1),\\ \dot{x}_2 &= f_2(x_1, x_2) = x_2(1 - 2x_1 - x_2).\end{aligned} \tag{4.91}$$

Let us first study the rest points of the system. The zeros of f_1 are $x_1 = 0$, $x_1 + 2x_1 - 1 = 0$, and those of f_2 are $x_2 = 0$, $1 - 2x_1 - x_2 = 0$. We conclude that there are four rest points:

1. $x_1 = x_2 = 0$
2. $x_1 = 0$, $x_2 = 1$,
3. $x_1 = 1$, $x_2 = 0$,
4. $x_1 = x_2 = \frac{1}{3}$.

The last one corresponds to the mixed NE of the underlying symmetric game. We now study the stability around these four rest points. We first do a linearized analysis. The Jacobian is

$$\mathcal{J}(x_1, x_2) = \begin{pmatrix} 2x_1 + 2x_2 - 1 & 2x_1 \\ -2x_2 & -2x_2 - 2x_1 + 1 \end{pmatrix} \tag{4.92}$$

Let us analyze it in the four rest points:

1. $$\mathcal{J}(0,0) = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix} \tag{4.93}$$

   This is an unstable hyperbolic point.

2. $$\mathcal{J}(0,1) = \begin{pmatrix} 1 & 0 \\ -2 & -1 \end{pmatrix} \tag{4.94}$$

   This is also an unstable hyperbolic point.

3. $$\mathcal{J}(1,0) = \begin{pmatrix} 1 & 2 \\ 0 & -1 \end{pmatrix} \tag{4.95}$$

   Again an unstable hyperbolic point.

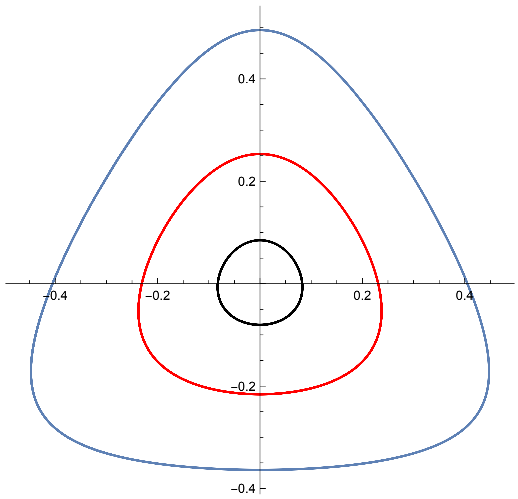

Figure 26. Three trajectories of the RSP RD, in barycentric coordinates. They correspond to the initial conditions $x_1(0) = 1/5$, $x_2(0) = 1/3$, $x_1(0) = 1/7$, $x_2(0) = 1/5$, and $x_1(0) = 10/35$, $x_2(0) = 1/3$.

4.
$$\mathcal{J}(1/3, 1/3) = \begin{pmatrix} 1/3 & 2/3 \\ -2/3 & -1/3 \end{pmatrix} \tag{4.96}$$

The characteristic polynomial is
$$\lambda^2 + \frac{1}{3} \tag{4.97}$$
with roots
$$\lambda = \pm \frac{\mathrm{i}}{\sqrt{3}}. \tag{4.98}$$
Therefore, the NE is not hyperbolic. The linear analysis does not determine the nature of the equilibrium.

In this case, the analysis of the stability of the NE can be performed as follows. It turns out that this dynamical system has a constant of motion, namely
$$\zeta = x_1 x_2 x_3, \tag{4.99}$$
which satisfies that
$$\frac{\mathrm{d}\zeta}{\mathrm{d}t} = 0. \tag{4.100}$$
This easily checked:
$$\dot{\zeta} = \dot{x}_1 x_2 x_3 + \dot{x}_2 x_1 x_3 + \dot{x}_3 x_1 x_2 = x_1 x_2 x_3 \left(x_2 - x_3 + x_3 - x_1 + x_1 - x_2 \right) = 0. \tag{4.101}$$
This implies that the trajectory of the dynamical system is included in the locus ζ =constant, where the constant is determined by the initial condition. It is easy to see from the existence of the constant of motion that, if

It is common to represent the motion of a dynamical system defined on the simplex Δ^2 in *barycentric coordinates*. Of course, any point in Δ^2 is given by three numbers $0 \le t_i \le 1$, $i = 1, 2, 3$, and such that
$$t_1 + t_2 + t_3 = 1. \tag{4.102}$$

Let us consider an equilateral triangle with center at the origin, and vertices given by the vectors

$$\begin{aligned}\mathbf{e}_1 &= \left(\frac{\sqrt{3}}{2}, -\frac{1}{2}\right), \\ \mathbf{e}_2 &= \left(-\frac{\sqrt{3}}{2}, -\frac{1}{2}\right), \\ \mathbf{e}_3 &= (0, 1),\end{aligned} \tag{4.103}$$

Let us now represent a point in the simplex Δ^2 by a point inside the triangle, given by

$$(x, y) = \sum_{i=1}^{3} t_i \mathbf{e}_i. \tag{4.104}$$

More concretely, we have

$$x = \frac{\sqrt{3}}{2}(t_1 - t_2), \qquad y = 1 - \frac{3}{2}(t_1 + t_2) \tag{4.105}$$

We represent in barycentric coordinates some of the trajectories for the replicator dynamics of RSP in Fig. 26. Clearly, the trajectories turn around the equilibrium point

$$t_i = \frac{1}{3}, \qquad i = 1, 2, 3. \tag{4.106}$$

This is the Nash equilibrium of the underlying game, and it corresponds to the origin $x = y = 0$ in barycentric coordinates. This point is stable, since trajectories which start arbitrarily close to it stay close to it. However, it is *not* asymptotically stable, and due to Theorem 4.19, it can not be an ESS, as we showed in the previous section. A phase portrait of the game, generated by the *Dynamo* software [26], is shown in Fig. 27.

□

5 Dynamic games

5.1 Games in extensive form

So far we have assumed that players in the game move simultaneously, but this is clearly not always the case. In many cases, players move *sequentially*. For example, in the game Bach or Stravinsky, it might happen that Gillian decides first where to go. She has two options, going to the Bach's concert or going to the Stravinsky's concert. Once she has decided, she tells Bob, who can join her where she is, or go to the other place. The resulting game can be represented by the tree diagram in Fig. 28.

The game in Fig. 28 is an example of a game in *extensive form*. Let us now formalize such games in some detail.

Definition 5.1. A game Γ in extensive form is given by the following items:

$$\Gamma = \left\{\mathcal{I}, \{K_i\}_{i=1,\cdots,I}, R, \{H_i\}_{i=1,\cdots,I}, \{A(x)\}_{x\in K\setminus Z}, \{\{u_i(z)\}_{i=1,\cdots,I}\}_{z\in Z}\right\}. \tag{5.1}$$

Let us make them explicit:

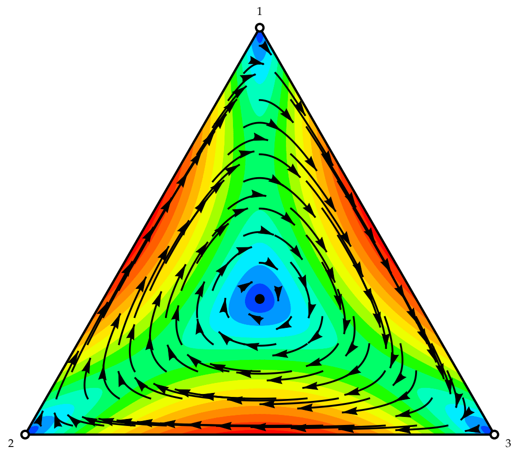

Figure 27. Phase space portrait of the dynamics of RSP, by using the *Dynamo* software.

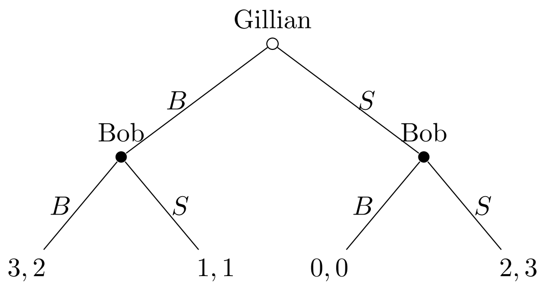

Figure 28. Bach or Stravinsky in extensive form: Gillian moves first.

1. First of all, we need a set $\mathcal{I} = \{1, \cdots, I\}$ of I players.

2. The structure of the game is encoded in a *game tree*. This tree consists of a set of *nodes* K, connected by a set of *edges*.

3. The set of nodes is ordered, and this determines the order of events in the game. Formally, we have a strict, partial ordering relation R of the nodes. If

$$x\,R\,y \tag{5.2}$$

we say that x *precedes* y in the game. In particular, for a given node, there will be nodes representing immediately prior events, which we call *predecessor nodes*, and nodes representing immedialtey posterior events, which we call *successor events*. Each edge connects a *parent node* (or *immediate predecessor*) to a *child node* (or *immediate successor*). The immediate predecessors of the node x form the set $P(x)$. The set of successors of the node x will be denoted by $P^{-1}(x)$. The game tree has the following tree property: first of all, there is a single *root node* x_0 which has no parent. This node represents the beginning of the game. There is a set of *terminal nodes* Z which have no children, and represent the

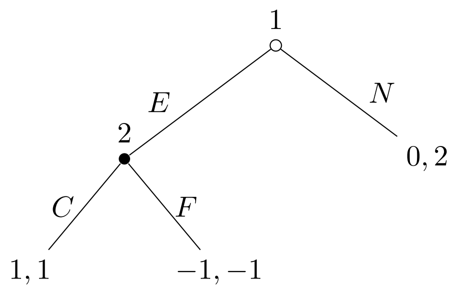

Figure 29. Entry game

final outcomes of the game. In addition, for each node $x \in K$ which is not the root node, there exist a unique, finite path of predecessors $x_1, \cdots, x_r$, joining x to the root node x_0. It follows from this tree property that each node has a single immediate predecessor, so $P(x)$ has a single element.

4. The set of intermediate nodes $K \backslash Z$ is partitioned into n subsets $K_1, \cdots, K_I$, associated to the players. This means that, if $x \in K_i$, it is player i who moves at the node x. A very important notion in extensive form games is the *information set*, which takes into account the possibility that a player does not know which decisions were previously taken into the game. To formalize this, we simply divide the set K_i in a set H_i of disjoint subsets h of K_i,

$$K_i = \bigcup_{h \in H_i} h, \qquad h \cap h' = \emptyset. \tag{5.3}$$

   Each of these sets is called an *information set.* The interpretation of an information set is that the player i cannot distinguish between the nodes of this set, i.e. she does not know (or did not observe) the choices which led to the different nodes in that set.

5. Let $x \in K_i$. At this node, there is a set of available actions for the player i, which we denote by $A(x)$. They are represented by the edges connecting x to $P^{-1}(x)$. Thee set of actions $A(x)$ is therefore isomorphic to the set $P^{-1}(x)$. Since nodes in the same information set are indistinguishable, it follows that, if $x, x' \in h$, we must have $A(x) = A(x')$, i.e. the possible actions for the player i in each of the nodes of an information set must be the same (otherwise she could distinguish the nodes!). We will denote $A(h) \equiv A(x)$ for any $x \in h$.

6. Finally, for each of the terminal nodes $z \in Z$ of the game, we have a set of n payoffs for the players, which we will denote by

$$u_i(z), \qquad i = 1, \cdots, I. \tag{5.4}$$

As in Fig. 28, we will represent a game in extensive form by a graphical description. At each node, we will indicate the index of the player who is playing, and the edges attached to a parent node x are labelled by the different actions in $A(x)$. Finally, nodes in the same information set will be joined or surrounded by a dashed line.

Example 5.2. *The entry game.* In the entry game, shown in Fig. 29, player 1 is a firm which is considering entering into a market where there is already another firm operating, represented

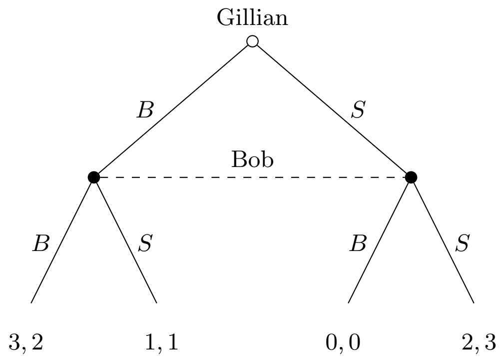

Figure 30. Bach or Stravinsky in extensive form, but with a non-trivial information set for Bob.

by player 2. The firm 1 has two options: enter (E) or not (N). If it does not enter, the game ends there, with payoff 0 for 1 and 2 for the existing firm 2. If it enters, firm 2 can fight (F) or concede (C), with the payoffs shown in Fig. 29. $\square$

Example 5.3. Let us now reconsider Bach or Stravinsky, but now with a non-trivial information set, as shown in Fig. 30. Bob has now only one information set, so when he has to move he does not know which decision has been taken by Gillian. Therefore, this game is at all practical levels identical to the simultaneous move game in Fig. 2. $\square$

Definition 5.4. An extensive form game in which every information set contains a single decision node is called a *game of perfect information.*

Before proceeding, let us think intuitively how players should play these games in order to maximize their payoffs. Let us first consider Bach or Stravinsky in extensive form, and let us look at the choices of Bob. In the first node, he will choose B, since this gives him a better payoff, and Gillian will get 3. In the second node, he will choose S, since this gives him a payoff of 3, while Gillian gets 2. Clearly, it is more convenient for Gillian to be at the first node, since this will give her a better payoff. We conclude that Gillian should play B, and Bob should then play B. This type of procedure to find the most "rational" outcome of the game is called *backward induction*: we have put ourselves in the last nodes of the game to see what are the best options for the player who moves last, and then we have gone backwards to make the best choice for the player who makes the previous move. Backward induction also allows us to "solve" for the Entry game. Let us go to the node where Player 2 chooses. Her best choice is C (since this gives her a payoff of 1). This gives a payoff of 1 to Player 1, as compared to the payoff 0 she would get by choosing N. Therefore, Player 1 should choose E, and player 2 should choose C.

We will now formalize these ideas, by introducing the notion of *subgame perfect equilibrium*, and of *generalized backward induction.* In order to do that, we have to extend the notion of Nash equilibrium to games in extensive form. This requires finding the normal or strategic form associated to a game in extensive form.

The first ingredient we need is the concept of strategy for a game in extensive form.

Definition 5.5. A *strategy* in a game in extensive form is a complete plan of action, i.e. a complete collection of choices for every possible situation. Formally, this goes as follows. Let H_i

be the partition of the decision nodes K_i of the i-th player. We will also denote

$$A_i = \bigcup_{h \in H_i} A(h). \tag{5.5}$$

A strategy for the player i is then a function

$$s_i : H_i \to A_i \tag{5.6}$$

such that

$$s_i(h) \in A(h), \tag{5.7}$$

for all $h \in H_i$.

Since any strategy is a complete path of action, the strategy profile

$$s = (s_1, \cdots, s_I) \tag{5.8}$$

determines a path in the game tree from the root node to a terminal node $\zeta(s) \in Z$. The payoff of the strategy for the i-th player is

$$u_i(s) = u_i(\zeta(s)). \tag{5.9}$$

In this way, we obtain a game in normal form with N players, the strategy sets S_i for the players, $i = 1, \cdots, I$, and the payoffs u_i. Once we obtain this game, we can search for Nash equilibria.

Remark 5.6. The concept of strategy that we have used corresponds to a pure strategy for a game in normal form. The analogue of a mixed strategy in a game in extensive form is sometimes called a *behavioral strategy*, which is a map

$$\gamma_i : H_i \to \Delta(A_i), \tag{5.10}$$

such that $\gamma_i(h) \in \Delta(A(h))$.

We will now consider various examples.

		Player 2	
		F	C
Player 1	E	(-1,-1)	(1,1)
	N	(0,2)	(0,2)

Figure 31. Entry game in normal form

Example 5.7. We start with the entry game of Fig. 29. Player 1 has two strategies, E and N, while player 2 has two strategies, C and F. The game in normal form can then be represented as in Fig. 31. Let us now look for Nash equilibria in pure strategies. The best-response correspondences are

$$BR_1(F) = \{N\}, \qquad BR_1(C) = \{E\}, \tag{5.11}$$

and

$$BR_2(E) = \{C\}, \qquad BR_2(N) = \{F, C\}. \tag{5.12}$$

Therefore, there are two Nash equilibria in pure strategies, namely,

$$(E, C), \qquad (N, F). \tag{5.13}$$

The first equilibrium is clear, since this is what we found from backward induction. The second equilibrium is more strange: the first player does not enter, and the second player fights. This should be interpreted as follows: the second player threatens to fight if player 1 enters, so the latter decides not to enter. In order words, the first player is dissuaded to enter due to the threat of the second player. However, this is clearly an example of an *incredible threat*: if player 1 decides to enter anyway, it is not in the very interest of player 2 to realize her threat, since fighting is not her best response. Therefore, the Nash equlibrium (N, F) is not a natural outcome for the game. □

		Player 2			
		(B,B)	(B,S)	(S,B)	(S,S)
Player 1	B	(3,2)	(3,2)	(1,1)	(1,1)
	S	(0,0)	(2,3)	(0,0)	(2,3)

Figure 32. Bach or Stravinsky, with Gillian moving first, in normal form.

Example 5.8. Let us now work out the normal form of the extensive form game in Fig. 28. Player 1, Gillian, has two strategies, B and S. Bob has *four* strategies, since he has two information sets $H_2 = \{h_1, h_2\}$, each one with one node, and two actions in each node. We will denote these strategies as (B, B), (B, S), (S, B) and (S, S). The strategy (B, B) means: if player 1 plays B, play B; if she plays S, play B. In terms of a function, we have:

$$s_2 : H_2 \to \{B, S\} \tag{5.14}$$

with

$$s_2(h_1) = s_2(h_2) = B. \tag{5.15}$$

The interpretation for the other strategies is similar. The resulting normal form is shown in Fig. 32. Let us calculate the Nash equilibria of the game. The best response correspondences are

$$BR_1((B, B)) = \{B\}, \quad BR_1((B, S)) = \{B\}, \quad BR_1((S, B)) = \{B\}, \quad BR_1((S, S)) = \{S\}, \tag{5.16}$$

for player 1, and

$$BR_2(B) = \{(B, B), (B, S)\}, \qquad BR_2(S) = \{(B, S), (S, S)\}, \tag{5.17}$$

for player 2. We conclude that $(B, (B, B))$, $(B, (B, S))$, and $(S, (S, S))$ are Nash equilibria. Note that $(B, (B, B))$ and $(S, (S, S))$ involve incredible threats. In the case of $(S, (S, S))$, Bob might threaten Gillian with going Stravinsky no matter what. But once Gillian has decided to go to

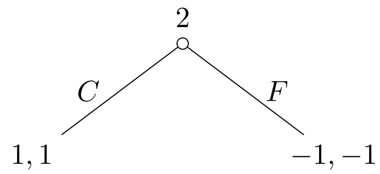

Figure 33. The proper subgame of the entry game.

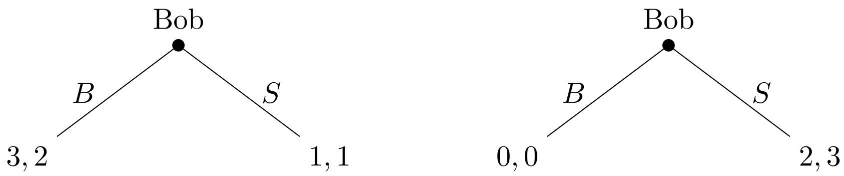

Figure 34. The two proper subgames of Bach or Stravinsky

the Bach's concert, the best option for Bob is to go with Gillian, so this is not a natural outcome of the game. In $(B, (B, B))$, Bob threatens with going to the Bach concert no matter what, but if Gillian decided to go to the Stravinsky concert, Bob would be better off by going to the Stravinsky concert as well. So the best strategy for playing this game should be $(B, (B, S))$. Since Gillian will choose B, the natural outcome of the game is that both go to the Bach concert. This is precisely what we found with the heuristic argument based on backward induction. □

5.2 Subgame perfect equilibrium

As the examples above show very clearly, the main problem of using the concept of Nash equilibrium for a game in extensive form is that many of the equilibria are not "natural". They involve incredible threats, and they are not the ones singled out by backward induction. A rational player will avoid strategies involving these threats. Let us now introduce the concept of subgame perfect equilibrium, which leads to a precise formulation of backward induction. To start with, we define what a subgame is.

Definition 5.9. A subgame of a game given in extensive form is a subset of nodes $\widehat{K} \subset K$, where K are the nodes of the original game, satisfying the following properties. First, there is an information set $\widehat{h}$ such that

$$\widehat{K} = \widehat{h} \cup \{x' \in K : x \in h,\, x\, R\, x'\}. \tag{5.18}$$

i.e. $\widehat{K}$ consists of the information set $\widehat{h}$, together with *every* node following the nodes in $\widehat{h}$. Second, any information set is wholly included in $\widehat{K}$ or in its complement (in other words, $\widehat{K}$ does not "cut" any information set). A subgame is *proper* if it does not coincide with the original game. A subgame is *admissible* if the initial information set $\widehat{h}$ consists of a single node.

As examples of subgames, we show in Fig. 33 the only proper subgame of the entry game, and in Fig. 34 the two proper subgames of Bach or Stravinsky in extensive form.

Clearly, any strategy profile of the original game leads to a strategy profile in any subgame by restriction. Namely, if h is a node in the information set $\widehat{H}_i$ of the subgame, then the induced profile

$$\widehat{s}_i = s_i|_{\widehat{\Gamma}} \tag{5.19}$$

is defined by

$$\widehat{s}_i(h) = s_i(h). \tag{5.20}$$

Definition 5.10. A profile s^* is a *subgame perfect equilibrium* of the game Γ if, for every admissible subgame $\widehat{\Gamma}$, $s^*|_{\widehat{\Gamma}}$ is a Nash equilibrium of $\widehat{\Gamma}$.

Remark 5.11. Note that a subgame perfect equilibrium of Γ must be in particular a Nash equilibrium of the normal form of Γ. One can extend the concept of subgame perfect equilibrium to behavioral strategies $\gamma = (\gamma_i, \cdots, \gamma_I)$ in the obvious way.

Example 5.12. *Subgame perfect equilibrium in the entry game.* In the proper subgame in Fig. 33, the Nash equilibrium is clearly C. The Nash equilibirum (N, F) of the original game induces the strategy F in the subgame, therefore it is not a subgame perfect equilibrium, while (E, C) is subgame perfect. The requirement of subgame perfection removes the incredible threat. □

Example 5.13. *Subgame perfect equilibrium in Bach or Stravinsky.* Let us now look at the subgames in Bach or Stravinsky, in Fig. 34. In the first one, corresponding to the node following B by Gillian, the Nash equilibrium is B, while in the second one, it is S. We conclude that the only subgame perfect equilibrium is $(B, (B, S))$. We have removed again the incredible threats. □

It is possible to show that any game in extensive form has a subgame perfect equilibrium in behavioral strategies (see, for example, [27]). In practice, the best way to find subgame perfect equilibria in pure strategies is to use the procedure of *generalized backward induction*. This formalizes precisely the intuitive arguments based on backward induction and presented before. Generalized backward induction goes as follows:

1. Start at the end of the game tree, and identify the Nash equilibria for each of the *final* subgames.

2. Select one Nash equilibrium in each of these final subgames, and derive the reduced extensive form game in which these final subgames are replaced by the payoffs that result in these subgames when players use these equilibrium strategies.

3. Repeat steps 1 and 2 for the reduced game. Continue the procedure until every move in the game is determined. This collection of moves at the various information sets constitutes a profile of subgame perfect equilibrium strategies.

4. If multiple equilibria are never encountered, this profile is the unique subgame perfect equilibrium. If multiple equilibria are encountered, the full set of subgame perfect equilibria is obtained by repeating the procedure for each possible equilibrium found in the subgames.

Note that this procedure guarantees that Nash equilibria are picked out in all subgames, by starting from the smallest ones to the larger ones. A proof that it indeed leads to all subgame perfect equilibria can be found in for example Chapter 9 of [21]. The main advantage of generalized backward induction is that it does not require us to use the normal form of the game, which sometimes is quite involved. Let us now apply it in some examples.

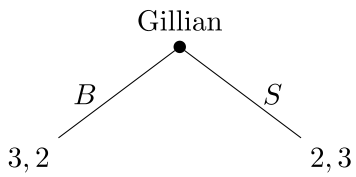

Figure 35. The game obtained in the first stage of generalized backward induction, applied to the game in Fig. 28.

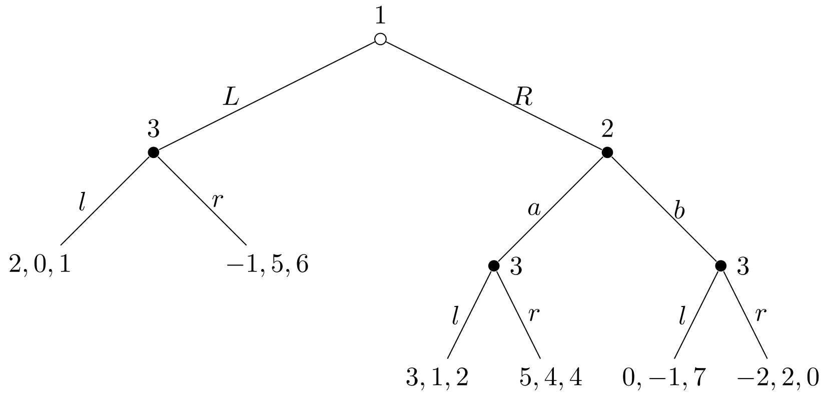

Figure 36. An extensive form game with three players.

Example 5.14. Let us consider the extensive form of Bach or Stravinsky. The final subgames are shown in Fig. 34. In the first one, Bob chooses B, with payoff $(3, 2)$ while in the second one he chooses S, with payoff $(2, 3)$. After replacing the payoffs, we obtain the game shown in Fig. 35. Now, Gillian will choose B. We conclude that the strategy obtained by generalized backward induction is $(B, (B, S))$, which is precisely the (only) subgame perfect equilibrium of the game. □

Example 5.15. Let us now consider the game shown in Fig. 36. The final subgames have Nash equilibria with payoffs $(-1, 5, 6)$, $(5, 4, 4)$, $(0, -1, 7)$ (from left to right) in which Player 3 chooses r, r, and l, respectively. After replacing the subgames with the equilibrium payoffs, we end up with the game shown in Fig. 37. We now repeat the procedure. The final subgame has a Nash equilibrium with payoffs $(5, 4, 4)$, in which 2 plays a. We end up with the game in Fig. 38, in which Player 1 chooses R, with a final payoff of $(5, 4, 4)$. We conclude that the game has a unique subgame perfect equilibrium, namely the strategy (R, a, rrl). □

As we have seen, subgame perfect equilibrium is a useful refinement of Nash equilibrium for extensive form games. However, it becomes less useful if the game has not many proper subgames. This tends to happen in games of incomplete information, since information sets can not be broken down when looking for subgames.

Another problem is that, although subgame perfect equilibria were introduced to eliminate "irrational" equilibria with incredible threats, they do not always describe correctly the way in which real humans interact. A famous example is the so-called *ultimatum game*. In this game

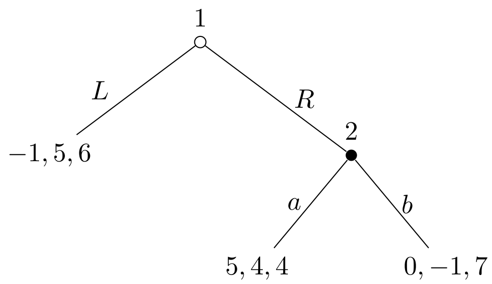

Figure 37. Applying generalized backward induction to the game in Fig. 36, first step.

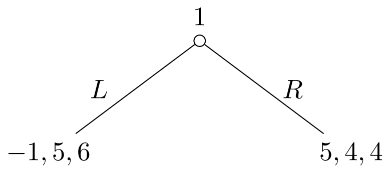

Figure 38. Applying generalized backward induction to the game in Fig. 36, second step.

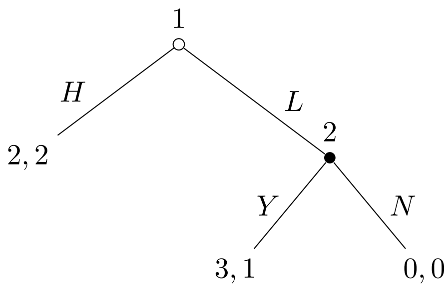

Figure 39. Ultimatum minigame

there are two players: the proposer and the responder. The proposer is given a certain amount of money x and she has to make a decision of how to split it. It then offers a share to the responder, who can accept or reject the offer. If he accepts the offer of, say y, then the payoffs to proposer and responder are $x - y$ and y, respectively. If the responder rejects the offer, both get nothing. By backward induction, it is clear that the best choice for the responder is to accept the offer y for *any* y, since otherwise she will get nothing. The proposer, knowing this, will propose the smallest minimal offer. A simple variant of the game, called the *ultimatum minigame* [11], is shown in Fig. 39. In the ultimatum minigame, the pie to share has 4 units. Player 1 can make a high offer (H) in which the pie is shared equally and both players get a payoff of 2. She can also make a low offer (L), in which player 1 gets 3 and player 2 gets one. Player 2 can accept (Y) or not (N) the low offer. If she rejects it, nobody gets anything. The subgame perfect equilibrium of the ultimatum minigame is clearly (L, Y), in which player 1 chooses the low offer and player 2 accepts it.

However, many experimental realizations of the ultimatum game indicate that unfair offers

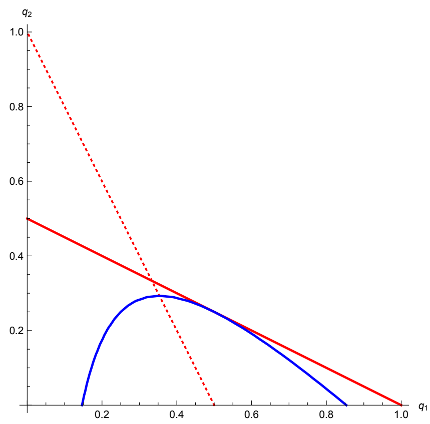

Figure 40. The subgame perfect equilibrium in the Stackelberg duopoly.

are often rejected, and anticipating this, people make fair offers! Moreover, the behaviour in this game depends sometimes on the cultural environment. The failure of "rational" game theory predictions in this type of games has led to many debates. For some people, this means that the view of players as selfish agents, only interested in maximizing their payoffs, is flawed (for a discussion of the ultimatum game according to this view, see [4], p. 111-116). For others [11], this just indicates that the subgame perfect equilibrium is not the most appropriate notion of equilibrium for this game.

5.3 Application: Stackelberg duopoly

The setting for Stackelberg model of duopoly is very similar to the setting for Cournot's model: there are two firms in the market, with cost functions given in (2.11), facing an inverse demand function (2.12), and with payoffs (2.37). However, in Stackelberg's model, firm 1 moves first, deciding its level of output q_1 before the second firm (the first firm is sometimes called the "leader", while firm 2 is the "follower."). What are the outputs characterizing the subgame perfect equilibrium? Let us apply backward induction. Given a level of output q_1 from firm 1, the second firm decides its level of output by maximizing its payoff, i.e. by determining

$$s_2^*(q_1) = \max_{q_2 \in \mathbb{R}_{>0}} u_2(q_1, q_2), \tag{5.21}$$

which is the best response function $B_2(q_1)$ of the firm 2 as computed in (2.17). In view of this response, firm 1 chooses its level of output, i.e. it maximizes the function

$$\upsilon(q_1) = u_1\left(q_1, s_2^*(q_1)\right) = \frac{q_1}{2}\left(a - c - bq_1\right), \tag{5.22}$$

as a function of q_1. The first order condition is

$$\frac{\mathrm{d}\upsilon}{\mathrm{d}q_1} = \frac{\partial u_1}{\partial q_1}\left(q_1, s_2^*(q_1)\right) + \frac{\partial u_1}{\partial q_2}\left(q_1, s_2^*(q_1)\right)\frac{\mathrm{d}s_2^*}{\mathrm{d}q_1} = 0, \tag{5.23}$$

with solution

$$q_1^* = \frac{a-c}{2b}. \tag{5.24}$$

Once the leader has decided to produce this output, the follower produces

$$s_2^*\left(\frac{a-c}{2b}\right) = \frac{a-c}{4b}. \tag{5.25}$$

Note that, in contrast to the Cournot duopoly, the equilibrium point is obtained by requiring the level curve of the payoff function of firm 1 to be tangent to the best response function of firm 2, as shown in Fig. 40. If we now compare to the equilibrium value obtained in the Cournot duopoly (2.19), we find that in the Stackelberg model the first firm has a larger output, while the second firm has a smaller output, as expected. This is also clearly seen in Fig. 40. Correspondingly, the payoff for firm 1 in the subgame perfect equilibrium,

$$u_1(q_1^*, s_2^*(q_1^*)) = \frac{(a-c)^2}{8b}, \tag{5.26}$$

is bigger than the payoff in the Cournot duopoly, while the payoff for firm 2,

$$u_2(q_1^*, s_2^*(q_1^*)) = \frac{(a-c)^2}{16b}, \tag{5.27}$$

is smaller. There are advantages in being the first mover!

6 Games and information

6.1 Incomplete information and bayesian games

So far, we have assumed that all players in the game have complete, albeit imperfect information. Namely, they know the payoffs of the players, even though they might not know at which node they are in an information set. However, in many games, players do not even know what the payoffs will be. Take for example poker. We usually don't know for sure what are the cards of the other players, therefore we don't know the payoffs that will be obtained in the game.

There is a clever way, due to Harsanyi, to model games with incomplete information as games in extensive form with imperfect information. To do this, we pretend that before the actual game starts, Nature "moves" and assigns a *type* to each player, with a certain probability distribution. For example, in a card game, the type of a player corresponds to the cards she gets. Typically, the type of a player is known to the player herself, but not to the others. This means that the different types for a given player should be regarded as different nodes in a single information set for the other players.

		Bob	
		B	S
Gillian	B	(3,2)	(2,1)
	S	(0,0)	(1,3)

		Bob	
		B	S
Gillian	B	(1,2)	(0,1)
	S	(2,0)	(3,3)

Figure 41. Bach or Stravinsky if Gillian is a Bach fan (left) and if Gillian is a Stravinsky fan (right).

To understand how this works, it is useful to work out a concrete example. We will then consider a version of Bach or Stravinsky in which Gillian can be of two types: a Bach fan, or a Stravinsky fan. If she is a Bach fan, the payoffs of the game are modified as in Fig. 41, left, and she is a Stravinsky fan, the payoffs are as in Fig. 41, right. Bob does not know which type Gillian is: she can be a Bach fan with probability p, or a Stravinsky fan with probability $1-p$. Following the idea of Harsanyi, we can model this as an extensive form game in which Nature

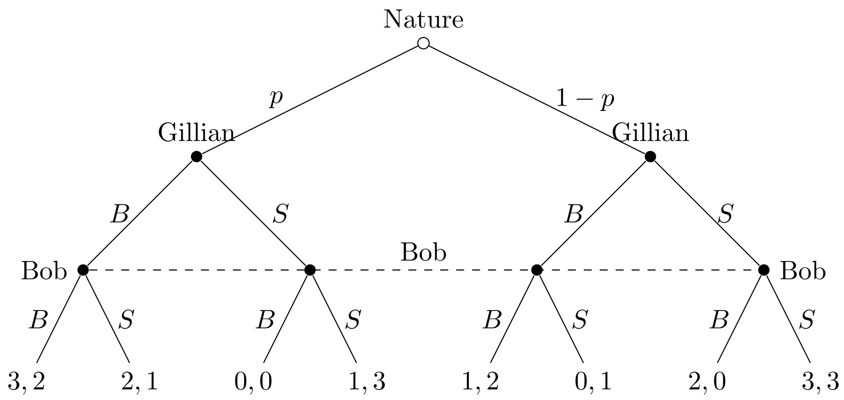

Figure 42. The Bach or Stravinsky in which Nature moves first.

moves first, assigning to Gillian the types “Bach fan” and ”Stravinsky fan” with probabilities p and $1-p$, respectively. Once Nature has moved, we have two simultaneous move games. We can therefore represent them in the form Fig. 42. Note that Bob has a single information set consisting of four nodes, since he does not know the move of Nature, nor the moves of Gillian.

Let us now represent the game in Fig. 42 in normal form, where the players are Gillian and Bob. We first note that Gillian has four strategies, which we will denote (B, B), (B, S), (S, B), (S, S). Thus, (B, B) means that Gillian chooses the Bach concert if she is a Bach fan, and if she is a Stravinsky fan. Bob has two stragegies, B and S. It remains to calculate the payoffs. In a game where Nature moves, a strategy profile s for the players determines various paths in the game tree. For example, in the game in Fig. 42, the strategy profile $((B, B), B)$ determines two different paths in the game. The first one, with probabililty p, ends up in the leftmost final node with payoffs $(3, 2)$, while the second one, with probability $1-p$, ends up in the node with payoffs $(1, 2)$. The total payoff is then

$$p(3, 2) + (1-p)(1, 2) = (2p + 1, 2). \tag{6.1}$$

In general, to calculate the payoff of a strategy profile s in a game where Nature moves, we consider all possible paths associated to s, $\mathcal{P}_s$. These are in one to one correspondence with terminal nodes t_γ. Given a path $\gamma \in \mathcal{P}_s$, we multiply the probabilities associated to the nodes of γ where Nature moves. This gives the quantity $p(s, \gamma)$. Then, given the payoff to the i-th player at the terminal node $u_i(t_\gamma)$, we define the payoff of the strategy as

$$u_i(s) = \sum_{\gamma \in \mathcal{P}_s} p(s, \gamma) u_i(t_\gamma). \tag{6.2}$$

With this principle, we can calculate the normal form of the game in Fig. 42, which we show in Fig. 43. If we assume that $0 < p < 1$, we see that (B, S) is a strictly dominating strategy for Gillian. To find the best strategy for Bob once Gillian has chosen (B, S), we compare $2p$ to $3 - 2p$. If $p > 3/4$, he should play B, while if $p < 3/4$, he should play S. If $p = 3/4$, he is indifferent between the two strategies. We have then found, if $0 < p < 1$ and $p \neq 3/4$, a unique Nash equilibrium for the game. Note that the game in Fig. 42 does not have proper subgames, so the concept of subgame perfect equilibrium does not apply.

	B	S
(B,B)	$(1+2p,2)$	$(2p,1)$
(B,S)	$(2+p,2p)$	$(3-p,3-2p)$
(S,B)	$(1-p,2-2p)$	$(p,1+2p)$
(S,S)	$(2-2p,0)$	$(3-2p,3)$

Figure 43. Bach or Stravinsky in which Gillian can be of two types, in normal form.

The game that we have just studied is an example of a *Bayesian game*. In these games, Nature moves first, choosing the type of the players, and then the players (who know only their own type) play a simultaneous move game. More formally, we have the following:

Definition 6.1. A *Bayesian game* consists of:

1. A set of I players, $\mathcal{I} = \{1, \cdots, I\}$.
2. For each player $i \in \mathcal{I}$, a space of types T_i, a set of pure actions A_i, and a payoff function
$$u_i : T \times A_1 \times \cdots \times A_I \to \mathbb{R}. \tag{6.3}$$
where
$$T = T_1 \times \cdots \times T_I \tag{6.4}$$
is the space of type profiles.
3. A probability distribution
$$p : T \to [0,1] \tag{6.5}$$
which specifies the probabilities that Nature selects the type profile $t = (t_1, \cdots, t_I)$.

A pure strategy for the i-th player in a Bayesian game is an application
$$s_i : T_i \to A_i \tag{6.6}$$
which specifies the action for a given type (mixed strategies are defined similarly, by considering applications into $\Delta(A_i)$).

The best way to analyze a Bayesian game is to first regard it as an extensive form game, and then write it in normal form, as we did with the Bach or Stravinsky. In the resulting game, we can look for Nash equilibria. A Nash equilibrium of the resulting normal form game is usually called a *Bayes–Nash equilibrium* of the Bayesian game. A Bayes–Nash equilibrium is a strategy profile s^* which satisfies
$$\sum_{t\in T} p(t) u_i(t, s_1^*(t_1), \cdots, s_i^*(t_i), \cdots, s_I^*(t_I)) \geq \sum_{t\in T} p(t) u_i(t, s_1^*(t_1), \cdots, s_i(t_i), \cdots, s_I^*(t_I)), \tag{6.7}$$
for any strategy s_i and for any player $i \in \mathcal{I}$. This is equivalent to finding a Nash equilibrium for the game in normal form, in which the payoffs of a strategy profile $s = (s_1, \cdots, s_I)$ are the expected payoffs
$$\sum_{t\in T} p(t) u_i(t, s_1(t_1), \cdots, s_i(t_i), \cdots, s_I(t_I)). \tag{6.8}$$

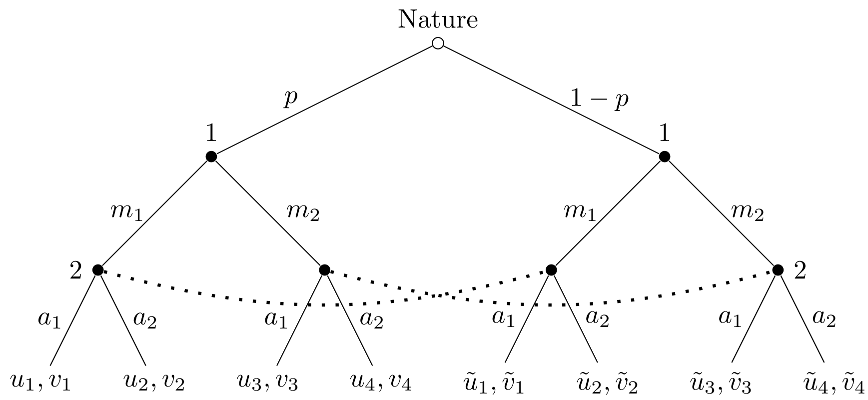

Figure 44. A typical signaling game.

6.2 Signaling games

We will now consider a type of games of incomplete information which are called *signaling games*. In these games, there are two agents moving in sequence. The first agent, usually called the *sender*, has some private information, which is relevant for the second agent, called the *receiver*. The first agent moves first and sends a *message* or signal, which can be observed by the receiver. In view of this message (which might be related to the underlying information), the receiver chooses an action, and the game concludes.

This simple scenario has many applications in real life. One of the first examples of a game of this type was proposed by M. Spence in 1973 and it describes signaling in the job market. The sender is in this case a worker, whose private information is her own ability. The receiver is a potential employer, and the message is her education level, which is in principle related to her ability. The employer observes this education level, but not the ability, and chooses a salary offer in view of this level. In the animal world, signals are everywhere, and we will address specific examples in biology in the next section.

Let us now formalize a signaling game:

1. First of all, Nature selects the type t of the sender from a set of possible types T, with probabilities $p(t)$ which are known to all players.

2. The sender, once her type is known, sends a message m belonging to some prescribed space of messages M.

3. The message is observed by the receiver, which then chooses an action a in a set A.

The payoffs of the players depend on the type of the sender, and the choices of message and action. We will denote by $u_R(t, m, a)$ and $u_S(t, m, a)$ the payoffs to the receiver and to the sender, respectively. A (pure) strategy for the sender is a specification of a message m for each of her information sets, i.e. for each type $t \in T$, so it is an application

$$\mu : T \to M. \tag{6.9}$$

A (pure) strategy for the receiver is a specification of an action a for each of her information sets, i.e. for each message $m \in M$, so it is an application

$$\alpha : M \to A. \tag{6.10}$$

Behavioural strategies are defined in a similar way by considering instead applications into $\Delta(M)$ and $\Delta(A)$, respectively.

A typical signaling game is shown in Fig. 44. In this simple game, player 1 can be of two types, t_1 and t_2, with probabilities p and $1-p$, respectively, and it can send two types of messages, m_1 and m_2. Player 2 can respond with two types of actions, a_1 and a_2. We have labelled the payoffs in a simplified way, so that, for example, u_1 stands for $u_S(t_1, m_1, a_1)$, and so on.

As for any other game in extensive form (with or without complete information), we can write it in normal form and find its (Bayes)–Nash equilibria. Let us spell in detail what are the conditions for such an equilibrium, focusing for simplicitiy on pure strategies. Given a strategy profile (μ, α), the expected payoff to the sender is

$$U_S(\mu, \alpha) = \sum_{t \in T} p(t)\, u_S\left(t, \mu(t), \alpha(\mu(t))\right), \tag{6.11}$$

while the expected payoff for the receiver is

$$U_R(\mu, \alpha) = \sum_{t \in T} p(t)\, u_R\left(t, \mu(t), \alpha(\mu(t))\right), \tag{6.12}$$

Following the general definition (6.7), we say that a strategy profile (μ^*, α^*) is a Bayes–Nash equilibrium if, for all possible pairs of applications (μ, α), we have

$$U_S(\mu^*, \alpha^*) \geq U_S(\mu, \alpha^*), \qquad U_R(\mu^*, \alpha^*) \geq U_R(\mu^*, \alpha). \tag{6.13}$$

Note that the equlibrium condition for the sender is equivalent to finding $m^* : T \to M$ which maximizes

$$\sum_{t \in T} p(t)\, u_S\left(t, m, \alpha^*(m)\right). \tag{6.14}$$

Since this is a sum of the payoffs $u_S(t, m, \alpha^*(m))$ over $t \in T$ with non-negative coefficients $p(t)$, maximizing this sum is equivalent to maximizing each summand, i.e. finding, for each $t \in T$, the message $\mu^*(t)$ which solves

$$\max_{m \in M} u_S\left(t, m, \alpha^*(m)\right). \tag{6.15}$$

We can rewrite a similar maximizing condition for the receiver, but we have to note that α is defined on the set of messages. However, not all messages are necessarily used by the sender in the equilibirum configuration. Let us define

$$T^*(m) = \left(\mu^*\right)^{-1}(m) = \{t \in T : \mu^*(t) = m\}. \tag{6.16}$$

Then, for each $m \in M$, we have to solve

$$\max_{a \in A} \sum_{t \in T^*(m)} p(t) u_S(t, m, a). \tag{6.17}$$

This maximization problem produces a function $\alpha^*(m)$.

Example 6.2. A famous example of a signaling game is "breakfast in the American West" [6][6]. In this game, there are two rival bands, the Montescos and the Capuletos. Player 1, which belongs to the Montescos, is sitting in a bar. She can belong to two different types: the tough Montescos

[6]Sometimes called "real men don't eat quiche" [13].

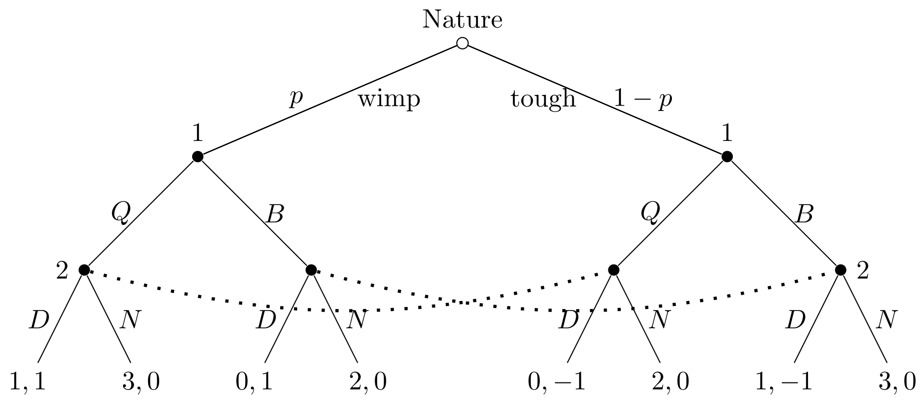

Figure 45. Breakfast in the American West.

(t_1), or the wimp Montescos (t_2), with probability p and $1-p$, respectively. She can also have two different dishes for breakfast: either beer ($m_1 = B$) or quiche ($m_2 = Q$). Tough Montescos prefer beer, while wimps prefer quiche, so that the breakfast plays the role of the signal for the type of player. Player 2 is a member of the Capuletos and enters the bar. He sees the Montesco having its breakfast. He can challenge her to duel ($a_1 = D$) or not ($a_2 = N$). However, the Capuletos are of medium type: they always win the duel to the wimp Montescos, but they always lose it with the tough Montescos. The structure of the game, with the corresponding payoffs, is shown in Fig. 45. Note that the payoffs of player 1 take into account, not only the result of the duel, but also whether she has taken his preferred breakfast (the worst thing that can happen to a wimp Montesco is indeed to pretend she is tough and having beer for breakfast, and then lose the duel to a Capuleto.)

	DD	DN	ND	NN
QQ	$(p, 2p-1)$	$(p, 2p-1)$	$(p+2, 0)$	$(p+2, 0)$
QB	$(1, 2p-1)$	$(3-2p, p)$	$(1+2p, p-1)$	$(3, 0)$
BQ	$(0, 2p-1)$	$(2p, p-1)$	$(2-2p, p)$	$(2, 0)$
BB	$(1-p, 2p-1)$	$(3-p, 0)$	$(1-p, 2p-1)$	$(3-p, 0)$

Figure 46. Breakfast in the American West, in normal form.

Let us first work out the normal form of the game. Player 1 has four strategies: QQ, QB, BQ and BB, where XY means: take X if wimp and Y if tough. Player 2 also has four strategies: DD, DN, ND, NN, where XY means do X if message is quiche, do Y if message is beer. The normal form of the game, together with its expected payoffs, is shown in Fig. 46. It is easy to compute the best-response correspondences. We have, for the first player,

$$BR_1(DD) = \{QB\}, \quad BR_1(DN) = \{BB\}, \quad BR_1(ND) = \{QQ\}, \quad BR_1(NN) = \{QB\}. \tag{6.18}$$

The best response correspondences of the second player depend on p. We have,

$$\begin{gathered} BR_2(QQ) = \begin{cases} \{DD, DN\}, & \text{if } p > 1/2, \\ \{DD, ND\}, & \text{if } p < 1/2, \end{cases} \quad BR_2(QB) = \{DN\}, \\ BR_2(BQ) = \{ND\}, \quad BR_2(BB) = \begin{cases} \{DD, ND\}, & \text{if } p > 1/2, \\ \{NN, DN\}, & \text{if } p < 1/2. \end{cases} \end{gathered} \tag{6.19}$$

We see that, if $p < 1/2$, the Nash equilibria in pure strategies are

$$(BB, DN), \qquad (QQ, ND). \tag{6.20}$$

This type of signaling equilibrium, in which the sender sends always the same message, independently of her type, is called *pooling equilibrium*. When different types of senders send different messages, we have a *separating equilibrium*. If $p > 1/2$, there are no Nash equilibria in pure strategies, and one has to search for mixed strategies. □

6.3 Costly signaling 1: the Sir Philip Sidney Game

Signaling games can be used to model communication between agents. In the animal world, communication strategies are pervasive. For example, many animals send signals of strength prior to a fight or contest. In courtship, animals signal their genetic quality with extravagant displays, like the peacock tail. In signaling games among animals (or humans) a crucial issue is the honesty of the signal. It used to be taken for granted that communication systems among animals are clean and honest. However, it was noted by Dawkins and Krebs in the 1970s that, if there is a conflict of interest between sender and receiver, signals could be distorted and manipulated by senders. This could start an "arms race" between senders and receivers, in which senders try to manipulate receivers, and receivers develop mechanisms to unmask senders. Eventually, an equilibrium must be reached where signals are honest, otherwise the communication system would not be stable. One important question in animal communication is then: what maintains the honesty of signals? (see [8], Chapter 14, for an excellent survey of these issues.)

One possible answer to this question is provided by the theory of *costly signaling*, originally proposed by Amotz Zahavi. Roughly speaking, the theory affirms that a signal is honest if it is costly for the sender. The receiver, seeing the cost incurred by the sender, can then trust the signal. This theory is also called the *handicap principle*, since according to costly signaling, honest signals are *handicaps* for the sender. We will develop two models illustrating the theory, both motivated by animal communication: the Sir Philip Sidney Game, due to John Maynard Smith [22], and Alan Grafen's model of the handicap principle [16].

The Sir Philip Sidney game is inspired by the story of the XVI-th century British statesman and poet Sir Philip Sidney, who was fatally wounded in a battle against the Spanish at Zutphen. According to legend, while lying wounded he gave his water to another wounded soldier, saying, "Thy necessity is yet greater than mine."[7] A typical situation which is modeled by this game occurs when chicks beg for food to their parents. The chick can be in two different states: very hungry, or slightly hungry, and it can choose between two different actions: signaling by chirping or not. The parents have two possible actions: they can keep the food or give it to the chick.

The Sir Philip Sidney game applies more generally to a situation in which a sender can be in two different states, t_1 and t_2, with probabilities $1-p$ and p. The sender can signal (S) or not

[7] This was described by Maynard Smith as "an unusual example of altruism by a member of the British upper classes" [22].

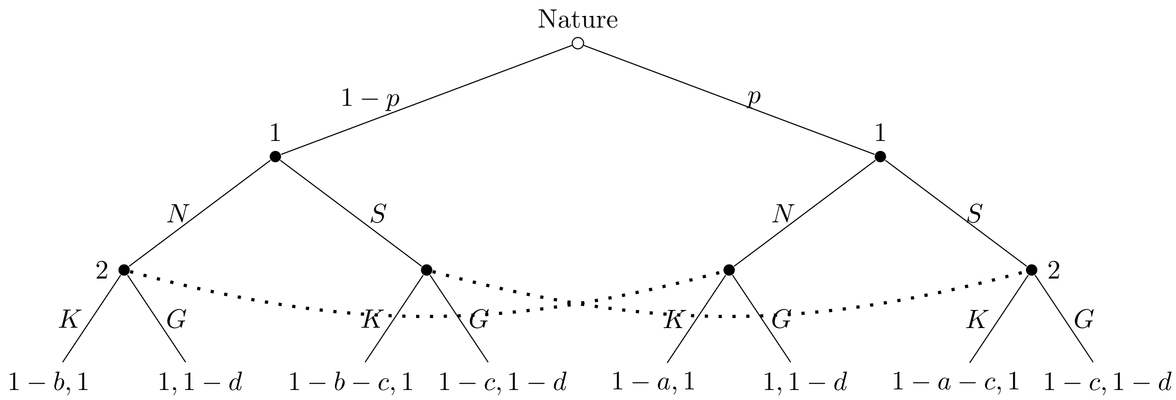

Figure 47. The Sir Philip Sidney game.

(N), and the receiver can keep (K) or give (G) a resource. We now have to specify the payoffs. If the receiver keeps the resource, she will get a payoff of 1, while if she gives it, the payoff will be $1-d$. In other words, d is the cost to the receiver of transferring the resource. On the sender's side, the payoff depends on its state, on its action, and on getting or not the resource. If she gets the resource, her payoff is 1, independently of its state. If she does not get the resource, she will have a payoff of $1-b$ if she is in the "healthy" state t_1, and $1-a$ if she is in the "needy" state t_2 (so that $a > b$). In addition, we will assume that signaling has a cost c (for example, when chicks chirp begging for food, they can attract predators).

We can represent this signaling game in the standard form of a signaling game, as in Fig. 47. There is however a small twist in the calculation of the payoffs. As in the example of the chicks and their parents, we will assume that the players in this game are genetically related. If we interpret the total payoffs as genetic fitness, we should calculate them by using *Hamilton's rule*. According to this rule, the total fitness of an individual (also called her *inclusive fitness*) is her own fitness, plus the fitness of her relatives, weighted by a coefficient measuring their genetic relatedness. For example, the genetic relatedness between parents and offspring is $1/2$, since they share half of their genes. Therefore, an outcome that enhances the fitness of the offspring is also enhancing the fitness of half of the genes of the parents, and should be taken into account in the total payoff. We conclude that in the Sir Philip Sidney game, the payoff of the i-th player is the "bare" payoff calculated above, plus the bare payoff of the other player, multiplied by the coefficient of genetic relatedness r.

We are now ready to analyze the emergence of honest signaling in this game. The honest signaling equilibrium is the stragegy in which player 1 plays NS (i.e. not signal if healthy, and signal if needy), while player 2 plays KG (keep if there is no signal, give if there is signal). We want to derive the conditions under which this is a Nash equilibrium of the game in normal form. For this to happen, each strategy has to be a best response.

Let us first study the conditions under which NS is a best response to KG. We first calculate the possible payoffs to player 1 from the different strategies. A simple calculation shows the following:

$$\begin{aligned} u_1(NS, KG) &= (1-p)(1-b) + r + p(1-c-rd), \\ u_1(SS, KG) &= 1-c+r(1-d), \\ u_1(NN, KG) &= 1+r-pa-(1-p)b, \\ u_1(SN, KG) &= p(1-a)+(1-p)(1-c-rd)+r. \end{aligned} \tag{6.21}$$

Let us first impose

$$u_1(NS, KG) \geq u_1(SS, KG). \tag{6.22}$$

This gives

$$1-b(1-p)-pc+r(1-dp) \geq 1-c+r(1-d) \tag{6.23}$$

which can be simplified to

$$c \geq b-rd. \tag{6.24}$$

Next, let us impose

$$u_1(NS, KG) \geq u_1(NN, KG). \tag{6.25}$$

This gives,

$$1+r-b(1-p)-pc-rdp \geq 1+r-pa-b(1-p) \tag{6.26}$$

which can be simplified to

$$a \geq c+rd. \tag{6.27}$$

Finally, the condition

$$u_1(NS, KG) \geq u_1(SN, KG). \tag{6.28}$$

can be simplified to

$$(1-p)(c+rd-b) \geq p(c+rd-a). \tag{6.29}$$

By the previous conditions, the l.h.s. is non-negative, while the r.h.s. is non-positive, so if (6.24) and (6.27) hold, (6.29) is satisfied to.

The conditions (6.24) and (6.27) have a very simple interpretation. The l.h.s. of (6.24) is cost of the signal, while the r.h.s. is the increase in inclusive fitness obtained when receiving the resource while in good condition (the sender gets b, but there is a decrease d in the receiver, so by Hamilton's rule the net effect is precisely $b - rd$). Therefore, a condition for honest signaling is that the signal should be sufficiently costly. Otherwise, it is better to lie, i.e. to signal always. On the other hand, condition (6.27) says that the signal shouldn't be too costly. The cost of the signal plus the cost induced by genetic relatedness to the receiver is $c + rd$. If the benefit obtained by getting the resource, a, is not bigger than this cost, it clearly does not pay to signal, and NN is a better response.

Let us now consider the conditions for KG to be a best response. To do this, we have to compare the payoffs of player 2 due to different strategies. We find the following:

$$\begin{aligned} u_2(NS, KK) &= 1+(1-p)r(1-b)+pr(1-a-c), \\ u_2(NS, GG) &= 1+(1-p)(r-d)+p\left(r(1-c)-d\right), \\ u_2(NS, GK) &= 1+(1-p)(r-d)+pr(1-a-c). \end{aligned} \tag{6.30}$$

We first require

$$u_2(NS, KG) \geq u_2(NS, KK). \tag{6.31}$$

This implies

$$1 - dp + r\,(1 - (1-p)b - pc) \geq 1 + r\,(1 - b(1-p) - pa - pc)\,, \tag{6.32}$$

which can be simplified to

$$ar \geq d. \tag{6.33}$$

The second condition is obtained when we require

$$u_2(NS, KG) \geq u_2(NS, GG)). \tag{6.34}$$

This implies

$$1 - dp + r\,(1 - (1-p)b - pc) \geq 1 + (1-p)(r-d) + pr(1-c) - pd, \tag{6.35}$$

which can be simplified to

$$d \geq rb. \tag{6.36}$$

Finally, we have to impose

$$u_2(NS, KG) \geq u_2(NS, GK), \tag{6.37}$$

which implies

$$1 - dp + r\,(1 - (1-p)b - pc) \geq 1 + (1-p)(r-d) + pr(1-a-c), \tag{6.38}$$

which can be written as

$$(1-p)(d - rb) \geq p(d - ar). \tag{6.39}$$

However, if (6.33) and (6.36) hold, the l.h.s. of (6.39) is non-negative, while the r.h.s. is non-positive, so (6.39) is automatically verified.

The condition (6.33) is easily interpreted: $ar - d$ is the increase in inclusive fitness of the sender when she gives the resource. Since (6.33) can be written as

$$r \geq \frac{d}{a}, \tag{6.40}$$

a minimum degree of genetic relatedness is needed for the transfer to be beneficial. The condition (6.36) means that the resource should be sufficiently valuable to the receiver, otherwise she would give it even when the sender is not in need.

The most important conclusion of the Sir Philip Sidney game is that, indeed, honest signaling can be a Nash equilibrium, if the signal is sufficiently costly and there is a a certain degree of common interest between sender and receiver, as measured by the coefficient of genetic relatedness r. In practice, this means that we should take seriously the complaints of our relatives when their signals are sufficiently costly.

6.4 Costly signaling 2: signals as handicaps

In Grafen's signaling model, the male type is chosen by Nature at the beginning of the game, and is given by a real number $t \in [t_{\min}, \infty) = T$. This number corresponds to the genetic quality of the male. It is a quantity of interest to females but cannot be observed directly by them. The higher the value of t, the better its quality. Each male knows its quality, and he can choose a signal $s \in S \subset \mathbb{R}$ (s could be for example the size of the tail in the peacock). A male strategy is a function

$$s = f(t), \qquad s \in [s_{\min}, \infty), \tag{6.41}$$

and all functions will be assumed to be smooth in their domains. Females use the value of s to infer the genetic quality of the male. The inferred value will be denoted by a, and it is a function of the signal s:

$$a = g(s). \tag{6.42}$$

The payoff or fitness of a male depends on his true quality t, his signal s, and his perceived value a, and will be denoted by

$$u(t, s, a). \tag{6.43}$$

We will now impose some natural conditions that guarantee a honest signaling equilibrium. First of all, we will suppose that signaling is costly, so that

$$u_s = \frac{\partial u}{\partial s} < 0. \tag{6.44}$$

Second, we will assume that the payoff of the male increases when the female's inferred value increases (a male who is regarded as a good male by females will have more fitness). In other words,

$$u_a = \frac{\partial u}{\partial a} > 0. \tag{6.45}$$

The payoff of a female is measured by the error she makes in her evaluation. We will denote by $d(t, a)$ the loss of fitness to a female assessing as a a male with true quality t. We have

$$d(t, a) = \begin{cases} > 0 & \text{if } t \neq a, \\ = 0, & \text{if } t = a, \end{cases} \tag{6.46}$$

so females are penalized if they do not assess male quality correctly. If $\rho(t)$ is the probability distribution of males of quality t in the population, we will take as the payoff of the female the average error she has made, in view of the strategies $a = g(s)$ and $s = f(t)$. In other words,

$$v(f, g) = -\int_{\mathbb{R}} d(t, g(f(t)))\rho(t)\mathrm{d}t. \tag{6.47}$$

Let us suppose that $s = f^*(t)$ and $a = g^*(s)$ is a Nash equilibrium for the game. Then, it must be the case that $f^*(t)$ maximizes the payoff of the male, given the strategy $g^*(s)$ of the female, as in (6.15). In other words, we must have

$$u\left(t, f^*(t), g^*(f^*(t))\right) \geq u\left(t, s, g^*(s)\right), \qquad \text{for all } s, t. \tag{6.48}$$

At the same time, given $s = f^*(t)$, the female must choose a strategy $g^*(s)$ which minimizes her error, i.e.

$$\int_{\mathbb{R}} d(t, g^*(f^*(t)))\rho(t)\mathrm{d}t \geq \int_{\mathbb{R}} d(t, g(f^*(t))\rho(t)\mathrm{d}t, \tag{6.49}$$

for all functions $g(s)$.

Let us now find the explicit Nash equilibrium, following Grafen's seminal analysis. The first order condition for maximizing the payoff of the male, given $a = g(s)$ and t, is that the function

$$U(s, t) = u(t, s, g(s)) \tag{6.50}$$

satisfies

$$\frac{\partial U}{\partial s} = \frac{\partial u}{\partial s} + \frac{\partial u}{\partial a} g'(s) = 0. \tag{6.51}$$

On the other hand, it is clear that the female minimizes her error if

$$g^*(f^*(t)) = t, \tag{6.52}$$

since in this case the l.h.s. of (6.49) vanishes. We will then assume that $g(s)$ is the inverse of $f(t)$. We now find the extremum of u under this assumption, which is equivalent to assuming that

$$t = g(s). \tag{6.53}$$

Substituting this into the first order condition (6.51), we find an equation for $g^*(s)$,

$$(g^*)'(s) = -\frac{u_s(g^*(s), s, g^*(s))}{u_a(g^*(s), s, g^*(s))}. \tag{6.54}$$

The equilibrium strategy of the male is then obtained from (6.52). Note that, since $u_s < 0$ and $u_a > 0$, the r.h.s. is positive and $g^*(s)$ is an increasing function of s. In particular, its inverse function $f(t)$ is well-defined and it is also increasing. This means that, in the equilibrium, higher quality males advertize more. In other words, signaling is honest. The integration constant is fixed by requiring

$$g^*(s_{\min}) = t_{\min}. \tag{6.55}$$

In order to show that the strategy of the male leads to a local maximum, we need slightly more than the first-order condition (6.51). We first note that, when (6.51) is evaluated at the equilibrium value of $g(s)$, we have

$$\begin{aligned} \left.\frac{\partial U}{\partial s}\right|_{a=g^*(s)} &= u_s(t, s, g^*(s)) - u_a(t, s, g^*(s)) \frac{u_s(g^*(s), s, g^*(s))}{u_a(g^*(s), s, g^*(s))} \\ &= u_a(t, s, g^*(s)) \left\{ \frac{u_s(t, s, g^*(s))}{u_a(t, s, g^*(s))} - \frac{u_s(g^*(s), s, g^*(s))}{u_a(g^*(s), s, g^*(s))} \right\}. \end{aligned} \tag{6.56}$$

Let us now assume that

$$\frac{u_s(t, s, a)}{u_a(t, s, a)} \tag{6.57}$$

is an strictly increasing function of t. If this is the case, it follows that $s = f^*(t)$ is a local maximum of $U(s, t)$. This is because, under this assumption,

$$\left.\frac{\partial U}{\partial s}\right|_{a=g^*(s)} \begin{cases} < 0 & \text{if } t < g^*(s), \\ = 0 & \text{if } t = g^*(s), \\ > 0 & \text{if } t > g^*(s). \end{cases} \tag{6.58}$$

In terms of the inverse function, this means that

$$\left.\frac{\partial U}{\partial s}\right|_{a=g^*(s)} \begin{cases} < 0 & \text{if } s > f^*(t), \\ = 0 & \text{if } s = f^*(t), \\ > 0 & \text{if } s < f^*(t). \end{cases} \tag{6.59}$$

This is of course saying that $s = f^*(t)$ is a local maximum. The condition that (6.57) is increasing means essentially that better males do better by advertising more. For example, if $u_{at} = 0$, the fact that (6.57) is an increasing function of t is equivalent to having $u_{st} > 0$, which means that the cost of advertising becomes smaller (in absolute value) for better males.

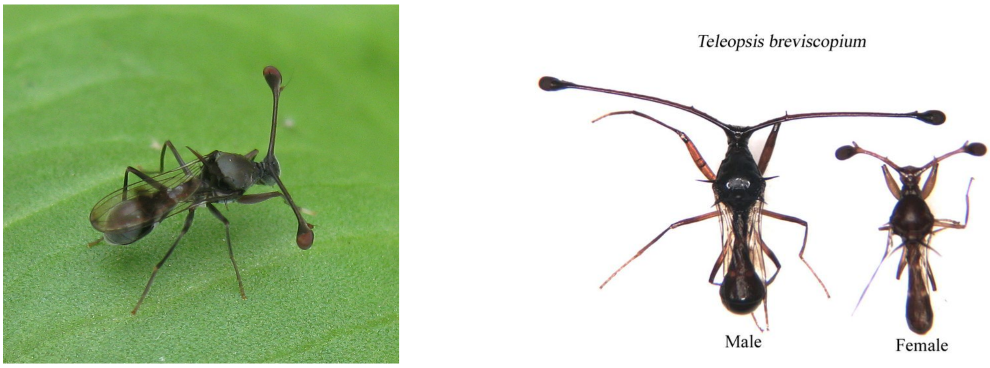

Figure 48. Stalk-eyed flies.

Example 6.3. In order to see how Grafen's model works in practice, let us consider the payoff function for the male,

$$u(t, s, a) = a^r t^s, \tag{6.60}$$

where $r > 0$ and $t \in [t_{\min}, 1) \subset (0, 1)$. We have,

$$u_s = a^r t^s \log t, \qquad u_a = r a^{r-1} t^s, \tag{6.61}$$

and we note that

$$\frac{u_s}{u_a} = \frac{1}{r} a \log t \tag{6.62}$$

is a strictly increasing function of t, as required. Since $t \in (0, 1)$, we have $u_s < 0$, $u_a > 0$. The condition for equilibrium (6.54) gives

$$(g^*)'(s) = -\frac{1}{r} g^* \log g^*, \tag{6.63}$$

which is equivalent to

$$\mathrm{d} \log g^* = -\frac{1}{r} \log g^*. \tag{6.64}$$

This can be easily integrated to

$$\log g^*(s) = C\mathrm{e}^{-s/r}. \tag{6.65}$$

The integration constant C is fixed by (6.55). We can now find $f^*(t)$ by requiring

$$\log g^*(f^*(t)) = C\mathrm{e}^{-f^*(t)/r} = \log t. \tag{6.66}$$

$\square$

The handicap principle has been experimentally tested in some animals. In the stalk-eyed fly, the eyes are located at the end of long stalks coming out of the head. It turns out that these stalks are much longer in males than in females (in some males, the total length of the stalks is longer than the length of the body). Such an extravagant trait is surely a handicap, but females mate preferably with males with longer stalks. It is natural to assume that the length of the stalks is a signal of quality. For the handicap principle to be at work, worse males should advertise less, and the cost of advertising should increase as the quality decreases. This has been tested in detail in [7] by controlling the amount of nutrient provided to larvae of flies. As the

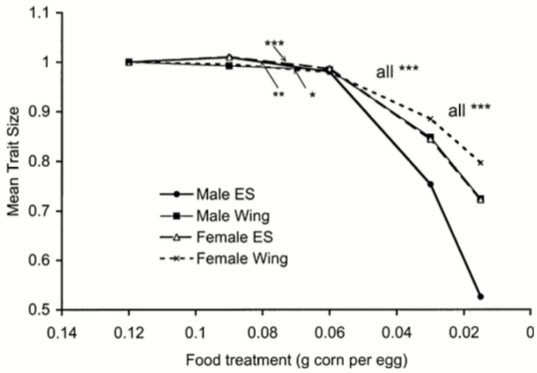

Figure 49. The decrease in eye stalk (ES) and in wing length, as a function of the (decreasing) amount of food provided to larvae. The former is much steeper than the latter, and it is steeper in males than in females. Graphic from [7].

quantity of nutrient decreased, leading to males of worse quality, the length of the eye stalks decreased more rapidly than, for example, wing length. In addition, the decrease was steeper in males than in females (see Fig. 49, taken from [7]). These and other results seem to vindicate the handicap principle as the key to understand some of the extravagant signals used in sexual selection[8].

References

[1] V. I. Arnold, *Équations différentielles ordinaires*, Ed. Mir, 1974.

[2] K. Binmore, *Playing for real. A text on game theory*, Oxford University Press, 2007.

[3] K. Binmore, *Game theory. A very short introduction*, Oxford University Press, 2007.

[4] S. Bowles, *Microeconomics. Behavior, institutions and evolution*, Princeton University Press, 2004.

[5] P. D. Taylor and M. G. Bulmer, "Local mate competition and the sex ratio," Journal of Theoretical Biology **86** (1980) 409-419.

[6] I.-K. Cho and D. M. Kreps, "Signaling games and stable equilibria," The Quarterly Journal of Economics **102** (1987) 179-221.

[7] S. Cotton, K. Fowler, and A. Pomiankowski, "Condition dependence of sexual ornament size and variation in the stalk-eyed fly Cyrtodiopsis dalmanni (Diptera: Diopsidae)," Evolution **58** (2004), 1038-1046.

[8] N.B. Davies, J. R. Krebs and S. A. West, *Principles of behavioural ecology*, Fourth edition, Wiley-Blackwell, 2012.

[9] R. Dawkins, *The extended phenotype*, Oxford University Press, 1999.

[10] J. Elster, *The cement of society*, Cambridge University Press, 1989.

[11] J. Gale, K. G. Binmore, and L. Samuelson, "Learning to be imperfect: the ultimatum game," *Games and Economic Behavior* **8**, 56-90 (1995).

[12] R. Gibbons, *A primer in game theory*, Prentice Hall, 1992.

[13] H. Gintis, *Game theory evolving*, Princeton University Press, 2000.

[14] S. J. Gould, *The panda's thumb. More reflections in natural history*, W.W. Norton, 1980.

[8] See however the recent criticism of the handicap principle in [24].

[15] S. J. Gould, *The flamingo's smile. Reflections in natural history.* W. W. Norton, 1985.

[16] A. Grafen, "Biological signals as handicaps," Journal of Theoretical Biology **144** (1990) 517-546.

[17] W. D. Hamilton, "Extraordinary sex ratios," Science **156** (1967): 477-488.

[18] G. Hardin, "The Tragedy of the Commons," Science **162** (1968), 1243-1248.

[19] J. Hofbauer and K. Sigmund, *Evolutionary games and population dynamics*, Cambridge University Press, 1998.

[20] E. Kalai and M. Smorodinsky, "Other solutions to Nash's bargaining problem," Econometrica **43** (1975) 513-518.

[21] A. Mas-Colell, M. D. Whinston and J. Green, *Microeconomic theory*, Oxford University Press, 1995.

[22] J. Maynard Smith, "Honest signalling: the Philip Sidney game," Animal Behaviour **42** (1991) 1034-1035.

[23] E. Olin Wright, *Class counts*, Student edition, Cambridge University Press, 2004.

[24] R. O. Prum, *The evolution of beauty*, Doubleday, 2017.

[25] J. Roemer, *Theories of distributive justice*, Harvard University Press, 1996.

[26] W. H. Sandholm, E. Dokumaci, and F. Franchetti, *Dynamo: Diagrams for Evolutionary Game Dynamics* (2012), `http://www.ssc.wisc.edu/~whs/dynamo`.

[27] F. Vega-Redondo, *Economics and the theory of games*, Cambridge University Press, 2003.

[28] S. A. West, *Sex allocation*, Princeton University Press, 2009.